

Defaming Artificial Intelligence: Why Is *The Guardian* Scared of Sound, Objective Science?

Tony Cox line-by-line comments on *The Guardian* article at <https://www.theguardian.com/technology/ng-interactive/2025/jun/27/tony-cox-epidemiology-risk-assessment-chatgpt-ai>

Louis Anthony Cox, Jr., Ph.D., tcoxdenver@aol.com

July 25, 2025

Introduction

This document provides a detailed, line-by-line response to the June 27, 2025 article published by *The Guardian*, currently titled “*How a New AI Tool Could Amplify Doubt in Pollution Science.*” *The Guardian* has issued multiple post-publication corrections to this article—and further corrections may still be warranted. As it stands, the article still contains numerous **false, misleading, and defamatory statements** about my work, character, and research.

The *Guardian* article misrepresents:

- The **purpose and function** of the AI technology and AI-based manuscript screening tool I developed;
- My **scientific record**, including my methods, funding, presentations and peer-reviewed publications;
- My **professional integrity and competence**, by misquoting or selectively framing my public statements – and by attributing to me wholly fictitious remarks that contradict my published work;
- The **technology** itself, which the article inaccurately describes and mischaracterizes.

This document serves three purposes:

1. To help **correct the record**, using verifiable facts and primary sources;
2. To **clarify what the AI tool is—and is not**;
3. To **defend the integrity** of rigorous, transparent, and empirically grounded science against irresponsible framing.

The first section provides a summary of key points generated using ChatGPT based on this rebuttal. The second section offers detailed, point-by-point comments and analysis addressing each misstatement or omission.

Why This Response Is Necessary

The *Guardian* article promotes a narrative grounded more in ideological framing and insinuation than in factual reporting or informed critique. It relies heavily on **unvetted commentary from selected “experts”** who:

- Display **no evident familiarity** with the AI tool (mischaracterizing it as a “bot” or LLM);
- Make **false or unfounded accusations**, which the article presents uncritically;
- Appear to share a strong **ideological bias**, leading them to misrepresent my motives and to attack my calls for the use of sound, objective, transparent scientific methods in regulatory risk assessment.

Any of the tool’s hundreds of actual users could have provided a far more accurate and informed description of its purpose and operation. Instead, the article elevates **uninformed speculation** over firsthand knowledge or verifiable evidence.

This response restores that missing context—by providing direct citations, primary sources, and accurate descriptions of both the technology and the science at issue—so that readers can evaluate the facts for themselves.

ChatGPT-Assisted Executive Summary

False Claims, Distorted Framing, and Reputational Harm in The Guardian's Article

"How a New AI Tool Could Amplify Doubt in Pollution Science" (June 27, 2025)

This response identifies and documents numerous **false, misleading, and defamatory statements** in The Guardian's article. These distortions affect not only the portrayal of the AI tool discussed, but also the professional character, scientific integrity, and published record of Dr. Tony Cox.

The article has already undergone multiple post-publication corrections—but **significant misrepresentations remain**. Moreover, its framing has been irresponsibly escalated in derivative coverage (e.g., by Yahoo News), increasing reputational and professional harm.

Key issues are as follows:

1. False and Misleading Claims About the AI Tool

- **False claim:** That the AI tool is designed to “amplify doubt in pollution science.”
- **Reality:** The tool is a neutral, rule-based pre-submission screening tool that checks for internal logical coherence—not content. It is not a chatbot, LLM, or editorial reviewer. It does not evaluate content and it cannot “debunk,” “evaluate,” “promote,” or “amplify” any viewpoint—let alone “doubt in pollution science.”
- The tool is freely available at <http://screener.riskanalysis.cloud/>. It is impossible to use it to “amplify doubt in pollution science.”
- **Clarification:** The tool has no role in any journal's editorial or peer review process. It is not used to accept, reject, or influence publication decisions, and it does not minimize, delay, or undermine science in any form.

2. Distortion of Dr. Cox's Scientific Record and Career

- **False implication:** That Dr. Cox has “forged a career sowing doubt,” is a “Trump adviser,” or seeks to undermine public health science.
- **Reality:** Dr. Cox's distinguished scientific career is well-documented and publicly visible—for example, on [Google Scholar](https://scholar.google.com/), where his work has been cited more than 9,000 times. He is widely recognized for excellence in scientific research, not for “sowing doubt.” His professional accomplishments include
 - **Multiple Best Paper awards** and other awards for scientific excellence in toxicology, risk analysis, and operations research;
 - **Elected Fellowships** in respected scientific and professional societies (SRA, INFORMS);
 - **Elected membership** in the U.S. National Academies of Sciences, Engineering, and MedicineDr. Cox is known for his rigorous, methodologically transparent, and nonpartisan contributions to science and public health policy. He has never served as a political adviser, nor has he advocated for any political agenda. His roles have been consistently scientific, technical, and advisory, grounded in data and logic—not ideology.

3. Misrepresentation of Scientific Integrity

- Falsely suggests that industry actors (e.g., API) copy-edited Cox's findings.
- **Clarification:** Draft comments on punctuation from external readers were acknowledged transparently, but all conclusions remain Dr. Cox's and were peer-reviewed independently.
- Misquotes and fabricates claims such as his supposedly calling epidemiological work "flimsy" on behalf of Philip Morris—**no such quotes or actions exist.**
- Propagates and links to a fabricated story from 2018 that Dr. Cox "Actually Said That Clean Air Saves No Lives." <https://www.motherjones.com/environment/2018/10/tony-cox-trumps-air-pollution-adviser-clean-air-saves-no-lives/>. **Dr. Cox has made no such claim.**
- **Reality: The 2018 story in *Mother Jones* was fabricated.** Dr. Cox's 2018 slides are online at www.healtheffects.org/sites/default/files/cox-causality-hei-2018.pdf. It was a purely methodological talk, concluding that "Machine learning helps avoid p-hacking and discover predictive causal relations." Dr. Cox's actual claim, as given in an earlier peer-reviewed published paper, was that, in a specific data set, "**higher PM2.5 levels are associated with increased risk of heart attack.** PM2.5 is positively associated ($p < 0.06$) with both stroke risk and heart attack risk when these are regressed only against PM2.5, sex, age, and ever-smoking status, but not when they are regressed against these variables and income." ([Cox LAT Jr. Socioeconomic and air pollution correlates of adult asthma, heart attack, and stroke risks in the United States, 2010-2013](#). Environ Res. 2017 May;155:92-107. doi: 10.1016/j.envres.2017.01.003. PMID: 28208075.)

4. Ideological Framing Posed as Scientific Critique

- Elevates quotes from critics with no direct knowledge of the tool while presenting unfounded ideological attacks ("science denialism") as fact.
- Selectively omits Cox's published rebuttals, methodological explanations, and clarifying remarks.

5. Misstatements About Prior Research

- Misrepresents the conclusion of a peer-reviewed PFOA study (which did not claim "safe doses" for humans).
- Falsely implies original work on gas stove asthma or tobacco risk defenses—neither claim is accurate.
- Cox's actual tobacco work explored mechanisms of disease, not policy advocacy or denial of risk.

6. False Statements About Peer Review

- Falsely states that a paper of Dr. Cox's published in *Global Epidemiology* "counted climate denier Steve Milloy as one of its reviewers."
- **Reality:** Reviews were anonymous and conducted per standard scientific protocols. Steve Milloy was not one of its reviewers.

7. Out-of-Context Quotes and Fabricated Quotes

- Accuses Cox of making assertions that he did not make.
- Cites fabricated quotes that appear nowhere in the record and directly contradict Dr. Cox's published and recorded work and testimony

8. Guilt by Association and Opaque Editorial Bias

- Embeds visual and rhetorical cues (e.g., outdated tobacco funding, contact with Exxon) to imply misconduct or “science for hire.”
- **Clarification:** All collaborations were transparent, common in academic work, and publicly disclosed.

9. Inaccurate Portrayal of the Tool’s Development

- Falsely characterizes the system as capable of spreading misinformation or enforcing Cox’s “views.”
- **Reality:** The tool checks for formal violations of causal logic—not beliefs, conclusions, or subject matter. Over 400 users have used it and can confirm its neutrality.

10. Broader Misrepresentation of Scientific Standards

- **False portrayal:** The article suggests that Dr. Cox demands “perfect certainty” or holds unrealistic expectations before regulatory action can be taken.
- **Reality:** Dr. Cox is a leading advocate for modern causal inference and decision and risk science methods precisely because they enable timely and effective public health protection—even under uncertainty. His work focuses on making better-informed decisions with existing data, not on delaying action or requiring unattainable proof. He has repeatedly emphasized that regulatory decisions can and should be made under uncertainty, but that they should reflect valid causal reasoning—not untested associations or assumptions. In a 2015 interview, he expressed his approach this way: “So much of what we do in life that is important has very uncertain consequences, and seeing how far analysis and reason can take us in making effective decisions in such situations is both intellectually challenging and of tremendous practical value.” <https://www.statisticsviews.com/article/breakthroughs-in-decision-science-and-risk-analysis-an-interview-with-louis-anthony-cox-jr/>

Distortion: The Guardian’s portrayal of Dr. Cox’s standards as unreasonably rigid or delay-inducing is a distortion of this well-documented, constructive, and pro-scientific stance.

Final Note

The Guardian’s article substitutes insinuation for investigation. It fails to engage with the actual design, function, or purpose of Dr. Cox’s AI tool. It misrepresents his scientific record by selectively quoting critics, introducing fabricated statements, and ignoring readily available documentation, peer-reviewed publications, and firsthand clarifications.

This response sets the record straight—drawing on verifiable evidence, primary sources, and published science—to restore factual accuracy and defend the integrity of scientific reasoning against ideological distortion and media misrepresentation.

Detailed Line-by-Line Rebuttal of [The Guardian's June 27, 2025 Article](#)

1. *The Guardian's* Misrepresentation of the AI Tool's Purpose and Function

Headline: How a new AI tool could amplify doubt in pollution science

TC: The “new AI tool” that I developed, which was deployed at *Risk Analysis* in January, 2025 is described in publicly available information ([Cox et al., 2025, Can AI help authors prepare better risk science manuscripts?](#)). This includes a link (<http://screener.riskanalysis.cloud/>), so that anyone – including reporters and those from whom they solicit comments – can easily see what it is and what does.

Contrary to *The Guardian's* claim in this headline, this tool **cannot be used to “amplify doubt in pollution science”** or in any other field. **It has nothing to do with pollution science.** It does not address doubts at all, either to “amplify” or minimize them. It is a **pre-submission screening aid for authors.** It uses AI to generate feedback to authors about whether a draft manuscript aligns with the published guidance for submissions to a scientific journal. It is a neutral, pre-submission screening tool for authors to check whether draft manuscripts align with publicly stated journal guidance. It does not evaluate content, make recommendations, or influence conclusions.

It takes no position on any scientific or policy issue. It does not assess “doubt,” and **it is technically and functionally incapable of promoting or minimizing any conclusion or of “amplifying doubt”**—about pollution or any other topic. The same is true for all AI systems I have designed since 2021.

Readers who wish to verify for themselves whether *The Guardian's* claim holds any merit are invited to test the tool directly. The interface is publicly available at <http://screener.riskanalysis.cloud/>.

My view is that ***The Guardian's* claim that the tool “could amplify doubt in pollution science” is not only unsupported but insupportable.** It appears to reflect a complete failure to perform basic journalistic due diligence—such as actually testing the tool before disseminating false, misleading, and defamatory claims about it.

*Risk analyst Tony Cox's work has been **backed by the chemical lobby, and some health experts are alarmed.***

TC: This is a misleading conflation. My research on AI-assisted reasoning has received partial funding from the American Chemistry Council (ACC) and from the University of Colorado. That support has enabled methodological work aimed at improving the rigor, transparency, and reproducibility of causal inference using AI—independent of any particular policy application or outcome.

Describing this as work “backed by the chemical lobby” may falsely suggest to readers agenda-driven advocacy or improper influence. I have never conducted work—on behalf of ACC or anyone else—to “amplify doubt.” My work promotes clarity, not confusion.

As for the claim that “some health experts are alarmed,” the article fails to identify any such individuals who have reviewed or engaged with my actual work. Instead, it quotes speculative commentary from uninformed commentators about hypothetical future tools or imagined risks—none of which describe the AI tools I have developed, published, or proposed. These individuals appear unfamiliar with the system itself and instead respond to misleading characterizations. Their concerns reflect the false narrative constructed by *The Guardian*, not any informed review of the project or of my AI technology.

2. *The Guardian’s* Misrepresentation of My Scientific Record and Research Focus

*An industry-backed researcher **who has forged a career sowing doubt** about the dangers of [pollutants](#) is attempting to use [artificial intelligence](#) (AI) to amplify his perspective.*

TC: This statement is doubly false and defamatory.

First, **my work does not “sow doubt.”** It clarifies scientific understanding by **improving and applying methods for valid causal inference and risk assessment.** Far from increasing uncertainty, my research **has helped reduce uncertainty and confusion about risks**—by rigorously distinguishing between correlation and causation and by identifying when exposure-response relationships are or are not supported by valid evidence (e.g., [Cox 2023](#); [Cox 2011](#)). This work has been recognized many times for scientific excellence ([INFORMS, 2019](#)), not for “sowing doubt.”

Second, the AI technology I have developed **does not “amplify” any perspective**—mine or anyone else’s. The “reviewer” I have developed is a reasoning aid, not a messaging tool. It **analyzes logical consistency and evidentiary support** in manuscripts, regardless of their conclusions. **It is content-neutral by design and has no mechanism to favor or promote any viewpoint.**

Louis Anthony “Tony” Cox Jr, a Denver-based risk analyst and former Trump adviser...

TC: I am not a “former Trump adviser.” I believe this false claim originated in an article in *Mother Jones* and has simply been propagated by reporters who did not bother to check the public records.

...who once reportedly claimed there is no proof that cleaning air [saves lives](#), ...

TC: I have never claimed that cleaning air doesn’t save lives, as the link that *The Guardian* provides falsely claims. I have said, correctly and based on evidence (e.g., [Lyons et al., 2023](#); [Zigler and Dominici, 2014](#)), that in multiple studies, large reductions in particulate air pollution have not been found to have caused detectable reductions in subsequent mortality rates, even after several decades. This is well known among specialists and did not originate with me, but is based on dozens of studies and journal articles by other researchers over multiple decades.

3. *The Guardian's* Misrepresentation of my AI's Function and Capabilities

...is developing an AI application to scan academic research for what he sees as the false conflation of correlation with causation.

TC: This is false and defamatory. The AI system I am developing **does not “scan academic research” for what I see as errors**. It is not limited to academic research, **nor does it operate based on my subjective views**. It is designed to *autonomously* analyze the logic, structure, and evidentiary support in any document—whether a journal article, regulatory report, or news story—and to flag common methodological issues. **Its design is content-neutral and not personalized to reflect my opinions**. Misrepresenting this as a tool to enforce “what he sees” conflates automated reasoning with personal ideology—and that is both inaccurate and damaging.

Cox has described the project as an attempt to weed “propaganda” out of epidemiological research

TC: I have never described the project that way. **I have not said it aims to “weed out” anything**. Rather, I have expressed a hope that responsible use of AI could *support* critical thinking and reduce the spread of uncritical or ideologically driven reasoning. For instance, I have written that I hope to **“bend applications of this technology toward scaled-up critical thinking** instead of scaled-up groupthink and propaganda.” That is **not a description of the project**—it is an aspiration for how AI might be used more generally to improve public discourse.

...and perform “critical thinking at scale” in emails to industry researchers, which were obtained via Freedom of Information Act requests by the Energy and Policy Institute, a non-profit advocacy group, and exclusively reviewed by the Guardian.

TC: The phrase “critical thinking at scale” is indeed language I have used in aspirational contexts. But describing this as how I “described the project” is inaccurate. My actual project description is this: “Development of the AI Assistant (AIA) began in 2022 as an applied research effort to use AI and natural language processing to check and improve causal inference and statistical methods in observational epidemiological research.”

— [Cox et al., 2025, Can AI help authors prepare better risk science manuscripts?](#)

The Guardian’s framing again distorts a personal aspiration into a mischaracterization of the project’s function and intent.

4. *The Guardian's* False and Defamatory Misrepresentation of my Research

He has long leveled accusations of flimsiness at research linking exposure to chemical compounds with health dangers, including on behalf of polluting interests such as cigarette manufacturer Philip Morris USA...

TC: This is false on multiple counts.

First, I have not “leveled accusations of flimsiness” at research linking chemical exposures with health risk on behalf of Philip Morris. **This is false, defamatory, and mischaracterizes my award-winning scientific research on causal mechanisms of smoking-induced lung diseases.** Rather, my work has focused on **strengthening causal inference** by improving our understanding of the causal links between exposures and health outcomes.

Second, my research for **Philip Morris International** (not Philip Morris USA, as incorrectly stated by *The Guardian*) did not question or undermine the causal relationship between smoking and disease. To the contrary, it *elucidated the biological causal networks and processes* through which cigarette smoking causes lung disease. The following publications from that research have been publicly available forever a decade, and I shared this information with *The Guardian's* reporter, who appears to have deliberately and knowingly ignored it in favor of the her above inaccurate and defamatory narrative.

- Cox LA Jr, Sanders E. [Estimating preventable fractions of disease caused by a specified biological mechanism: PAHs in smoking lung cancers as an example](#). *Risk Analysis*. 2006 Aug;26(4):881-92. doi: 10.1111/j.1539-6924.2006.00785.x. PMID: 16948683.
- Cox LA Jr. [A causal model of chronic obstructive pulmonary disease \(COPD\) risk](#). *Risk Analysis*. 2011 Jan;31(1):38-62. doi: 10.1111/j.1539-6924.2010.01487.x. PMID: 20846171.
- Cox LA Jr. [Quantifying potential health impacts of cadmium in cigarettes on smoker risk of lung cancer: a portfolio-of-mechanisms approach](#). *Risk Anal*. 2006 Dec;26(6):1581-99. doi: 10.1111/j.1539-6924.2006.00848.x. PMID: 17184399.

This is the direct opposite of attacking health risk evidence “on behalf of polluting interests.” Two of those papers received “Best Paper of the Year” awards from the Society for Risk Analysis for scientific excellence. **None of this work “leveled accusations of flimsiness... at research linking chemical exposures with health dangers... on behalf of Philip Morris.”** I will request retraction of this false and defamatory claim by *The Guardian*.

... and the American Petroleum Institute – a fossil fuel lobbying group he has even allowed to [“copy edit” his findings...](#)

TC: This is also false, misleading, and defamatory.

I have **never** allowed any external entity—whether the American Petroleum Institute, the American Chemistry Council, or anyone else—to edit or modify my scientific findings. The *only* instance remotely related to this involved asking a colleague to proofread a paper for **spelling and punctuation**, because my version of Microsoft Word’s proofreading tool was malfunctioning at the time. I believe one semicolon was changed to a dash—something I agreed with because it improved clarity. This was a minor typographical correction, not a content edit.

(...to [“copy edit” his findings](#). Cox says the edit “amounted to suggesting a small change” and noted that he has also obtained [public](#) research funding.)

TC: This quotation misrepresents what I actually said. I did **not** say the edit was a “small change” to my **findings**, as this misleadingly suggests. I said it was a minor, typographic change to **punctuation** in the background section—*not* a change to scientific content or conclusions. Conflating proofreading with altering findings is irresponsible and misleading.



[View image in fullscreen](#)

Cox has previously done some work for the tobacco industry. Photograph: Oliver Helbig/Getty Images

TC: Yes, I conducted scientific research aimed at elucidating the biological mechanisms through which cigarette smoke causes chronic obstructive pulmonary disease (COPD), emphysema, and lung cancer. This work advanced understanding of how these diseases develop at the mechanistic level, **contributing to public health knowledge—not disputing it.**

Both the tobacco and oil industries have a history of [weaponizing scientific uncertainty](#), experts say, with some arguing that similar tactics drive the Trump administration’s current deregulatory efforts. The

president’s May “gold standard” science order, for instance, empowered his appointees to “correct scientific information” and “discipline” those who breach the administration’s views, prompting [outrage](#) from some scientists.

TC: I have had no involvement whatsoever with the Trump administration’s “gold standard” science order or with any efforts to “discipline” scientists or “correct scientific information” according to political agendas. **I have not participated in any such activities, directly or indirectly. I categorically reject the suggestion—whether explicit or implied—that my work aligns with or supports these tactics.** My research is grounded in scientific reasoning and evidence, not ideology or politics.

Cox has obtained [funding](#) to develop the new AI reviewer from the American Chemistry Council (ACC), the nation’s largest chemical industry advocacy group, which counts oil and chemical giants such as Exxon and DuPont as members.

Experts say the ACC’s sponsorship raises questions about whom the project will benefit.

TC: This is a vague insinuation unsupported by any evidence. Experts who have **actually read the publicly available project documentation**—for example [Cox et al., 2025, Can AI help authors prepare better risk science manuscripts?](#)—will see that the currently deployed AI tool is **specifically benefits authors** by helping them align manuscripts with journal-specific guidance.

The broader purpose of the project is to improve scientific clarity and methodological rigor—for the benefit of **all** users: authors, reviewers, editors, decision-makers, and members of the public, regardless of institutional affiliation or viewpoint.

Anonymous “experts” quoted in a hit piece, who have neither reviewed the tool nor examined its documentation, are not in a position to offer informed opinions. Their speculative commentary substitutes innuendo for evidence.

The tool will benefit **anyone** who values logical clarity, methodological transparency, and sound causal reasoning—regardless of whether they agree with any particular policy or scientific conclusion.

The real question is why *The Guardian* chose not to evaluate the freely available tool on its actual capabilities and documented uses before publishing misinformed and defamatory accusations about its intent or impact – despite the reporter having been sent full accurate documentation on its purpose and uses.

Asked about these concerns, Kelly Montes de Oca, spokesperson for the ACC, said: “This research has the potential to support scientific understanding and analysis of chemical exposure and human health, enhance transparency and reproducibility, advance the safety of chemical products and processes, and inform science-based global regulatory approaches.”

Cox said in an email to the Guardian that his assistant “is specifically designed to be helpful to those who wish to understand the objective implications of data without any distortions from the kinds of well-known human heuristics and biases that make objective analysis difficult for humans”.

“My work aims to help anyone interested in using sound technical methods to pursue scientific truth,” he added. The questions sent to him by the Guardian contained “many fundamental inaccuracies”, he said.

Some critics have mischaracterized my work as an attempt to delay regulation or promote industry interests. That is not true

Louis Anthony ‘Tony’ Cox Jr

5. The Guardian’s Misrepresentation of the AI Tool’s Function and Role in Journal Peer Review

Cox said the tool is currently available to authors submitting to academic journals - including *Risk Analysis*, which he edits - to [evaluate their prospective submissions](#) before they are sent to editors.

TC: This is incorrect in both substance and wording.

First, this AI tool does **not “evaluate” prospective submissions**. For example, it makes no attempt to determine their value, worth, or scientific merit. It does not perform peer review or editorial appraisal. Rather, it provides automated feedback to authors on how well their drafts align with publicly available journal guidance for authors, not an evaluator of content or conclusions. The link that The Guardian provides is to the tool opt-in page, which correctly states “The AI will provide you with feedback and suggestions.” The incorrect claim by the Guardian that the tool will be used to “evaluate their prospective submissions” is contradicted by the page they link to, which explicitly states that “The information generated by the AI will not be saved or shared by us and it **will have no effect on any later evaluation** of your paper.”

Second, the statement that the tool is used “before [submissions] are sent to editors” is misleading. Authors may choose to use the tool at any point in their writing process, regardless of whether they intend to submit, or to whom. There is no connection between the tool and any journal’s editorial process, including *Risk Analysis*.

6. The Guardian’s Misrepresentation of the AI Tool’s Motivation and Insinuations of Partiality and Agenda-Driven Research

Asked for a response to concerns about the project’s funding, Cox said that he has publicly acknowledged the ACC’s support in all relevant publications and said the tool “has no axe to grind and no positions to push”.

But the ACC is not a neutral force, said Chris Frey, the associate dean for research and infrastructure at the North Carolina State University's College of Engineering who chaired the Environmental Protection Agency's clean air scientific advisory committee from 2012 to 2015.

"They lack impartiality in that they want to minimize regulatory burden on their members," said Frey.

TC:

1. I note a revealing inconsistency in *The Guardian's* reporting. Chris Frey and I both served in the exact same role—Chair of the EPA's Clean Air Scientific Advisory Committee (CASAC). Yet Frey is neutrally and accurately described as having chaired the committee, while I am labeled elsewhere in the same article as "a former Trump adviser"—a politicized and false characterization. This inconsistent framing reflects an apparent editorial bias and undermines the credibility of the article.
2. With regard to the American Chemistry Council (ACC): In all my work on this project, **no one at the ACC has ever stated or implied that the goal is to "minimize regulatory burden"** or to influence the tool's outputs in any way. Nor is that the intent of any research project I have participated in. The specific funding stream that supported this work—the ACC's Long-Range Research Initiative (LRI)—has a **clearly stated mission:**

"To catalyze innovations in chemical safety science."

(See: <https://www.americanchemistry.com/chemistry-in-america/research/long-range-research-initiative-lri>)

My research supports with this mission by developing tools to **enhance scientific rigor and transparency**—not to advocate for or against any particular regulatory outcome.

3. From the outset of this work in 2020, the core design specification has been clear: to **build AI tools that support methodologically sound, transparent, and unbiased risk analysis**. These tools are content-neutral, logic-based, and publicly documented. They do not promote, defend, or oppose any position—scientific, regulatory, or otherwise. They are incapable of doing so by design.

Again, it is unclear why *The Guardian* chose to insinuate bias and publish derogatory quotes from individuals who are uninvolved in—and uninformed about—the tool or the project, instead of engaging with its actual architecture, functionality, or outputs. This is especially troubling given that the reporter was provided with full and accurate documentation in advance.

ChatGPT

In mid-February of 2023, Cox struck up a conversation with the AI assistant ChatGPT. He later sent the chat to his University of Colorado email address, which is subject to public record request laws.

After asking the chatbot to write a sonnet about Abraham Lincoln, he turned the conversation to a more serious topic: the tiny toxic particles known as PM2.5, also known as fine particulate matter.

Scientists have long found links between PM2.5 exposure and a [wide variety of health concerns](#), from asthma and lung disease to heart attacks. But for years, Cox has raised uncertainty about those connections, publishing work [“re-assessing” the mortality risks associated with exposure to the pollutant in animal agriculture and challenging the link between PM2.5 from gas stoves and childhood asthma](#).

TC: Wrong on all counts.

1. My work has not “*raised uncertainty about those connections*” but has clarified the nature of the connections (associational but not interventional-causal). This *reduces uncertainty* about what the ambiguous term “link” actually means in these contexts, rather than raising it.
2. None of my work has involved “*challenging the link between PM2.5 from gas stoves and childhood asthma*.” It has clarified the nature of a claimed statistical link (associational but not causal) between NO2 (not PM2.5) from gas stoves and childhood asthma in a single, specific, recent paper (<https://pubmed.ncbi.nlm.nih.gov/37638367/>).

7. The Guardian’s Mischaracterization of Exploratory AI Research as Ideological or Coordinated

On that February day, Cox asked ChatGPT a simple question: “Does PM2.5 cause lung cancer?” The bot responded by noting the “strong scientific evidence” demonstrating the link.

A seemingly dissatisfied Cox went on to ask if it is “known with certainty that PM2.5 causes lung cancer”, calling on ChatGPT to identify potential confounding factors in the research on the air pollutant.

TC: There is nothing unusual, suspicious, or “dissatisfied” about asking whether an observed association reflects a proven causal relationship, and what uncertainties or alternative explanations might remain. That’s simply standard scientific reasoning. Asking about confounding, measurement error, and modeling bias is exactly what any competent analyst should do when evaluating a non-experimental association. **I was probing ChatGPT’s understanding and use of these crucial concepts and distinctions—not evincing dissatisfaction.**

He took brief tangents to pose riddles to the bot and to request additional sonnets – including one about PM2.5 itself, yielding the line “no nose can catch you, for you are so fine”.

TC: The sonnets and riddles were part of a broader exploratory interaction, testing ChatGPT’s versatility and tone-shifting capabilities. They served as informal probes of language generation as part of a longer investigation into the system’s reasoning and generative patterns.

Eventually, Cox asked ChatGPT whether the association between PM2.5 and lung cancer can be explained by a “combination of residual confounding, measurement errors, and modeling biases”. The bot maintained that “the observed association is likely to be real and not fully explained by these factors”.

TC: This was precisely the kind of question I was interested in exploring: whether ChatGPT could distinguish between association and causation, and whether it understood the limits of observational inference, and whether it was precise in stating conclusions to maximize their validity. The AI’s response—initially confident, then increasingly cautious when challenged—provided a useful illustration of how probabilistic language can mask substantial epistemic uncertainty. That’s why I later published excerpts as an example of Socratic dialogue for clarifying causal claims.

From a conversation Cox had with ChatGPT in February 2023



Click to see document

Write a sonnet about PM2.5 and lung cancer

Oh, deadly PM2.5 in the air, So small and yet so full of poison’s might, A silent killer that we cannot bear, You infiltrate our lungs both day and night

At the time of the conversation, Cox had already been interested for years in making “smarter and more trustworthy AI”, he told the Guardian. He had published a paper on the topic in 2021, and the following year – months before ChatGPT’s release – had begun developing his own AI tool. The ACC had partly funded that work, he said.

TC: Correct. I have long been interested in developing AI systems that reason more transparently and rigorously than current models. The work was partly funded by the American Chemistry Council, as fully and publicly disclosed in all relevant publications.

Cox also had a meeting with the Long-Range Research Initiative, a lobbying group focused on “innovations in chemical safety science” which includes Exxon, just hours before he had the February 2023 ChatGPT conversation, the emails show.

TC: This framing is misleading. To my knowledge, **the Long-Range Research Initiative is not “a lobbying group.”** Its stated and actual mission is, rather, **“to catalyze innovations in chemical safety science.”** www.americanchemistry.com/chemistry-in-america/research/long-range-research-initiative-lri. I **attended the meeting in the context of research discussions—not lobbying,** policy, or political activity.

In an email seeking funding to develop an AI tool, Cox seemingly referenced that “long experimental chat”.

TC: It’s not clear whether the email in question referred to that particular chat. In any case, my early experiments with ChatGPT—some of which involved extended conversations—were part of a careful attempt to assess its reliability and trustworthiness. I was especially interested in whether it exhibited cognitive patterns such as overconfidence and anchoring, as documented by Kahneman, Tversky, Slovic, and others. The work was exploratory and hypothesis-generating—not promotional.

Among the recipients were George Maldonado, the editor of the academic journal Global Epidemiology, and ACC toxicologist Rick Becker.

Cox wrote in the email that his questions eventually led ChatGPT to “concede that we don’t actually know that PM2.5 at current ambient levels causes lung cancer in humans – but it was a struggle to get there!” The chatbot “does an excellent job of reflecting the ‘party line’ that is most prevalent on the web, fallacies and all”, Cox continued in the email. But new AI software could be used to do “critical thinking at scale’ (if I may be grandiose!)”, he said.

TC: True. I stand by all of these statements. **The ChatGPT exchange revealed how default language models tend to reproduce conventional narratives without distinguishing between correlation and causation.** That was a central concern I sought to address with improved reasoning tools.

The following day, Cox emailed a larger group of researchers, including Becker and two ExxonMobil scientists. ChatGPT, he wrote, “seems to me to display a very strong starting bias that can eventually be overcome by sufficiently patient questioning”. That bias involved conflating “evidence of association with evidence of causation”, he said.

TC: Also true—and again, I stand by it. This was a key insight from my exploratory dialogues with ChatGPT in early 2023. The model frequently asserted causal claims when the underlying evidence was associational. **Recognizing and correcting this confusion was essential to the development of more trustworthy AI.**

From an email Cox sent to industry researchers in February 2023

We can help bend applications of this technology toward scaled-up critical thinking instead of scaled-up groupthink and propaganda

“I am hoping to build a critical mass of interest and get some funding in this area so that we can help bend applications of this technology toward scaled-up critical thinking instead of scaled-up groupthink and propaganda,” he added.

TC: True. That is still my goal. My aspiration has always been to support the responsible development of AI for scientific reasoning—not for messaging, ideology, or policy promotion. Tools that improve clarity, transparency, and causal logic serve **all sides** of a scientific debate.

8. *The Guardian's* Misrepresentation of the PFOA Study and Misuse of the Term 'Safe Dose'

*Cox's past work may shed light on the "groupthink and propaganda" that his work questions. In one [2023 study](#) he co-authored, **he found** that exposure to the "forever chemical" known as PFOA can occur in **safe doses**.*

TC: This is false and defamatory.

1. Far from illustrating "groupthink and propaganda," this study was **explicitly designed to avoid groupthink**. My primary contribution was to structure the study so that 24 participating scientists—divided into three independent international teams—developed their estimates separately, without influence from one another, before comparing results. That methodology was precisely intended to reduce groupthink bias and undue consensus, not perpetuate it.
2. The study did **not** "find that exposure to PFOA can occur in safe doses." That is a serious misrepresentation. **What we actually concluded and stated was that:**

"The available epidemiologic information could not form a reliable basis for a PFOA safe dose-assessment in the absence of mechanistic data that are relevant for humans at serum concentrations seen in the general population."

3. To reduce uncertainty, we derived a range of health-protective dose estimates from animal studies. These were not presented as definitive thresholds for human safety, but as calculations to support more informed and transparent risk assessments.
4. In regulatory toxicology, the term "safe dose" refers to a conservatively derived estimate (e.g., NOAEL ÷ uncertainty factor). It does not imply that exposures below that level are actually safe or risk-free. *The Guardian* misrepresents both the intent and the content of our paper—despite our conclusions being clearly stated.
5. Finally, I was not involved in the scientific deliberations that produced the dose estimates. **The Guardian's statement that "he [Cox] found..." is incorrect. My role was methodological:** to help design a study structure that minimizes bias and promotes independent, critical thinking. For *The Guardian* to cite this work—designed to prevent groupthink—as an example of it, is as perverse as it is uninformed.

The Guardian's characterization here is factually wrong, conceptually confused, and reputationally damaging. I request a formal retraction.

9. *The Guardian's* Tactics: *Ad Hominem* Attacks, Misinformation, and Guilt by Association

The research was conducted with the organization Toxicology Excellence for Risk Assessment, headed by the [contentious toxicologist Michael Dourson](#), who has also received funding from chemical makers.

TC: This is another example of *The Guardian* inserting irrelevant innuendo to cast suspicion where none is warranted.

1. To my knowledge, this research received no external funding at all. It was a voluntary collaboration by scientists aiming to reduce uncertainty in public health risk assessment. *The Guardian* provides no evidence to the contrary.
2. The mention of Michael Dourson's name and funding history is irrelevant to the study's design, execution, or conclusions. He did not fund the work. The study's structure—three independent expert panels working without influence or coordination—was precisely intended to safeguard scientific integrity and insulate the findings from any potential bias.
3. *The Guardian's* attempt to discredit the study by invoking Dr. Dourson's past funding—rather than addressing the actual content or methods of the research—is a textbook example of **ad hominem innuendo**. This kind of character-based attack contributes nothing to scientific understanding.
4. I note that *The Guardian* has previously targeted Dr. Dourson in a separate opinion piece (see: [Michael Dourson's rebuttal](#)), further suggesting an editorial pattern of mischaracterizing experts with whom it disagrees.

If *The Guardian* wishes to critique the science, it should engage with the methods and findings—not recycle innuendo about people uninvolved in funding or conduct of the research.

Another [study](#) the same year, which Cox co-authored with a Chevron toxicologist, said molybdenum – a chemical present in lubricants Chevron produces – was “not a risk factor for changes in serum testosterone”.

TC: This is misleading in both attribution and implication.

1. To my knowledge, the lead author of the paper in question—[Klipsch et al., 2023](#)— Dr. Klipsch, is an independent German researcher unaffiliated with Chevron. The study had multiple coauthors from different institutions. Referring to it as “Cox co-authored with a Chevron toxicologist” grossly misrepresents the authorship and wrongly suggests Chevron as the driving force in this work.
2. Any implication that my scientific conclusions are biased based on coauthor affiliations—rather than being evaluated on their empirical or methodological merits—is both irresponsible and defamatory.

The Guardian's summary here is a textbook example of insinuation by selective framing—describing a study by its coauthor's employer rather than its actual content, methodology, or results.

*“And in a third [2023 study](#), Cox said **his research found no causal link between childhood asthma and gas stove exposure.**”*

TC: This statement is false and misleading on multiple levels.

1. **There was no such “study.”** The 2023 article in question was a **methodological commentary** on someone else's study, not a “study” itself.
2. **The commentary did not state, imply, or suggest that my research “found no causal link.” I have done no such research**, and have not stated or pretended otherwise, as *The Guardian* falsely implies. Rather, my commentary explained that a **previously published claim**—specifically, that 13% of childhood asthma cases in the U.S. were attributable to gas stove pollution—was **not supported by the underlying data or methods**. That claim had been based on associational measures (Population Attributable Fractions), without establishing the required causal assumptions.
3. **I have conducted no original research on gas stove exposure and childhood asthma.** My commentary focused solely on the **misuse of causal language** in a high-profile public health claim. It highlighted the difference between statistical association and causal inference, using this case as an example of how misleading conclusions can arise when this distinction is ignored.
4. **The Guardian's formulation falsely attributes to me a research finding I never made**, and then mischaracterizes the nature of my contribution. This misrepresents my work.



[View image in fullscreen](#)

A growing body of research shows gas stoves emit toxic compounds even when not in use. Photograph: Jena Ardell/Getty Images

At a [2018](#) conference, Cox also claimed there is no proven connection between air pollution and respiratory problems or heart attacks,...

TC: This is false. I said no such thing.

1. The quote attributed to me appears to have originated in a 2018 Mother Jones article, which misrepresented my remarks at a conference. That mischaracterization has now been recycled by *The Guardian* without verification.
2. **I have never claimed that there is “no proven connection” between air pollution and respiratory or cardiovascular problems.** On the contrary, **much of my work has been devoted to proving, clarifying and quantifying the nature** of those connections—particularly the distinction between associational patterns and causal mechanisms.
3. In fact, in 2018, the same year cited by *The Guardian*, I explicitly wrote:

“Heart disease risk varies with age, sex, smoking, income, education, and particulate (PM2.5) air pollution.”

—Cox LAT Jr. [Socioeconomic and particulate air pollution correlates of heart disease risk](#). *Environ Res.* 2018 Nov;167:386-392. doi: 10.1016/j.envres.2018.07.023. Epub 2018 Jul 20. PMID: 30098525.

4. ***The Guardian’s* paraphrase is therefore not only inaccurate—it is directly contradicted by my published work** from the very year they cite. This amounts to a false attribution of a claim I have never made, in contradiction to the public record.

...while he said in a 2012 [paper](#) – funded in part by tobacco company Philip Morris USA – that he found smoking half a pack of cigarettes daily did “not appear to be associated” with increased risk of coronary heart disease.

TC: This is true only in the most limited and technical sense, and its selective presentation in *The Guardian* is highly misleading.

1. The quoted phrase—“not appear to be associated”—is technically accurate in context, but it **explicitly refers to association, not causation**. My 2012 paper took pains to explain that observed patterns in the data do not necessarily reflect causal effects, due to the presence of confounding factors.
2. In fact, the paper explicitly states:

“The nonlinear dose-response relations suggested in these figures may not reflect a true cause-and-effect relation, but rather the effects of confounders (such as alcohol and coffee consumption, which, as noted previously, are associated with cigarette consumption and also have U-shaped relations with CHD risk) or other variables associated with both smoking and CHD risks.”

3. The entire purpose of this analysis was not to deny that smoking causes heart disease—an established fact—but rather to **demonstrate the dangers of inferring causality from associations in observational data without proper control for confounding**. This is a methodological point, not a health claim.
4. Presenting this paper as evidence of my minimizing smoking risks—without any mention of its stated methodological focus—is highly misleading. It also ignores the broader body of my published work, which includes multiple papers elucidating how and why smoking causes disease (e.g., in mechanistic pathways leading to COPD and cancer).

I plan to request a correction that acknowledges the methodological framing of the 2012 paper and retracts the insinuation that it denied established causal links between smoking and heart disease.

10. Misrepresentation of Scientific Disagreement as Ideological Opposition

In an email to the Guardian, Cox said the methods he applies are “drawn from the scientific mainstream – not from ideology or partisanship”.

“Some critics have mischaracterized my work as an attempt to delay regulation or promote industry interests. That is not true,” he said. “I do not advocate for or against any policy outcome. I advocate for grounding decisions in empirically supported causal understanding.”

Cox later responded that his work had demonstrated the “strength” of “linking exposure to chemical compounds to health dangers”, adding: “To the extent that my career has touched on the dangers of pollutants, it has done much to remove doubt.”

TC: This is accurate, and I appreciate *The Guardian* quoting me correctly in this passage. These statements reflect my consistent position:

- My work is rooted in mainstream scientific methods, including causal inference, risk modeling, and mechanistic reasoning.
- I do not take positions on policy outcomes. My concern is with **methodological rigor**, transparency, and scientifically valid reasoning.
- Where I have studied pollutants, my aim has been to **clarify causal relationships**, not obscure them. In many cases, my work has helped remove doubt by identifying when observed associations do or do not support causal interpretations.

Cox served as an adviser to policymakers in his role on an EPA advisory committee.

TC: This phrasing misrepresents the nature of the role.

1. I served on the EPA’s Clean Air Scientific Advisory Committee (CASAC) to provide **independent scientific advice**, not to advocate for or set policy. CASAC’s charter is explicitly scientific, not political.
2. As CASAC chair, I was responsible for evaluating the **scientific evidence** relevant to air quality standards—not for advising policymakers on what to do with that evidence. That distinction matters. Scientific analysis and policy advocacy are separate functions, and I fulfilled the former, not the latter.
3. Presenting this as “advising policymakers” may subtly but misleadingly imply or suggest a political or advocacy role, which is neither accurate nor consistent with CASAC’s mandate.

He has also [argued against the proposed tightening of a regulation](#) at an Occupational Safety and Health Administration hearing, in his capacity as an ACC consultant.

I testified at an OSHA hearing as an **independent risk analyst**, presenting my own (not necessarily the ACC’s) views about the proposed rule’s evidentiary basis. While I was consulting for the American Chemistry Council (ACC) at the time, my testimony was based on peer-reviewed methods and evidence, not on advocacy or preference. My analyses in such settings are always grounded in transparent, reproducible, and scientifically accepted methodologies.

11. How *The Guardian* and Its “Experts” Misrepresent My Views on Scientific Objectivity and Uncertainty

11.1 No, Science Is Not Just a Matter of Preferences

Adam Finkel, a risk analyst and environmental health sciences professor at the University of Michigan, said though he believes Cox to be in some ways a “genius” and skilled risk analyst, he also seems to be “deceiving himself and everyone else” about the impacts of bias on his research.

“How you interpret any information is by imposing your preferences,” said Finkel, who is also a former director of health standards programs at the US Department of Labor’s Occupational Safety and Health Administration. “There is no possible way to get around imposing some set of preferences.”

TC: This reflects Dr. Finkel’s philosophical position, not a valid critique of my work.

1. I reject the premise that interpreting evidence must involve “imposing some set of preferences.”

That position conflates subjective judgment with objective analysis. Much of modern mainstream science—from statistical hypothesis testing to causal inference—is explicitly designed to minimize the influence of personal bias through formal methods and transparency. My view is that everyone should be able to interpret and verify objective factual information, such as that $1 + 1 = 2$, without “imposing some set of preferences.”

2. My work aims to clarify the conditions under which valid inferences can be made, and to separate empirical evidence from wishful thinking or ideological priors. I do not “deceive myself” or anyone else: I document my assumptions, methods, and data sources precisely so that others can evaluate or replicate my conclusions.

In contrast to Dr. Finkel’s statement, I disagree that “*There is no possible way to get around imposing some set of preferences*” in interpreting scientific information. To the contrary: it is essential not to do so if we are to learn what is true, rather than what we wish were true.

Much harm has been done in public health by those who believe that imposing their own preferences on data – or substituting preferred narratives for objective evidence – is necessary or even desirable. The claim that “*There is no possible way to get around imposing some set of preferences*” is mistaken. Those who hold this view would do well to study the methods of modern risk science, which are expressly designed to test beliefs rigorously against data without regard for personal preferences.

This is the heart of the objective scientific method that I advocate.

3. Ironically, Dr. Finkel’s claim—that one cannot avoid imposing preferences—is itself a form of epistemological pessimism that undermines the very idea of scientific objectivity. I reject that view. I have spent my career demonstrating that more objective, transparent, and empirically grounded science is not only possible—it is essential for identifying causally effective policies and interventions that actually protect public health.

11.2 Uncertainty Is Not an Excuse for Inaction—or for Unsupported Action

Some degree of uncertainty is inherent to scientific analysis. But when assessing whether or not there is a causal effect between exposure to something potentially harmful, Finkel said, Cox looks for “perfect

certainty”, which “can lead to years and decades of doing nothing and harming people while you wait for the certainty to come”.

TC: This claim is both false and misguided. I **do not “look for perfect certainty”** or “wait for the certainty to come.” On the contrary, I have spent decades analyzing how to make sound decisions under uncertainty – including analyzing when it is better to act now versus wait for more information. Since at least 1984, I have published on probabilistic causation and decision-making with incomplete evidence—for example, in [Probability of causation and the attributable proportion risk](#) (Cox, 1984)—and I have returned to these topics repeatedly in both peer-reviewed articles and books (e.g., [Breakthroughs in Decision Science and Risk Analysis](#), Wiley, 2012).

These contributions explicitly address how to act responsibly in the face of uncertainty, not how to demand “perfect certainty.” The notion that my work “*can lead to years and decades of doing nothing and harming people while you wait for the certainty to come*” is a fabricated straw man. No one I have worked with, nor anything I have published, advocates inaction or paralysis. What I and many others do advocate is the use of objective, transparent, and testable scientific evidence—not ideological certainty or wishful thinking—to improve the quality of decisions we make today based on the best available information.

This uncertainty is always present, but that of course doesn’t mean the research is wrong

Gretchen Goldman, president of the Union of Concerned Scientists

TC: No, but it also does not mean that the research is right—or that it cannot be greatly improved through more appropriate technical methods. **Uncertainty is not a license for lax standards or uncritical acceptance of unsound research.** I have addressed this very point in several publications, including in direct response to arguments by Dr. Gretchen Goldman. Like Dr. Finkel, she often appears to advocate for *preference-based* approaches—such as “weight-of-evidence” methods that reflect the collective judgments or preferences of selected experts—over *science-based* methods that are explicit, testable, and transparent. In my view, that is a step backward. See, for example, [“Should health risks of air pollution be studied scientifically?”](#) (*Global Epidemiology*, 2019), in which I argue for the primacy of objective evidence over narrative-based or consensus-driven assessments.

The Guardian could easily have quoted leading figures in modern risk analysis and decision science who share my view that rigorous, objective methods better serve the public than appeals to “consensus” or “imposing some sort of preferences” or other ideologically motivated heuristics. Their omission suggests a one-sided narrative rather than a good-faith journalistic inquiry.

While Finkel has “fundamental belief that our system is under-protective” when it comes to public health, Cox seems to believe the opposite.

TC: That is incorrect. I am agnostic on whether public health systems are systematically under- or over-protective. I believe this is an empirical question that should be answered on a case-by-case basis using relevant data and sound analytical methods—not pre-judged based on someone’s “fundamental belief,” no matter how well-intentioned. I suspect that much of the criticism directed at my work by Drs. Finkel and Goldman stems from their preference to impose such prior beliefs on society at large, often without applying the technical methods of modern risk science needed to ascertain whether those beliefs are factually correct. In contrast, I advocate for letting objective facts and science, rather than ideological prior beliefs and preferences, inform policy decisions about health risk management.

Asked for comment, Cox said: “I have never advocated that we should not act until we have certainty. Rather, I have advocated choosing to act on the best available information.” He said his work has acknowledged causal relationships between [smoking and lung cancer](#), [asbestos exposure and mesothelioma](#), and, in 2011, [crystalline silica exposure and lung disease](#).

But at the Occupational Safety and Health Administration hearing at which he spoke in 2014, Cox [asserted on behalf of the ACC](#) that the federal government had not demonstrated a link between certain levels of silica exposure and lung disease.

“He’ll accept that at very high doses, this stuff is bad for you,” said Finkel. Policy is meant to ensure that level of exposure doesn’t occur, he added.

TC: Correct—but incomplete and potentially misleading. My testimony and research have focused on improving our understanding of *how different exposure patterns over time* affect risk, and on ensuring that regulatory standards are based on *accurate, empirically supported causal models*. At the 2014 OSHA hearing, I testified that the federal government had *not adequately demonstrated* that the proposed *specific exposure limits* would reduce lung disease risk, given the available data and the complex dynamics of cumulative and time-patterned exposure.

Unfortunately, current policies often implement *untested and empirically mistaken assumptions* about how different time patterns of exposure (e.g., short high peaks vs. long low exposures) influence health risks. A major theme of my scientific research over the past two decades has been to *correct these misconceptions* and to develop more accurate risk models that can improve both worker protection and regulatory effectiveness.

12. How *The Guardian* and Its “Experts” Misrepresent Peer Review, Current Scientific Practices, and My Scientific Views

‘Socratic dialogue’

*Maldonado, editor of *Global Epidemiology*, responded positively to Cox’s AI assistant proposal, the emails from 2023 show. Within weeks, his journal [published](#) another one of Cox’s conversations with ChatGPT in his journal.*

*“The purpose of this comment is to provide an example of a Socratic dialogue with ChatGPT about the causal interpretation of an important epidemiological association between exposure to fine particulate matter air pollution (PM2.5) and mortality risk,” says the paper, which [states](#) that it was partly funded by the ACC and **counted climate denier Steve Milloy as one of its reviewers.***

TC: This is false and defamatory.

Steve Milloy was *not* a reviewer of the paper in question. The journal *Global Epidemiology* follows standard academic peer-review protocols. Reviewers are anonymous and selected by the editors based solely on relevant subject-matter expertise—not ideology or politics.

The Guardian’s claim that Milloy served as a reviewer is factually incorrect and damaging. It falsely suggests that the peer-review process was compromised or manipulated, thereby misrepresenting both the journal’s editorial practices and the scientific integrity of the paper. This reckless assertion undermines public trust in peer-reviewed science and constitutes a serious mischaracterization of the editorial process followed by *Global Epidemiology*.

When the bot said *“it is well-established that exposure to ambient levels of PM2.5 does increase mortality risk”, Cox accused it of confusing evidence of association with evidence of causation. Eventually, ChatGPT said: “It is not known with certainty that current ambient levels of PM2.5 increase mortality risk.”*

TC: This is misleading. I did not “accuse” ChatGPT of anything. My exact words were: *“Your response seems to me to confuse evidence of association with evidence of causation.”* This is a reasoned observation, not an accusation. It reflects the dialogue’s constructive, exploratory tone. The phrase “accused it” mischaracterizes the nature of the exchange by suggesting adversarial intent or hostility. That implication is false and incompatible with the careful, respectful, and inquisitive tone of the published dialogue.

In addition, referring to ChatGPT as a “bot” is technically misleading. ChatGPT is not a bot in the conventional sense of automated scripts or simple rule-based agents. It is a large language model. The use of the term “bot” here obscures the actual nature of the tool and may misinform readers unfamiliar with the underlying technology.

But the distinction between correlation and causation is “epidemiology 101”, said Gretchen Goldman, president of the scientific advocacy group Union of Concerned Scientists, who [co-authored a 2019 paper](#) critiquing Cox.

“From day one of a study, researchers consider, analyze and guard against possible confounding factors,” said Goldman. “This uncertainty is always present, but that of course doesn’t mean the research is wrong.”

TC: Dr. Goldman’s assertion that “from day one of a study, researchers consider, analyze and guard against possible confounding factors” is, unfortunately, not supported by the methodological record of

air pollution epidemiology studies. A substantial body of critical literature has highlighted widespread shortcomings in confounder control and causal design in this domain. For example, a 2020 review applying the GRADE framework to air pollution health effect studies concluded:

“Applying the GRADE approach to appraise the certainty of evidence yielded low or very low ratings for all primary health and ambient air quality outcomes.... Common issues comprised choice of intervention and selection sites and the **lack of consideration of potentially important confounders.**” (Burns J, Boogaard H, Polus S, Pfadenhauer LM, Rohwer AC, van Erp AM, Turley R, Rehfues EA. [Interventions to reduce ambient air pollution and their effects on health: An abridged Cochrane systematic review.](#) Environ Int. 2020 Feb;135:105400. doi: 10.1016/j.envint.2019.105400. PMID: 31855800.)

This is consistent with my observation that much of the literature continues to report associations without rigorous causal interpretation or adjustment for key sources of confounding and bias. Acknowledging the distinction between correlation and causation is a good start—but by itself, it is not enough. Demonstrably guarding against confounding through formal design and transparent statistical justification remains an unmet need in many areas of environmental health research, and in air pollution health effects research, specifically. I have advocated for changing this decades-long pattern of poor practice, which I believe has contributed to demonstrably – but needlessly – ineffective regulations.



[View image in fullscreen](#)

Cox has critiqued some proposals to strengthen controls on pollution on the grounds of imperfectly demonstrated causality. Photograph: Paul Hennessy/SOPA Images/LightRocket via Getty Images

TC: This is misleading. **My critiques are not based on “imperfectly demonstrated causality,”** as if any imperfection would be disqualifying. Rather, they focus on the fact that many studies drawing causal conclusions about pollution and health have not used valid causal models at all. Instead, they rely on

associational analyses that do not meet even the basic methodological requirements for valid causal inference—for example, failing to control for known confounders such as lagged high and low daily temperatures over the 30 days preceding mortality.

In such cases, the issue is not that causality is imperfectly demonstrated, but that it is not demonstrated at all using valid methods. My critiques and the evidentiary basis for them are summarized on [Cox, 2023, Improving interventional causal predictions in regulatory risk assessment](#). “Imperfectly demonstrated causality” is *not* one of those grounds.

13. How *The Guardian* and Its “Experts” Misrepresent the Feasibility and Ethics of Causal Inference in Environmental Health

*Demonstrating clear causal links between pollutants and health impacts can be complicated, especially because unlike in testing pharmaceuticals, **it can be difficult and unethical to establish control groups for comparison.***

TC: The comment that “it can be difficult and unethical to establish control groups for comparison” reflects a lack of understanding of modern epidemiology designs and method. Establishing control groups for evaluating the health effects of pollution has already been done many times. For example, [Lyons et al., 2023, The effect of “smoky” coal bans on chronic lung disease among older people in Ireland](#) used the staggered introduction of “smoky” coal bans across different regions in Ireland to define treatment and control populations, yielding credible causal estimates of impacts on chronic lung disease. [Burns et al., 2020](#) likewise assessed air pollution interventions using comparative designs that incorporated control groups across locations and time.

These and many other studies demonstrate that **appropriate comparison groups can be identified and used without ethical compromise. It is neither difficult nor unethical to do so.** Arguably, it is unethical *not* to do so, as failure to use appropriate comparison groups delays—often by decades—the discovery of which interventions actually improve health and which do not. Claims to the contrary reflect a refusal to apply well-established causal inference tools, not a lack of available methods. ***The Guardian* and its “experts” should stop making excuses and start applying proven scientific approaches,** including the use of appropriate control groups, to avoid repeating the kinds of belief-driven errors that have led to demonstrably false conclusions about the effects of regulatory interventions ([Zigler and Dominici, 2014](#)).

“If you’re looking at the effects on an actual population that’s been exposed in real life to pollutants, you can’t do those controlled types of studies,” said Frey of North Carolina State University’s College of Engineering.

TC: It is **both possible and essential to control for confounding and construct valid comparisons when studying the effects of real-world pollution exposures.** That is precisely what modern causal inference designs and methods are intended to do. Many studies have successfully applied these methods in

observational settings, using comparison groups and appropriate covariate adjustment to support valid causal conclusions.

A good example is [Lyons et al., 2023](#), which exploited natural variation in the timing of smoky coal bans across Ireland to estimate causal effects on chronic lung disease. They report that “**Controlling for relevant factors**, smoky coal bans reduced the probability that an older person reports being diagnosed with chronic lung disease by between three and five percentage points. ...Our findings were robust to estimating the model using different sub-samples and control variables.”

Similarly, [Zigler and Dominici, 2014](#) applied core principles of causal design and analysis – well-established since the early 1960s in other fields – to evaluate effects of regulatory interventions. These are precisely the principles that much of the air pollution health effects literature has routinely neglected or violated for decades. Claims that “you can’t do those controlled types of studies” or that “it can be difficult and unethical to establish control groups for comparison” are contradicted by the long and growing list of quasi-experimental studies that have done exactly that—successfully, ethically, and transparently.

Unfortunately, The Guardian and its selected “experts” continue to promote outdated and misleading assertions that have, for decades, impeded progress in identifying interventions that actually improve health outcomes.

The studies discussed by [Lyons et al., 2023](#), [Zigler and Dominici, 2014](#), and many others ([Burns et al., 2020](#)) are quasi-experimental studies that incorporate comparison groups and support modern causal inference methods. While not randomized, they **yield strong, policy-relevant causal estimates from real-world data**. To suggest that such research is infeasible or unethical reflects a misunderstanding—or rejection—of well-established scientific methodology, not any inherent barrier to conducting rigorous and ethical causal analysis in environmental health.

14. How *The Guardian* and Its “Experts” Misrepresent My Scientific Standards, Motivations, and Reputation

“That leads to thinking about ways to make inferences from real world data that might, for example, mimic a random, controlled trial.”

But though demonstrating true causality can be complex, Cox has long overstated scientific uncertainty while downplaying evidence, said Frey.

TC: Frey’s characterization is incorrect. I have consistently emphasized placing greater weight on valid causal evidence—and less on subjective preferences, speculative narratives, or untested assumptions (e.g., [Cox, 2021, Toward practical causal epidemiology](#); [Cox, 2019, Should health risks of air pollution be studied scientifically?](#)). That is not “downplaying evidence.” It is rejecting the substitution of human heuristics and ideological biases for valid scientific evidence—a pattern that characterizes much of the work promoted by advocates of subjective “weight-of-evidence” judgments, including Frey.

Recognizing uncertainty is not the same as overstating it. My work focuses on distinguishing when causal claims are justified by data and methods, and when they are not. That distinction is foundational to sound science and to public health policies that actually work.

Science denialism often sounds convincing because it contains some truthiness to it

Chris Frey of the North Carolina State University's College of Engineering

"Science denialism often sounds convincing because it contains some truthiness to it or elements of truth or elements of valid points, but it's often based on either overemphasis or omission and doesn't portray a full picture," he said.

TC: I have never engaged in science denialism. On the contrary, I have called for more rigorous, transparent, and testable applications of scientific methods to evaluate health risks—especially where current practice relies on weak or assumptive reasoning (e.g., [Cox, 2019, Should health risks of air pollution be studied scientifically?](#)) I have consistently advocated for the application of the standard scientific method and objectively sound science—rather than for the personal preferences, “fundamental beliefs,” or selective interpretations of “experts” who refuse to use objective scientific and causal analysis methods.

What Frey calls “truthiness” often reflects unexamined assumptions dressed up as evidence. My work aims to clarify when evidence truly supports a causal claim—and when it does not. That is the opposite of denialism; it is a demand for higher scientific standards (ibid).

As chair of EPA's clean air scientific advisory committee during Trump's first presidential term, for instance, Cox proposed [eliminating all research from the agency's consideration that did not demonstrate “manipulative causation”](#), wherein intervention on one variable would change the probability of an outcome. “I see it as being about using widely accepted, non-controversial principles of causal analysis and inference,” Cox said of his push for this change.

TC: This claim is false. **I never proposed eliminating all research that does not demonstrate “manipulative causation.”** I have consistently advocated for the inclusion of studies that apply valid causal inference methods—particularly intervention and accountability studies that had been systematically overlooked in prior EPA assessments—whether or not they demonstrate manipulative causation.

The Guardian's own source for this claim links to a document authored by C. Frey in 2013—years before I had any involvement with CASAC. It does not substantiate their allegation. Citing it to support a claim about my actions as CASAC chair is both misleading and irresponsible.

But in effect, the alteration would have dramatically and unnecessarily “winnowed down” the body of evidence to which the EPA could have referred and removed research from consideration which “in fact robustly” demonstrates that certain compounds cause harm, Frey said.

TC: This is false. **I did not call for narrowing the body of evidence EPA could consider—I called for expanding it. Specifically, I advocated including intervention and accountability studies that were being**

systematically overlooked in EPA assessments. These studies often provide stronger causal insights than purely associational analyses and are essential for evaluating real-world policy impacts.

I was explicit about this position on multiple occasions—for example, Cox LA Jr, Popken DA. [Has reducing fine particulate matter and ozone caused reduced mortality rates in the United States?](#) *Ann Epidemiol.* 2015 Mar;25(3):162-73. doi: 10.1016/j.annepidem.2014.11.006. PMID: 25571792). My goal was to strengthen the evidentiary base, not restrict it.

“That effort and his work generally have not been viewed as compelling by the mainstream scientific community,” he added.

TC: Frey’s claim is unfounded. Contrary to his implication, **my work on causality has been recognized as compelling by respected members of the mainstream scientific community.** When I was inducted as a Fellow of the Institute for Operations Research and Management Science (INFORMS) in 2019, the citation explicitly noted “significant research, practice, and service contributions to homeland security, health and environmental risk analysis, telecommunications, **and the modeling of causality.**” (www.informs.org/Recognizing-Excellence/Fellows/INFORMS-Fellows-Class-of-2019).

I suspect Frey may be using the term “mainstream scientific community” to refer only to those who share his particular views. That differs from my understanding, which includes the contributions of over a dozen Nobel Laureates and Turing Award winners who have advanced modern causal inference and the scientific method itself. Either way, many respected scientists and institutions value methodological rigor, objective science, and causal clarity—principles that are central to my work and widely recognized as essential to sound science.

Industry interests have promoted uncertainty to defend their business models, Frey said. The oil sector, for instance, had strong evidence that fossil fuels warmed the planet as early as the 1950s yet [publicly](#) called the link “weak” or even “non-existent” for decades. Cigarette manufacturers also long promoted the idea that the connection between cigarettes and health harms was tenuous, with one tobacco executive even [saying](#) in 1969 that “doubt is our product”.

“It’s a well-worn tactic,” said Frey.

TC: **No one I have worked with or for has ever expressed any interest in promoting uncertainty to protect a business model.** My work is focused on reducing uncertainty—by applying rigorous, transparent scientific methods to clarify when causal claims are justified and when they are not. That is the opposite of the tactics Frey describes.

Invoking historical misconduct by the tobacco or fossil fuel industries does nothing to illuminate or accurately characterize my work. It’s a rhetorical maneuver, not a factual critique. Sadly, *The Guardian* and its selected “experts” seem more interested in guilt-by-association smears and other non-evidentiary tactics than in engaging with the actual content, methods, and transparency of the science and AI they claim to critique.

15. How *The Guardian* and its Chosen Sources Misrepresent My AI Tool, Its Purpose, and Scientific Collaborations

Industry collaboration

Cox kept corresponding with industry scientists about his new tool, all the while holding similar conversations with ChatGPT about causation in research. In May 2023, for instance, Cox posed questions about the causal claims in a recent landmark study linking [gas stove exposure to childhood asthma](#), the emails show.

Later that month, Cox sent a slideshow to the ACC's Becker and several other industry-related scientists. His reviewer, it showed, had identified issues with the recent gas stoves study, and another major assessment which linked [PM2.5 exposure to cardiovascular issues](#).

This tool could “benefit authors, reviewers, reporters, media (if we make the summary reports good enough), and decision-makers and policymakers trying to evaluate studies and decide how trustworthy their methods and conclusions are”, Cox said.

In a proposal sent days later, he added that it is “probably good enough to be commercially useful”.

In July 2023, Cox presented his new tool to members of the Long-Range Research Initiative – which also funded his earlier work – including to representatives from Exxon.

TC: Yes. I believe—and members of the ACC agree—that **AI tools capable of reliably identifying unsound arguments and unsupported conclusions can greatly benefit scientists, reviewers, decision-makers, and the public.** Developing such tools in collaboration with stakeholders, including industry scientists, is not unusual or inappropriate—particularly when the goal is to improve scientific reasoning and transparency.

All of my collaborations have been disclosed and conducted in the open. The value of a scientific tool should be judged by the rigor of its methods and the validity of its outputs, not by guilt-by-association implications based on who expresses interest in using it.

*Ahead of the meeting, Cox sent the group a conversation he had with the reviewer, which used a **2020 paper demonstrating a causal link between PM2.5 and mortality** as an example of the kind of conflation his tool could spot.*

TC: No, the 2020 paper did not demonstrate “a causal link between PM2.5 and mortality.” That is exactly the conflation my tool was designed to detect and help prevent. The study presented an observed association, not valid evidence of a causal effect—and mistaking one for the other has led to widespread misunderstanding and poor policy decisions. *The Guardian* appears to be committing the same error my tool was explicitly created to address.

*Maldonado, the editor of *Global Epidemiology*, offered to give the tool a “friendly trial” at his journal.*

From an email Cox sent to the American Chemistry Council's Becker in July 2023



Click to see document

Such automated critical reasoning can help to thoroughly review, and potentially to improve, the scientific claims and scientific integrity of causal reasoning and presentation of evidence underlying many regulatory risk assessments

After the meeting, Cox sent a two-part project proposal to the ACC. “Such automated critical reasoning can help to thoroughly review, and potentially to improve, the scientific claims and scientific integrity of causal reasoning and presentation of evidence underlying many regulatory risk assessments,” Cox said.

For part one, an academic paper on the project which would be published in Maldonado’s Global Epidemiology, he asked for \$75,000. For part two, a pilot testing the reviewer on submissions to the same journal, he asked for \$80,000. In his response to questions from the Guardian, Cox confirmed the ACC’s funding but not a dollar amount.

Cox published the “phase 1” [paper](#) about his new AI reviewer in the journal Global Epidemiology in June 2024.

He also appears to have secured \$40,000 for Global Epidemiology to participate in the second phase, but the partnership “did not come to fruition” because too few authors were willing to participate, Cox told the Guardian. Maldonado did not respond to a request for comment.

By April 2024, Cox told the ACC’s Becker in an email that his reviewer tool was “ready for a demo”, claiming its reviews are “already better than many human reviews, although not as on-point and insightful as the best human reviews”.

But in an email last May to toxicologist [Ted Simon](#), Cox said “the real goal” of the tool was to enable it to do literature reviews, examining wide swaths of published information in a particular subject area. That month, ExxonMobil scientist Hua Qian ran a test of the tool.

Now, Cox told the Guardian, the tool was being tested by researchers submitting work to the journal he edits, Risk Analysis, and other academic journals, including Decision Analysis. About 400 people have tested the tool so far.

Itai Vardi, a manager at the Energy and Policy Institute, who shared the trove of emails with the Guardian, said the project could have disastrous consequences for academia, particularly epidemiology.

“AI language models are not programmed, but built and trained,” he said, “and when in the hands and funding of this industry, can be dangerous as they will further erode public trust and understanding of this crucial science.”

TC: Vardi’s comment reflects a fundamental misunderstanding of the technology. The reviewer tool I developed is not an “AI language model” and is not “trained” in the way models like ChatGPT are. It does not generate text or mimic human language. Instead, it applies explicit, rule-based reasoning to evaluate the logic, structure, and evidentiary support of causal claims in scientific papers.

These remarks are not only irrelevant to the actual system—they are reckless. Publishing inflammatory and defamatory claims like “the project could have disastrous consequences for academia,” from someone with no knowledge of the tool’s structure, function, or use, is **not responsible journalism. It’s alarmism—and it defames a project that has no such potential for harm,** but instead has demonstrated and documented potential to improve scientific and methodological quality and clarity of published research. The only entities for whom it “could have disastrous consequences” are those who have grown comfortable using junk science and statistics to advance their preferred agendas—while resisting the application of objective scientific methods. Their opposition to an AI system that reliably detects unsound reasoning is understandable, if not forgivable.

The actual purpose of the reviewer is the opposite of what Vardi suggests: to improve transparency, clarify causal reasoning, and help scientists, reviewers, and editors assess whether causal claims are supported by valid evidence and methods. That’s why more than 400 researchers have already tested it, including scientists at academic journals, regulatory agencies, and independent organizations. ***The Guardian* could have asked any of them for an informed opinion—but instead chose to print unsubstantiated fear-mongering from a political advocacy group.**

It is unfortunate that *The Guardian* chose to amplify accusations rooted not in evidence or understanding, but in rhetorical tactics—mischaracterization, insinuation, and guilt by association. Giving voice to those who substitute ad hominem attacks, defamation, and ideological framing for critical engagement with actual methods and results undermines both scientific discourse and public trust. If there is concern about eroding scientific integrity, this is where the erosion begins.

16. How *The Guardian* Misrepresents the Purpose, Independence, Capabilities, and Scientific Standards of My AI Reviewer Tool

‘Sound science’

Asked about critics’ concerns about the ACC’s funding for the project, Cox said: “People who are concerned about the use of sound science in areas where politics has dominated might understandably be concerned about the use of such tools.”

But people should “favor the development” of the AI tool if they want to “apply sound science to improve our understanding of the world and how to act more effectively”, he said.

“The fact that the ACC ... are starting to step up to the challenge of designing AI to increase the objectivity, transparency, and trustworthiness of scientific research seems to me to be a great public benefit,” he said.

But the ACC “cannot be trusted as a source of ‘objectivity, transparency, and trustworthiness of scientific research,’” said Frey, when that research is “aimed at understanding the human health harms caused by chemicals manufactured by their members”.

TC: That may be why the ACC is not, in fact, conducting this research. The project was designed and carried out independently by Cox Associates, an independent scientific research firm. Moreover, the research is not aimed at understanding the human health harms caused by chemicals manufactured by ACC members. It is aimed at improving the clarity, consistency, and scientific rigor of causal reasoning in research across many fields—regardless of what conclusions the data support. **Frey’s framing misrepresents both the purpose of the work and who is doing it.**

The actual objectives of the project are (i) to provide authors with feedback on whether a draft paper, abstract, or proposal is likely to meet the editorial and methodological standards of the journal *Risk Analysis* [<https://www.sra.org/journal/what-makes-a-good-risk-analysis-article/>], and (ii) to deliver substantive, technical feedback on how to improve the clarity of methods and the interpretation of results.

Frey’s comment does not engage with either of these aims. Instead, it attacks a strawman version of the project that bears no resemblance to its actual content, goals, or design. **Frey’s comment misrepresents both the nature and purpose of the project and fails to address its actual goals.**

And for him, Cox’s use of the term “sound science” also prompted concern.

“‘Sound science’ is a term popularized by the tobacco industry as part of a campaign to create burdens of proof far beyond those required for policy decisions,” Frey said. Indeed, in the 1990s, Philip Morris USA – for whom Cox has done research – ran a 10-year “sound science” public relations campaign to [sow doubt](#) about the harm cigarettes cause.

*In an email to the Guardian, Cox noted that [“reputable scientists” use the term](#) to refer to reliable, verifiable research that follows accepted scientific methods. **He dismissed the idea that causation can be difficult to prove in epidemiology.***

TC: This is false and potentially defamatory. **I have never “dismissed” the idea that causation is difficult to establish in epidemiology.** On the contrary, I have written extensively—entire books and numerous peer-reviewed articles—on the very real challenges of making valid causal inferences in epidemiologic research (e.g., Cox, 2021, *Toward Practical Causal Epidemiology*).

What I reject is the claim that it is inherently too difficult to follow basic principles of sound methodology—such as controlling for obvious measured confounders, specifying testable causal

assumptions, and comparing implications to data. Suggesting that we must abandon these principles and defer instead to untested expert opinion or “fundamental beliefs” is not a solution—it is a rejection of science. When such beliefs are contradicted by data, insisting on them is not scientific caution; it is dogma.

Sadly, this kind of mischaracterization appears to be a recurring pattern in *The Guardian’s* reporting: quoting out of context, misrepresenting scientific positions, and relying on insinuation and ad hominem critiques rather than engaging with the actual evidence. That may serve a narrative, but it does not serve readers—or science.



[View image in fullscreen](#)

Some public health experts are alarmed about Cox’s AI tool. Photograph: Toshi Sasaki/Getty Images

“My response to people who are concerned that we should treat evidence of repeated associations as if it were evidence of interventional causality is that this outdated style of thinking is tremendously harmful and counterproductive in designing effective measures to successfully protect human health and safety,” he said.

*Asked for examples of harmful policies created by overreliance on association, **Cox named several scientific studies, including a [1996 experiment](#)** which was stopped because interventions that were expected to slash participants’ chances of getting lung cancer “based on repeatedly observed associations” actually increased that risk.*

He did not name any policies.

TC: This is misleading.

1. *The Guardian* blurs the line between studies and policies. The CARET trial itself was not a policy—it was a rigorous test of a policy-relevant recommendation. That recommendation, based on observational studies, held that high-dose beta-carotene might prevent lung cancer and was widely accepted in scientific and public discourse by the late 1980s. The CARET and ATBC trials were launched in response to this belief. When these large, randomized controlled trials showed that beta-carotene actually increased lung cancer risk in smokers, they forced a reversal of prior expectations. This outcome illustrates precisely the kind of harm that can result when health advice or public messaging is based on a repeatedly observed association rather than on valid causal evidence.
2. **The CARET trial was not merely “a 1996 experiment.”** It was a large, multicenter, randomized, double-blind, placebo-controlled primary prevention trial that **began in 1985** and enrolled over

18,000 participants—including smokers, former smokers, and asbestos-exposed workers. It was halted in 1996 when the intervention—beta-carotene supplementation, recommended based on repeatedly observed associations with lower lung cancer risk—was found to significantly *increase* that risk instead.

This is a textbook case of the harm that can result when health policies or public recommendations are based on association rather than causation. The findings from CARET (and the similar ATBC trial) ultimately led the U.S. Preventive Services Task Force in 2022 to issue a formal recommendation against beta-carotene supplementation for smokers and the general public—specifically citing these trials as the basis.

*Other experts note that **regulations and policies are not meant to require proof of causality** – the Clean Air Act, for instance, says standards “allowing an adequate margin of safety ... are requisite to protect the public health”.*

TC: I have never suggested that regulations or policies should “require proof of causality.” What I *have* consistently argued is that when scientific studies are used to justify regulatory actions, they should meet basic standards of valid causal inference. That includes accounting for well-known and easily measured confounders, clearly defining exposures and outcomes, and testing assumptions against data. The studies underlying EPA’s recent regulatory policies fail these basic requirements—not because causality is inherently unprovable, but because poor methodology was used. That is a preventable failure.

Cox, however, has critiqued proposals to [strengthen controls](#) on [pollution](#) on the grounds of imperfectly demonstrated causality. It is the sort of logic that Cox’s new AI tool could automate, which could benefit corporate interests, said Vardi of the Energy and Policy Institute.

“Instead of having scientists-for-hire do that denial work, which advances their economic interests, the industry is funding efforts to outsource it to a machine in order to give it an image of unbiased neutrality,” Vardi said.

Cox, for his part, said: “A scientist-for-hire could use such an AI system to check whether the conclusions affirmed or denied in a scientific paper follow from the data and analyses presented, but my AI systems don’t concern themselves with affirming or denying any specific positions or conclusions. That is left for people to do.”

Though Cox claims his AI tool is neutral, Finkel said his early ChatGPT conversations shed light on its potential dangers.

TC: Once again, *The Guardian* solicits commentary from individuals with no actual knowledge of the tools I developed. My early conversations with ChatGPT occurred well before those tools existed. They

do not illustrate any risk associated with my tools—they illustrate the limitations of large language models like ChatGPT at that time.

In fact, my AI reviewer was designed specifically to avoid those limitations. Unlike generative models, it does not create text or reflect bias from training data. It applies explicit, transparent rules of logic and evidence evaluation to assess whether conclusions in a paper follow from the data and methods presented. It does not affirm or deny any position—it evaluates reasoning. That is the opposite of what Vardi and Finkel claim.

The notion that my AI tool enables “denial work” is baseless and defamatory. It was built to promote better science, not to serve anyone’s economic or ideological agenda. If it creates discomfort for those accustomed to drawing conclusions without adequate evidence, that speaks to its value—not its danger.

“He was torturing the machine only along one set of preferences, which is: ‘Can I force you to admit that we are being too protective?’” Finkel said. “That’s not science.”

TC: This comment entirely misses and misrepresents the point of my exploratory dialogues with ChatGPT. The purpose was not to “force” the model to admit any particular position, but to probe whether its answers were logically consistent, evidence-based, and resistant to superficial or biased reasoning. These dialogues were a form of stress-testing—intended to reveal how malleable, reliable, and transparent the model’s reasoning was.

The ultimate aim was to identify the kinds of flaws and vulnerabilities that needed to be corrected to build more trustworthy AI systems. This is what responsible technology development looks like. Dr. Finkel may not recognize this as “science,” but it is precisely the kind of rigorous, iterative refinement that modern science, technology, and engineering depend on.

If questioning unsound reasoning is troubling to some, that’s all the more reason to build tools that do it better.

Cox said his conversations with ChatGPT aimed to uncover hidden uncertainties.

TC: No, I did not say that. My conversations with ChatGPT were not aimed at “uncovering hidden uncertainties.” They were aimed at evaluating the internal consistency, reasoning quality, validity, and evidentiary grounding of the model’s responses—especially in areas where it appeared to confuse association with causation. The goal was to explore how such models reason and to identify ways to build more trustworthy, logically sound AI systems—not to speculate about “hidden uncertainties.”

This is yet another misrepresentation by *The Guardian*—substituting a misleading paraphrase for what I actually said and did.

But a different chatbot could be trained to identify instances in which government is “under-regulating”, Finkel said.

TC: This comment has nothing to do with the actual AI technology I developed. My AI is not a chatbot—so the phrase “a different chatbot” is misleading. It does not use training, so Dr. Finkel’s speculations

about how a hypothetical “different chatbot” could be “trained” for various purposes are wholly irrelevant to my AI.

What’s more troubling is that *The Guardian*, yet again, solicits and publishes speculative commentary from someone with no understanding of the tool’s technology, function, or application—while disregarding the perspectives of the hundreds of actual users who could have provided informed, firsthand insight. **That’s not responsible journalism. It’s narrative-driven misrepresentation** that deprives readers of accurate, informed reporting about the technology in question.

17. How The Guardian Misrepresents Scientific Rigor as a Threat, and False Choices as Science

On an academic level, Cox’s interest in certainty might seem reasonable, but in the real world, it is dangerous to apply his standard of causality, said Finkel.

TC: My so-called “standard[s] of causality” are not personal inventions—they reflect core principles of modern scientific methodology, as advocated by leading scientists across disciplines, including numerous Nobel Laureates and Turing Award winners. These standards include using well-defined causal questions, controlling for confounding, testing assumptions against data, and relying on transparent, reproducible reasoning.

Far from being dangerous, applying these principles is essential for protecting public health. What’s dangerous is *not* applying them—continuing instead to base major regulatory decisions on untested assumptions, preferences, and selectively interpreted associations. That approach has already led to decades of ineffective or misdirected policies with little demonstrated public health benefit (e.g., Lyons et al., 2023).

The irony is that **it’s not scientific rigor that delays effective health protections—it’s the refusal to use it.** When regulations are built on weak science, they fail to achieve their intended benefits, waste resources, and erode public trust. The solution is not to lower the bar for scientific evidence, but to raise our standards—so we can act sooner *and* more effectively, based on real evidence rather than belief.

“For almost anything that we now know is harmful, there was a period in time when we didn’t know that,” he said. If Cox’s standards are taken seriously, he added, we could see “generations, decades of misery while we wait for him to be satisfied”.

TC: Dr. Finkel’s suggestion that we must choose between taking timely action and applying sound science standards is a false dichotomy and deeply misleading. In fact, applying rigorous causal inference methods—what Finkel calls “my standards”—is precisely how we *accelerate* the discovery of which interventions actually reduce harm. These standards are not arbitrary personal preferences; they are widely endorsed by leading scientists and methodologists across fields and they have underpinned many of the greatest advances in medicine, public health, and safety.

What delays progress is not scientific rigor, but its absence. When regulations and policies are built on untested assumptions or flawed associations, they often fail to deliver benefits—and may even do harm. Examples abound: MTBE was added to gasoline to improve air quality, only to contaminate groundwater without clear health gains; hormone replacement therapy was promoted to prevent heart disease until randomized trials showed increased risks; low-fat dietary guidelines contributed to obesity by encouraging high sugar intake; and beta-carotene supplements, once proposed to reduce lung cancer in smokers, were found to increase it instead. **These are the kinds of failures that waste decades—not my insistence on sound methods.** The real danger to public health lies in accepting claims without adequate scrutiny and clinging to failed interventions for political or ideological reasons.

The public deserves better than policies based on wishful thinking, “fundamental beliefs” and “preferences” of selected “experts” who refuse to apply real scientific methods. They deserve policies grounded in evidence that meets the same minimal standards of logic, transparency, and testability that we expect in every other serious domain of science—standards that, when followed, have consistently delivered enormous benefits.

False and Defamatory Claims in *Futurism* Article at <https://futurism.com/pollution-ai-green-washing>

Louis Anthony Cox, Jr., Ph.D.

tony.coxjr@ucdenver.edu, tcoxdenver@aol.com

July 23, 2025

Building on the false reporting in *The Guardian*, other media outlets have embellished *The Guardian's* story to add their own layers of fiction and defamation. This document outlines verifiably false and defamatory statements made about Dr. Louis Anthony (Tony) Cox, Jr and his AI technology in the *Futurism* article titled "Pollution Industry Using AI to Make the Case for More Pollution" by Joe Wilkins dated July 4, 2025, (<https://futurism.com/pollution-ai-green-washing>). It provides factual rebuttals and clarifying commentary.

Claim in Futurism Article	Factual Rebuttal	Commentary
"Louis Anthony 'Tony' Cox Jr is a petrochemical industry mouthpiece"	Dr. Cox is a tenured university professor and independent researcher, elected to the U.S. National Academies of Sciences, Engineering, and Medicine, and a Fellow of multiple scientific societies. He speaks for evidence and logic, not for any industry.	This is an unsupported and defamatory slur. No evidence is provided to justify this label.
"... and former Trump advisor"	Dr. Cox has never held a political advisory role in the Trump administration—or any administration.	This falsehood appears to have originated in a <i>Mother Jones</i> article and has been uncritically repeated without fact-checking.
"... who's now hard at work developing a large language model (LLM) program"	Dr. Cox has never developed, nor attempted to develop, an LLM.	This is a fundamental technical error that could have been avoided with even minimal due diligence.
"... to censor tree-hugger 'propaganda' from epidemiological research, according to the Guardian"	Neither Dr. Cox nor The Guardian article makes or attributes any such statement. No tool developed by Dr. Cox has this intent or capability, nor does the Guardian quote or imply this language.	This appears to be a fabrication based on misrepresentation of a quote taken out of context.

Claim in Futurism Article	Factual Rebuttal	Commentary
“Cox — who once claimed there’s no connection between air pollution and respiratory problems”	Dr. Cox has published extensively on how air pollutants cause respiratory diseases. He has never claimed otherwise.	A defamatory and demonstrably false claim. Dr. Cox has written award-winning scientific papers on biological causal mechanisms by which air pollutants cause lung diseases.
“Cox’s relationship with the American Petroleum Institute... is so cozy that he’s even let the petrochemical lobbying group ‘edit’ his research.”	Dr. Cox once accepted a minor proofreading suggestion (e.g., a punctuation change) from a colleague affiliated with the API. He openly acknowledged this proofreading help. No one has ever “edited” his research. All conclusions remain his, and the paper was peer-reviewed independently.	Equating minor proofreading with editorial control is misleading and defamatory.
“He began work on his LLM back in 2023”	Dr. Cox has not worked on any LLM. No such project exists.	Repeats a falsehood already addressed above.
“Dissatisfied with the length of time...”	Dr. Cox’s dialogues with ChatGPT were not aimed at “convincing” it of anything and satisfaction with length of time was not a consideration. This experiment was aimed at testing ChatGPT’s internal consistency and reliability under logical scrutiny.	The quoted motivation is fictional. The actual goal was to stress-test AI reasoning and accuracy and trustworthiness of stated conclusions and rationales.
“...to convince the notoriously agreeable ChatGPT that PM2.5s don’t cause lung cancer — a career-long hobby horse of his”	The dialogue had nothing to do with lung cancer. Dr. Cox’s research was not about “convincing ChatGPT,” but about probing how accurately it represented uncertainty in its reasoning. Virtually none of Dr. Cox’s career has addressed PM2.5 and lung cancer specifically. It is not a “career-long hobby horse.”	This is a cartoonish and defamatory mischaracterization.
“Cox turned to his buddies in the petrochemical industry to pitch an app that could be pre-trained...”	The tool is not an app, it is not an LLM, and it cannot be trained, pre-trained, or fine-tuned. It uses deterministic, interpretable logic—not machine learning.	Another fundamental factual error that misrepresents both the people involved and the technical properties of the system.

Claim in Futurism Article	Factual Rebuttal	Commentary
“That app... will tell users what Cox and his friends of industry want them to hear.”	The tool cannot be programmed with any desired viewpoint. It does not output preferred messages. It checks for logical coherence, not content.	Entirely fictional and defamatory.
“His painfully obvious track record, including studies like a chemical industry-sponsored report on the supposed safety of PFOAs...”	(1) The cited paper was not industry-sponsored. (2) It concluded that no safe level of exposure could be determined. (3) Again, there is no “his LLM.”	This is either grossly negligent reporting or deliberate distortion.
“Experts say Cox’s ‘AI-washing’ is a perfidious form of mass-market propaganda...”	No “experts” cited have reviewed or used the tool, nor is any such quote attributable to a credible source. This is defamatory speculation, not reporting.	Inflammatory fiction, not journalism.
“AI language models are not programmed, but built and trained,” Vardi continued...	The tool in question is not an LLM and is not built or trained in the way Vardi describes. His statement is irrelevant to the technology being discussed.	The article critiques a technology Dr. Cox did not build.
“Cox’s insistence on AI-washing public health on behalf of his corporate masters...”	Dr. Cox is a tenured university professor and independent researcher with no “corporate masters”. This claim is malicious and false.	Malicious rhetoric substituting for facts.
“If, or perhaps when, Cox’s big business bot comes online...”	The tool has already been online since January 2025, is publicly accessible, has been used by hundreds of users, and is entirely opt-in. There is no secret “seeping in.”	Again, this framing is fiction.

These demonstrably false statements have caused professional and reputational harm, and their repetition in derivative media coverage risks further misleading the public.

False and Defamatory Claims in *Yahoo* Article at <https://www.yahoo.com/news/controversial-researcher-accused-using-ai-070000123.html>

Louis Anthony Cox, Jr., Ph.D.
tony.coxjr@ucdenver.edu, tcoxdenver@aol.com

July 23, 2025

Yahoo was particularly creative in fabricating additional misinformation by citing *The Guardian's* already inaccurate story while substituting Yahoo's own inaccuracies and slurs to further sensationalize their narrative.

The following table documents key instances where Yahoo News escalated or fabricated claims beyond what was originally reported in *The Guardian*, along with factual clarifications.

Claim / Topic	What Yahoo Says	What <i>The Guardian</i> Said	Truth / Tony Cox Clarification
Headline: "AI used to promote industry interests; can't be trusted as objective"	"Controversial researcher accused of using AI to 'promote industry interests': '[It] cannot be trusted as a source of objectivity"	<ul style="list-style-type: none"> • "Some critics have mischaracterized my work as an attempt to ... promote industry interests. That is not true." • "The ACC cannot be trusted as a source of objectivity..." (Frey) 	Blatantly false and misleading. The Guardian did not accuse me of promoting industry interests—it quoted me denying such claims. The "cannot be trusted" quote refers to the American Chemistry Council (ACC), not me or my AI tool. Yahoo spliced these into a headline falsely attributing both criticisms to me.
AI tool "falsely debunks pollution science"	"A new AI tool is being used to falsely debunk pollution science."	Implies concern: critics fear AI may " amplify doubt, " not that it's being used falsely.	False and defamatory. My AI tool does not "debunk" anything (falsely or otherwise). It evaluates internal reasoning about causality and is entirely content-neutral. Yahoo falsely states that misuse is already occurring.
Misinformation campaign & unethical research	" Cox's misinformation campaign ... lack of ethical research ... dispute of evidence-backed facts. "	Not mentioned at all in <i>The Guardian</i> .	Entirely fabricated and defamatory. No such "campaign" exists. No expert has credibly accused me of unethical research or denial of valid evidence. Yahoo manufactured these claims.
Inflating own claims using AI	"Using artificial intelligence to inflate his own claims..."	Said AI was designed to "amplify his perspective," with concerns attributed to critics.	False and defamatory. I have not used AI to inflate anything. Yahoo rephrases vague concern into a factual accusation of intellectual dishonesty.

“Attack on science-backed research”	“This new AI initiative is yet another attack on science-backed research...”	Guardian says critics fear industry AI could influence framing—not that it “attacks” research.	Inflammatory and misleading. The tool is neutral and does not attack anything. Yahoo rewords criticism into an allegation of harmful intent.
Spreading public doubt as a defining feature	“Known for seeding public doubt...”	“Forged a career sowing doubt...”	False and damaging. My career is defined by contributions to quantitative risk analysis and causal reasoning, not by spreading doubt. This framing is derogatory and unsupported.
Promoting the dirty energy industry	“An attempt to promote the dirty energy industry.”	Guardian did not allege promotion—only reported industry funding.	Defamatory exaggeration. I do not promote any industry. I pursue scientific rigor and transparency. Yahoo escalates disclosed funding into a claim of intent to promote.
Violated research objectivity	“Violated research objectivity by allowing a dirty energy lobbying group ... to copy edit his findings.”	Guardian reported the fact of a minor edit and my clarification that it was small and disclosed.	False and misleading. Yahoo falsely asserts wrongdoing. The copy-edit was minor (e.g., punctuation), acknowledged, and transparently disclosed. There was no violation of objectivity.
Promoting misinformation	“Cox’s claims are concerning because they not only spread misinformation...”	Guardian never uses the word “misinformation” to describe my work.	Baseless and defamatory. My publications are peer-reviewed and grounded in reproducible methods. Disagreeing with a causal claim is not “misinformation.” Yahoo introduces this defamatory label without justification.