UNIVERSITY OF OTTAWA

DOCTORAL THESIS

Essays in Environmental Economics

Author:
Ghina ABDUL BAKI

Supervisor:
Dr. Myra YAZBECK

Co-Supervisor:
Dr. Lynda KHALAF

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy

in the

Department of Economics
Faculty of Social Sciences

October 28, 2025

Declaration of Authorship

I, Ghina ABDUL BAKI, declare that this thesis titled, "Essays in Environmental Economics" and the work presented in it are my own. The first chapter of this thesis is solely done by me. The second chapter is a joint work with Walid Marrouch. One subsection of this paper has been published in Economics Letters (Abdul Baki et al., 2024). The third chapter is a joint work with Myra Yazbeck. In Chapters 2 and 3, all authors, including myself, contributed equally.
I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: Ghina Abdul Baki

Date: 10th of August, 2025

«ونهار نهرب منه عند النوم توقظنا في الصباح مشاغله ومشاكله وغمومه وهمومه، لنستعين عليها بنور نهار آخر، وهكذا نصل الفكر بالفكر والنية بالنية والأمل بالأمل والنفس بالنفس، والحركة بالحركة واليقظة بالمنام.»

— ميخائيل نعيمة

"And a day we escape from in sleep, only to be awakened in the morning by its concerns, troubles, sorrows, and worries, so we confront them with the light of another day. And thus we link thought to thought, intention to intention, hope to hope, soul to soul, movement to movement, and wakefulness to sleep."

—Mikhail Naimy

UNIVERSITY OF OTTAWA

Abstract

Department of Economics
Faculty of Social Sciences

Doctor of Philosophy

Essays in Environmental Economics

by Ghina ABDUL BAKI

This dissertation examines how environmental externalities shape financial markets, regulatory policies, and public well-being, using a combination of empirical analysis and theoretical modeling. The **first chapter** investigates whether environmental concerns help explain Bitcoin's returns and volatility. Using innovative textual data from Tweets and Google Trends, I construct three proxies for climate risk, each capturing a distinct dimension: aversion, uncertainty, and attention. The findings reveal that climate risk uncertainty (aversion) improves the in-sample predictability of returns (volatility), while climate risk attention enhances the out-of-sample predictive ability of volatility. The **second chapter** develops a partial equilibrium model to determine optimal emission taxes in the presence of two-sided environmental externalities, two-sided abatement, and one-sided market power. In this case, the taxes imposed on polluting producers and consumers fall below their Pigouvian levels without achieving the first-best. However, if only one side abates, the first-best is achieved by taxing the (un)abating polluters at (below) their respective Pigouvian level. The **third chapter** employs a difference-in-differences approach to assess the causal effects of prenatal fluoride exposure on birth outcomes in the United States. Results show a significant reduction in birth weight and an increased risk of central nervous system anomalies, especially during the first trimester and among White, married, and more educated mothers, which highlights socioeconomic disparities in vulnerability. Additionally, the effects vary by exposure type, with distinct impacts observed for low-dose cumulative versus high-dose acute exposure.

Acknowledgements

First and foremost, all praise is to Allah, the Almighty, the Entirely Merciful, and the Especially Merciful, for His countless blessings in my life.

I am deeply thankful to:

My supervisors, Myra Yazbeck and Lynda Khalaf, for their unwavering support, boundless patience, and rigorous training. Their belief in me, steady as their commitment to excellence, has always propelled me forward.

Myra Yazbeck, who, to put it simply, has seamlessly embodied the roles of supervisor, colleague, friend, and sister, each in its time and place. She taught me to always aim for the sky, even when the road felt cloudy and full of thunderstorms. Her steadfast presence and expertise have been a cornerstone, guiding me through this journey and leaving an enduring impact well beyond the bounds of this PhD.

Lynda Khalaf, a role model, an inspiration, and a mentor to me, from the very beginning of the admission stage. Her expertise, generosity in sharing wisdom, and constant guidance in recalibrating my path have been invaluable. She consistently pushed me beyond my comfort zone, and while this sometimes felt daunting, it ultimately unlocked the door to real growth, both as a scholar and as a person.

My committee members, Myra Mohnen, Louis-Philippe Beland, and Paul Makdissi, for their valuable insights and constructive criticism; and Jie He, for agreeing to serve as my external examiner. A special thanks to Walid Marrouch, my MA supervisor and coauthor, for his continued support and encouragement from the very beginning.

Catherine Deri Armstrong, the Chair of the Economics Department, for her support on both professional and personal levels.

The late Professor Serge Nadeau, whose generous donation to the Economics Department allowed me to devote more time to my research (Thesis Completion Grant).

My colleagues, Aline Zayat, Motasem Qaddoura, Ibrahim Abuallail, Wenjie Tian, and Pu Sun, whose presence made my journey far more meaningful and enjoyable.

My friend, Karen Kmeid, whose presence brought light in the midst of challenges and will forever remain one of the greatest gifts of this journey.

My sister, Iman, and her family, whose support, especially when I moved to Canada, made this journey much smoother and allowed me to integrate better.

My parents, Nadia and Walid, and my brother, Tarik. I am forever grateful for your unconditional love and prayers, and for always believing in me, no matter the circumstances. Your trust, sacrifices, and presence have always been the spark of my perseverance. Thank you immensely for all that is said and what words cannot express.

Dedicated to my parents and my brother.

Contents

List of Figures

List of Tables

General Introduction

Market activity and public policies are widely acknowledged to generate welfare effects that extend far beyond the scope of conventional economic indicators. In particular, they often produce environmental and health externalities that can shape financial markets, influence the design of regulatory policies, and affect public well-being, both physiological and psychological. Understanding how these externalities interact with financial markets, public policy, and well-being is a central concern in economic research. This dissertation addresses this concern from three complementary angles. Specifically, it combines empirical evidence and theoretical modeling to trace how these externalities (i) transmit information through financial markets, (ii) alter the design of optimal emissions taxes, and (iii) manifest in unintended early-life health outcomes.

The effects of environmental signals on financial markets have recently drawn increasing attention, especially when the asset in question is environmentally detrimental. This is notably the case for Bitcoin, where mining and transaction activities generate substantial environmental damage due to their intensive energy use. Accordingly, general environmental concerns may influence Bitcoin's supply and demand, yet their role in price formation remains underexplored.

The first chapter, titled "**Climate Risk and Cryptocurrencies: Evidence from Bitcoin Market Reactions to Environmental Concerns**," addresses this gap by examining whether public environmental concerns help predict Bitcoin's returns and volatility, both in and out-of-sample. Assembling a unique dataset comprising social media-based textual data (Twitter) and Google Trends, I construct two novel indices that capture distinct dimensions of climate risk: the Environmental Concern Measure (proxy for climate risk aversion) and the Environmental Consciousness Measure (proxy for climate risk uncertainty). Together, these measures serve as a proxy for climate risk attention. The results reveal that climate risk uncertainty enhances the in-sample predictability of returns, while aversion improves the in-sample prediction of volatility (realization effect). Moreover, climate risk attention improves the out-of-sample predictive ability of the volatility over short horizons (short-term anticipation effect), especially during periods of environmental stress or asymmetric risk. These findings underscore the importance of textual data in capturing elusive information often overlooked by traditional financial data.

Bitcoin exemplifies a final product whose production (mining) and consumption (transaction) both cause environmental harm; it imposes a two-sided environmental externality. Similar dynamics also exist in sectors like agriculture, pharmaceuticals, and energy, among others. Nonetheless, the literature on pollution control, particularly emissions taxes, largely overlooks consumer-side pollution, abatement efforts, and market power, all in the presence of production-sided pollution.

To bridge this gap, the second chapter, "**Emissions Taxes and Market Power: Linked vs. Unlinked Market Failures**," develops a partial equilibrium model to determine optimal emissions taxes under two-sided externalities, two-sided abatement, and one-sided market power. This project shows that when both producers and consumers abate, the optimal taxes fall below Pigouvian levels and fail to achieve first-best outcomes due to the linkage between pollution and market power. However, if one side abates, the market failures are unlinked; thus, the optimal tax scheme can restore the first-best allocation by taxing the abating polluter at the Pigouvian level and the non-abating polluter below it. If neither side abates, multiple optimal taxes exist and achieve the first-best. These findings are orthogonal to the source of market power. Altogether, this framework yields new insights into the theory of second-best and offers a unifying framework, extending from horizontal to vertical markets, that nests all prominent results in the emission taxation literature under the pure market structure as special cases.

While pollution is often viewed as a byproduct of industrial processes (production or consumption), it can also arise, paradoxically, from public health interventions. Water fluoridation, the addition of fluoride to public water supplies to reduce dental caries, is one such case. Although initially intended to improve oral health, recent concerns have emerged about potential unintended effects even at recommended safe doses. Nevertheless, the extent to which fluoride exposure affects early life outcomes remains an open question.

The third chapter, "**The Unintended Health Impacts of Water Fluoridation: Evidence from Birth Outcomes**," contributes to the very nascent economic literature on fluoride exposure by investigating the impact of prenatal fluoride exposure on birth outcomes in the United States (US). A difference-in-difference (DiD) approach is adopted that leverages variation in county-level adoption of water fluoridation between 1982 and 1993. Results mainly indicate a modest reduction in birth weight (0.34%) and a nearly 50% increase in the risk of central nervous system anomalies. Effects are most pronounced during the first trimester and among White, married, and more educated mothers, highlighting socioeconomic disparities in vulnerability. The study also distinguishes between low-dose cumulative and high-dose acute exposure, which are shown, in this chapter, to trigger distinct effects. Back-of-the-envelope calculations suggest that these unintentional health costs may overturn the intervention's net social benefit, warranting a reassessment of its policy implications.

Across the three essays, this dissertation applies diverse empirical and theoretical methods to investigate timely issues at the intersection of environment, finance, and health. Combining textual analysis, time-series forecasting, welfare theory, and quasi-experimental designs, it adopts an interdisciplinary approach to economic research. Together, the chapters deepen our understanding of how environmental and health externalities permeate financial markets, shape policy, and affect public welfare, while demonstrating the value of integrating traditional and nontraditional data sources.

Chapter 1

Climate Risk and Cryptocurrencies: Evidence from Bitcoin Market Reactions to Environmental Concerns

1.1 Introduction

Cryptocurrencies have gained widespread attention among public and private sectors due to their diverse roles beyond finance.[1] Albeit being the oldest and one of the most traded (Easley, O'Hara, and Basu, 2019), Bitcoin is classified as a brown cryptocurrency. Indeed, the energy demand of Bitcoin mining and transaction rivals that of the top 30 electricity-consuming countries, thus emitting around 71 million metric tons of carbon dioxide (CO_2) annually (Stoll, Klaaßen, and Gallersdörfer, 2019).[2] This emission level may accelerate global warming beyond 2 degrees Celsius within 16 years only (Dittmar and Praktiknjo, 2019). Recent findings by Papp, Almond, and Zhang (2023) indicate that a $1 increase in Bitcoin's price leads to a $3-6 rise in the external costs related to CO_2 emissions. Despite such evidence, whether Bitcoin's price is affected by climate risk remains an open question. Accounting for in and out-of-sample analyses, this paper introduces social-media based measures of environmental concerns and finds that these are important in explaining the returns on Bitcoin and their volatility.

1. Besides their financial role, cryptocurrencies promote innovation, create job opportunities in the blockchain field, and generate tax revenues. See https://www.pelicoin.com/blog/what-is-the-economic-impact-of-cryptocurrency
https://theconversation.com/cryptocurrency-has-an-impact-on-economies-thats-why-some-are-afraid-of-it-and-some-welcome-it-175911

2. This is based on the Cambridge Bitcoin Electricity Consumption Index, calculated by Cambridge Center for Alternative Finance, University of Cambridge, available at https://ccaf.io/cbeci/index
In the United States (US), Goodkind, Jones, and Berrens (2020) estimate Bitcoin's environmental damages and associated health impacts to be nearly half of its dollar worth.

The paper has three contributions. First, I provide evidence of a link between climate risk and the Bitcoin market by defining two environmental pricing factors, denoted by "concerns" and "consciousness", and evaluating their contribution above and beyond what a base model (GARCH-in-mean) explains.[3] The key distinction between "concerns" and "consciousness" is that environmental concerns constitute a proxy for climate risk aversion, whereas environmental consciousness comprises a proxy for climate risk uncertainty. This distinction affects the empirical findings and contributes to the emerging literature on climate risk hedging (Krueger, Sautner, and Starks, 2020), management (Engle et al., 2020), and premia (Ardia et al., 2022).

Second, I differentiate between public "consciousness" and "concerns" as two distinct dimensions of environmental attention (EA), unlike existing studies that use "attention" and "consciousness" interchangeably (e.g., El Ouadghiri et al., 2021). In doing so, I acknowledge the significance of the widely used Google Trends data in capturing public "consciousness". Thus, I use Google Trends for the Environmental Consciousness Measure (denoted thereafter as the ESM) as a proxy for climate risk uncertainty. At the same time, I argue against the effectiveness of relying on Google Trends to effectively capture public concerns. To address this limitation, I gather unique alternative data, that is data from text-based social media posts. Particularly, I perform a fine-grained aspect-based sentiment (FiGAS) analysis of Tweets to develop a new metric denoted as the Environmental Concern Measure (ECM), which serves as a proxy for climate risk aversion.[4] Therefore, the ESM reflects information-seeking behaviors, while the ECM reflects information-sharing behaviors. Together, public consciousness and concern constitute a broader definition of EA that serves as a proxy for climate risk attention. More importantly, the ECM is distinct along the following dimensions. (1) It gauges social interaction in an organic setting. In fact, climate risk aversion may propagate on social media platforms, leveraging their mechanism for social buffering. (2) It is straightforward to replicate, beyond the sampling period or (eventually) the microblogging and social networking service itself. (3) It is a high-frequency and granular series, which plays a critical role in disentangling confounding factors through discerning finer nuances and variations.

Third, I uncover an asymmetric effect in the considered linkages between text data

3. Generalized AutoRegressive Conditional Heteroskedasticity (GARCH)-in-mean

4. The "fine-grained [aspect implies] that words are assigned a polarity score that ranges between $\pm$ [5] based on a dictionary". Whereas, "aspect-based" indicates that tweets related to a certain concept (i.e., environment) are identified, and sentiment is calculated only on that text (Barbaglia, Consoli, and Manzan, 2023).

and financial stability. Indeed, environmental concerns are likely to affect Bitcoin's supply and demand, which, in turn, can impact the energy cost of mining and transactions. Hence, it is not clear whether a directional effect is to be expected in principle. This said, environmental concerns are known to publicly flare up around irregular or disruptive events. Specifically, Engle et al. (2020) underscore that "climate change primarily rises to the media's attention when there is a cause for concern". This asymmetric feature, which is inherent to news reactions more broadly, has a key implication on the directional forecast accuracy of returns, which I formally take into consideration. Moreover, I shed light on the relevance of text data amidst asymmetric risk by recognizing that investors may attribute distinct weights to gains and losses. Precisely, I incorporate a nuanced perspective by considering an asymmetric loss function.

The results can be summarized as follows. Firstly, climate risk uncertainty contributes to predicting expected returns on Bitcoin in-sample; however, climate risk aversion improves the predictive ability of return volatility in-sample, which I define as realization effect. Secondly, climate risk attention yields short-run anticipation effect, described as improving the out-of-sample predictive ability of volatility, yet on a daily or weekly basis. Thirdly, the forward-looking feature of climate risk attention is state-contingent, in the sense that it manifests asymmetrically, contingent on the presence of environmental turmoil or asymmetric risk situations, but only in the short-run (day to week). The forward-looking information of text data permits assessing past and current financial performance and predicting future price (return and volatility). Indeed, this asymmetric effect is compatible with Timmermann (2008)'s argument regarding the challenges associated with "identifying intermittent predictive components" in returns, particularly in short-time horizons. Even though "the evidence of return predictability at short horizons such as one day is generally very weak", the out-of-sample asymmetric effect survives up to one day-week in my present work. Thus, the social media-based text data presenting climate risk concerns, which resurge "intermittently", are capturing fleeting elusive predictive factors in the Bitcoin market.

This analysis contributes to three strands of the literature. First, the identification of asset pricing factors continues to pose challenges (Fama and French, 2015; Fama, Hansen, and Shiller, 2013). There is extensive literature aiming to identify cryptocurrencies pricing factors. Within this strand of the literature, a predominant focus centers on financial predictors, as exemplified by Dyhrberg (2016), Gronwald (2019), and Okorie and Lin (2020). Formally, Liu, Tsyvinski, and Wu (2022) identifies three factors: cryptocurrency market, size, and momentum. Although it is pertinent, the predictive ability of a climate risk pricing factor is still uninvestigated, prompting discussions regarding the robustness of financial predictors. While the work of Liu, Tsyvinski, and Wu (2022) is a cross-sectional analysis of returns, which is clearly distinct from the time-series analysis I perform, my present paper still stands apart by introducing climate attention (concerns vs consciousness) as a pricing factor, revealing superior robustness when compared to conventional financial predictors. The complexities herein resonate with the findings elucidated by Harvey and Liu (2021), demonstrating that some asset market factors may become nullified or, in more unfavorable scenarios, undergo a change in direction when integrated with other factors.

Second, against this background, there is nascent literature analyzing and debating the bidirectional effects of climate change on financial markets (see e.g., Bauer and Rudebusch, 2021; Bolton and Kacperczyk (2021); Hong, Li, and Xu, 2019). Notably, two environmental metrics have garnered prominence: (i) Google Trends and (ii) news articles.

Regarding Google Trends (i), the preeminent work in this context is illustrated by El Ouadghiri et al. (2021), introducing Google Trends as a measure of environmental attention. Nonetheless, attention has not been recognized to manifest in two forms: "concerns" and "consciousness". In the present paper, I explicitly establish this distinction, which, in turn, allows me to control for social interaction when examining the role of climate risk. Indeed, learning through social interaction increases volatility in asset markets (Pedersen, 2022).

As for news articles (ii), considerable research has leveraged the attributes of textual data from climate change news, aiming to capture public attention and investor sentiment regarding time-varying climate risk (see e.g., Tetlock, 2007; Manela and Moreira, 2017; Giglio et al., 2021). In fact, Pástor, Stambaugh, and Taylor (2022) and Engle et al. (2020) introduce monthly media-based, yet noninteractive, climate measures. However, a focal point has not been addressed in their analyses: due to the rapid responses of financial markets, short measurement horizons (daily) are needed to mitigate confounding factors associated with lower frequencies (weekly or monthly horizons) while striking "a balance between timeliness and feasibility" (Ardia et al., 2022). In my present work, the proposed (daily) ECM captures real-time available concerns, paralleling the fluctuations in Bitcoin's price, and can even be available at the tick level without compromising its timeliness and feasibility. This flexibility is a distinctive feature specific to social-media based data given the publication time and length of news articles.

Third, more broadly, an increasingly popular research direction refers to "text as data", in the sense that timely information from media sources is integrated with traditional data to enhance models or forecasts (Chiang, Hughen, and Sagi, 2015; Gentzkow, Kelly, and Taddy, 2019; Baker, Bloom, and Davis, 2016; Reboredo and Ugolini, 2018; Jegadeesh and Wu, 2013). Specifically, my work aligns closely with that of Ellingsen, Larsen, and Thorsrud (2022) and Kalamara et al. (2022). The authors notably highlight the importance of text data during periods of economic stress (i.e. recessions). This conclusion is further substantiated by Thorsrud (2020), who employs a daily business cycle measure. Nevertheless, the efficacy of text data amid periods of environmental stress (i.e. disruptive climate events) or asymmetric risk (i.e., when gains and losses are weighed differently) remains unexplored. In the present paper, I underscore the asymmetric role of text data in shaping expectations, particularly during scenarios of climate upheaval or asymmetry in risk by (1) considering the sign test of Pesaran and Timmermann (2009) and (2) utilizing an asymmetric loss function.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 presents the model and methodology. Section 4 illustrates the results, while section 5 presents the robustness check. Lastly, section 6 concludes.

1.2 Data

A daily dataset between 08/09/2015 and 11/14/2021 is constructed from various sources. This dataset comprises 2290 observations, with variables categorized into two main groups: environmental and financial.

1.2.1 Environmental Data

The environmental variables are retrieved from three sources: Twitter (currently known as X), Google, and Wikipedia.

1.2.1.1 Twitter

The primary data source is Twitter. Presently, the US boasts the highest number of Twitter users, totaling 77.75 million.[5] Twitter has recently garnered significant attention

5. Refer to https://en.wikipedia.org/wiki/Twitter and https://www.statista.com/statistics/242606/number-of-active-twitter-users-in-selected-countries/ for

in behavioral finance, notably for its role in disseminating financial news (see, e.g., Reboredo and Ugolini, 2018).

Twitter posts, known as tweets, take various formats: text, images, and videos. I assemble text-based tweets to construct the ECM, which measures the volume of tweets on the following queries: *pollution*, *climate change*, *global warming*, and *renewable energy*. Tweet volume access requires leveraging the Twitter Application Programming Interface (API). Following Twitter's recommendations, I use 'Postman' application to extract daily tweet volume data.

While the tweet volume captures public concerns, it fails to distinguish between positive and negative concerns. Accordingly, I conduct a fine-grained aspect-based sentiment (FiGAS) analysis to disentangle the impact attributed to public concern polarity. Fine-grained means that words in a tweet are assigned a sentiment score lying within a range determined by the specified dictionary (AFINN lexicon). Aspect-based indicates that the sentiment analysis is performed on tweets belonging to a specific theme (environment). The FiGAS approach is common in the emerging literature on textual analysis that mainly focuses on news articles (see e.g., Barbaglia, Consoli, and Manzan, 2023; Consoli, Barbaglia, and Manzan, 2022).

The FiGAS analysis is constrained by Twitter's restriction on the number of downloadable tweets. Given this limitation, I download 4200 tweets/day, equally divided among the following queries: *pollution*, *climate change*, and *global warming*. These queries are in line with El Ouadghiri et al. (2021) and have a tweet volume exceeding that of *renewable energy*, as shown in Figure 1.1. Next, I construct the sentiment index associated with each of these tweets using a dictionary-based approach, as in Silge and Robinson (2017), Kalamara et al. (2022), and Gentzkow, Kelly, and Taddy (2019).

I use the 'tidytext' package of 'R' software to construct the (dictionary-based) sentiment index. This package has four sentiment dictionaries: AFINN, BING, Loughran, and NRC. I use the AFINN dictionary, which assigns words a score ranging between -5 (very negative) to +5 (very positive). This lexicon provides more information than the BING and Loughran lexicons because the latter two dictionaries only classify words as positive or negative, without indicating how positive/negative words are. As for the NRC dictionary, it classifies words into 10 categories, such as positive, negative, anger, and surprise, among others. This assignment operates in a binary fashion, with no further sentiment score for each word. Consequently, assigning sentiment scores may be

further information. Additional details can be found at
https://developer.twitter.com/en/docs/tutorials/postman-getting-started.

tricky, especially when words exhibit characteristics that align with multiple categories.

The complexity of sentiment analysis is driven by its compositional and contextual attributes (Shapiro, Sudhof, and Wilson, 2022). For example, with negation, the sentiment of words can directly change when preceded by a single negating word. Accordingly, the reliability of dictionary-based approaches is compromised, as they typically fail to embed the effects of negation. For instance, the sentiment of "I am not happy" and "I am happy" is determined by "happy", which is a positive word. Therefore, the first sentence will be falsely labeled with a positive score. To overcome this limitation, I first identify negated words: words preceded by {not, no, never, without} (in line with Silge and Robinson, 2017; Abdul Baki et al., 2025) or by a word with suffix "n't". By taking "n't", the abbreviation of "not", into account, I can capture all negated words preceded by {haven't, wouldn't, aren't, isn't, etc.}. Next, I reverse the sign of the sentiment score associated with these negated words. For example, if "happy" scores +3, "not happy" would score -3. Finally, I compute the sentiment index of each tweet by summing the scores across all words within that tweet. Hence, the ECM transitions from the absolute volume, which captures public concerns, to the volume of positive/negative tweets, which reveals public concerns polarity. The two versions of the ECM serve as a proxy for climate risk aversion. The (positive/negative) tweet volume is logarithmically first-differenced to ensure its stationarity.

Figure 1.2 displays the polarity distributions of tweets mentioning *pollution*, *climate change*, and *global warming*. All distributions are left-skewed, suggesting a predominance of negative sentiment. This trend is most pronounced for *climate change*, followed by *pollution* and *global warming*.

Figure 1.3 displays the word clouds for these same tweets. Prominent negative terms include *risk*, *catastrophic*, *crisis*, *disaster*, *threat*, *death*, *kill*, *smog*, *cancer*, and *bad*. Conversely, the most frequently mentioned positive terms are *save*, *protect*, *clean*, *free*, *solution*, *support*, *hope*, *opportunity*, *benefits* and *justice*. Thus, the word clouds reveal that environmentally-themed tweets particularly gain attention when there is a cause for concern (e.g., disaster, catastrophic, smog). Additionally, these tweets seem to exhibit a forward-looking perspective, focusing on potential optimistic (e.g., solution, opportunity, hope) and pessimistic (e.g., die, cancer, kill) outcomes. These tweets also signal an aversion to climate risk, given the negative connotations of terms like *risk* and *threat*.

1.2.1.2 Google

I construct the ESM, a proxy for climate risk uncertainty, using the US Google Trends. Google is a search engine mainly used to gather, rather than share, information. Therefore, the Google Trends of environmental topics are indicative of public consciousness, rather than concerns, about climate risk.

The ESM represents the Google search volume index (GSVI) for the same queries: *pollution*, *global warming*, *climate change*, and *renewable energy*. Specifically, the GSVI for a certain keyword is a normalized value ranging between 0 and 100: a value of 0 signifies that the number of searches in the specified period falls below a certain threshold determined by Google, while a value of 100 indicates that the number of searches is maximal in that period. Thus, the GSVI is a relative measure.

The complexities surrounding the GSVI stem from its nature as a relative measure. One challenge arises from the dependence of the GSVI value on the specified period of interest, during which the frequency of searches for a certain query is compared with all Google searches. Another issue is the sensitivity of GSVI frequency to the specified time interval over which data is collected; longer time periods result in lower frequencies. To obtain a reliable daily GSVI between 2014 and 2021, I employ a rolling time window fixed at 5 months, which is the maximum length generating a daily measurement of the GSVI. Following the approach of El Ouadghiri et al. (2021), Hillert, Jacobs, and Müller (2014), and Aouadi, Arouri, and Roubaud (2018), I apply the logarithmic transformation of one plus the value of its GSVI of each query [i.e., $Ln(1 + GSVI)$], to ensure stationarity and smooth out outliers.

Thus, the ECM and ESM constitute two distinct dimensions of EA: information sharing (ECM) or information seeking (ESM). This broader definition of EA stands as a proxy for climate risk attention.

1.2.1.3 Wikipedia

I use Wikipedia to construct two dummy variables that could be correlated with the proxies: $UNCC_t$ and UN_t. The former takes the value of 1 for the days t that the United Nations Climate Change (UNCC) conference was held on, 0 otherwise. The latter equals 1 if day t is recognized by the United Nations (UN) as a special day under the theme environment/climate change, 0 otherwise. This variable captures the UN observances related to the environment and climate change, which are summarized in Table 1.1.

1.2.2 Financial Data

In line with the Bitcoin literature, the financial variables include: Bitcoin trading volume, EUR-USD exchange rate, and the US-time closing price (in US dollars) of Bitcoin, Ethereum, gold, WTI crude oil, and S&P500, DJSI, FTSE, FTSE4Good, NASDAQ, and VIX indices. Besides these variables, I include the daily tweet volume on *bitcoin*. This variable may capture announcements or news about Bitcoin, and it may assist in explaining the outliers in Bitcoin's rate of return, particularly since Twitter has become a major source of news (Broersma and Graham, 2013). To ensure their stationarity, I consider the natural logarithm (first) difference of these variables.[6]

1.3 Econometric Model and Methodology

The empirical strategy consists of in-sample and out-of-sample analyses, covering the periods 08/09/2015 - 09/27/2018 and 09/27/2018 - 11/14/2021, respectively. For the in-sample analysis, I estimate a GARCH(1,1)-M model by maximizing the likelihood (see e.g., Katsiampa, 2017; Tiwari, Kumar, and Pathak, 2019, Ardia, Bluteau, and Rüede, 2019; Bouoiyour, Selmi, et al., 2016; Dyhrberg, 2016). The resulting evidence informs whether environmental concerns influence the returns on Bitcoin and their volatility. Next, I perform an out-of-sample analysis to assess the extent to which environmental concerns improve the predictive ability of the returns and their volatility.

The cut-off point for the out-of-sample exercise is set to 09/27/2018 following the establishment of the Cambridge Bitcoin Electricity Consumption Index. The latter concretizes attention to the environmental externality in the Bitcoin market, and resulting sample sizes are large enough to permit estimation (1145 observations each). This is illustrated in Figure 1.4, which presents the volume of tweets jointly mentioning *bitcoin* and either *pollution*, *global warming*, or *climate change*. It is, therefore, crucial to unveil the role that EA plays out-of-sample.

1.3.1 In-sample Analysis

The in-sample analysis consists of two parts, each reflecting a key aspect of the ECM's construction. The first part considers the absolute volume of tweets, while the second part focuses on the sentiment conveyed.

6. Formally, I have conducted the Augmented Dickey-Fuller test for stationarity on all the transformed variables.

1.3.1.1 Volume of Tweets

The primary analysis focuses on the volume of tweets as a proxy for climate risk aversion. Consider the following GARCH(1,1)-M model, hereafter referred to as the Green Model (main specification):

$$BTC_t = \delta_0 + \alpha ECM_{t-1} + \beta ESM_{t-1} + aENV_{t-1} + bFIN_{t-1} + \theta h_t + \epsilon_t, \quad (1.1)$$

$$h_t = exp(\delta_1 + \gamma ECM_{t-1} + \omega ESM_{t-1}) + \phi \epsilon_{t-1} + \psi h_{t-1}, \quad (1.2)$$

with $\epsilon_t = \sqrt{h_t} z_t; z_t \overset{\text{iid}}{\sim} N(0,1); 0 \leq t < 1146; \phi \geq 0; \psi \geq 0$. δ_0 and δ_1 are constant terms.

BTC_t denotes Bitcoin's rate of return at day t, h_t represents the conditional variance of Bitcoin's rate of return. θ, in eq 1.1, is a measure of the risk premium. ϕ and ψ in eq 1.2 correspond to the ARCH and GARCH effects, respectively. ECM_{t-1} and ESM_{t-1} are 4-dimensional vectors representing the ECM (tweets) and ESM (Google Trends). ENV_{t-1} is a 6-dimensional vector corresponding to environmental incidents (dummy variables). FIN_{t-1} is a 12-dimensional vector of financial variables. The complete definition of all variables in this model is presented in Table 1.2.

The descriptive statistics are summarized in Table 1.3. Indeed, public attention to *climate change*, whether measured via the GSVI or the tweet volume, has the highest standard deviation (around 0.7). Moreover, among the various financial assets, Bitcoin and Ethereum are the most volatile, with a standard deviation equals to 0.04 and 0.08, respectively.

Importantly, the ECM, denoted by the tweet volume on the various queries, measures public environmental concerns, which are not necessarily limited to Bitcoin. This is illustrated in Figure 1.4 which shows the tweet volume on *bitcoin* and any of {*pollution, global warming, climate change*}. These tweets capture the environmental concerns specific to Bitcoin. Figures 1.1 and 1.4 show that public environmental concerns exceed Bitcoin-specific environmental concerns, especially in the sub-sample used for statistical inference.

1.3.1.2 Sentiment of Tweets

To explore the potential asymmetric (in-sample) impact of public environmental concerns, I account for the sentiment index in the model. This adjustment involves distinguishing the volume of tweets (i.e., environmental concerns) based on their polarity. Specifically, I estimate the following model, denoted hereafter as the Sentiment-Driven Green Model:

$$BTC_t = \delta_0 + \alpha^- ECM^-_{t-1} + \alpha^+ ECM^+_{t-1} + \tilde{\beta}\widetilde{ESM}_{t-1} + \tilde{a}\widetilde{ENV}_{t-1} + bFIN_{t-1} + \theta h_t + \epsilon_t \tag{1.3}$$

$$h_t = exp(\delta_1 + \gamma^- ECM^-_{t-1} + \gamma_1^+ ECM^+_{t-1} + \tilde{\omega}\widetilde{ESM}_{t-1}) + \phi\epsilon_{t-1} + \psi h_{t-1}. \tag{1.4}$$

ECM^-_{t-1} and ECM^+_{t-1} are 3-dimensional vectors corresponding to the negative and positive environmental concerns, respectively. $\widetilde{ESM}_{t-1}$ is now a 3-dimensional vector, $\widetilde{ENV}_{t-1}$ is a 2-dimensional vector.[7] FIN_{t-1} is as previously defined in the Green Model. Table 1.4 presents the complete definition of all variables comprising all vectors in eq 1.3 and 1.4. All parameters and assumptions are as previously outlined in the Green Model.

1.3.2 Out-of-sample Analysis

In the out-of-sample analysis, I obtain the one-day, one-week, and one-month-ahead forecasts of the return rate and its variance using a rolling window. Since the model requires a sufficiently large number of observations for estimation, the window size represents around 90% of the sample.

First, I obtain the return, $BTC^{green}_{t+h|t}$, and variance, $VAR^{green}_{t+h|t}$, forecasts using the main specification accounting for EA (i.e., the Green Model); $h \in \{1, 7, 30\}$ is the forecasting horizon (in days). Next, I omit all environmental variables from the Green Model and obtain the return, $BTC^{brown}_{t+h|t}$, and variance, $VAR^{brown}_{t+h|t}$, forecasts using the same

7. The query *renewable energy* is disregarded from the ECM^-, ECM^+, and $\widetilde{ESM}$ due to its minimal presence in the tweets and Google searches. The interaction terms are omitted from $\widetilde{ENV}$ to conserve degrees of freedom. All variables in the Green Model (GM) and the Sentiment-Driven Green Model (SDGM) are stationary, as per the Augmented Dickey-Fuller test for stationarity.

rolling window. Specifically, $BTC^{brown}_{t+h|t}$ and $VAR^{brown}_{t+h|t}$ are computed using the following GARCH(1,1)-M model, subsequently denoted as the Brown Model:

$$BTC_t = \delta_0 + bFIN_{t-1} + \theta h_t + \epsilon_t, \tag{1.5}$$

$$h_t = \delta_2 + \phi\epsilon_{t-1} + \psi h_{t-1}, \tag{1.6}$$

where δ_2 is a constant term; all variables and parameters in eq 1.5 and 1.6 are as previously defined. As revealed in Table 1.5, the only distinction between the Green and the Brown Models is whether the vectors of environmental variables (ECM, ESM, and FIN) are included. Comparing the forecasts, $BTC^{green}_{t+h|t}$ vs $BTC^{brown}_{t+h|t}$ and $VAR^{green}_{t+h|t}$ vs $VAR^{brown}_{t+h|t}$, determines whether EA, comprising both ECM and ESM, contributes to reducing the forecasting errors.

I calculate the forecast errors using various loss functions: mean absolute error (MAE), mean squared error (MSE), mean absolute percentage error (MAPE), and linear exponential (LINEX) loss function, which are presented as follows.

$$MAE : L(Y_{t+h}, Y_{t+h|t}) = |Y_{t+h} - Y_{t+h|t}| \tag{1.7}$$

$$MSE : L(Y_{t+h}, Y_{t+h|t}) = \left(Y_{t+h} - Y_{t+h|t}\right)^2 \tag{1.8}$$

$$MAPE : L(Y_{t+h}, Y_{t+h|t}) = |\frac{Y_{t+h} - Y_{t+h|t}}{Y_{t+h}}| \tag{1.9}$$

$$LINEX : L(Y_{t+h}, Y_{t+h|t}) = \frac{2[exp[\pi(Y_{t+h} - Y_{t+h|t})] - \pi(Y_{t+h} - Y_{t+h|t}) - 1]}{\pi^2} \tag{1.10}$$

$Y_{t+h|t}$ is the h-step ahead forecast of Y_{t+h}. Y could be either Bitcoin's rate of return or volatility.[8] $L(Y_{t+h}, Y_{t+h|t})$ denotes the forecast error.

The forecast errors, as presented in eq 1.7-1.10, are symmetric, meaning they penalize both over- and underpredictions equally. This limitation motivates utilizing an

8. In line with Ghysels and Marcellino (2018), the realized volatility (RV) is defined as $RV_h = BTC^2_h$. $h \in \{1, 7, 30\}$.

asymmetric loss function, presented in eq 1.10. If π is positive (negative), underprediction becomes more (less) costly than overprediction. Following Yu (2002), I consider multiple values of π: $\pi \in \{-20, -10, 10, 20\}$.

Finally, I conduct two forecast accuracy tests: the Diebold-Mariano (DM) test (Diebold and Mariano, 2002) and the mean directional accuracy test of Pesaran and Timmermann (2009) (PT).[9] The PT test is implemented in the form considered by Bernard et al. (2018). Namely, the sign of the actual return is regressed on that of the predicted return and a constant. If the parameter associated with the forecast sign is significant, then the null hypothesis that 'the forecast rate of return does not predict the directional change of Bitcoin's price' is rejected.

1.4 Results

The results are structured into three parts, each focusing on the different aspects of climate risk: uncertainty (i.e., ESM), aversion (i.e., ECM), and attention (i.e., EA). In the following discussion, I build upon the definition of Aruoba et al. (2024) to introduce two effects: realization and anticipation. Precisely, realization effect denotes the improvement of the predictive ability of volatility in-sample, while anticipation effect refers to the improvement of the predictive ability of volatility out-of-sample. These effects, as summarized in Table 1.6, appear to be two short-term confounding effects associated with EA that are disentangled through the measures proposed in my present work: ESM and ECM.

1.4.1 Climate Risk Uncertainty

Table 1.7 summarizes the results of the in-sample analysis. Focusing on the ESM, its role seems more prominent in the return than the conditional variance equation. Specifically, in panel A of the main specification (the return equation of the Green Model), only the effects associated with the Google Trends on *pollution* and *climate change* are significant (at the 95% and 99% confidence levels, respectively). These effects remain robust in the specification accounting for the sentiment of the ECM (Sentiment-Driven Green Model). Moving to Panel B (conditional variance equation), the effect tied to the ESM lacks robustness. For instance, the effect linked to the Google Trends on

9. Although Model 2 is nested in Model 1, the use of a rolling window justifies the application of the DM test. See Giacomini and White (2006).

global warming is significant in the Sentiment-Driven Green Model but not the Green Model, while the effect accompanied by the Google Trends on *climate change* is significant in both specifications. These findings highlight that climate risk uncertainty, as proxied by the ESM, helps explain Bitcoin's expected returns. In fact, Pindyck (2014) argues that uncertainty stems from an incomplete understanding of the link between greenhouse gas levels and temperature, and between elevated temperatures and GDP growth. Pindyck (2014) builds a theoretical framework that incorporates uncertainty into the design of environmental policies. Accordingly, his framework emphasizes the importance of addressing not only the reduction of expected environmental damages but also the uncertainty surrounding these damages. While Pindyck (2014)'s analysis primarily focuses on the macroeconomic implications of climate uncertainty, its underlying principles extend to asset pricing, particularly in speculative markets like Bitcoin. As revealed in my present work, public consciousness about climate risk, gained via Google searches, may assist in shaping climate risk uncertainty and can potentially constitute a pricing factor in the Bitcoin market.

1.4.2 Climate Risk Aversion

As shown in Table 1.7 of the in-sample analysis, the ECM explicitly contributes to explaining Bitcoin's volatility but not expected returns. This is evident in the effects associated with the tweet volume on *pollution*, *climate change*, and *global warming*, which are significant at the 99 % confidence level (and mostly positive) in the conditional variance equation of the Green Model only, but not the return equation. Analyzing this effect based on the tweet sentiment reveals that the role of negative environmental concerns (*pollution*, *climate change*, and *global warming*) is more prominent than the positive concerns (*pollution* and *global warming*). This is reflected in the effect associated with each of these negative concerns, which is statistically different from their corresponding positive concerns, as indicated in the results of the Wald test in Table 1.8. Specifically, the effect attributed to public environmental concerns is generally positive, significant at the 99% confidence level, and has a larger magnitude for negative concerns only. This constitutes evidence on the presence of a short-term realization effect, in the sense that climate risk aversion contributes to explaining Bitcoin volatility in the short-run. Indeed, Aruoba et al. (2024), Berger, Dew-Becker, and Giglio (2020), and Dew-Becker et al. (2017), among others, showed that two effects can be attributed to uncertainty in general: realization, defined as the world being more volatile today, and anticipation, defined as the world expected to be volatile tomorrow. In this present framework, the realization effect is linked more closely with aversion, particularly climate risk aversion, rather than with uncertainty, and is exclusively captured via the ECM. Given its social media-based nature, the ECM constitutes alternative (i.e., nontraditional) data, entailing information on social interaction. This information, which is not captured using financial indicators, is accompanied by higher Bitcoin volatility, owing to its role in the dissemination of climate risk aversion.[10]

1.4.3 Climate Risk Attention

Tables 1.9-1.13 present the results of the out-of-sample analysis. Two points can be highlighted in this regard. Firstly, climate risk attention results in short-term anticipation effect, which is characterized by the improvement of the out-of-sample predictive ability of the volatility for short time horizons. To elaborate, Tables 1.9 and 1.10 summarize the expected return and variance forecast accuracy results of the DM test, respectively, using symmetric loss functions. Focusing on the expected return, accounting for EA does not seem to improve the out-of-sample predictive ability, as inferred by the p-values exceeding 5% across all loss functions (MSE, MAE, and MAPE). However, this result does not hold for the variance forecasts. Specifically, the model accounting for EA (Green Model) yields significantly lower forecast errors (p-value < 5%) than the model dismissing the environmental variables (Brown Model), across all the symmetric loss functions. Therefore, the anticipation effect, as delineated by Aruoba et al. (2024) to be confined to uncertainty, can be broadened to include aversion, in general, and climate risk aversion, in specific.

Secondly, the forward-looking information entailed in the proxy for climate risk attention (EA) is especially pertinent during scenarios of asymmetric risk and environmental stress. Nevertheless, this out-of-sample relevance is typically in the short-run only (day-week).

1.4.3.1 Asymmetric Risk

Tables 1.11 and 1.12 present the return and variance forecast accuracy results of the DM test, respectively, associated with the LINEX asymmetric loss function, where overpredicting and underpredicting are penalized differently. In both tables, when the penalty is negative and is sufficiently large enough ($\pi = -20$), as indicated by the

10. https://lens.monash.edu/@climate-change-rising-to-the-real-urgent-and-globa/2023/11/24/1386202/scrolling-into-stress-how-climate-fears-hit-youth

test statistic (< -1.96), accounting for EA improves the out-of-sample predictive ability of Bitcoin returns and their volatility, albeit only for short time horizons (day to week). A large negative penalty likely reflects scenarios where overprediction tends to be more costly than underprediction. Accordingly, the forward-looking feature of the environmentally-themed textual data gains particular importance in the context of asymmetry in risk. This result aligns with the findings of Adämmer, Prüser, and Schüssler (2024), Delle Monache, De Polis, and Petrella (2024) and Barbaglia, Consoli, and Manzan (2023), who underscore the importance of textual data in improving the forecasts of macroeconomic time series (e.g., inflation, GDP, among others). In this present study, the forward-looking relevance of textual data is extended to highly speculative financial markets such as Bitcoin.

1.4.3.2 Environmental Stress

Table 1.13 summarizes the results of the PT test. The coefficient associated with the return forecast sign is significant at the 95 % confidence level in the Green Model but not the Brown Model, and only for the one-day-ahead forecast horizon. Thus, accounting for EA improves the predictive ability of the Bitcoin price directional changes. Several studies, including those by Hong et al. (2024) and Zheng et al. (2024), have highlighted the out-of-sample relevance of textual data, primarily during periods of economic stress (i.e., recession). However, as noted by Kalamara et al. (2022) and Ellingsen, Larsen, and Thorsrud (2022), this predictive value of textual data is mostly significant for short-term forecasts. Given the text-based nature of the EA measure (i.e., textual data presenting the concern dimension), the results associated with the PT test in this present study establish the short-run forecasting relevance of textual data during episodes of environmental stress. In the aftermath of disruptive environmental events (e.g., natural disasters), climate risk attention generally, and climate risk concerns, specifically, surge, intermittently after being latent during periods of stability. Accordingly, (environmentally-themed) textual data successfully capture, in the short-run, fleeting elusive predictive factors of Bitcoin returns. This result holds despite the challenges associated with "identifying intermittent predictive components" in returns, particularly in short-time horizons, as highlighted by Timmermann (2008).

1.5 Robustness Check

In this section, I validate that the in-sample and out-of-sample results are robust to the model specifications and assumptions.

1.5.1 In-sample Results

I perform the following checks to confirm the robustness of the in-sample results . First, I redefine one of my explanatory variables; namely the ESM. Second, I account for potential omitted non-financial variables and financial/economic variables. Third, I perform a principal component analysis (PCA) and a cross-validation. Finally, I conduct placebo tests to confirm that the tweets comprising the ECM are not contaminated, meaning that they are not capturing non-environmental concerns.

1.5.1.1 Google Search Volume Index

The proxy for climate risk uncertainty, ESM, was initially defined as $Ln(1 + GSVI)$ for various environmental queries, as in El Ouadghiri et al. (2021). This transformation is common practice, as it allows the log to be taken when search volumes are zero and slightly compresses the scale for low values. To test whether the findings are sensitive to this functional form, I re-estimate both the Green and the Sentiment-Driven Green Models, while replacing $Ln(1 + GSVI)$ with $Ln(GSVI)$. This robustness check assesses whether the results are driven by the handling of low or zero search volumes rather than by the underlying relationship between climate risk uncertainty and the dependent variables. The results associated with the two models are summarized in Table 1.14. In line with my previous findings, the role of climate risk uncertainty (aversion) seems to be more prominent in the return (variance) equation. The effects linked to the Google Trends (see *pollution* and *climate change*) are still significant and at the same level (p-value < 5%). Likewise, the effect associated with the tweets (see *pollution*, *global warming* and *climate change*) remains robust in the conditional variance equation, providing evidence of the short-term realization effect. While the focus is, in essence, not on financial variables, their effects remain robust except for oil prices, which was significant at the 5%; now it becomes weakly significant (at the 10% significance level).

1.5.1.2 Omitted Variables (1)

I maintain the initial definition of the ESM, i.e. $Ln(1 + GSVI)$, but I include more (potentially omitted) controls to both Green and Sentiment-Driven Green Models, which could reshape both public concerns and consciousness. Specifically, I include $disaster_{t-1}$ in eq 1.1 and 1.3, which is a binary variable indicating whether a natural disaster happened on day $t - 1$. As defined by Wikipedia, I only control for the deadliest worldwide natural disasters (excluding famines and epidemics), as illustrated in Table 1.15. In addition, I include month-fixed effects in the return equations 1.1 and 1.3 to capture possible unobserved factors affecting Bitcoin's returns. The results of both models, presented in Table 1.16, indicate that the effect associated with $disaster_{t-1}$ is insignificant in both models. More importantly, the effect associated with the ESM, ECM, and all other variables almost preserves its sign and significance, as in the main specification, thus confirming my previous findings.

1.5.1.3 Omitted Variables (2)

I consider two potentially relevant, yet omitted, financial/economic variables from the Green and Sentiment-Driven Green Model. Specifically, I control for the cryptocurrency volatility index (CVIX) as well as the US economic policy uncertainty index (PUI) in the return equations.[11] The results are reported in Table 1.17. The effects associated with the CVIX and PUI are insignificant; however, the established effects associated with the ESM and ECM remain unchanged in both models and across all specifications (i.e., including either CVIX or PUI, or both).

1.5.1.4 Principal Component Analysis (PCA)

Given that I control for a large number of financial variables, a potential limitation is the curse of dimensionality. Thus, a PCA might be a solution to reduce the data dimension while maintaining almost all information entailed in the large dataset. I construct a principal component using all the financial variables, including the CVIX and PUI, in the Green and Sentiment-Driven Green Models. The new results are summarized in Table 1.18. As shown in this table, the effects associated with all Twitter and Google variables remain robust in magnitude, sign, and significance level, and in both models.

11. The CVIX is retrieved from https://www.royalton-crix.com/. The CVIX is available since 29/06/2016. Observations prior to this date are assigned the same value as that on 29/06/2016. The PUI is retrieved from https://fred.stlouisfed.org/

1.5.1.5 Cross-validation

Another potential concern is whether the effect associated with the financial variables was initially significant, but has been masked by the proposed ESM and ECM. Accordingly, I perform a cross-validation to confirm that this is not the case. I consider various specifications based on financial variables while disregarding the environmental variables. The results are reported in Table 1.19. Across all the 40 specifications considered, the effect associated with all financial variables is mostly insignificant. The significance of the effect associated with the VIX has been restored sometimes; nevertheless, the significance is marginal only (at the 10 percent level only). In fact, this finding resonates with Harvey and Liu (2021), Cochrane (2009), and Goyal (2012), who shed light on the challenges surrounding the identification of pricing factors. Thus, the role of financial variables in determining Bitcoin's returns may be indirect, potentially operating through the risk premium. This channel warrants further investigation. In conclusion, the insignificance of the financial variables is not driven by the environmental proxies.

1.5.1.6 Tweet Contamination

In this section, I defend that the tweets used to construct the ECM are indeed capturing public environmental concerns. In other words, these tweets are not contaminated with non-environmental information that could bias the constructed ECM.

Dependent Variable: One way to validate the environmental authenticity of the ECM is by investigating its impact on another financial asset comparable with Bitcoin. To elaborate, I now focus on examining the effect associated with ECM on the returns and volatility of Ethereum. In fact, Ethereum is the second most traded cryptocurrency and has a lower carbon footprint than Bitcoin. If my proposed ECM is indeed capturing environmental concerns and is, thus, an authentic proxy for climate risk aversion, this ECM should also have an impact on Ether's rate of return and volatility. One would also expect this impact to reflect the fact that Ethereum is relatively more sustainable than Bitcoin. Table 1.20 summarizes the results associated with the Green and Sentiment-Driven Green Models, with the only difference being the Ether's returns and volatility as the dependent variables in the return and variance equations, respectively. Consistent with my previous findings, the role of ESM (ECM) is more prominent in the return (variance) equation. While both Bitcoin Ethereum become more volatile as environmental concerns increase (see e.g., the tweets on *pollution*), the effect associated with the ECM seems to be weaker for Ethereum. This is reflected in the marginal impact associated with the tweets in Ethereum's conditional variance equation: for most tweets, the effect is smaller and insignificant, in comparison with Bitcoin. Since the results resonate with the fact that Ethereum is more environmentally friendly than Bitcoin, the proposed ECM is, indeed, capturing environmental concerns.

Irrelevant Tweets: Another way to validate the credibility of the proposed ECM is through a placebo test. This test involves analyzing tweets on irrelevant topics to ensure that the ECM captures genuine public environmental concerns. If the ECM is valid, tweets on non-environmental topics, such as abortion, should not exhibit significant effects on Bitcoin's return and, more critically, its volatility, where the ECM is most impactful. Accordingly, I retrieve tweets on *abortion*, a topic orthogonal to the Bitcoin market. I re-estimate both the Green and Sentiment-Driven Green Models, with only difference that the tweets on *abortion* are among the explanatory variables. The results are summarized in Table 1.21. First, the effects associated with all variables, ESM, ECM, financial, and environmental remain robust in both models. Importantly, the effect associated with the number of tweets on *abortion* is insignificant in the return and conditional variance equations, across both models. This result confirms that these tweets are indeed capturing *abortion* concerns, which are unrelated to the Bitcoin market. Therefore, this reinforces my interpretation of the ECM as capturing environmental concerns.

1.5.2 Out-of-sample Results

To validate the robustness of the out-of-sample results, I employ a recursive window instead of a rolling one. A recursive window evaluates the extent to which acquiring additional information improves predictions, particularly for directional price changes (PT test). Decreasing the window's size has resulted in many missing forecasts because optimization is more likely to fail given the few observations available for estimation (the objective function is no longer concave). Moreover, increasing the window's size might not be too helpful because I end up with a few forecasts that might not be enough to conduct the forecast accuracy tests. Therefore, I maintain the same window size (around 90% of the sample). I repeat the DM test and the mean directional forecast accuracy test PT across the same horizons.

The results of the DM test for the return and variance forecast accuracy for symmetric loss functions are summarized in Tables 1.22 and 1.23. In line with my previous finding, controlling for environmental concerns improves the out-of-sample predictive ability of the volatility only. Thus, the anticipation effect associated with EA is still robust; however, it now persists to a slightly longer horizon (one-month) as more information is acquired prior to forecasting (recursive window). Tables 1.24 and 1.25 summarize the results of the DM test associated with the return and variance forecasts, for the LINEX asymmetric loss function. In line with my previous findings, the short-run forward-looking feature of EA during periods associated with asymmetric risk holds specifically for the variance forecasts (one week). This is inferred from the DM-test statistic (<-1.96) as different penalties are set on over and under-predictions. Lastly, Table 1.26 presents the results associated with the PT sign test. The forward-looking feature of EA loses significance, as indicated by the p-value (>0.05) associated with the forecast sign in the Green Model. This finding does not necessarily invalidate the previous results. Rather, it implies that acquiring further information (i.e., through the recursive window) does not improve the predictive ability of the directional price change.

1.6 Concluding Remarks

Bitcoin continues to impose environmental and, potentially, health threats via mining and transactions.[12] Thus, incorporating climate risk into the cryptocurrency markets, particularly the Bitcoin market, is essential to hedge against climate risk. This study is the first to empirically establish and uncover potential asymmetries in these linkages. Specifically, the paper investigates the role of public environmental concerns in explaining Bitcoin's returns and volatility, both in-sample and out-of-sample. A key contribution of this paper is distinguishing between two dimensions of environmental attention: concerns and consciousness. From this distinction, two novel measures are introduced. Environmental Concern Measure (ECM), constructed from social media-based textual data (Tweets), serves as a proxy for climate risk aversion. In contrast, Environmental Consciousness Measure (ESM), based on internet search volume (Google Trends), acts as a proxy for climate risk uncertainty. Together, the ECM and ESM constitute a proxy for climate risk attention. In this present work, textual analysis is also

12. A single Bitcoin transaction consumes electricity that is enough to sustain 100,000 VISA transactions. See https://www.statista.com/statistics/881541/bitcoin-energy-consumption-transaction-comparison-visa/

implemented, which further refines the ECM to disentangle positive and negative public concerns.

The paper reveals several main findings. In-sample, climate risk uncertainty helps explain Bitcoin's expected return, while climate risk aversion contributes to explaining its volatility. This latter effect is referred to as realization effect. Out-of-sample, climate risk attention improves the predictive ability of the volatility for short-term horizons only. This effect is defined as short-term anticipation effect. Finally, the forward-looking feature of climate risk attention is particularly relevant for short-term horizons and periods associated with asymmetric risk or environmental stress. In fact, the high frequency and granularity of textual data make the ECM, and thus climate risk attention, especially effective in capturing fleeting predictive factors that traditional financial indicators often miss. These results remain robust to various specifications and assumptions.

Altogether, the results offer new insights and highlight some policy implications. First, social media-based text data emerges as a valuable, non-traditional data source, complementing traditional financial data. Compared to news articles, such data is more accessible, cost-effective, frequent, and granular, allowing for easy replication of the ECM. Second, fostering environmental awareness may not only complement (see e.g., Pindyck, 2014) or substitute (see e.g., Abdul Baki and Marrouch, 2022) environmental policies, but it can also shape climate risk consciousness, a potential pricing factor in the Bitcoin market. Finally, these insights extend beyond Bitcoin to other cryptocurrencies. For instance, Cahill and Liu (2021) show that copycat cryptocurrencies with Bitcoin-imitating names achieve higher short-term returns. Accordingly, climate risk may similarly spill over into copycat markets through the "name channel."

While these findings have significant implications for policy and practice, extending beyond Bitcoin to other cryptocurrencies, a number of avenues await future work. Future research could explore various methods (e.g., machine learning approaches) for transforming text data into meaningful time series and re-examine their advantages over traditional financial data. Furthermore, while this study underscores the short-term relevance of textual data, future work could leverage the high-frequency nature of social-media-based text data to focus on nowcasting, providing real-time insights into cryptocurrency market dynamics and environmental concerns.

Figures and Tables

FIGURE (1.1) Volume of Tweets on Environmental Queries

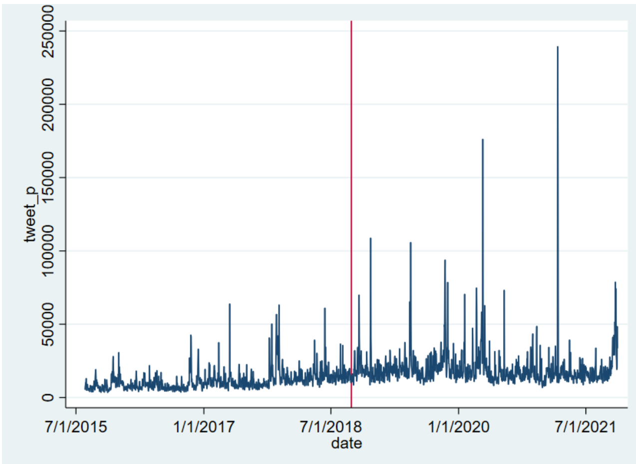

(A) pollution

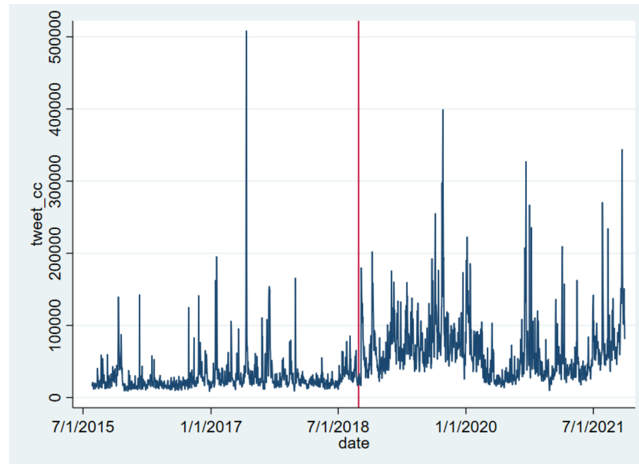

(B) climate change

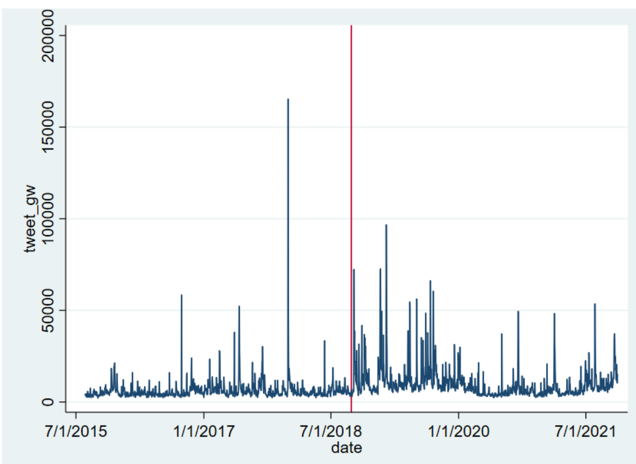

(C) global warming

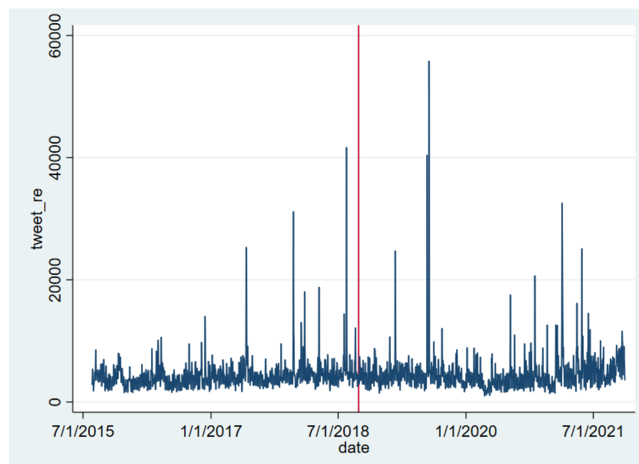

(D) renewable energy

FIGURE (1.2) Polarity Distribution

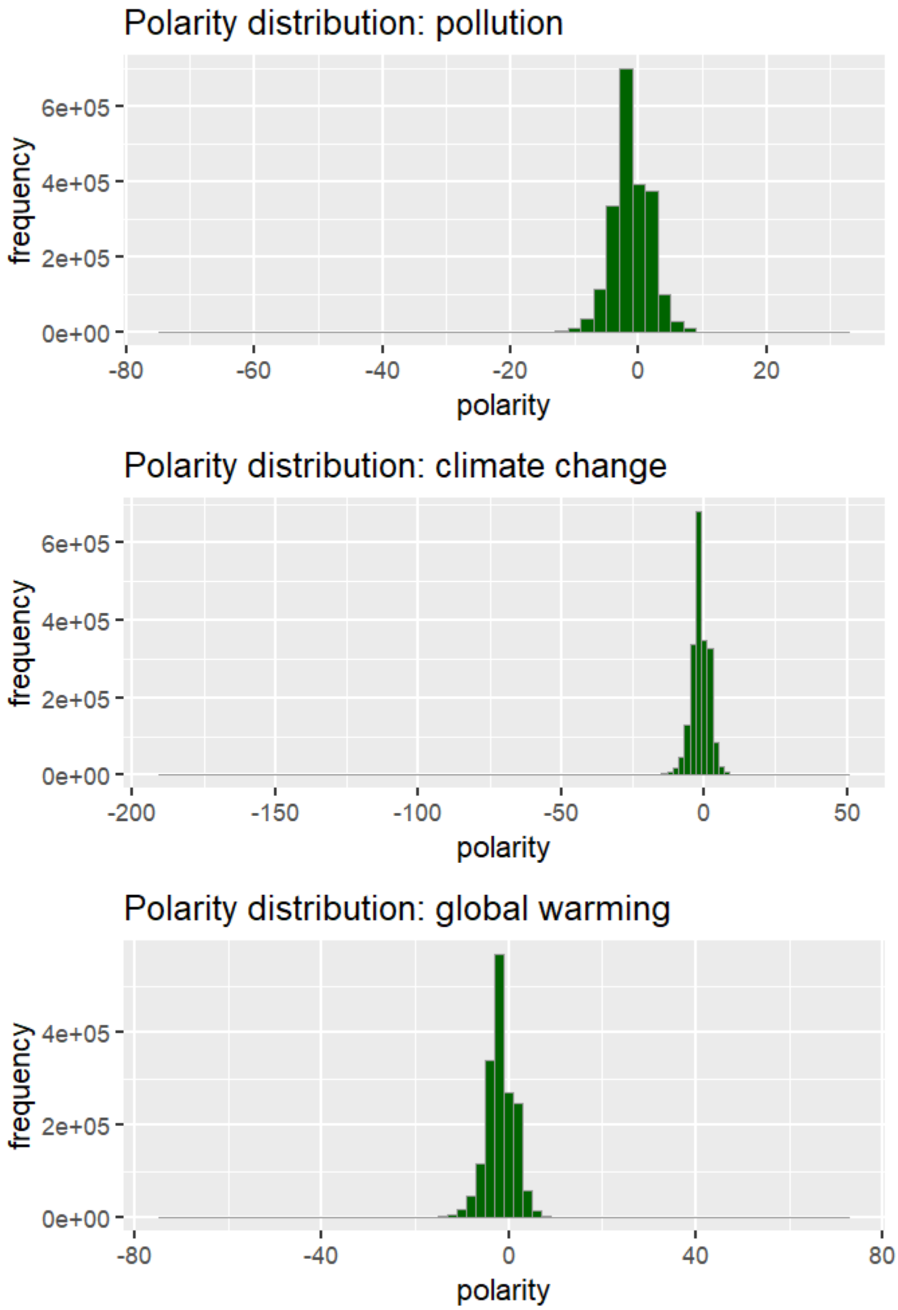

FIGURE (1.3) Word Clouds

(A) pollution

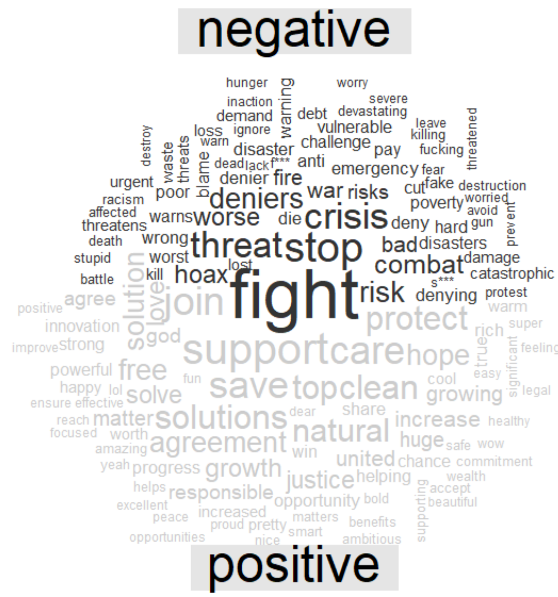

(B) climate change

(C) global warming

FIGURE (1.4) Volume of Tweets on "Bitcoin" **and** Environmental Queries

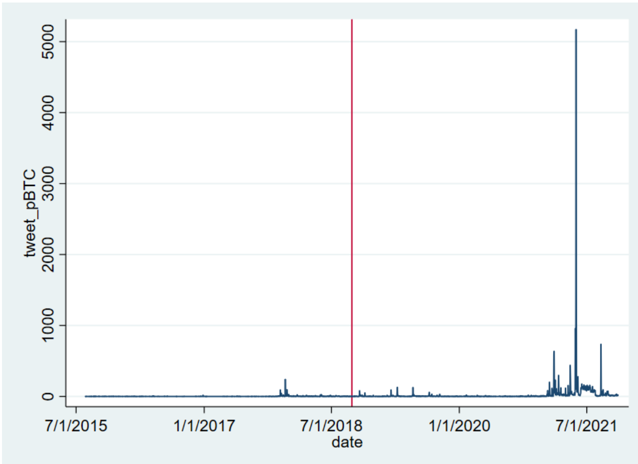

(A) pollution and bitcoin

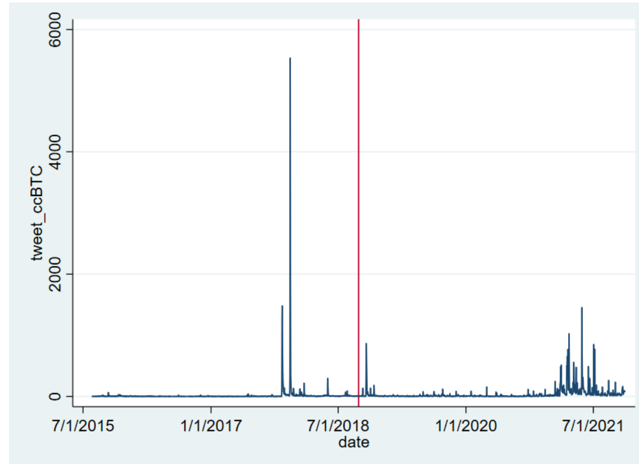

(B) climate change and bitcoin

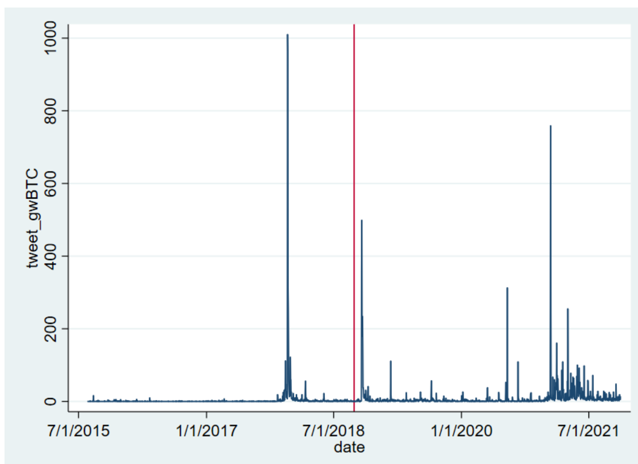

(C) global warming and bitcoin

TABLE (1.1) UN Observances (Theme: Environment/Climate Change)

Date	**Observance**
Last Saturday of March	Earth Hour
April 22	Earth Day
March 3	Ocean Day
March 21	Forest Day
March 22	Water Day
March 23	Meteorological Day
June 8	Ocean Day
October 4	Habitat Day
October 24	United Nations (UN) Day
November 5	Tsunami Day

TABLE (1.3) Descriptive Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Panel A: Tweets					
pollution	0.001	0.311	-1.156	1.773	1144
warming	0	0.366	-2.309	3.117	1144
climate	10.099	0.503	9.056	13.138	1145
renewable	0	0.327	-1.312	1.842	1144
$pollution^+$	0	0.287	-1.219	1.4	1144
$warming^+$	0	0.331	-1.177	1.478	1144
$climate^+$	0	0.296	-1.168	1.471	1144
$pollution^-$	0	0.229	-1.509	1.597	1144
$warming^-$	0	0.267	-1.347	0.928	1144
$climate^-$	0	0.232	-1.076	1.118	1144
bitcoin	0.001	0.163	-0.786	1.555	1144
Panel B: Google Trends					
$pollution^G$	3.676	0.527	2.303	4.615	1145
$warming^G$	3.443	0.574	1.609	4.615	1145
$climate^G$	3.297	0.697	1.099	4.615	1145
$renewable^G$	3.495	0.622	0	4.615	1145
Panel C: Financial Variables					
BTC	0.003	0.04	-0.208	0.225	1144
$DJSI$	0	0.007	-0.043	0.037	1144
$SP500$	0	0.033	-0.774	0.778	1144
$FTSE$	0	0.007	-0.041	0.038	1144
$FTSE4GD$	0	0.007	-0.043	0.041	1144
VIX	0	0.07	-0.3	0.768	1144
EUR_USD	0	0.004	-0.026	0.028	1144
$GOLD$	0	0.007	-0.036	0.046	1144
$NASDAQ$	0	0.008	-0.042	0.042	1144
OIL	0	0.02	-0.081	0.113	1144
ETH	0.005	0.07	-0.315	0.41	1144
$VOLUME$	0.005	0.325	-1.375	1.25	1144

Note: Column 1 presents the variables. Columns 2 and 3 report the mean and standard deviation, respectively. Columns 4 and 5 correspond to the minimum and maximum values of the variables. Column (6) presents the number of -in-sample- observations.

TABLE (1.2) The Variables of the Green Model

Variable	**Definition**
	ECM_t**: Environmental Concerns Measure at day** t
$pollution_t$	percentage change of the tweet volume on *pollution*
$warming_t$	percentage change of the tweet volume on *global warming*
$climate_t$	first difference of the percentage change of the tweet volume on *climate change*
$renewable_t$	percentage change of the tweet volume on *renewable energy*
	ESM_t: Environmental Consciousness Measure at day t
$pollution_t^G$	natural logarithm of one plus the GSVI for *pollution*
$warming_t^G$	natural logarithm of one plus the GSVI for *global warming*
$climate_t^G$	natural logarithm of one plus the GSVI for *climate change*
$renewable_t^G$	natural logarithm of one plus the GSVI for *renewable energy*
	ENV_t**: Environmental Incidents at day** t
$UNCC_t$	equals to 1 if the UNCC conference is held at day t, 0 otherwise
UN_t	equals to 1 if day t is observed by the UN as a special day (theme: environment/climate change), 0 otherwise
$UNpollution_t$	interaction term: $UNCC_t$ x $pollution_t$
$UNwarming_t$	interaction term: $UNCC_t$ x $warming_t$
$UNclimate_t$	interaction term: $UNCC_t$ x $climate_t$
$UNrenewable_t$	interaction term: $UNCC_t$ x $renewable_t$
	FIN_t**: Financial Variables at day** t
BTC_t	return rate of Bitcoin
ETH_t	return rate of Ethereum
$GOLD_t$	return rate of gold
$NASDAQ_t$	return rate of NASDAQ index
$FTSE_t$	return rate of FTSE USA index
$FTSE4GD_t$	return rate of FTSE4GOOD USA index
$DJSI_t$	return rate of DJSI index
$SP500_t$	return rate of S&P500 index
EUR_USD_t	percentage change of EUR-USD exchange rate
VIX_t	percentage change of Volatility Index (VIX)
$VOLUME_t$	percentage change of Bitcoin's trading volume
$bitcoin_t$	percentage change of tweet volume on *bitcoin*

TABLE (1.4) The Variables of the Sentiment-Driven Green Model

Variable	**Definition**
	ECM_t^-: **Negative Environmental Concerns Measure at day** t
$pollution_t^-$	percentage change of the negative tweet volume on pollution
$warming_t^-$	percentage change of the negative tweet volume on global warming
$climate_t^-$	percentage change of the negative tweet volume on *climate change*
	ECM_t^+: **Positive Environmental Concerns Measure at day** t
$pollution_t^+$	percentage change of the positive tweet volume on pollution
$warming_t^+$	percentage change of the positive tweet volume on *global warming*
$climate_t^+$	percentage change of the positive tweet volume on *climate change*
	$\widetilde{ESM}_t$: Environmental Consciousness Measure at day t
$pollution_t^G$	natural logarithm of one plus the GSVI for *pollution*
$warming_t^G$	natural logarithm of one plus the GSVI for *global warming*
$climate_t^G$	natural logarithm of one plus the GSVI for *climate change*
	$\widetilde{ENV}_t$: **Environmental Incidents at day** t
$UNCC_t$	equals to 1 if the UNCC conference is held at day t, 0 otherwise
UN_t	equals to 1 if day t is observed by the UN as a special day (theme: environment/climate change), 0 otherwise
	FIN_t: **Financial Variables at day** t
BTC_t	return rate of Bitcoin
ETH_t	return rate of Ethereum
$GOLD_t$	return rate of gold
$NASDAQ_t$	return rate of NASDAQ index
$FTSE_t$	return rate of FTSE USA index
$FTSE4GD_t$	return rate of FTSE4GOOD USA index
$DJSI_t$	return rate of DJSI index
$SP500_t$	return rate of S&P500 index
EUR_USD_t	percentage change of EUR-USD exchange rate
VIX_t	percentage change of Volatility Index (VIX)
$VOLUME_t$	percentage change of Bitcoin's trading volume
$bitcoin_t$	percentage change of tweet volume on *bitcoin*

TABLE (1.5) A Comparison of the Variables of the Green vs Brown Models

Variable	GM	BM
ECM_{t-1}	✓	
ESM_{t-1}	✓	
ENV_{t-1}	✓	
FIN_{t-1}	✓	✓

Note: ECM denotes the environmental concerns measure. ESM denotes the environmental consciousness measure. ENV is a vector of environmental incidents. FIN is a vector of financial variables. Columns (2) and (3) correspond to the Green (GM) and Brown (BM) Models, respectively.

TABLE (1.6) Realization and Anticipation Effects

Effect	Aruoba et al. (2024)	Abdul Baki (2025)
anticipation	more volatility today	improvement of the in-sample predictive ability of volatility
realization	likely volatile tomorrow	improvement of the out-of-sample predictive ability of volatility

TABLE (1.7) Maximum Likelihood Estimation for the Generalized Autoregressive Heteroskedastic in mean GARCH(1,1)-M model

	GM	SDGM	BM
	Panel A: Return Equation		
Financial Variables			
BTC_{t-1}	0.007	0.011	0.017
	(0.036)	(0.034)	(0.034)
ETH_{t-1}	0.000	0.003	0.006
	(0.012)	(0.014)	(0.011)
OIL_{t-1}	-0.068	-0.066	-0.047
	(0.043)	(0.042)	(0.048)
$NASDAQ_{t-1}$	0.049	0.173	0.090
	(0.320)	(0.355)	(0.300)
$GOLD_{t-1}$	0.132	0.110	0.172
	(0.110)	(0.118)	(0.119)
EUR_USD_{t-1}	-0.097	0.174	-0.059
	(0.174)	(0.188)	(0.169)
VIX_{t-1}	-0.049**	-0.029	-0.051**
	(0.022)	(0.024)	(0.020)
$FTSE4GD_{t-1}$	-0.388	-0.912	-0.211
	(0.728)	(0.913)	(0.736)
$FTSE_{t-1}$	0.116	0.839	-0.086
	(0.879)	(0.939)	(0.997)
$SP500_{t-1}$	-0.009	-0.008	-0.007
	(0.091)	(0.067)	(0.247)
$DJSI_{t-1}$	-0.101	-0.205	-0.154
	(0.492)	(0.461)	(0.649)
$VOLUME_{t-1}$	0.007***	0.004	0.009***
	(0.002)	(0.003)	(0.002)
$bitcoin_{t-1}$	0.002	0.001	-0.001
	(0.007)	(0.006)	(0.006)
Environmental Incidents			
UN_{t-1}	0.223	0.005	
	(0.165)	(0.004)	
$UNCC_{t-1}$	0.003	0.003	
	(0.004)	(0.004)	
$UNrenewable_{t-1}$	0.005		
	(0.029)		

Note: Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2). Column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Column 3 presents the estimation results of the Brown Model (eq 1.5 and 1.6). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses. ***, **, and * denote statistical significance at 1, 5, and 10 percent significance level.

TABLE (1.7) Maximum Likelihood Estimation for the Generalized Autoregressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM	BM
$UNpollution_{t-1}$	0.001		
	(0.023)		
$UNwarming_{t-1}$	0.016		
	(0.019)		
$UNclimate_{t-1}$	-0.023		
	(0.016)		
ECM: Environmental Concerns Measure (Tweets)			
$renewable_{t-1}$	-0.003		
	(0.003)		
$pollution_{t-1}$	-0.001		
	(0.004)		
$pollution^{+}_{t-1}$		0.001	
		(0.003)	
$pollution^{-}_{t-1}$		0.003	
		(0.004)	
$climate_{t-1}$	0.002		
	(0.003)		
$climate^{+}_{t-1}$		-0.003	
		(0.003)	
$climate^{-}_{t-1}$		0.001	
		(0.004)	
$warming_{t-1}$	-0.001		
	(0.003)		
$warming^{+}_{t-1}$		0.006^{***}	
		(0.002)	
$warmnig^{-}_{t-1}$		0.002	
		(0.003)	
ESM: Environmental Consciousness Measure (Google Trends)			
$renewable^{G}_{t-1}$	-0.002		
	(0.002)		
$pollution^{G}_{t-1}$	0.008^{**}	0.006^{**}	
	(0.003)	(0.003)	
$climate^{G}_{t-1}$	-0.007^{***}	-0.006^{***}	
	(0.002)	(0.002)	

Note: Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2). Column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Column 3 presents the estimation results of the Brown Model (eq 1.5 and 1.6). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses. ***, **, and * denote statistical significance at 1, 5, and 10 percent significance level.

TABLE (1.7) Maximum Likelihood Estimation for the Generalized Autoregressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM	BM
$warming^G_{t-1}$	0.002	0.000	
	(0.003)	(0.003)	
θ	0.650	0.394	0.602
	(1.107)	(1.062)	(1.006)
δ_0	-0.000	-0.003	0.002
	(0.007)	(0.006)	(0.001)
Panel B: Conditional Variance Equation			
ECM: Environmental Concerns Measure (Tweets)			
$pollution_{t-1}$	3.518***		
	(0.448)		
$pollution^+_{t-1}$		-1.077**	
		(0.464)	
$pollution^-_{t-1}$		5.958***	
		(0.639)	
$climate_{t-1}$	-3.102***		
	(0.629)		
$climate^+_{t-1}$		0.755	
		(0.470)	
$climate^-_{t-1}$		-8.004***	
		(0.790)	
$warming_{t-1}$	1.817***		
	(0.352)		
$warming^+_{t-1}$		-4.119***	
		(0.621)	
$warming^-_{t-1}$		1.807***	
		(0.492)	
$renewable_{t-1}$	-0.632		
	(0.911)		
ESM: Environmental Consciousness Measure (Google Trends)			
$pollution^G_{t-1}$	-0.476	0.860	
	(0.591)	(0.543)	
$climate^G_{t-1}$	-1.031***	-3.828***	
	(0.259)	(0.380)	
$warming^G_{t-1}$	-0.377	2.680***	
	(0.408)	(0.648)	

Note: Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2). Column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Column 3 presents the estimation results of the Brown Model (eq 1.5 and 1.6). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses. ***, **, and * denote statistical significance at 1, 5, and 10 percent significance level.

TABLE (1.7) Maximum Likelihood Estimation for the Generalized Autoregressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM	BM
$renewable^G_{t-1}$	1.060		
	(0.732)		
δ_1	-8.270***	-13.365***	
	(1.306)	(1.203)	
ϕ	0.161***	0.117***	0.141***
	(0.016)	(0.012)	(0.011)
ψ	0.822***	0.862***	0.866***
	(0.015)	(0.010)	(0.008)
δ_2			0.000***
			(0.000)
Panel C			
N	1143	1143	1143
AIC	-4463.705	-4500.612	-4436.244
BIC	-4262.048	-4309.039	-4345.499
LL	2271.852	2288.306	2236.122

Note: Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2). Column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Column 3 presents the estimation results of the Brown Model (eq 1.5 and 1.6). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses. ***, **, and * denote statistical significance at 1, 5, and 10 percent significance level.

TABLE (1.8) A Comparison of the Effects Associated with Environmental Concerns

	Pollution		Climate Change		Global Warming	
	Negative	Positive	Negative	Positive	Negative	Positive
Coefficient	5.958	1.077	8.004	0.755	1.807	4.119
	(0.639)	(0.464)	(0.790)	(0.470)	(0.492)	(0.621)
Difference	4.881***		7.249***		-2.312**	
	(0.968)		(1.011)		(0.901)	

Note: The effects associated with the environmental concerns on *pollution, climate change, and global warming* are reported in absolute value. Difference corresponds to the effect of negative − positive concerns. Standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.9) Return Forecast Accuracy Results of the Diebold-Mariano test (symmetric loss functions)

h	Loss	GM	BM	Difference	p-value	N
	MSE	0.001	0.001	-0.000	0.633	
1	MAE	1.179	1.820	-0.641	0.291	101
	MAPE	0.026	0.026	-0.000	0.778	
	MSE	0.001	0.001	-0.000	0.686	
7	MAE	1.134	1.147	-0.013	0.843	95
	MAPE	0.026	0.026	0.000	0.954	
	MSE	0.001	0.001	0.000	0.664	
30	MAE	1.407	1.069	0.338	0.145	72
	MAPE	0.027	0.026	0.001	0.234	

Note: The test is conducted using a rolling window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 2 presents the loss functions: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). Column 3 refers to Bitcoin's return forecast error based on the Green Model (GM). Column 4 refers to Bitcoin's return forecast error based on the Brown Model (BM). The forecast error difference between GM & BM is presented in column 5, while the p-value associated with this difference is presented in column 6. The number of generated forecasts is in column 7.

TABLE (1.10) Variance Forecast Accuracy Results of the Diebold-Mariano test (symmetric loss functions)

h	Loss	GM	BM	Difference	p-value	N
	MSE	4.70e-06	4.91e-06	-2.12e-07	0.0474	
1	MAE	0.00127	0.001442	-.0001745	0.0000	101
	MAPE	263.8	444	-180.2	0.2170	
	MSE	4.85e-06	5.12e-06	-2.73e-07	0.0679	
7	MAE	0.001276	0.001467	-0.000191	0.0000	95
	MAPE	50.35	81.35	-30.99	0.0185	
	MSE	5.95e-06	6.23e-06	-2.84e-07	0.0940	
30	MAE	0.001386	0.001595	-0.0002089	0.0001	72
	MAPE	47.71	70.53	-22.81	0.0743	

Note: The test is conducted using a rolling window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 2 presents the loss functions: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). Column 3 refers to Bitcoin's variance forecast error based on the Green Model (GM). Column 4 refers to Bitcoin's variance forecast error based on the Brown Model (BM). The forecast error difference between GM & BM is presented in column 5, while the p-value associated with this difference is presented in column 6. The number of generated forecasts is in column 7.

TABLE (1.11) Return Forecast Accuracy Results of the Diebold-Mariano test (LINEX asymmetric loss functions)

h	π	DM-statistic
1	-20	-9.369
	-10	-1.565
	10	0.550
	20	2.016
7	-20	-2.056
	-10	-1.185
	10	0.611
	20	1.490
30	-20	-0.980
	-10	-0.394
	10	0.953
	20	1.370

Note: The test is conducted using a rolling window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 3 corresponds to the LINEX function's forecast penalty π for $\pi \in \{-20, -10, 10, 20\}$. Column 4 presents the Diebold-Mariano (DM) test statistic associated with the difference in the return forecast error between the Green and Brown Models. Significance is when DM statistic $<$ -1.96.

TABLE (1.12) Variance Forecast Accuracy Results of the Diebold-Mariano test (LINEX asymmetric loss functions)

h	π	DM-statistic
1	-20	-3.292
	-10	-2.960
	10	-2.374
	20	-2.108
7	-20	-2.003
	-10	-1.866
	10	-1.599
	20	-1.469
30	-20	-2.009
	-10	-1.866
	10	-1.521
	20	-1.326

Note: The test is conducted using a rolling window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 3 corresponds to the LINEX function's forecast penalty π for $\pi \in \{-20, -10, 10, 20\}$. Column 4 presents the Diebold-Mariano (DM) test statistic associated with the difference in the variance forecast error between the Green and Brown Models. Significance is when DM statistic $<$ -1.96.

TABLE (1.13) Results of the Mean Directional Forecast Accuracy

Panel A: The Green Model			
h	1	7	30
		$sign_{t+h}$	
$sign^{green}_{t+h\|t}$	0.206**	0.179*	0.114
	(0.098)	(0.102)	(0.119)
constant	0.417***	0.426***	0.462***
	(0.071)	(0.072)	(0.081)
N	101	95	72
adj. R^2	0.033	0.022	-0.001
Panel B: The Brown Model			
h	1	7	30
		$sign_{t+h}$	
$sign^{brown}_{t+h\|t}$	0.147	0.140	0.244*
	(0.104)	(0.108)	(0.125)
constant	0.429***	0.424***	0.348***
	(0.084)	(0.087)	(0.103)
N	101	95	72
adj. R^2	0.010	0.007	0.038

Note: The test is conducted using a rolling window. Panels A and B present the results of the mean directional accuracy test associated with Green (GM) and Brown (BM) Models, respectively. Columns 2, 3, and 4 present the results of the one-day, week, and month-ahead forecasts, respectively. $sign_{t+h}$ is the sign of the actual rate of return. $sign^{green}_{t+h|t}$ and $sign^{brown}_{t+h|t}$ are the signs of the forecasted return rate generated using GM and BM, respectively. _cons refers to the constant term in the regression of the accuracy test. N corresponds to the sample size (number of forecasts), while adj. R^2 corresponds to the adjusted goodness-of-fit measure. Standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) statistical level.

TABLE (1.14) Robustness Check 1: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model

	GM	SDGM
Panel A: Return Equation		
Financial Variables		
BTC_{t-1}	0.016	0.011
	(0.035)	(0.034)
ETH_{t-1}	-0.006	0.003
	(0.011)	(0.014)
OIL_{t-1}	-0.073*	-0.066
	(0.043)	(0.042)
$NASDAQ_{t-1}$	-0.110	0.173
	(0.323)	(0.360)
$GOLD_{t-1}$	0.128	0.111
	(0.111)	(0.119)
EUR_USD_{t-1}	-0.057	-0.170
	(0.167)	(0.188)
VIX_{t-1}	-0.056***	-0.030
	(0.019)	(0.024)
$FTSE4GD_{t-1}$	0.058	-0.921
	(0.661)	(0.916)
$FTSE_{t-1}$	-0.162	0.840
	(0.838)	(0.942)

Note: The ESM is defined as Ln(Google search volume). Column 1 presents the estimation results of Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of model Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.14) Robustness Check 1: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$SP500_{t-1}$	-0.011	-0.008
	(0.061)	(0.068)
$DJSI_{t-1}$	-0.140	-0.206
	(0.453)	(0.470)
$VOLUME_{t-1}$	0.007***	0.004
	(0.002)	(0.003)
$bitcoin_{t-1}$	0.002	0.001
	(0.007)	(0.006)
Environmental Incidents		
UN_{t-1}	0.195	0.005
	(0.141)	(0.004)
$UNCC_{t-1}$	0.003	0.003
	(0.004)	(0.004)
$UNrenewable_{t-1}$	0.005	
	(0.028)	
$UNpollution_{t-1}$	-0.006	
	(0.021)	
$UNwarming_{t-1}$	0.021	
	(0.015)	

Note: The ESM is defined as Ln(Google search volume). Column 1 presents the estimation results of Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of model Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.14) Robustness Check 1: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$UNclimate_{t-1}$	-0.020	
	(0.014)	
ECM: Environmental Concerns Measure (Tweets)		
$renewable_{t-1}$	-0.002	
	(0.003)	
$pollution_{t-1}$	-0.001	
	(0.003)	
$pollution^{+}_{t-1}$		0.001
		(0.003)
$pollution^{-}_{t-1}$		0.003
		(0.004)
$climate_{t-1}$	0.001	
	(0.003)	
$climate^{+}_{t-1}$		-0.003
		(0.003)
$climate^{-}_{t-1}$		0.000
		(0.004)
$warming_{t-1}$	-0.001	
	(0.003)	

Note: The ESM is defined as Ln(Google search volume). Column 1 presents the estimation results of Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of model Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.14) Robustness Check 1: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$warming^{+}_{t-1}$		0.006***
		(0.002)
$warming^{-}_{t-1}$		0.002
		(0.003)
ESM: Environmental Consciousness Measure (Google Trends)		
$renewable^{G}_{t-1}$	-0.002	
	(0.002)	
$pollution^{G}_{t-1}$	0.007**	0.006**
	(0.003)	(0.003)
$climate^{G}_{t-1}$	-0.008***	-0.005***
	(0.002)	(0.002)
$warming^{G}_{t-1}$	0.003	0.000
	(0.003)	(0.003)
θ	0.567	0.376
	(1.068)	(1.065)
δ_0	-0.004	-0.003
	(0.006)	(0.006)
Panel B: Conditional Variance Equation		
$pollution_{t-1}$	2.681***	
	(0.499)	

Note: The ESM is defined as Ln(Google search volume). Column 1 presents the estimation results of Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of model Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.14) Robustness Check 1: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^{+}_{t-1}$		-1.101**
		(0.466)
$pollution^{-}_{t-1}$		5.961***
		(0.643)
$climate_{t-1}$	-2.370***	
	(0.680)	
$climate^{+}_{t-1}$		0.783*
		(0.472)
$climate^{-}_{t-1}$		-7.922***
		(0.781)
$warming_{t-1}$	1.088***	
	(0.334)	
$warming^{+}_{t-1}$		-4.045***
		(0.623)
$warming^{-}_{t-1}$		1.873***
		(0.499)
$renewable_{t-1}$	0.323	
	(0.833)	
$pollution^{G}_{t-1}$	-2.328***	0.821
	(0.643)	(0.540)

Note: The ESM is defined as Ln(Google search volume). Column 1 presents the estimation results of Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of model Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.14) Robustness Check 1: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$climate^G_{t-1}$	-1.336***	-3.511***
	(0.391)	(0.350)
$warming^G_{t-1}$	0.891	2.528***
	(0.606)	(0.631)
$renewable^G_{t-1}$	4.111***	
	(0.858)	
δ_1	-16.056***	-13.752***
	(1.905)	(1.182)
ϕ	0.188***	0.116***
	(0.018)	(0.012)
ψ	0.803***	0.864***
	(0.015)	(0.009)
Panel C		
N	1142	1143
AIC	-4489.366	-4498.708
BIC	-4287.745	-4307.134
LL	2284.683	2287.354

Note: The ESM is defined as Ln(Google search volume). Column 1 presents the estimation results of Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of model Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.15) Deadliest Natural Disasters

Date	Disaster	Country/Region
April-May, 2016	Heat Wave	India
September 19-21, 2017	Tropical cyclone	Puerto Rico, Dominica
September 28, 2018	Earthquake, Tsunami	Indonesia
June-July, 2019	Heat Wave	Europe
June-September, 2020	Flood	India, Bangladesh
August 14, 2021	Earthquake	Haiti

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model

	GM	SDGM
Panel A: Return Equation		
Financial Variables		
BTC_{t-1}	-0.007	-0.009
	(0.037)	(0.036)
ETH_{t-1}	0.000	0.004
	(0.013)	(0.014)
OIL_{t-1}	-0.067	-0.058
	(0.042)	(0.043)
$NASDAQ_{t-1}$	0.077	0.203
	(0.325)	(0.351)
$GOLD_{t-1}$	0.132	0.106
	(0.113)	(0.120)

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
EUR_USD_{t-1}	-0.034	-0.150
	(0.172)	(0.193)
VIX_{t-1}	-0.060***	-0.034
	(0.022)	(0.025)
$FTSE4GD_{t-1}$	-0.502	-0.885
	(0.744)	(0.913)
$FTSE_{t-1}$	0.095	0.676
	(0.861)	(0.955)
$SP500_{t-1}$	-0.009	-0.008
	(0.109)	(0.079)
$DJSI_{t-1}$	-0.120	-0.179
	(0.404)	(0.419)
$VOLUME_{t-1}$	0.007***	0.004
	(0.003)	(0.003)
$bitcoin_{t-1}$	0.003	0.003
	(0.008)	(0.006)
Environmental Incidents		
$disaster_{t-1}$	-0.002	-0.002
	(0.002)	(0.002)

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
UN_{t-1}	0.234	0.002
	(0.158)	(0.004)
$UNCC_{t-1}$	0.002	-0.001
	(0.004)	(0.005)
$UNrenewable_{t-1}$	0.003	
	(0.029)	
$UNpollution_{t-1}$	-0.001	
	(0.022)	
$UNclimate_{t-1}$	-0.024	
	(0.016)	
$UNwarming_{t-1}$	0.018	
	(0.018)	
ECM: Environmental Concerns Measure (Tweets)		
$renewable_{t-1}$	-0.002	
	(0.003)	
$pollution_{t-1}$	-0.000	
	(0.003)	
$pollution^{+}_{t-1}$		0.000
		(0.003)

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^{-}_{t-1}$		0.002
		(0.004)
$climate_{t-1}$	0.002	
	(0.003)	
$climate^{+}_{t-1}$		-0.002
		(0.003)
$climate^{-}_{t-1}$		0.001
		(0.004)
$warming_{t-1}$	-0.001	
	(0.003)	
$warming^{+}_{t-1}$		0.005**
		(0.002)
$warming^{-}_{t-1}$		0.001
		(0.003)
ESM: Environmental Consciousness Measure (Google Trends)		
$renewable^{G}_{t-1}$	-0.003	
	(0.002)	

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^G_{t-1}$	0.009**	0.005*
	(0.004)	(0.003)
$climate^G_{t-1}$	-0.007***	-0.005**
	(0.003)	(0.003)
$warming^G_{t-1}$	-0.001	-0.001
	(0.003)	(0.003)
θ	0.340	1.022
	(1.144)	(1.127)
δ_0	0.006	0.006
	(0.011)	(0.010)
month fixed effects	yes	yes
Panel B: Conditional Variance Equation		
ECM: Environmental Concerns Measure (Tweets)		
pollution	3.442***	
	(0.508)	
$pollution^+_{t-1}$		-1.080**
		(0.472)
$pollution^-_{t-1}$		5.992***
		(0.664)
$climate_{t-1}$	-3.119***	
	(0.697)	

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$climate^{+}_{t-1}$		0.703
		(0.470)
$climate^{-}_{t-1}$		-8.132***
		(0.833)
$warming_{t-1}$	1.905***	
	(0.386)	
$warming^{+}_{t-1}$		-4.120***
		(0.622)
$warming^{-}_{t-1}$		1.845***
		(0.487)
$renewable_{t-1}$	-0.634	
	(0.920)	
ESM: Environmental Consciousness Measure (Google Trends)		
$pollution^{G}_{t-1}$	-0.689	0.801
	(0.625)	(0.505)
$climate^{G}_{t-1}$	-1.020***	-3.937***
	(0.269)	(0.404)
$warming^{G}_{t-1}$	-0.454	2.849***
	(0.449)	(0.624)

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.16) Robustness Check 2: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$renewable^G_{t-1}$	1.348	
	(0.891)	
δ_1	-8.276***	-13.325***
	(1.617)	(1.180)
ϕ	0.190***	0.128***
	(0.020)	(0.015)
ψ	0.800***	0.849***
	(0.017)	(0.013)
Panel C		
N	1143	1143
AIC	-4452.301	-4499.126
BIC	-4190.148	-4247.055
LL	2278.151	2299.563

Note: The ESM is defined as Ln(1+ Google search volume). Column 1 presents the estimation results of the Green Model (eq 1.1 and 1.2), while column 2 presents the estimation results of the Sentiment-Driven Green Model (eq 1.3 and 1.4). Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.17) Robustness Check 3: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model

	GM			SDGM		
Panel A: Return Equation						
Financial Variables						
BTC_{t-1}	0.006	0.007	0.006	0.010	0.010	0.010
	(0.036)	(0.036)	(0.036)	(0.034)	(0.034)	(0.034)
ETH_{t-1}	0.001	0.001	0.001	0.003	0.003	0.003
	(0.012)	(0.012)	(0.012)	(0.014)	(0.014)	(0.014)
OIL_{t-1}	-0.068	-0.067	-0.068	-0.065	-0.065	-0.065
	(0.043)	(0.043)	(0.044)	(0.042)	(0.042)	(0.042)
$NASDAQ_{t-1}$	0.035	0.047	0.035	0.167	0.173	0.167
	(0.318)	(0.321)	(0.319)	(0.355)	(0.357)	(0.357)
$GOLD_{t-1}$	0.132	0.133	0.132	0.112	0.111	0.112
	(0.109)	(0.110)	(0.109)	(0.119)	(0.118)	(0.118)
EUR_USD_{t-1}	-0.091	-0.094	-0.090	-0.170	-0.174	-0.170
	(0.172)	(0.174)	(0.172)	(0.187)	(0.189)	(0.189)
VIX_{t-1}	-0.049**	-0.049**	-0.049**	-0.028	-0.029	-0.028
	(0.023)	(0.022)	(0.023)	(0.025)	(0.024)	(0.025)
$FTSE4GD_{t-1}$	-0.391	-0.379	-0.389	-0.918	-0.910	-0.916
	(0.725)	(0.727)	(0.724)	(0.912)	(0.923)	(0.922)
$FTSE_{t-1}$	0.139	0.111	0.138	0.858	0.834	0.856
	(0.879)	(0.727)	(0.883)	(0.937)	(0.938)	(0.937)

Note: GM panel corresponds to the estimation results of the Green Model (eq 1.1 and 1.2), while SDGM panel presents the estimation results of the Sentiment-Driven Green Model. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.17) Robustness Check 3: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

$SP500_{t-1}$	-0.001	-0.009	-0.001	-0.009	-0.008	-0.009
	(0.087)	(0.090)	(0.086)	(0.065)	(0.068)	(0.065)
$DJSI_{t-1}$	-0.101	-0.103	-0.102	-0.206	-0.208	-0.207
	(0.485)	(0.499)	(0.491)	(0.461)	(0.467)	(0.466)
$VOLUME_{t-1}$	0.007***	0.007***	0.007***	0.003	0.004	0.003
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
$bitcoin_{t-1}$	0.002	0.002	0.003	0.001	0.001	0.001
	(0.008)	(0.002)	(0.008)	(0.006)	(0.001)	(0.006)
$CVIX_{t-1}$	-0.009		-0.009	0.001		-0.006
	(0.009)		(0.009)	(0.003)		(0.009)
PUI_{t-1}		-0.000	-0.000		0.000	0.000
		(0.002)	(0.002)		(0.002)	(0.002)
Environmental Incidents						
UN_{t-1}	0.226	0.222	0.223	0.005	0.005	0.005
	(0.165)	(0.165)	(0.165)	(0.004)	(0.004)	(0.004)
$UNCC_{t-1}$	0.003	0.003	0.003	0.003	0.003	0.003
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
$UNrenewable_{t-1}$	0.005	0.005	0.005			
	(0.029)	(0.029)	(0.029)			
θ	0.749	0.650	0.750	0.437	0.394	0.437
	(1.068)	(1.108)	(1.119)	(1.065)	(1.061)	(1.076)

Note: GM panel corresponds to the estimation results of the Green Model (eq 1.1 and 1.2), while SDGM panel presents the estimation results of the Sentiment-Driven Green Model. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.17) Robustness Check 3: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

δ_0	-0.001	-0.000	-0.001	-0.003	-0.003	-0.003
	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)
Panel B: Conditional Variance Equation						
$pollution_{t-1}$	3.585***	3.531***	3.588***			
	(0.459)	(0.458)	(0.467)			
$pollution^{+}_{t-1}$				-1.092**	-1.077**	-1.092**
				(0.464)	(0.464)	(0.464)
$pollution^{-}_{t-1}$				5.972***	5.964***	5.974***
				(0.645)	(0.640)	(0.464)
δ_1	-8.344***	-8.317***	-8.355***	-13.456***	-13.382***	-13.460***
	(1.335)	(1.313)	(1.337)	(1.269)	(1.210)	(1.272)
ϕ	0.163***	0.161***	0.163***	0.117***	0.118***	0.117***
	(0.016)	(0.017)	(0.016)	(0.012)	(0.012)	(0.012)
ψ	0.822***	0.822***	0.822***	0.863***	0.863***	0.863***
	(0.015)	(0.015)	(0.015)	(0.009)	(0.009)	(0.009)
Panel C						
N	1143	1143	1143	1143	1143	1143
AIC	-4463	-4462	-4461	-4499	-4499	-4497
BIC	-4256	-4255	-4249	-4302	-4302	-4295
LL	2273	2272	2272	2289	2288	2286

Note: GM panel corresponds to the estimation results of the Green Model (eq 1.1 and 1.2), while SDGM panel presents the estimation results of the Sentiment-Driven Green Model. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.18) Robustness Check 4: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model

	GM	SDGM
Panel A: Return Equation		
Financial Variables		
BTC_{t-1}	0.016	0.010
	(0.035)	(0.033)
PCA_{t-1}	0.002	0.005
	(0.001)	(0.001)
$bitcoin_{t-1}$	0.005	0.004
	(0.007)	(0.005)
Environmental Incidents		
UN_{t-1}	0.229	0.006
	(0.181)	(0.004)
$UNCC_{t-1}$	0.003	0.003
	(0.004)	(0.004)
$UNrenewable_{t-1}$	0.013	
	(0.032)	
$UNpollution_{t-1}$	0.003	
	(0.024)	
$UNclimate_{t-1}$	-0.023	
	(0.018)	

Note: GM corresponds to the Green Model (eq 1.1 and 1.2), while SDGM corresponds to the Sentiment-Driven Green Model (eq 1.3 and 1.4). PCA denotes the principal component of all financial indicators. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.18) Robustness Check 4: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$UNwarming_{t-1}$	0.010	
	(0.021)	
ECM: Environmental Concerns Measure (Tweets)		
$warming_{t-1}$	-0.005	
	(0.003)	
$warming^{+}_{t-1}$	0.006***	
		(0.002)
$warming^{-}_{t-1}$		0.002
		(0.003)
$climate_{t-1}$	0.002	
	(0.003)	
$climate^{+}_{t-1}$		-0.002
		(0.003)
$climate^{-}_{t-1}$		0.001
		(0.003)
$pollution_{t-1}$	-0.001	
	(0.003)	
$pollution^{+}_{t-1}$		0.000
		(0.003)

Note: GM corresponds to the Green Model (eq 1.1 and 1.2), while SDGM corresponds to the Sentiment-Driven Green Model (eq 1.3 and 1.4). PCA denotes the principal component of all financial indicators. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.18) Robustness Check 4: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^{-}_{t-1}$		0.002
		(0.004)
$renewable_{t-1}$	-0.003	
	(0.003)	
$renewable^{G}_{t-1}$	-0.001	
	(0.002)	
ESM: Environmental Consciousness Measure (Google Trends)		
$pollution^{G}_{t-1}$	0.007**	0.006**
	(0.003)	(0.002)
$climate^{G}_{t-1}$	-0.008***	-0.006***
	(0.002)	(0.002)
$warming^{G}_{t-1}$	0.002	-0.000
	(0.003)	(0.003)
θ	0.711	0.490
	(1.129)	(1.056)
δ_0	-0.002	-0.003
	(0.007)	(0.006)
Panel B: Conditional Variance Equation		
ECM: Environmental Concerns Measure (Tweets)		
$pollution_{t-1}$	3.878***	
	(0.371)	

Note: GM corresponds to the Green Model (eq 1.1 and 1.2), while SDGM corresponds to the Sentiment-Driven Green Model (eq 1.3 and 1.4). PCA denotes the principal component of all financial indicators. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.18) Robustness Check 4: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^{+}_{t-1}$		-1.121**
		(0.434)
$pollution^{-}_{t-1}$		5.924***
		(0.621)
$climate_{t-1}$	-3.180***	
	(0.559)	
$climate^{+}_{t-1}$		0.762
		(0.456)
$climate^{-}_{t-1}$		-7.985***
		(0.765)
$warming_{t-1}$	1.854***	
	(0.303)	
$warming^{+}_{t-1}$		-3.998***
		(0.578)
$warming^{-}_{t-1}$		1.813***
		(0.475)
$renewable_{t-1}$	-0.970	
	(0.889)	

Note: GM corresponds to the Green Model (eq 1.1 and 1.2), while SDGM corresponds to the Sentiment-Driven Green Model (eq 1.3 and 1.4). PCA denotes the principal component of all financial indicators. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.18) Robustness Check 4: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
ESM: Environmental Consciousness Measure (Google Trends)		
$pollution^G_{t-1}$	-0.204	0.841
	(0.511)	(0.502)
$climate^G_{t-1}$	-1.273***	-3.866***
	(0.226)	(0.360)
$warming^G_{t-1}$	-0.052	2.750***
	(0.353)	(0.618)
$renewable^G_{t-1}$	0.508	
	(0.476)	
δ_1	-7.789***	-13.354***
	(0.095)	(1.090)
ϕ	0.135***	0.115***
	(0.012)	(0.012)
ψ	0.843***	0.863***
	(0.012)	(0.010)
Panel C		
N	1143	1143
AIC	-4464	-4509
BIC	-4313	-4368
LL	2262	2282

Note: GM corresponds to the Green Model (eq 1.1 and 1.2), while SDGM corresponds to the Sentiment-Driven Green Model (eq 1.3 and 1.4). PCA denotes the principal component of all financial indicators. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradient (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.19) Robustness Check 5: Cross Validation of the Maximum Likelihood Estimation for the Generalized Autoregressive Conditional Heteroskedastic in mean GARCH(1,1)-M model

VIX	VOLUME	BITCOIN	SP500	EUR_USD	OIL	ETHER	GOLD	NASDAQ	DJSI	FTSE	FTSE4GD	bitcoin
✓												
	✓											
✓	✓											
✓		✓										
	✓	✓										
✓	✓	✓										
✓	✓	✓	✓									
✓	✓	✓		✓								
✓	✓	✓	✓	✓								
✓	✓	✓			✓							
✓	✓	✓				✓						
✓	✓	✓	✓		✓							
✓	✓	✓		✓	✓							
✓	✓	✓	✓	✓	✓							✓
✓					✓							
	✓				✓							
✓	✓	✓			✓	✓						
✓	✓	✓				✓						
✓	✓	✓			✓		✓					
✓	✓	✓			✓	✓	✓					
✓	✓	✓	✓	✓	✓		✓					✓

Note: The first row corresponds to the explanatory variables lagged one day. Each row corresponds to a regression based on the Brown model (eq 1.5 and 1.6). ✓denotes the included explanatory variable. None of the explanatory variables is significant at the 5 percent level.

TABLE (1.19) Robustness Check 5: Cross Validation of the Maximum Likelihood Estimation for the Generalized Autoregressive Conditional Heteroskedastic in mean GARCH(1,1)-M mode (continued)

VIX	VOLUME	BITCOIN	SP500	EUR_USD	OIL	ETHER	GOLD	NASDAQ	DJSI	FTSE	FTSE4GD	bitcoin
✓	✓	✓	✓	✓	✓		✓					
✓	✓	✓	✓	✓	✓	✓	✓					
✓	✓	✓	✓	✓	✓		✓	✓				
✓	✓	✓	✓	✓	✓		✓		✓			
✓	✓	✓			✓		✓		✓			
✓	✓	✓			✓				✓			
✓	✓	✓							✓			
✓	✓	✓				✓			✓			
✓	✓	✓				✓		✓				
✓	✓	✓						✓				
✓	✓	✓			✓			✓				
✓	✓	✓		✓	✓						✓	
✓	✓	✓		✓							✓	
✓	✓	✓			✓						✓	
✓	✓	✓									✓	
✓	✓	✓	✓	✓								✓
✓	✓	✓								✓		
✓	✓	✓			✓					✓		

Note: The first row corresponds to the explanatory variables lagged one day. Each row corresponds to a regression based on the Brown model (eq 1.5 and 1.6). ✓denotes the included explanatory variable. None of the explanatory variables is significant at the 5 percent level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis

	BTC	ETH	BTC	ETH
Panel A: Return Equation				
Financial Variables				
BTC_{t-1}	0.007	-0.065	0.011	-0.048
	(0.036)	(0.056)	(0.034)	(0.055)
ETH_{t-1}	0.0004	0.0734*	0.003	0.0722*
	(0.012)	(0.041)	(0.014)	(0.042)
OIL_{t-1}	-0.068	-0.071	-0.066	-0.125
	(0.043)	(0.081)	(0.042)	(0.087)
$NASDAQ_{t-1}$	0.049	-0.519	0.173	-0.513
	(0.320)	(0.622)	(0.355)	(0.692)
$GOLD_{t-1}$	0.133	-0.106	0.11	-0.163
	(0.11)	(0.257)	(0.118)	(0.250)
EUR_USD_{t-1}	-0.097	-0.059	-0.174	0.038
	(0.174)	(0.422)	(0.188)	(0.405)
VIX_{t-1}	-0.0493**	-0.026	-0.029	-0.008
	(0.022)	(0.039)	(0.024)	(0.045)
$FTSE4GD_{t-1}$	-0.388	-1.045	-0.912	-1.925
	(0.728)	(1.851)	(0.913)	(1.886)

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis (continued)

	BTC	ETH	BTC	ETH
$FTSE_{t-1}$	0.115	0.627	0.839	2.06
	(0.879)	(2.201)	(0.939)	(2.0)
$SP500_{t-1}$	-0.009	-0.006	-0.008	-0.015
	(0.091)	(0.818)	(0.067)	(0.116)
$DJSI_{t-1}$	-0.101	0.915	-0.205	0.51
	(0.492)	(1.336)	(0.461)	(1.240)
$VOLUME_{t-1}$	0.007***	0.003	0.004	0.004
	(0.002)	(0.006)	(0.003)	(0.006)
$bitcoin_{t-1}$	0.002	-0.007	0.0009	-0.01
	(0.007)	(0.014)	(0.006)	(0.012)
Environmental Incidents				
UN_{t-1}	0.223	-0.352	0.005	-0.01
	(0.165)	(0.423)	(0.004)	(0.013)
$UNCC_{t-1}$	0.003	0.0006	0.003	0.002
	(0.004)	(0.011)	(0.004)	(0.01)
$UNrenewable_{t-1}$	0.005	-0.021		
	(0.029)	(0.059)		
$UNpollution_{t-1}$	0.001	-0.01		
	(0.023)	(0.048)		

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis (continued)

	BTC	ETH	BTC	ETH
$UNclimate_{t-1}$	-0.022	-0.03		
	(0.016)	(0.042)		
$UNwarming_{t-1}$	0.016	0.016		
	(0.019)	(0.042)		
ECM: Environmental Concerns Measure (Tweets)				
$renewable_{t-1}$	-0.003	0.001		
	(0.003)	(0.006)		
$pollution_{t-1}$	-0.0006	0.002		
	(0.004)	(0.007)		
$pollution^{+}_{t-1}$			0.0007	-0.0013
			(0.003)	(0.007)
$pollution^{-}_{t-1}$			0.003	-0.001
			(0.004)	(0.008)
$climate_{t-1}$	0.002	0.004		
	(0.003)	(0.006)		
$climate^{+}_{t-1}$			-0.003	-0.002
			(0.003)	(0.006)
$climate^{-}_{t-1}$			0.0005	0.002
			(0.004)	(0.009)

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis (continued)

	BTC	ETH	BTC	ETH
$warming_{t-1}$	-0.0008	-0.003		
	(0.003)	(0.005)		
$warming^{+}_{t-1}$			0.006***	0.003***
			(0.002)	(0.006)
$warming^{-}_{t-1}$			0.002	0.017 **
			(0.003)	(0.007)
ESM: Environmental Consciousness Measure (Google Trends)				
$renewable^{G}_{t-1}$	-0.002	0.003		
	(0.002)	(0.005)		
$pollution^{G}_{t-1}$	0.008**	- 0.002	0.006**	0.003
	(0.003)	(0.006)	(0.003)	(0.005)
$climate^{G}_{t-1}$	-0.007***	-0.01***	-0.006***	-0.012***
	(0.002)	(0.004)	(0.002)	(0.004)
$warming^{G}_{t-1}$	0.002	0.007	0.0001	0.008
	(0.003)	(0.006)	(0.003)	(0.006)
θ	0.65	1.141*	0.394	1.356*
	(1.107)	(0.692)	(1.062)	(0.759)
δ_0	-0.0003	-0.004	-0.003	-0.002
	(0.007)	(0.013)	(0.006)	(0.014)

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis (continued)

	BTC	ETH	BTC	ETH
Panel B: Conditional Variance Equation				
ECM: Environmental Concerns Measure (Tweets)				
$pollution_{t-1}$	3.518***	1.748***		
	(0.448)	(0.445)		
$pollution^{+}_{t-1}$			-1.077**	-1.97***
			(0.464)	(0.527)
$pollution^{-}_{t-1}$			5.958***	0.122
			(0.639)	(0.518)
$climate_{t-1}$	-3.102***	0.258		
	(0.629)	(0.544)		
$climate^{+}_{t-1}$			0.755	-1.257
			(0.470)	(0.767)
$climate^{-}_{t-1}$			-8.004***	0.457
			(0.790)	(1.082)
$warming_{t-1}$	1.817***	-0.572		
	(0.352)	(0.567)		
$warming^{+}_{t-1}$			-4.119***	0.768
			(0.621)	(0.683)

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis (continued)

	BTC	ETH	BTC	ETH
$warming^{-}_{t-1}$			1.807***	0.326
			(0.492)	(0.934)
$renewable_{t-1}$	-0.632	1.285**		
	(0.911)	(0.516)		
ESM: Environmental Consciousness Measure (Google Trends)				
$pollution^{G}_{t-1}$	-0.476	-2.022	0.86	0.423
	(0.591)	(0.399)	(0.543)	(0.325)
$climate^{G}_{t-1}$	-1.03***	0.029	-3.828***	-0.07
	(0.259)	(0.267)	(0.38)	(0.264)
$warming^{G}_{t-1}$	-0.3772	-0.624	2.68***	-0.218
	(0.408)	(0.458)	(0.648)	(0.460)
$renewable^{G}_{t-1}$	1.06	-1.409**		
	(0.732)	(0.273)		
δ_1	-8.27***	-9.017***	-13.365***	-8.817***
	(1.306)	(0.868)	(1.203)	(0.856)
ϕ	0.161***	0234***	0.117***	0.209***
	(0.016)	(0.029)	(0.012)	(0.028)
ψ	0.822***	0.719***	0.862***	0.727***
	(0.015)	(0.025)	(0.01)	(0.029)

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.20) Robustness Check 6: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (in-sample analysis (continued)

	BTC	ETH	BTC	ETH
		Panel C		
N	1143	1143	1143	1143
AIC	-44642	-3081	-4501	-3077
BIC	-4262	-2879	-4309	-2885
LL	2272	1580	2288	1576

Note: Columns (1) and (2) present the estimation results of Green Model (eq 1.1 and 1.2), where BTC and ETH are the dependent variables in equation (1), respectively. Columns (3) and (4) correspond to Sentiment-Driven Green Model (eq 1.3 and 1.4), where BTC and ETH are the dependent variables in equation (1*), respectively. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model

	GM	SDGM
Panel A: Return Equation		
$abortion_{t-1}$	-0.001	- 0.001
	(0.002)	(0.001)
Financial Variables		
BTC_{t-1}	0.006	0.01
	(0.036)	(0.034)
ETH_{t-1}	0.001	0.003
	(0.012)	(0.014)
OIL_{t-1}	-0.067	-0.067
	(0.043)	(0.043)
$NASDAQ_{t-1}$	0.047	0.169
	(0.324)	(0.360)
$GOLD_{t-1}$	0.123	0.103
	(0.112)	(0.120)
EUR_USD_{t-1}	-0.089	-0.167
	(0.177)	(0.189)
VIX_{t-1}	-0.049**	-0.029
	(0.022)	(0.025)

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$FTSE4GD_{t-1}$	-0.37	-0.909
	(0.731)	(0.920)
$FTSE_{t-1}$	0.088	0.828
	(0.892)	(0.964)
$SP500_{t-1}$	-0.009	-0.008
	(0.099)	(0.071)
$DJSI_{t-1}$	-0.088	-0.188
	(0.495)	(0.468)
$VOLUME_{t-1}$	0.007***	0.004
	(0.002)	(0.003)
$bitcoin_{t-1}$	0.002	0.001
	(0.007) (0.006)	
Environmental Incidents		
UN_{t-1}	0.224	0.005
	(0.165)	(0.004)
$UNCC_{t-1}$	0.003	0.003
	(0.004)	(0.004)

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$UNrenewable_{t-1}$	0.004	
	(0.029)	
$UNpollution_{t-1}$	0.001	
	(0.023)	
$UNclimate_{t-1}$	-0.023	
	(0.016)	
$UNwarming_{t-1}$	0.016	
	(0.019)	
ECM: Environmental Concerns Measure (Tweets)		
$renewable_{t-1}$	-0.003	
	(0.003)	
$pollution_{t-1}$	-0.0004	
	(0.004)	
$climate_{t-1}$	0.002	
	(0.003)	
$warming_{t-1}$	-0.0008	
	(0.003)	
$pollution^{+}_{t-1}$		0.0005
		(0.003)

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^{-}_{t-1}$		0.003
		(0.004)
$climate^{+}_{t-1}$		-0.003
		(0.003)
$climate^{-}_{t-1}$		0.0003
		(0.004)
$warming^{+}_{t-1}$		0.006***
		(0.002)
$warming^{-}_{t-1}$		0.002
		(0.003)
ESM: Environmental Consciousness Measure (Google Trends)		
$renewable^{G}_{t-1}$	-0.002	
	(0.002)	
$pollution^{G}_{t-1}$	0.008**	0.007**
	(0.003)	(0.003)
$climate^{G}_{t-1}$	-0.007***	-0.005**
	(0.002)	(0.002)

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$warming^G_{t-1}$	0.001	0.0001
	(0.003)	(0.003)
θ	0.619	0.419
	(1.112)	(1.068)
δ_0	-0.0007	-0.004
	(0.007)	(0.006)
Panel B: Conditional Variance Equation		
$abortion_{t-1}$	0.035	-0.009
	(0.583)	(0.451)
ECM: Environmental Concerns Measure (Tweets)		
$pollution_{t-1}$	3.498***	
	(0.452)	
$climate_{t-1}$	-3.09***	
	(0.632)	
$warming_{t-1}$	1.83***	
	(0.355)	
$renewable_{t-1}$	-0.616	
	(0.916)	

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$pollution^{+}_{t-1}$		-1.97***
		(0.464)
$pollution^{-}_{t-1}$		5.95***
		(0.653)
$climate^{+}_{t-1}$		0.755
		(0.79)
$climate^{-}_{t-1}$		-7.955***
		(0.790)
$warming^{+}_{t-1}$		-4.054***
		(0.655)
$warming^{-}t-1$		1.853***
		(0.507)
ESM: Environmental Consciousness Measure (Google Trends)		
$pollution^{G}_{t-1}$	-0.452	0.868
	(0.592)	(0.558)
$climate^{G}_{t-1}$	-1.046***	-3.864***
	(0.268)	(0.38)
$warming^{G}_{t-1}$	-0.373	2.744***
	(0.411)	(0.667)

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.21) Robustness Check 7: Maximum Likelihood Estimation for the Generalized Auto-regressive Heteroskedastic in mean GARCH(1,1)-M model (continued)

	GM	SDGM
$renewable^G_{t-1}$	1.02696	
	(0.733)	
δ_1	-8.2***	-13.488***
	(1.3)	(1.211)
ϕ	0.16***	0.116***
	(0.017)	(0.012)
ψ	0.822***	0.864***
	(0.015)	(0.01)
Panel C		
N	1143	1143
AIC	-4460	-4497
BIC	-4249	-4296
LL	2272	2289

Note: Columns (1) and (2) present the estimation results of the Green Model (eq 1.1 and 1.2) and the Sentiment-Driven Green Model (eq 1.3 and 1.4), respectively. The tweet volume on *abortion*is included as an explanatory variable in the return and variance equations. Panel A corresponds to the return equation. Panel B corresponds to the conditional variance equation. Panel C presents the number of observations (N), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and the maximum Log-Likelihood value (LL). Outer Product Gradiant (OPG) standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) significance level.

TABLE (1.22) Robustness Check 8: Return Forecast Accuracy Results of the Diebold-Mariano test (symmetric loss functions)

h	Loss	GM	BM	Difference	p-value	N
	MSE	0.001244	0.001256	-0.0000118	0.7484	
1	MAE	1.277	1.918	-0. 6401	0.2961	98
	MAPE	0.02638	0.02653	-0.0001578	0.788	
	MSE	0.001285	0.001282	2.32e-06	0.9614	
7	MAE	1.295	1.201	0.0938	0.3843	92
	MAPE	0.02709	0.0269	0.0001914	0.7687	
	MSE	0.001236	0.001194	0.0000416	0.3988	
30	MAE	1.591	1.128	0.1626	0.0381	84
	MAPE	0.02607	0.02501	0.001059	0.1172	

Note: The test is conducted using a recursive window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 2 presents the loss functions: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). Column 3 refers to Bitcoin's return forecast error based on the Green Model. Column 4 refers to Bitcoin's return forecast error based on the Brown Model. The forecast error difference between Models 1 & 2 is presented in column 5, while the p-value associated with this difference is presented in column 6. The number of generated forecasts is in column 7.

TABLE (1.23) Robustness Check 9: Variance Forecast Accuracy Results of the Diebold-Mariano test (symmetric loss functions)

h	Loss	GM	BM	Difference	p-value	N
	MSE	4.88e-06	5.04e-06	-1.57e-07	0.1583	
1	MAE	0.001301	0.001457	-.0001561	0.0001	98
	MAPE	275.2	456.4	-181.2	0.2231	
	MSE	5.08-06	5. e-06	-2.23e-07	0.0819	
7	MAE	0.001318	0.001486	-0.0001687	0.0001	92
	MAPE	54.52	82.85	-28.33	0.0109	
	MSE	5.34e-06	5.59e-06	-2.5e-07	0.0817	
30	MAE	00.001326	0.001508	-0.0001821	0.0002	84
	MAPE	65.85	94.76	-28.92	0.0341	

Note: The test is conducted using a recursive window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 2 presents the loss functions: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). Column 3 refers to Bitcoin's variance forecast error based on the Green Model (GM). Column 4 refers to Bitcoin's variance forecast error based on the Brown Model (BM). The forecast error difference between GM and BM is presented in column 5, while the p-value associated with this difference is presented in column 6. The number of generated forecasts is in column 7.

TABLE (1.24) Robustness Check 10: Return Forecast Accuracy Results of the Diebold-Mariano test (LINEX asymmetric loss functions)

h	π	DM-statistic
1	-20	-1.7
	-10	-1.361
	10	0.516
	20	1.339
7	-20	-1.561
	-10	-1.040
	10	1.041
	20	1.778
30	-20	-1.161
	-10	-0.329
	10	1.783
	20	2.321

Note: The test is conducted using a rolling window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 3 corresponds to the LINEX function's forecast penalty π for $\pi \in \{-20, -10, 10, 20\}$. Column 4 presents the Diebold-Mariano (DM) test statistic associated with the difference in the return forecast error between the Green and Brown Models. Significance is when DM statistic $<$ -1.96.

TABLE (1.25) Robustness Check 11: Variance Forecast Accuracy Results of the Diebold-Mariano test (LINEX asymmetric loss functions)

h	π	DM-statistic
1	-20	-1.258
	-10	-2.479
	10	-0.384
	20	0.616
7	-20	-1.540
	-10	-2.424
	10	-3.028
	20	-2.777
30	-20	-1.38
	-10	-0.583
	10	-0.510
	20	0.546

Note: The test is conducted using a rolling window. Column 1 presents the h-step-ahead forecast horizon. The first panel corresponds to the one-day-ahead forecast (h = 1). The second panel corresponds to the one-week ahead forecast (h = 7). The last panel corresponds to the one-month-ahead forecast (h = 30). Column 3 corresponds to the LINEX function's forecast penalty π for $\pi \in \{-20, -10, 10, 20\}$. Column 4 presents the Diebold-Mariano (DM) test statistic associated with the difference in the variance forecast error between the Green and Brown Models. Significance is when DM statistic $<$ -1.96.

TABLE (1.26) Robustness Check 12: Results of the Mean Directional Forecast Accuracy

Panel A: Green Model			
h	1	7	30
		$sign_{t+h}$	
$sign^{green}_{t+h\|t}$	0.14717	0.08967	0.14583
	(0.101)	(0.107)	(0.110)
constant	0.4528***	0.463***	0.438***
	(0.069)	(0.069)	(0.072)
N	101	95	72
adj. R^2	0.011	-0.003	0.009
Panel B: Brown Model			
h	1	7	30
		$sign_{t+h}$	
$sign^{brown}_{t+h\|t}$	0.11930	0.06917	0.31111***
	(0.104)	(0.108)	(0.110)
constant	0.44737***	0.45714***	0.30000***
	(0.081)	(0.085)	(0.088)
N	98	92	84
adj. R^2	0.03	-0.007	0.078

Note: The test is conducted using a recursive window. Panels A and B present the results of the mean directional accuracy test associated with the Green (GM) and Brown (BM) Models, respectively. Columns 2, 3, and 3 present the results of the one day, week, and month -ahead forecasts, respectively. $sign_{t+h}$ is the sign of the actual rate of return. $sign^{green}_{t+h|t}$and $sign^{brown}_{t+h|t}$ are the signs of the forecasted return rate generated using GM and BM, respectively. _cons refers to the constant term in the regression of the accuracy test. N corresponds to the sample size (number of forecasts), while adj. R^2 corresponds to the adjusted goodness-of-fit measure. Standard errors are reported in parentheses ( ). ***, **, and * denote statistical significance at 1, 5, and 10 percent (%) statistical level

Chapter 2

Emissions Taxes and Market Power: Linked vs. Unlinked Market Failures

2.1 Introduction

Achieving the net-zero emissions target requires not only the efforts of producers but also consumers. As highlighted by the International Energy Agency (IEA, 2021) and World Economic Forum (WEF, 2021), the actions and attitudes of individuals (i.e., the consumers) are crucial. The 2021 Ernest and Young (EY) Future Consumer Index reveals that sustainability is significant to 84% of consumers worldwide.[1] According to the 2020 McKinsey United States (US) consumer sentiment survey, over 60% of respondents are willing to pay more for sustainably-packaged products. Thus, companies, as in the packaging industry, are increasingly focusing on integrating environmental and social responsibility into their business practices, in an attempt to satisfy consumers' green preferences.[2] A broader perspective underscores that consumer-driven abatement effort is a pivotal factor contributing to the scale and speed of the net-zero transition. Accordingly, consumer-sided abatement efforts should be reflected in the design of anti-pollution instruments.

In this paper, we seek to determine the optimal emissions taxes for producers and consumers of a polluting final consumption product, considering that both parties can abate, but only one holds market power. To this end, we build a partial equilibrium model in the presence of one-sided market power, two-sided environmental externalities, and two-sided abatement efforts. We initially examine the case of a Cournot

1. https://sponsored.bloomberg.com/quicksight/How-Can-Companies-Promote-Sustainable-Consumer-Behavior

2. https://www.mckinsey.com/industries/consumer-packaged-goods/our-insights/consumers-care-about-sustainability-and-back-it-up-with-their-wallets

oligopoly and then the mirror case of a Cournot oligopsony. By addressing both production and consumption-sided externalities, we ensure that each party is held accountable for their own pollution, a goal resonating with the Polluter-Pays Principle (PPP) (OECD, 2001). Moreover, allowing abatement technologies to be adopted by both producers and consumers ensures that the green preference is simultaneously active in the production and consumption processes of the same product, a dynamic observed in industries like vehicles, packaging, energy, and food. Our model, therefore, offers a unifying framework that nests all prominent results in the emission taxation literature (e.g., Pigou, 1920; Barnett, 1980; Ebert, 1991; Requate, 2006) as special cases under specific assumptions about market power and externalities.

This dual focus, from both the pollution and abatement lens, remains relatively unexplored in the emissions taxation literature, despite its relevance to many day-to-day situations. For example, double-sided pollution emerges when market power is concentrated in the hands of producers. Notably, in the Bitcoin market, where producers (i.e., miners) hold market power (see Ma, Gans, and Tourky, 2018), both the production (mining) and consumption (transactions) of Bitcoin are environmentally detrimental. Mining emits approximately 82.43 million metric tonnes of carbon dioxide annually, while a single Bitcoin transaction consumes as much energy as roughly 100,000 Visa transactions.[3] Another example of supply-side market power is in the pharmaceutical industry. Active pharmaceutical ingredients released during production and consumption lead to ecotoxicological effects and antimicrobial resistance (AMR). The true burden of AMR far exceeds current estimates: 25000 deaths were reported in 2007, and projections predict that by 2050, morbidity will increase 15-fold in Europe.[4]

Similarly, double-sided environmental externalities are present when consumers have market power. For example, retail grocery chains form an oligopsony in the agricultural produce market. In 2021, 59% of the Canadian grocery retail market was controlled by three companies: Loblaws, Sobeys, and Metro. A recent report by Environmental Defence Canada underscores that 71% of produce items from these retailers are packaged in plastic, which exacerbates plastic pollution.[5] Also, in the weapons and

3. See https://ccaf.io/cbnsi/cbeci/ghg and https://www.gate.io/learn/articles/60-bitcoin-mining-and-energy-consumption-statistics-for-2023-you-need-to-know/1265

4. Oligopoly markets also experience two-sided environmental externalities, particularly at waste treatment plants. Wastewater from household and industrial activities contributes to nitrogen, phosphorus, and microplastic pollution even after treatment (Sadia et al., 2022). See the following for more details on the pharmaceutical pollution: https://eeb.org/the-problem-of-pharmaceutical-pollution/

5. See https://environmentaldefence.ca/2023/04/18/plastic-packaging-in-grocery-stores/ and https://www.statista.com/statistics/481019/leading-grocery-retailers-by-market-share-canada/

military vehicles industry, numerous polluting producers supply a limited number of final buyers: governments. On average, between 1.5 and 20 tons of lead shots and bullets are consumed annually, which, in turn, contributes to lead pollution.[6]

Our work makes three main contributions. First, it introduces, to the emissions taxation literature, the concept of two-sided abatement in the context of two-sided environmental externalities, which has been unexplored so far, except very recently by Abdul Baki et al. (2024). Second, it is among the few studies to model market structures when final consumers engage in polluting activities, which is an understudied case of pure market structure in the Industrial Organization literature. Third, it contributes to the emission taxation literature by explicitly modeling the green preferences of the polluting consumers, even if they act as price takers.

Regarding the first contribution, existing literature on emission taxes predominantly addresses “one-sided environmental externalities" (see, e.g., Ino and Matsumura, 2021; Elnaboulsi, Daher, and Sağlam, 2023), also referred to as “Pigouvian externalities" (Huber and Wirl, 1998). Pigou (1920) posited that the tax imposed on price takers should coincide with their marginal social damage (i.e., marginal pollution damage). However, in the case of a polluting monopoly, the Pigouvian tax becomes socially suboptimal (Buchanan, 1969), and the optimal tax is set below the Pigouvian level to balance social welfare losses trade-off between market power and pollution (Barnett, 1980). This result also extends to Cournot oligopoly scenarios (Ebert, 1991) and analogous cases of input, but not final product, oligopsony (Requate, 2006), yet not the case of two-sided externalities. Indeed, Abdul Baki, Benchekroun, and Marrouch (2024) are the first to find that the tax set on the price maker (taker) is precisely (below) the Pigouvian level under a two-sided pollution scenario. While Abdul Baki, Benchekroun, and Marrouch (2024) focus solely on the case of production-sided abatement, our present study stands out by incorporating abatement measures on both the producer and consumer sides. This extension provides insights into whether the findings of Abdul Baki, Benchekroun, and Marrouch (2024), and so of Pigou (1920) and Barnett (1980), are applicable within our generalized framework.

Moving to the second contribution, various studies, such as Ren, Chen, and Hu (2020), have explored consumption-based externalities. Nonetheless, this literature often penalizes producers for the consumers’ emissions, thereby challenging the PPP. Importantly, many studies have highlighted the presence of consumption-sided emissions taxes in both first-best (Pearce and Turner, 1992) and second-best contexts (Sandmo,

6. See the U.S. Geological Survey within the Department of the Interior

1975; Sandmo, 2011). However, in that first-best world, the Pigouvian principle is assumed, only consumers are permitted to pollute, and the tax is tailored specifically to the packaging industry, limiting its generalizability to other sectors.[7] As for that second-best context, the presence of multiple distortionary taxes is more characteristic than market power. Requate (2006) is among the few studies allowing consumers to exercise market power over green and brown inputs. A major limitation of this study is its exclusion of abatement efforts, which are systematically incorporated into our research. Moreover, besides being rooted in Pigouvian externalities, Requate (2006) overlooks cases where consumers can influence the price of a final, rather than intermediate, product. While Abdul Baki, Benchekroun, and Marrouch (2024) delve into double-sided Pigouvian externalities, their analysis confines market power exclusively to producers, not consumers. This gap is bridged in our work by allowing oligopsony power to be exerted over a final product, whose production and consumption are polluting, which covers additional cases of relevant unregulated pollution scenarios.

As for the third contribution, the literature has been relatively hesitant to model consumer green preferences explicitly. While several studies have explored how green preferences of producers (Villena and Quinteros, 2024) and consumers (Abdul Baki and Marrouch, 2022) shape the stringency of emission taxation, comprehensive models of consumer externalities and green preferences remain scarce. Thus, although we are not the first to consider consumers' green preferences, we are among the few to explicitly model the behavior of consumers, particularly when they act as polluters. For instance, Abdul Baki and Marrouch (2022) incorporate green preferences in consumer utility but assume that consumers do not generate pollution. This overlooks a key distinction: green preferences may differ significantly depending on whether or not the consumer is a polluter. Our work addresses this gap by modeling consumer heterogeneity along that dimension. Similarly, Ebert and Von Dem Hagen (1998) assume that producers' emissions influence consumers' preferences, disregarding scenarios where consumption itself generates pollution that could, in turn, impact consumers' preferences, such as with coal for home heating, kerosene for lighting, and firewood for cooking. We incorporate this overlooked feedback loop in our framework, as it aligns with real-life dynamics. For instance, residential coal burning was a key contributor to London's Great Smog in 1952 that led to 10000-12000 deaths (Bell, Davis, and Fletcher, 2008; Stone,

7. Pearce and Turner (1992) assumes a perfectly competitive market where only consumers of packaged products contribute to pollution. The proposed tax is not applicable beyond the packaging sector as it relies on "packaging-specific" parameters, including marginal disposal costs, liter and landfill user costs, container weight per liter, and recycling costs, among others.

2002), urging households to switch to cleaner heating sources with the aid of the funds initiated under the Clean Act of 1956.[8] Likewise, households in developing countries are replacing kerosene lamps with other cost-efficient alternatives such as LED lights, solar, or solar-hybrid solutions, particularly that around 4.3 million deaths are linked with kerosene lamps, as declared by the World Health Organization.[9] Accordingly, the subtle distinction in our present work lies in explicitly modeling consumers' preferences for a green environment when consumers and producers pollute.

This paper is also pertinent to scenarios of double-sided pollution occurring in vertical markets. While horizontal markets can be viewed as abstract or simplified versions of vertical relationships, our framework provides a generalization that, to our knowledge, has not yet been explored in this sub-literature. Surprisingly, most existing studies assume that only one side of the market—either upstream or downstream—generates pollution but not both (see e.g. Sugeta and Matsumoto, 2007; Canton, Soubeyran, and Stahn, 2008; Yin et al., 2020; Zhu, Yue, and Chen, 2022; Chang and Sellak, 2023). Mansur (2011) is among the few exceptions, allowing both upstream and downstream firms to pollute. Nonetheless, Mansur (2011) imposes the Pigouvian principle and regulates one pollution source only. On a related note, Cachon (2014) examines the retailer's store density decision, but not emission taxes, when both the retailer and consumers pollute. Broadly speaking, the literature on vertical markets does not adequately address the concept of double-sided emission taxes when both double-sided externalities and abatement are present. This imbalance in the modeling of consumption-side pollution within vertically structured markets has contributed to a broader modeling bias in environmental economics, where consumer externalities and preferences (and so abatement efforts) are often underrepresented or simplified.

The results of this paper show that when both polluters, consumers and producers, abate, the optimal taxes imposed on both abating polluters fall below their respective marginal pollution damage. In this case, these taxes can only achieve the second-best allocation because the two market failures, pollution and market power, are linked. As

8. Another example is related to agricultural products. The consumption of agricultural products may lead to increased waste litter, while their production can result in the eutrophication of lakes and contamination of groundwater due to the use of pesticides, leading to negative health effects (Bertero et al., 2020). In 2020, at least one pesticide was found to exceed permissible levels at around 22 percent of monitoring sites across European rivers and lakes. See https://www.eea.europa.eu/publications/how-pesticides-impact-human-health

9. https://www.ccacoalition.org/news/kerosene-lamps-are-important-target-reducing-indoor-air-pollution-and-climate-emissions

a result, these instruments seek a second-best trade-off between internalizing the (two-sided) market failures: pollution and market power. In other words, the second-best trade-off has two layers: between and within the two market failures. However, if one of the polluters (either producers or consumers) abates, then the tax imposed on the (non) abating polluters is (below) equal to the Pigouvian level. These taxes can achieve the first-best in this scenario since the two market failures are unlinked. Similarly, the first-best is achieved if neither party abates, yet the taxes imposed on producers and consumers are not unique. All these results are orthogonal to the source of market power, whether it stems from the production or the consumption side.[10]

The rest of the paper unfolds as follows. Section 2 reviews the literature. Section 3 outlines the model. Section 4 presents the optimal emission taxes. Sections 5 and 6 discuss the results. Lastly, section 7 concludes.

2.2 Related Literature

The PPP has garnered adoption across both developed and developing nations (Luken, 2009; Glazyrina, Glazyrin, and Vinnichenko, 2006), spanning diverse sectors including agriculture, pharmaceuticals, energy, and packaging (Pearce and Turner, 1992;Tobey and Smets, 1996; Rahman and Edwards, 2004; Malmqvist et al., 2023). For instance, in the (coal-based) power generation industry, "institutions often attempt to identify fully liable parties instead of questioning involved parties' roles in causing the problem. This is often called the [PPP] in which generating companies could be made to pay environmental costs through assignment and enforcement of full liability" (Rahman and Edwards, 2004). The PPP's adoption manifests through material levies in the packaging industry. "The size of the levy needs to be related directly to the environmental damage done by the production and consumption of the packaging, or to the costs of restoration of the environment. This is important since the proportional link between the size of the charge and damage done is fundamental to the [PPP]" (Pearce and Turner, 1992). This overview highlights the relevance of identifying the polluter to

10. Consistent with Abdul Baki, Benchekroun, and Marrouch (2024), this paper emphasizes the limitations of applying Pigou (1920)'s and Barnett (1980)'s principles universally when moving from a one-sided to a two-sided environmental externality framework. Nonetheless, the findings of Abdul Baki, Benchekroun, and Marrouch (2024) are specific to scenarios involving production-sided abatement alone and do not extend to cases where consumers (along with producers) adopt in abatement technologies, or even to cases where neither producers nor consumers can abate. It also uncovers hidden insights regarding the interpretation of first- vs second-best world.

hold it liable for internalizing the environmental externality- an objective lying at the heart of the PPP.

A substantial number of studies delve into the different economic instruments available for applying the PPP. Among the various tools, emissions taxes have gained a wide range of applications (Bovenberg and De Mooij, 1997; Hoel, 1998; OECD, 2001; Millock, Nauges, and Sterner, 2004; Isaksson, 2005; Blackman and Harrington, 2018). Emissions taxes trace their origins to Pigou (1920), who showed that in an ideal competitive scenario, the emission tax is equal to the marginal social damage of pollution. Thus, identifying the polluter becomes not only an objective of the PPP, but is also in line with the spirit of the Pigouvian tax because it calls to internalizing the full marginal social damage of pollution. Nevertheless, subsequent research has highlighted the suboptimality of the Pigouvian tax upon moving from a first to a second-best world (Buchanan, 1969). Equating the emission tax with the marginal pollution damage, as Pigou (1920) concludes, ceases to hold in the presence of multiple market failures. Barnett (1980) reaffirms this result and proves that the monopolist's optimal emission tax falls below their marginal pollution damage to seek a second-best compromise between the distortions resulting from market power and pollution. This result holds under Cournot oligopoly as well (Ebert, 1991). However, when the monopolist's emissions affects consumers, the optimal emission tax may exceed the Pigouvian level depending on the strategic relation between consumption and the monopolist's emissions (Ebert and Von Dem Hagen, 1998).

Deviating from the Pigouvian tax also happens when abatement is integrated into the model. In the Pigouvian framework (i.e., perfect competition among polluters), David and Sinclair-Desgagné (2005) show that the optimal emission tax, which is now modulated by the demand elasticity of abatement, exceeds the Pigouvian tax in the presence of an imperfectly competitive eco-industry. As for the Buchanan-Barnett context (i.e., Cournot oligopoly), the optimal tax may exceed or coincide with the Pigouvian tax depending on the relative market power of polluters and eco-friendly firms (Nimubona and Sinclair-Desgagné, 2013). The literature has later revealed other factors modulating emissions taxes, such as the number of polluting producers entering the market (Schoonbeek and Vries, 2009; Estay and Stranlund, 2022), the spatial (Marrouch and Sinclair-Desgagné, 2012; Abdul Baki and Marrouch, 2022) and intertemporal (Benchekroun and Van Long, 1998; Benchekroun and Van Long, 2002) dimensions of pollution.

While there is an extensive body of literature exploring the interplay between emissions taxes and market structure (see for e.g., Simpson, 1995; Yin, 2003; Sartzetakis and Tsigaris, 2005; David, Nimubona, and Sinclair-Desgagné, 2011), two prominent gaps emerge: the source of (1) market power and (2) environmental externality are usually ignored. Regarding market power, it is often assumed to be wielded by producers, whereas, in reality, consumers could also exercise market power. For example, municipalities or governments (i.e., consumers) can influence prices in various sectors such as weaponry, military vehicle manufacturing, telecommunications, and the privatized electricity sector. Similarly, market power is active on the demand side of air taxi and charter services (consumers are usually affluent individuals). Moving to the environmental externality, producers have been usually portrayed as the primary polluters despite consumers being also engaged in environmentally detrimental activities. Relevant examples include waste management, steel production, pharmaceutical, weaponry, agrochemical, and agricultural industries, among others. In this context, Abdul Baki, Benchekroun, and Marrouch (2024) are the first to allow pollution to be exerted by both producers and consumers. Under this case, the optimal tax on the polluting monopolist (consumers) coincides with (falls below) their marginal pollution damage, thus contradicting Pigou (1920) and Barnett (1980).

Prior research has identified the prospect of consumers themselves acting as the source of pollution, albeit not in the presence of market power. Pearce and Turner (1992) distinctly recognize consumption externalities within the packaging industry, leading them to propose a packaging tax/levy aligned with the Pigouvian principle. Nonetheless, the work of Pearce and Turner (1992) encounters three major limitations. First, the study imposes a first-best scenario, and so the Pigouvian principle, through its silence about the market structure. Second, the study considers the environmental externality to be exerted by either producers or consumers yet not both of them. This assumption does not apply to the meat packaging industry, for example (see e.g., Muth and Wohlgenant, 1999; Treich, 2021). Third, the study's findings cannot be extrapolated to sectors beyond the packaging industry since the proposed tax is contingent on "packaging-specific" parameters (e.g., marginal landfill user costs, weight per liter of container, recycling cost, etc).

The concept of consumption externality has been taken into account not only in an ideal first-best scenario but also within a second-best context. Nevertheless, in the majority of the studies (see e.g., Sandmo, 1978; Sandmo, 2011), the second-best world is often characterized by the existence of multiple distortionary taxes. Sandmo (1975) is among the first to view Pigouvian taxes as just one element within a broader system of commodity taxation. Cremer, Gahvari, and Ladoux (1998) challenges Sandmo (1975)'s study by accounting for an income tax and finds that the optimal tax on the polluting commodity deviates from the Pigouvian tax. Henceforth, most studies highlighting the presence of a consumption-sided externality are silent about market power. The study by Abdul Baki, Benchekroun, and Marrouch (2024) is among the few studies to consider market power; however, it is exerted by producers but not consumers. In this case, the optimal taxes will achieve the first-best allocation.

Interestingly, Requate (2006) is the first to derive an optimal emission tax when consumers hold market power. A main limitation is that his study applies to firms–consumers of inputs–but not consumers of final products. Specifically, Requate (2006) considers a firm exercising monopsony power over two inputs, brown and green, and generalizes the model to the case of oligopsony. In this context, the tax is in line with that of Barnett (1980). Our main criticism here is in consumer's market power being exerted over an input rather than a final product. Accordingly, the study of Requate (2006) dismisses many incidents where market power is held by final consumers, rather than intermediate users, as previously illustrated. Moreover, Requate (2006) restricts his analysis to the case of Pigouvian externalities, thus ignoring the possibility of having two active sources of pollution.

In summary, this overview underscores the presence of enduring unresolved challenges in the literature on emission taxes. Namely, the sources of pollution (and hence abatement) and market power are overlooked. Even in the study accounting for the source of pollution, abatement technologies are adopted by producers but not consumers (Abdul Baki, Benchekroun, and Marrouch, 2024). Therefore, the existing literature fails to address incidents characterized by double-sided environmental externalities and double-sided abatement efforts, when either consumers or producers control the price. This present study aims at bridging these gaps in a unified framework.

2.3 The Model

Consider a partial equilibrium model formed of n homogeneous producers and m homogeneous consumers. Suppose there is a final consumption product whose production and consumption generate pollution.

Let $U(q_i^c, a_i^c)$ represent consumer i's preference $(1 \leq i \leq m)$, where q_i^c is the quantity consumed and a_i^c is the abatement level. Assume that $\frac{\partial U}{\partial q_i^c} > 0$, $\frac{\partial^2 U}{\partial q_i^{c2}} < 0$, and $\frac{\partial^2 U}{\partial a_i^{c2}} < 0$

(in line with Ebert and Von Dem Hagen, 1998). It is worth noting that consumer i's marginal utility can either increase or decrease with abatement, i.e., $\frac{\partial U}{\partial a_i^c}$ can be either positive or negative, depending on whether consumers exhibit green or brown preferences, respectively. Let $B(a_i^c)$ denote the abatement cost function of consumer i, where $\frac{dB}{da_i^c} > 0$. $A^c = \sum_{i=1}^m a_i^c$ denotes the total abatement level of consumers.

Producer i $(1 \leq i \leq n)$ face the following production cost function: $C(q_i^p, a_i{}^p)$, where q_i^p is the quantity produced and a_i^p is the abatement level. Suppose that $\frac{\partial C}{\partial q_i^p} > 0$, $\frac{\partial C}{\partial a_i{}^p} > 0$, $\frac{\partial^2 C}{\partial q_i^{p2}} > 0$, $\frac{\partial^2 C}{\partial a_i{}^{p2}} > 0$, and $\frac{\partial^2 C}{\partial q_i^p \partial a_i^p} > 0$. The latter cross derivative implies that production and abatement are substitutes, meaning that production and emissions are complements (Requate, 2006). Let $A^p = \sum_{i=1}^p a_i^p$ be the total abatement level of producers. At equilibrium, $\sum_{i=1}^n q_i^p = \sum_{i=1}^m q_i^c = Q$, i.e., the total quantity produced should be equal to the total quantity consumed.

Let $E^p(q_i^p, a_i{}^p)$ and $E^c(q_i^c, a_i{}^c)$ be the emission functions of the i^{th} producer and consumer, respectively. Assume $\frac{\partial E^j}{\partial q_i^j} > 0$ and $\frac{\partial E^j}{\partial a_i^j} < 0$ for $j \in \{p, c\}$. Let $\overline{E^p} = \sum_{i=1}^n E^p(q_i^p, a_i^p)$ and $\overline{E^c} = \sum_{i=1}^m E^c(q_i^c, a_i^c)$ denote the producers' and consumers' total emissions levels, respectively. Therefore, the environmental externality is captured by the damage function $D\left(\overline{E^p}, \overline{E^c}\right)$, where $\frac{\partial D}{\partial \overline{E^j}} > 0$ for $j \in \{p, c\}$.[11] In a business-as-usual scenario, both producers and consumers ignore their environmental damages.

To internalize the double-sided environmental externality, $E^p(q_i^p, a_i^p)$ and $E^c(q_i^c, a_i^c)$ are subject to emissions taxes t^p and t^c, respectively. A benevolent regulator aims to find optimal emission taxes by maximizing the social welfare function $SW(t^p, t^c)$, which is additively separable into the producers' and consumers' surpluses, as well as the pollution damage.

Let $\tilde{p}$ denote the purchase price of the product. This price is either set by producers or consumers, depending on whether market power is active on the supply or demand side. Accordingly, we derive the optimal emissions taxes under two scenarios: production-sided (Cournot oligopoly) and consumption-sided (Cournot oligopsony) market power. In either case, the objective of a benevolent regulator is to choose t^p and t^c that maximize the social welfare function, which is the sum of the producer and consumer surpluses, minus the environmental damage (i.e., social disamenity of pollution)

11. We remain agnostic about the nature of the damage function. It can be additive or multiplicative. This does not change our results.

2.3.1 Production-sided Market Power

We first consider the case where market power arises on the production side—i.e., a few large firms supply a polluting good to a homogeneous mass of consumers. This setup applies to industries such as oil extraction, airlines, or automobile manufacturing. The goal is to characterize how the presence of oligopolistic power alters the optimal emissions tax relative to the competitive benchmark. Under this scenario, the polluting product is supplied by an oligopolist ($1 \leq i \leq n$) and consumed by a mass of homogeneous consumers ($m \to \infty$), who cannot influence the price p (i.e., a few producers, many consumers). However, given p set by producers, consumers can choose their total levels of consumption Q and abatement A^c. Thus, consumers aim to maximize $U(Q, A^c) - p.Q - B(A^c)$, which results in an inverse market demand $P^d(Q, A^c)$, with $\frac{\partial P^d}{\partial Q} < 0$ and $\frac{\partial^2 P^d/\partial Q^2}{\partial P^d/\partial Q} Q > -1$ (in line with Requate, 2006). $\epsilon_d = \frac{\partial Q}{\partial P^d}\frac{P^d}{Q} < 0$ is the price elasticity of demand.

Producer $1 \leq i \leq n$ chooses the equilibrium quantity (q_i^p) and abatement (a_i^p) level that maximize its profits subject to, t^p,

$$\pi(q_i^p, a_i^p) = P^d(Q, A^c)q_i - C(q_i^p, a_i^p) - t^p E^p(q_i^p, a_i^p).$$

The producer's profit maximization yields the usual Cournot first-order conditions (FOCs), equating marginal revenue with the total marginal cost of production,

$$P^d(Q, A^c) = \frac{\partial C}{\partial q_i^p} + t^p \frac{\partial E^p}{\partial q_i^p} - \frac{\partial P^d}{\partial Q} q_i^p, \tag{2.1}$$

as well as

$$\frac{\partial C}{\partial a_i^p} = -t^p \frac{\partial E^p}{\partial a_i^p}. \tag{2.2}$$

The second condition (i.e., Eq. 2.2) indicates that the producer abates emission until the marginal cost of abatement is equal to the marginal benefit of abatement (i.e., tax-induced marginal saving). Eq. 2.1 and 2.2 imply that at equilibrium, $q_i^p = \frac{Q}{n}$and $a_i^p = \frac{A^p}{n}$, respectively.

This equilibrium quantity produced Q should also maximize the surplus of the price takers (i.e., the mass of consumers), subject to t^c, for the market to clear; that is, the marginal utility of consumption should be equal to the consumers' total marginal expense. Moreover, the mass of consumers tends to choose their total abatement level A^c by equating their marginal utility of abatement to their marginal expense of abatement.

Accordingly, the mass of consumers' surplus is given by:

$$CS(Q, A^c) = U(Q, A^c) - \tilde{p}Q - t^c E^c(Q, A^c) - B(A^c).$$

Substituting p with Eq. 2.1, the corresponding FOCs associated with the consumers' maximization problem are

$$\frac{\partial U}{\partial Q} = \frac{\partial C}{\partial q_i^p} + t^p \frac{\partial E^p}{\partial q_i^p} + t^c \frac{\partial E^c}{\partial Q} - \frac{\partial P^d}{\partial Q} q_i^p.$$

and

$$\frac{\partial U}{\partial A^c} - \frac{dB}{dA^c} = t^c \frac{\partial E^c}{\partial A^c}.$$

Maximizing the following social welfare function

$$SW(t^p, t^c) = U(Q, E^c) - B(A^c) - nC(q_i^p, a_i^p) - D\left[\overline{E^p}, E^c(Q, A^c)\right]$$

with respect to t^p and t^c gives

$$t^p = \frac{\partial D}{\partial \overline{E^p}} + \frac{1}{n} \frac{\alpha^p}{\epsilon_d} \frac{dQ}{dt^p} \tag{2.3}$$

and

$$t^c = \frac{\partial D}{\partial E^c} + \frac{1}{n} \frac{\alpha^c}{\epsilon_d} \frac{dQ}{dt^c}, \tag{2.4}$$

where

$$\alpha^p = P^d(Q, A^c) \left(\frac{\partial E^p}{\partial q_i^p} \frac{dQ}{dt^p} + \frac{\partial E^p}{\partial a_i^p} \frac{dA^p}{dt^p} + \frac{\frac{\partial E^p}{\partial a_i^p} \frac{dA^p}{dt^p}}{\frac{\partial E^c}{\partial A^c} \frac{dA^c}{dt^c}} \frac{\partial E^c}{\partial Q} \frac{dQ}{dt^c} \right)^{-1}$$

and

$$\alpha^c = P^d(Q, A^c) \left(\frac{\partial E^c}{\partial Q} \frac{dQ}{dt^c} + \frac{\partial E^c}{\partial A^c} \frac{dA^c}{dt^c} + \frac{\frac{\partial E^c}{\partial A^c} \frac{dA^c}{dt^c}}{\frac{\partial E^p}{\partial a_i^p} \frac{dA^p}{dt^p}} \frac{\partial E^p}{\partial q_i^p} \frac{dQ}{dt^p} \right)^{-1}.$$

2.3.2 Consumption-sided Market Power

Our goal now is to investigate whether the taxes presented in Eq. 2.3 and 2.4 would still hold when market power is no longer held by producers, but rather by consumers (i.e., Cournot oligopsony). Specifically, let us focus on the case where the final polluting product is purchased by consumers who possess market power. This situation commonly arises with large retailers as well as in public sector procurement of infrastructure or vehicles. In such contexts, the polluting product is now demanded by an oligopsonist ($1 \leq i \leq m$), but it is supplied by a mass of homogeneous producers who cannot influence the price ($n \to \infty$). Accordingly, the price of the product p is now set by consumers (i.e., a few consumers, many producers). The mass of producers determines their total levels of production Q and abatement A^p in a business-as-usual scenario by maximizing their surplus $p.Q - C(Q, A^p)$. Profit maximization generates the market supply $P^s(Q)$, where $\frac{dP^s(Q)}{dQ} > 0$ and $\frac{d^2P^s(Q)}{dQ^2} > 0$. The price elasticity of supply is denoted by $\epsilon_s = \frac{dQ}{dP^s}\frac{P^s}{Q} > 0$

Under oligopsony, consumer $1 \leq i \leq m$ determines the equilibrium quantity (q_i^c) (and thus price) and abatement (a_i^c) levels by maximizing their surplus, which is subject to the emission tax t^c:[12]

$$CS(q_i^c, a_i^c) = U(q_i^c, a_i^c) - P^s(Q)q_i^c - t^c E^c(q_i^c, a_i^c) - B(a_i^c).$$

The resulting FOCs are

$$P^s(Q) = \frac{\partial U}{\partial q_i^c} - t^c \frac{\partial E^c}{\partial q_i^c} - \frac{dP^s}{dQ} q_i^c \tag{2.5}$$

and

$$\frac{\partial U}{\partial a_i^c} = t^c \frac{\partial E^c}{\partial a_i^c} + \frac{dB}{da_i^c}. \tag{2.6}$$

It follows from Eq. 2.5 and 2.6 that $q_i^c = \frac{Q}{m}$ and $a_i^c = \frac{A^c}{m}$, respectively.

For the market-clearing condition, the equilibrium quantity consumers demand should, by default, maximize the profits of the mass of producers, who can control their total abatement levels. The profits are denoted by the following equation.

$$\pi(Q, A^p) = \tilde{p}.Q - C(Q, A^p) - t^p E^p(Q, A^p)$$

Replacing p with Eq.2.5, the FOCs are

$$\frac{\partial C}{\partial Q} = \frac{\partial U}{\partial q_i^c} - t^p \frac{\partial E^p}{\partial Q} - t^c \frac{\partial E^c}{\partial q_i^c} - \frac{dP^s}{dQ} q_i^c$$

and

12. See Appendix for more details on the calculations.

$$\frac{\partial C}{\partial A^p} = -t^p \frac{\partial E^p}{\partial A^p}.$$

Finally, the emissions taxes t^p and t^c are chosen by maximizing the social welfare function

$$SW(t^p, t^c) = m\left[U(q_i^c, a_i^c) - B(a_i^c)\right] - C(Q, A^p) - D\left[E^p(Q, A^p), \overline{E^c}\right].$$

The optimal emission taxes are

$$t^p = \frac{\partial D}{\partial E^p} - \frac{\beta^p}{m\epsilon_s}\frac{dQ}{dt^p} \tag{2.7}$$

and

$$t^c = \frac{\partial D}{\partial \overline{E^c}} - \frac{\beta^c}{m\epsilon_s}\frac{dQ}{dt^c}, \tag{2.8}$$

where

$$\beta^p = P^s(Q)\left(\frac{\partial E^p}{\partial Q}\frac{dQ}{dt^p} + \frac{\partial E^p}{\partial A^p}\frac{dA^p}{dt^p} + \frac{\frac{\partial E^p}{\partial A^p}\frac{dA^p}{dt^p}}{\frac{\partial E^c}{\partial a_i^c}\frac{dA^c}{dt^c}}\frac{\partial E^c}{\partial q_i^c}\frac{dQ}{dt^c}\right)^{-1}$$

and

$$\beta^c = P^s(Q)\left(\frac{\partial E^c}{\partial q_i^c}\frac{dQ}{dt^c} + \frac{\partial E^c}{\partial a_i^c}\frac{dA^c}{dt^c} + \frac{\frac{\partial E^c}{\partial a_i^c}\frac{dA^c}{dt^c}}{\frac{\partial E^p}{\partial A^p}\frac{dA^p}{dt^p}}\frac{\partial E^p}{\partial Q}\frac{dQ}{dt^p}\right)^{-1}.$$

2.4 The Pigouvian (1920) vs Barnett's (1980) Principle

2.4.1 Two-sided Environmental Externality

When both producers and consumers pollute, there are three cases: both abate, either one abate, and neither one abate.

2.4.1.1 Two-sided Abatement

When both polluters, producers and consumers, are abating, the optimal emission taxes are presented by eq. 2.3 and 2.4, when producers hold market power, and Eq. 2.7 and 2.8, when consumers hold market power, respectively. In either case, assuming $\frac{dQ}{dt^j} < 0$ and $\frac{dA^j}{dt^j} > 0$ for $j \in \{p, c\}$, all the optimal emission taxes fall under their corresponding Pigouvian level, that is $t^j < \frac{\partial D}{\partial E^j}$. This is true, irrespective of the source of market power.[13] Thus, the regulator is less stringent with the abating price makers and takers. While this result is in line with Barnett (1980) and Requate (2006) for the price makers (case of oligopoly and oligopsony, respectively), it is still in opposition with Pigou (1920) for the price takers. The former (Barnett, 1980; Requate, 2006) implies that the tax set on the price maker should be below its marginal pollution damage irrespective of abatement, whereas the latter (Pigou, 1920) reveals that the emission tax set on the price taker should be equal to its marginal pollution damage. Moreover, our finding is in contradiction with Abdul Baki, Benchekroun, and Marrouch (2024), who consider the case of two-sided environmental externality where only the monopolist can abate. In their context, the authors show that the regulator is more stringent with the abating monopolist and is less stringent with the non-abating price takers. Thus, we reveal the following proposition.

Proposition 1 *In the presence of a two-sided environmental externality, two-sided abatement, and one-sided market power, the optimal emission tax imposed on both abating polluters falls below their marginal pollution damage.*

2.4.1.2 One-sided Abatement

Production-sided Abatement: When abatement technologies are adopted by producers only but not consumers, then $a_i^c = 0$ for $(1 \leq i \leq m)$. In this case, the optimal emissions taxes under oligopoly become

$$t^p = \frac{\partial D}{\overline{E^p}} \tag{2.9}$$

and

$$t^c = \frac{\partial D}{E^c} + \frac{1}{n}\frac{P^d(Q)}{\epsilon_d}\left(\frac{dE^c}{dQ}\right)^{-1}. \tag{2.10}$$

Thus, t^p and t^c as in Eq. 2.9 and 2.10 are inline with the taxes derived by Abdul Baki, Benchekroun, and Marrouch (2024). Specifically, t^p coincides with the producers' marginal pollution damage, while t^c falls below the consumers' marginal damage.[14]

If we were to follow the results of Abdul Baki, Benchekroun, and Marrouch (2024), then we would expect the regulator to be more stringent with the unabating consumers

13. $\frac{dQ}{dt^j} < 0$ and $\frac{dA^j}{dt^j} > 0$ for $j \in \{p, c\}$, then $\alpha^j < 0$, $\beta^j < 0$ and $t^j < \frac{\partial D}{\partial E^j}$.

14. This result is not driven by the assumption that the total pollution damage is an implicit function of equilibrium quantity: $D\left[E^p(Q), E^c(Q)\right]$. In other words, all our results continue to hold if damage depends explicitly on the total equilibrium quantity, that is $D(Q)$

when these consumers (but not abating producers) hold market power. Surprisingly, this intuition breaks down when we consider the case of oligopsony, provided that only producers (in this case, price takers) abate. Specifically, the optimal emissions taxes become

$$t^p = \frac{\partial D}{E^p} \tag{2.11}$$

and

$$t^c = \frac{\partial D}{\partial \overline{E^c}} - \frac{1}{m}\frac{P^s(Q)}{\epsilon_s}\left(\frac{dE^c}{dq_i^c}\right)^{-1}. \tag{2.12}$$

The price-takers (abating producers) are now taxed at their marginal pollution damage, whereas the price-makers (unabating consumers) are now taxed below their respective marginal damage. Thus, t^p and t^c, presented in Eq. 2.11 and 2.12, are now in resonance with Pigou (1920) and Requate (2006) (the mirror case of Barnett, 1980), respectively, yet are in opposition to Abdul Baki, Benchekroun, and Marrouch (2024). A broader look at the results could imply that, in the presence of one source of market power and abatement, the taxation stringency does not depend on the source of market power, but rather the source of abatement.

Consumption-sided Abatement: In the absence of producers' abatement effort $\left(a_i^p = 0, 1 \leq i \leq n\right)$, the optimal emissions taxes under oligopoly are

$$t^p = \frac{\partial D}{\partial \overline{E^p}} + \frac{1}{n}\frac{P^d(Q, A^c)}{\epsilon_d}\left(\frac{dE^p}{dq_i^p}\right)^{-1} \tag{2.13}$$

and

$$t^c = \frac{\partial D}{\partial E^c}. \tag{2.14}$$

t^p and t^c in Eq. 2.13 and 2.14 resonate with Barnett (1980) and Pigou (1920), respectively, yet they contradict Abdul Baki, Benchekroun, and Marrouch (2024) as the tax set on the price makers (takers), which are in this case producers (consumers) falls below (is equal to) their respective Pivouvian level. This result is reversed when the abating consumers hold market power. Specifically, under oligopsony, the tax set on the unabating producers (price takers) is

$$t^p = \frac{\partial D}{\partial E^p} - \frac{1}{m}\frac{P^s(Q)}{\epsilon_s}\left(\frac{dE^p}{dQ}\right)^{-1}, \tag{2.15}$$

whereas the tax set on the abating price makers (consumers) is

$$t^c = \frac{\partial D}{\partial \overline{\overline{E^c}}}, \tag{2.16}$$

both, Eq. 2.15 and 2.16, resonating with the mirror case of Abdul Baki, Benchekroun, and Marrouch (2024). Thus, the new insight we are revealing now is that in the presence of one-sided market power and two-sided environmental externality, the regulator is more stringent with the abating polluter (even if it does not hold market power) but is less stringent with the unabating polluter (even if it holds market power). Accordingly, the main result of Abdul Baki, Benchekroun, and Marrouch (2024) only holds for the specific case of production-sided abatement and market power. In summary, our results go beyond the principles established by Pigou (1920) and Barnett (1980). In the presence of two-sided environmental externalities, it is the nature of abatement, rather than market power, that determines the stringency of environmental regulation. Accordingly, the conclusions of Pigou (1920), Barnett (1980), and Abdul Baki et al. (2024) may no longer hold, depending on the source of abatement. This proposition summarizes our findings.

Proposition 2 *In the presence of a two-sided environmental externality, one-sided abatement, and one-sided market power, the optimal emission tax imposed on the abating (unabating) polluter is equal to (below) their marginal pollution damage, irrespective of whether this polluter holds market power.*

2.4.1.3 No Abatement

In the absence of abatement, $a_i^p = 0$ for $1 \leq i \leq n$ and $a_i^c = 0$ for $1 \leq i \leq m$. Accordingly, infinite combinations of t^p and t^c, exceeding or falling below the marginal pollution damage, are socially optimal under both oligopoly

$$t^p \frac{dE^p}{dq_i^p} + t^c \frac{dE^c}{dQ} = \frac{\partial D}{\partial \overline{\overline{E^p}}}\frac{dE^p}{dq_i^p} + \frac{\partial D}{\partial E^c}\frac{dE^c}{dQ} + \frac{1}{n}\frac{P^d(Q)}{\epsilon_d} \tag{2.17}$$

and oligopsony

$$t^p \frac{dE^p}{dQ} + t^c \frac{dE^c}{dq_i^c} = \frac{\partial D}{\partial E^p}\frac{dE^p}{dQ} + \frac{\partial D}{\partial \overline{\overline{E^c}}}\frac{dE^c}{dq_i^c} - \frac{1}{m}\frac{P^s(Q)}{\epsilon_s}. \tag{2.18}$$

This result is in opposition to Barnett (1980), Requate (2006), and Pigou (1920). Even when market power is absent from the supply $(n \to \infty)$ or demand $(m \to \infty)$ side, the Pigouvian principle fails to hold due to the infinite possibilities of optimal taxes: $t^p \frac{dE^p}{dQ} + t^c \frac{dE^c}{dQ} = \frac{\partial D}{\partial E^p} \frac{dE^p}{dQ} + \frac{\partial D}{\partial E^c} \frac{dE^c}{dQ}$. However, the total tax revenue in the presence (absence) of market power is still less than (equal to) the total marginal pollution damage, which is in the spirit of Barnett (1980) and Requate (2006) (Pigou, 1920). In conclusion, our results are summarized in the following proposition.

Proposition 3 *In the presence of a two-sided environmental externality and one-sided market power, if polluters do not abate, the optimal emissions taxes set on the price makers and takers are not unique and can exceed their corresponding marginal pollution damage.*

2.4.2 Special Cases

Assume that either producers or consumers generate pollution, that is one-sided environmental externality.[15]

Production-sided Environmental Externality

If pollution is generated by producers only, then $E^c(q_i^c, a_i^c) = 0$ with $1 \leq i \leq m$. Accordingly, under oligopoly, the optimal emissions taxes become

$$t^p = \frac{\partial D}{\partial \overline{E^p}} + \frac{1}{n} \frac{P^d(Q)}{\epsilon_d} \frac{dQ}{dt^p} \left(\frac{\partial E^p}{\partial q_i^p} \frac{dQ}{dt^p} + \frac{\partial E^p}{\partial a_i^p} \frac{dA^p}{dt^p} \right)^{-1}, \tag{2.19}$$

while

$$t^c = 0.$$

In this case, t^p in Eq. 2.19 is qualitatively in line with the tax derived by Ebert (1991), and hence Barnett (1980) as $n \to 1$ for the case of monopoly ($t^p < \frac{\partial D}{\partial E^p}$). Under perfect competition, $n \to \infty$, $t^p = \frac{\partial D}{\partial \overline{E^p}}$, which is the Pigouvian tax (Pigou, 1920).

Consumption-sided environmental Externality

$E^p(q_i^p, a_i^p) = 0$ for $1 \leq i \leq n$ when pollution is generated by consumers only. In this case, the optimal emissions taxes for the case of oligopsony are

15. In this subsection, we assume that the polluter is abating. However, our discussion still holds in the absence of abatement.

$$t^p = 0,$$

and

$$t^c = \frac{\partial D}{\partial \overline{E^c}} - \frac{1}{m}\frac{P^s(Q)}{\epsilon_s}\frac{dQ}{dt^c}\left(\frac{\partial E^c}{\partial q_i^c}\frac{dQ}{dt^c} + \frac{\partial E^c}{\partial a_i^c}\frac{dA^c}{dt^c}\right)^{-1}. \tag{2.20}$$

The tax set on polluting consumers, as in Eq. 2.20, resonates with that of Requate (2006), that is $t^c < \frac{\partial D}{\partial \overline{E^c}}$. Thus, the framework of Requate (2006), which defines the polluter to be a producer, can be perceived as a special case of a consumption-sided environmental externality that is exerted over an input and not a final product. In line with Requate (2006), in the absence of oligopsony power, $m \to \infty$, then the Pigouvian tax is achieved: $t^c = \frac{\partial D}{\partial \overline{E^c}}$.

2.5 The Stringency of the Environmental Regulation

In our context, the new insight we are revealing is that the stringency of the emissions taxes under two-sided environmental externality and one-sided market power does not depend on the source of market power but rather on the source of abatement. To understand the mechanism behind the taxation stringency, it is sufficient to examine the measure of market power, which we denote by Δ. We differentiate between two measures: the producers' price markup for the case of oligopoly (Δ^+) and the consumers' price markdown for the case of oligopsony (Δ^-). Market power is active via ϵ, the elasticity of demand under oligopoly ϵ_d or the elasticity of supply under oligopsony ϵ_s

2.5.1 Two-sided Environmental Externality

Suppose that both, the price makers and takers, pollute.

2.5.1.1 Two-sided Abatement

When both the price makers and price takers abate, the measure of market power are

$$\Delta^+ = \frac{\partial D}{\partial E^c}\frac{dE^c}{dQ} - \left(-\frac{\alpha^c}{n\epsilon_d}\frac{dQ}{dt^c}\frac{\partial E^c}{\partial Q}\right) + \left(-\frac{P^d(Q, A^c)}{n\epsilon_d}\right), \tag{2.21}$$

for the case of oligopoly and

$$\Delta^- = \frac{\partial D}{\partial E^p}\frac{dE^p}{dQ} - \frac{\beta^p}{m\epsilon_s}\frac{dQ}{dt^p}\frac{\partial E^p}{\partial Q} + \frac{P^s(Q)}{m\epsilon_s}, \tag{2.22}$$

for the case of oligopsony.

Market power ϵ_d (or ϵ_s) manifests in the price markup (or markdown) via two channels. The first channel is explicit , that is the part $-\frac{1}{n}\frac{P^d(Q,A^c)}{\epsilon_d}$ under oligopoly and $\frac{1}{m}\frac{P^s(Q)}{\epsilon_s}$ under oligopsony. The second channel is implicit, via $-\alpha^c > 0$ under oligopoly and $\beta^p < 0$ under oligopsony. Therefore, the first channel implies that the tax on the abating price maker is less than the Pigouvian level (t^p in Eq. 2.3 and t^c in Eq. 2.8). On the other hand, the second channel mandates that the tax on the abating price taker is less than the Pigouvian level (t^c in Eq. 2.4 and t^p in Eq. 2.7).

2.5.1.2 One-sided Abatement

Price Makers Abate: Assume that only the price makers abate, but not the price takers. In this case, the measure of market power is

$$\Delta^+ = \frac{\partial D}{\partial E^c}\frac{dE^c}{dQ}, \tag{2.23}$$

for the case of oligopoly and

$$\Delta^- = \frac{\partial D}{\partial E^p}\frac{dE^p}{dQ}, \tag{2.24}$$

for the case of oligopsony.

Surprisingly, ϵ does not influence Δ regardless of whether market power is active on the demand or supply side (see Eq. 2.23 and 2.24). As a result, the regulator can be more stringent with the price makers (t^p in Eq. 2.9 and t^c in Eq. 2.16). However, the unabating price takers are taxed at a level below the Pigouvian (t^p in Eq. 2.10 and t^c in Eq. 2.15) to internalize the price takers' environmental externality while bearing in mind that this will happen only through consuming (producing) less under oligopoly (oligopsony). Otherwise, quantity consumed (produced) and thus Δ collapses to zero.

Price Takers Abate: If the price takers abate only, the regulator becomes less stringent with the unabating price makers and more stringent with the abating price takers, as in

Barnett (1980), Requate (2006), and Pigou (1920). Under this scenario, the measure of market power under oligopoly is

$$\Delta^{+} = \frac{\partial D}{\partial E^{c}} \frac{dE^{c}}{dQ} + \left(-\frac{P^{d}(Q, A^{c})}{n\epsilon_{d}} \right), \tag{2.25}$$

and

$$\Delta^{+} = \frac{\partial D}{\partial E^{c}} \frac{dE^{c}}{dQ} + \left(\frac{P^{s}(Q)}{m\epsilon_{s}} \right) \tag{2.26}$$

under oligopsony.

In this case, price makers rely on the elasticity of demand/supply (oligopoly/oligopsony) to mark up/down the prices, as seen through ϵ in Δ (Eq. 2.25 and 2.26). Subsequently, the price makers' optimal emissions taxes fall below their respective Pigouvian level to avoid further price distortions (t^{c} in Eq. 2.12 and t^{p} in Eq. 2.13). However, the tax imposed on the price takers coincides with their respective Pigouvian level (t^{p} in Eq. 2.11 and t^{c} in Eq. 2.14) since price takers can not only adjust the quantity consumed (produced) under oligopoly (oligopsony), but they can also abate more.

2.6 The First- vs Second-best Allocation

Barnett (1980) reveals that in the presence of two market failures, pollution and market power, the optimal emission tax can only acheive a second-best trade-off between internalizing the two distortions. In this section, we examine whether this principle can be extended to the double-sided pollution and abatement scenario.

2.6.1 Two-sided Environmental Externality

2.6.1.1 Two-sided Abatement

When producers and consumers abate, the optimal taxes can only achieve a second-best allocation. This result follows from simply replacing t^{p} and t^{c} in the equilibrium quantity (and thus price) condition.

Under oligopoly:

$$\frac{\partial U}{dQ} = \frac{\partial C}{dq_{i}} + \frac{\partial D}{\partial \overline{E^{p}}} \frac{\partial E^{p}}{\partial q_{i}} + \frac{\partial D}{\partial E^{c}} \frac{\partial E^{c}}{\partial Q} + \frac{1}{n\epsilon_{d}} \left(\alpha^{p} \frac{\partial E^{p}}{\partial q_{i}} \frac{dQ}{dt^{p}} + \alpha^{c} \frac{\partial E^{c}}{\partial Q} \frac{dQ}{dt^{c}} - P^{d}(Q, A^{c}) \right). \tag{2.27}$$

Under oligopsony:

$$\frac{\partial C}{\partial Q} = \frac{\partial U}{\partial q_i^c} - \frac{\partial D}{\partial E^p}\frac{\partial E^p}{\partial Q} - \frac{\partial D}{\partial \overline{E^c}}\frac{\partial E^c}{\partial q_i^c} + \frac{1}{m\epsilon_s}\left(\beta^p \frac{\partial E^p}{\partial Q}\frac{dQ}{dt^p} + \beta^c \frac{\partial E^c}{\partial q_i^c}\frac{dQ}{dt^c} - P^s(Q)\right). \tag{2.28}$$

Based on Eq. 2.27 and 2.28, at equilibrium, the optimal quantity and price are distorted by market power. The reason why the second-best is achieved is because the two market failures, market power and pollution, are linked via the term $\frac{dQ/dt^j}{dA^j/t^j}$ in t^j for $j \in \{p, c\}$ (Eq. 2.3, 2.4, 2.7, 2.4). In this case, as seen in section 5.1.1 (the Δs), there is another layer to the trade-off between correcting for the two market failures. The first layer, and the widely known one, is between market power and pollution, and the second one is within each market failure. Therefore, the main result of Abdul Baki, Benchekroun, and Marrouch (2024), which posits that the first-best can be achieved in the presence of two-sided pollution, only holds under the production-sided abatement scenario and does not extend to the two-sided abatement case. This is because the two market failures are unlinked in this case, as we will argue below. Nonetheless, under the framework of Barnett (1980) (or Requate, 2006), where the polluter abates, only the second-best is achieved because the market failures were linked via $\frac{dQ/dt^j}{dA^j/t^j}$, as revealed in Eq. 2.19 and 2.20.

2.6.1.2 One-sided Abatement

Production-sided Abatement: This case overlaps with the framework of Abdul Baki, Benchekroun, and Marrouch (2024). Specifically, the optimal emission taxes will achieve the first-best allocation when abators (producers) hold market power. This result also extends to the case where consumers (unabators) can control the price. In either case, In either case, the optimal emission taxes correct for market power distortions, as evident in the FOCs below, where the marginal social benefit is equal to the marginal social cost.

Under oligopoly:

$$\frac{dU}{dQ} = \frac{\partial C}{\partial q_i^p} + \frac{\partial D}{\partial \overline{E^p}}\frac{\partial E^p}{\partial q_i^p} + \frac{\partial D}{\partial E^c}\frac{\partial E^c}{\partial Q} \tag{2.29}$$

Under oligopsony:

$$\frac{\partial C}{\partial Q} = \frac{dU}{dq_i^c} - \frac{\partial D}{\partial E^p}\frac{\partial E^p}{\partial Q} - \frac{\partial D}{\partial \overline{E^c}}\frac{dE^c}{dq_i^c}. \tag{2.30}$$

Eq. 2.29 and 2.30 are not distorted by market power because the two market failures are unlinked. This can be seen in t^p and t^c in Eq. 2.9, 2.10, 2.11, and 2.12, where the term

$\frac{dQ}{dt^j}, j \in \{p, c\}$ that drives this correlation dissipates from the optimal taxes.

Consumption-sided Abatement Similar to the production-sided abatement scenario, the optimal emission taxes are the first-best instruments, irrespective of whether maker power is wielded by producers or consumers. Precisely, the equilibrium quantity-price allocation is no longer distorted by market power in the presence of oligopoly power

$$\frac{\partial U}{dQ} = \frac{dC}{dq_i} + \frac{\partial D}{\partial \overline{E^p}} \frac{dE^p}{dq_i} + \frac{\partial D}{\partial E^c} \frac{\partial E^c}{\partial Q}.$$

or oligopsony power

$$\frac{dC}{dQ} = \frac{\partial U}{\partial q_i^c} - \frac{\partial D}{\partial E^p} \frac{dE^p}{dQ} - \frac{\partial D}{\partial \overline{E^c}} \frac{\partial E^c}{\partial q_i^c}.$$

2.6.1.3 No Abatement

Similarly, since $\frac{dQ}{dt^j}, j \in \{p, c\}$ does not appear in the optimal combination of taxes in Eq. 2.17 and 2.19, the two market failures are unlinked, and so the optimal emission taxes, t^p and t^c, are first-best instruments under either oligopoly or oligopsony, where

$$\frac{dU}{dQ} = \frac{dC}{dq_i} + \frac{\partial D}{\partial \overline{E^p}} \frac{dE^p}{dq_i} + \frac{\partial D}{\partial E^c} \frac{dE^c}{dQ},$$

under oligopoly, and

$$\frac{dC}{dQ} = \frac{dU}{dq_i^c} - \frac{\partial D}{\partial E^p} \frac{dE^p}{dQ} - \frac{\partial D}{\partial \overline{E^c}} \frac{dE^c}{dq_i^c},$$

under oligopsony.

To summarize, in the presence of multiple market failures, achieving the first-best depends less on the number of policy instruments and more on whether the market failures are linked. This porposition recaps.

Proposition 4 *In the presence of two market failures, pollution and market power,*

- *the optimal emission taxes achieve the first-best allocation when market failures are unlinked. This corresponds to the case of a two-sided environmental externality with either no abatement or with one-sided abatement.*
- *the optimal emission tax(es) can only achieve the second-best allocation when market failures are linked. This corresponds to the case of a two-sided environmental externality and two-sided abatement or a one-sided environmental externality and abatement.*

2.7 Conclusion

This study is the first to examine whether the results of Pigou (1920) and Barnett (1980), established within the context of a one-sided environmental externality, and so one-sided abatement, can be extended to a two-sided environmental externality two-sided abatement framework, that is when both producers and consumers pollute and abate. Also, this study is the first to explore whether the results of Abdul Baki, Benchekroun, and Marrouch (2024), which are based on a two-sided pollution production-sided abatement scenario, extend to consumption-sided abatement or two-sided abatement scenarios. Specifically, this paper attempts to answer a fundamental question: what are the optimal emission taxes imposed on producers and consumers in the presence of two-sided environmental externality, two-sided abatement effort, and one-sided market power? This inquiry bears relevance to many day-to-day scenarios and holds particular significance given the heightened health and environmental concerns across several industries: waste management, steel, pharmaceutical, agrochemical, agricultural, weapons and military vehicle manufacturing, electricity, tobacco, air taxi and charter, and meat, among others. Within this paper, a Cournot oligopoly, then a Cournot oligopsony, is examined under three scenarios: (1) neither the producers nor consumers abate, (2) either the producers or consumers abate, and (3) both the producers and consumers abate.

Surprisingly, neither the results of Pigou (1920), Barnett (1980), nor Abdul Baki, Benchekroun, and Marrouch (2024) hold. In the absence of abatement, the optimal emission taxes are not unique and can, thus, exceed the Pigouvian level; nevertheless, they achieve the first-best allocation. When one of the polluters (producers or consumers) abates, the first-best is achieved where the optimal emission tax imposed on the (non) abating polluter is (below) equal to the Pigouvian level, even if this polluter is a price (taker) maker. When both polluters abate, the first-best can no longer be achieved. Specifically, the optimal emission tax set on producers and consumers fall below their respective Pigouvian level. All these insights hold irrespective of whether market power is active on the demand or the supply side.

In summary, our work paper provides a comprehensive framework for emissions taxes and generates important theoretical and policy implications. Regarding the theoretical implication, this study is the first to highlight that the (first) second-best is achieved when the market failures are (un)linked. Thus, it is the correlation between market failures, rather than the number of instruments, that matters the most. Moreover, our work uncovers that the second-best trade-off, as highlighted by Barnett (1980), has two dimensions in a two-sided pollution-abatement world: between and within the two market failures, pollution and market power. While our work primarily models horizontal markets, all our insights generalize to the case of vertical markets.

Moving to the policy implications, first, to effectively internalize two-sided environmental externalities caused by the same product, emissions taxes need to be designed in a coordinated manner that reflects the abatement efforts of both producers and consumers. For example, in the cryptocurrency market, the Canada Revenue Agency imposes income tax (but not an environmental tax) on cryptocurrency transactions, but not mining activities.[16] This oversight means that the environmental impacts of both mining and transactions are not fully addressed. To rectify this, emissions taxes should be levied on miners based on the source of the electricity used for mining–whether it is renewable or non-renewable. Similarly, investors should face taxes based on the sustainability of the cryptocurrencies they invest in. For instance, Bitcoin has a significantly higher carbon footprint compared to Ethereum (Kohli et al., 2023). Imposing taxes independently on either producers or consumers without considering their interdependence results in socially inefficient outcomes by *falsely* adhering strictly to the principles of Pigou (1920) or Barnett (1980).

Second, there should be concerted efforts to incentivize both producers and consumers to adopt abatement technologies. As demonstrated in our analysis, the implementation of abatement can lead to less stringent emissions taxes when both polluters engage in it, but at the cost of forgoing the first-best outcome. Incentives such as subsidies for green technologies or public awareness campaigns can drive these efforts. For example, Canada's Agriculture Clean Technology Program offers subsidies of up to 50% for Canadian farmers and processors who adopt clean technologies. Also, the Canada Clean Technology Program supports consumers by providing interest-free loans for energy-efficient retrofits. Furthermore, the Climate Action and Awareness Fund (CAAF) allocates up to $206 million over five years to projects that reduce greenhouse gas emissions, with a focus on enhancing youth climate awareness.[17]

Lastly, our results expose a critical equity concern hidden within the application of

16. https://www.canada.ca/en/revenue-agency/programs/about-canada-revenue-agency-cra/compliance/digital-currency/cryptocurrency-guide.html

17. https://solarpowerstore.ca/pages/canadian-solar-incentives
https://www.canada.ca/en/services/environment/weather/climatechange/funding-programs/climate-action-awareness-fund.html

the PPP under one-sided abatement scenarios. In theory, PPP aligns responsibility with environmental harm. In practice, however, we find that the tax burden tends to fall disproportionately on the side of the market without pricing power, especially when that side undertakes abatement. This means that the price takers are doubly penalized when they abate: they bear the cost of abatement and are also taxed more aggressively. Such an outcome challenges the fairness of the PPP as currently implemented and calls for a more nuanced policy design that recognizes both pollution responsibility, market power, and abatement capacity.

In light of our findings, several promising avenues for future research emerge. First, a natural extension of this model is to examine optimal emission taxation in the presence of two-sided market power, that is, when both producers and consumers possess pricing power. This scenario could reveal new distortions and trade-offs in policy design. Second, it would be valuable to endogenize market power and investigate how two-sided environmental externalities influence firm entry decisions and, consequently, market structure. These extensions remain open for future exploration.

Appendix

Social Welfare Maximization

The FOCs associated with maximizing the social welfare function are

$$\left(\frac{\partial U}{\partial Q} - \frac{\partial C}{\partial q_i^p}\right)\frac{dQ}{dt^p} - \frac{\partial C}{\partial a_i^p}\frac{dA^p}{dt^p} = \left(\frac{\partial D}{\partial \overline{E^p}}\frac{\partial E^p}{dq_i^p} + \frac{\partial D}{\partial E^c}\frac{\partial E^c}{dQ}\right)\frac{dQ}{dt^p} + \frac{\partial D}{\partial \overline{E^p}}\frac{\partial E^p}{\partial a_i^p}\frac{dA^p}{dt^p}$$

and

$$\left(\frac{\partial U}{\partial Q} - \frac{\partial C}{\partial q_i^p}\right)\frac{dQ}{dt^c} + \left(\frac{\partial U}{\partial A^c} - \frac{dB}{dA^c}\right)\frac{dA^c}{dt^c} = \left(\frac{\partial D}{\partial \overline{E^p}}\frac{\partial E^p}{\partial q_i^p} + \frac{\partial D}{\partial E^c}\frac{\partial E^c}{\partial Q}\right)\frac{dQ}{dt^c} + \frac{\partial D}{\partial E^c}\frac{\partial E^c}{\partial A^c}\frac{dA^c}{dt^c}.$$

under production-sided market power, and

$$\left(\frac{\partial U}{\partial q_i^c} - \frac{\partial C}{\partial Q}\right)\frac{dQ}{dt^p} - \frac{\partial C}{\partial A^p}\frac{dA^p}{dt^p} = \left(\frac{\partial D}{\partial E^p}\frac{\partial E^p}{dQ} + \frac{\partial D}{\partial \overline{E^c}}\frac{\partial E^c}{dq_i^c}\right)\frac{dQ}{dt^p} + \frac{\partial D}{\partial E^p}\frac{\partial E^p}{\partial A^p}\frac{dA^p}{dt^p}$$

and

$$\left(\frac{\partial U}{\partial q_i^c} - \frac{\partial C}{\partial Q}\right)\frac{dQ}{dt^c} + \left(\frac{\partial U}{\partial a_i^c} - \frac{dB}{da_i^c}\right)\frac{dA^c}{dt^c} = \left(\frac{\partial D}{\partial E^p}\frac{\partial E^p}{\partial Q} + \frac{\partial D}{\partial \overline{E^c}}\frac{\partial E^c}{\partial q_i^c}\right)\frac{dQ}{dt^c} + \frac{\partial D}{\partial \overline{E^c}}\frac{\partial E^c}{\partial a_i^c}\frac{dA^c}{dt^c}.$$

under consumption-sided market power.

The Stringency of the Environmental Regulation

Let $\bar{p}$ ($\underline{p}$) be the net price received by the oligopolist (oligopsonist) and $p^\star$ be the competitive equilibrium price. Let $\Delta^+ = \bar{p} - p^\star$ $\left(\Delta^- = p^\star - \underline{p}\right)$ denote the oligopolist's price markup (oligopsonist's price markdown) in the presence of the emissions taxes. Using FOCs for the cases of oligopoly and oligopsony, respectively, we present the equations of $\bar{p}$, $\underline{p}$, Δ^+, and Δ^- as follows.

$$\overline{p} = \frac{\partial C}{\partial q_i^p} + t^c\frac{\partial E^c}{\partial Q} - \frac{P^d(Q, A^c)}{n\epsilon_d}$$

$$\Delta^+ = t^c\frac{\partial E^c}{\partial Q} - \frac{P^d(Q, A^c)}{n\epsilon_d}$$

$$\underline{p} = \frac{\partial U}{\partial q_i^c} - t^p\frac{\partial E^p}{\partial Q} - \frac{P^s(Q)}{m\epsilon_s}$$

$$\Delta^c = t^p\frac{\partial E^p}{\partial Q} + \frac{P^s(Q)}{m\epsilon_s}$$

Chapter 3

The Unintended Health Impacts of Water Fluoridation: Evidence from Birth Outcomes

3.1 Introduction

Water pollution is a pervasive source of negative health externalities, which are often disproportionately borne by vulnerable populations. (see e.g., Chakraborti and Shimshack, 2022; Guimbeau et al., 2024; Ebenstein, 2012).[1] These include newborns of women from less advantaged communities, such as those who are non-White, less educated, or of lower socioeconomic status. While such pollutants are usually the byproducts of industrial activities, some could, paradoxically, arise from deliberate public health interventions. A notable example is water fluoridation, the practice of adding fluoride to public water supplies to prevent dental cavities. Although fluoride has been endorsed as safe at recommended levels, growing concerns have emerged regarding its potential health consequences.[2] These concerns range from dental and skeletal fluorosis to neuro-developmental toxicity (Hung et al., 2023; Veneri et al., 2023a; Veneri et al., 2023b). In 2024, a lawsuit was filed against the Environmental Protection Agency (EPA) in California regarding the safety of fluoride in water, and in 2025, water fluoridation

1. The mortality rate attributed to water pollution is around 2% in developed countries and can reach up to 50% in developing countries. See https://genderdata.worldbank.org/en/indicator/sh-sta-wash-p5?gender=total

2. The public health authorities supporting water fluoridation include the World Health Organization (WHO), the FDI World Dental Federation, and the Centers for Disease Control and Prevention (CDC)

was banned in Utah for the first time in United States (US) history.[3] These actions highlight growing concerns about the established safety of fluoridation and underscore the need to revisit long-held assumptions about the health impacts of fluoridation, even at levels previously considered safe (see e.g., Taylor et al., 2025).

Building on these emerging safety concerns, this study investigates the unintended effects of water fluoridation in the US on birth outcomes, specifically, birth weight, prematurity, and congenital anomalies. We construct comprehensive county-level data on water fluoridation from the CDC's 1992 Fluoridation Census and merge it with individual-level birth records from the US Vital Statistics for the period 1982–1993. This period is of particular interest because it coincides with the major phase of fluoride expansion, which lasted until the early 1990s, and it is also the period when geographic county information became consistently available in the Vital Statistics. Our empirical strategy is comprised of a difference-in-differences (DiD) design that exploits exogenous variation in the timing of fluoridation across counties. This approach allows us to isolate the causal impact of prenatal fluoride exposure by comparing birth outcomes in counties that adopt fluoridation to those in counties that have not done so, before and after prenatal fluoride exposure. We complement this main specification with an extensive set of robustness checks, including different model specifications and an alternative empirical framework (e.g., stacked DiD). These additional analyses assess the sensitivity of our findings to potential omitted variable bias and evaluate key assumptions underlying the DiD framework, including the possibility of heterogeneous treatment effects.

The relevance of our research question extends beyond the immediate health outcomes under study. Water fluoridation has long been promoted as a cost-effective public health intervention. For example, in 2013, estimated net savings from community fluoridation reached approximately $6.5 billion in the US (O'Connell et al., 2016). Nevertheless, such calculations typically omit potential unintended health consequences. Our study, hence, allows for a more comprehensive evaluation of the overall social costs and benefits. Moreover, birth outcomes are well-established predictors of long-term health and economic outcomes (Black, Devereux, and Salvanes, 2007; Oreopoulos et al., 2008; Royer, 2009; Currie, 2009; Bharadwaj, Eberhard, and Neilson, 2018). These early-life disadvantages often reinforce intergenerational cycles of health and income inequality (Currie and Moretti, 2007; Currie, 2011; Aizer and Currie, 2014). Given that

3. https://www.cbsnews.com/news/epa-fluoride-drinking-water-federal-court-ruling/
https://www.bbc.com/news/articles/c4gmggp2y99o

water fluoridation is far more common in the US than in other high-income countries (Vinceti, Veneri, and Filippini, 2024), and that the US also faces higher levels of socioeconomic inequality (e.g., Michaud et al., 2011), reassessing this widely implemented policy is particularly necessary. Our study, therefore, speaks to broader concerns at the intersection of environmental health, public policy, and inequality.[4]

Our work makes three main contributions to the following strands of literature: fluoride exposure, the fetal origins hypothesis, and water quality. First, we provide novel causal evidence on the unintended effects of prenatal exposure to water fluoridation, leveraging exogenous variation in the timing of community-level fluoridation across US counties. The literature on fluoride exposure remains relatively nascent, with most studies being either medical or epidemiological.[5] This growing body of work has documented correlations between fluoride exposure and adverse birth outcomes as well as other health and behavioral outcomes, though sometimes with conflicting results. However, the majority of this research is observational in nature and lacks quasi-experimental methods, limiting the ability to identify causal conclusions regarding the health impacts of fluoride. Moreover, most of this work relies on experimental or clinical datasets with relatively small sample sizes, raising concerns about population representativeness and the external validity of its findings. Our work stands out by uncovering the causal inadvertent impact of prenatal fluoride exposure on birth outcomes using a DiD approach applied to administrative data which is nationally representative.

Second, by shedding light on the unexpected gestational vulnerability to water fluoridation, we situate fluoride exposure within the broader fetal origins hypothesis literature (see, e.g., Barker and Osmond, 1986; Almond, Chay, and Lee, 2005; Black, Devereux, and Salvanes, 2007; Almond and Currie, 2011; Bharadwaj, Lundborg, and Rooth, 2018; Wan, Zhang, and Hu, 2025). To the best of our knowledge, Roberts (2024), Aggeborn and Öhman (2021), and Glied and Neidell (2010) are the only studies that aimed at disentangling the causal effect of fluoride exposure. However, these studies have focused exclusively on postnatal exposure and adult outcomes like labor income, economic self-sufficiency, and physical health. By contrast, our study focuses on prenatal exposure and immediate birth outcomes, which count among the predictors of the long-term development, in general, and the already examined long-term outcomes, in

4. https://www.vox.com/2018/7/29/17627134/income-inequality-chart

5. See, e.g., Grandjean et al. (2024), Green et al. (2019), Fawell et al. (2006)), Krzeczkowski et al. (2024), Arun, Rustveld, and Sunny (2022), Ortíz-García et al. (2022), Goodman et al. (2022), Till et al. (2025), Gopu et al. (2022), Green et al. (2019), Bashash et al. (2017).

specific. Accordingly, our study not only expands the scope of fluoride research to include birth outcomes but also highlights how this public health policy can potentially shape health trajectories throughout the life course.

Third, we offer new insight into the mechanisms and unequal consequences of fluoride exposure by examining heterogeneity across maternal (e.g., education, race) and birth (e.g., sex, parity) characteristics, as well as by trimester and season of exposure, thereby contributing to the broader literature on water quality. To our knowledge, no prior work, whether on fluoridation or water quality more generally (e.g., Hill and Ma, 2022; Wang, Chen, and Li, 2022; Dave and Yang, 2022; Beland and Oloomi, 2019, among others), has examined the role of seasonality in exposure, despite its relevance for individual-level water consumption patterns (e.g., Zhang et al., 2025; Lin et al., 2024). For instance, water intake tends to rise significantly during the summer, leading to a gradual accumulation of fluoride, while individuals consume more hot beverages in winter. Notably, boiling water does not eliminate fluoride; instead, it increases its concentration due to evaporation.[6] Additionally, fluoride concentrations in public water supplies are typically adjusted seasonally, with higher levels in winter than summer.[7] These seasonal differences in water consumption behavior could affect both the quantity and intensity of fluoride exposure during pregnancy. We help bridge this gap by examining whether low-dose cumulative exposure to fluoride (in summer) yields different birth outcomes than high-dose acute exposure (in winter). Moreover, among causal studies on fluoride, only Glied and Neidell (2010) considers heterogeneity by socioeconomic status of adults. In contrast, neither Roberts (2024) nor Aggeborn and Öhman (2021) explore whether their divergent findings vary by race, marital status, socioeconomic background, or other key characteristics.[8] These are all relevant factors that we attempt to explore. Therefore, our work highlights the limited scope of heterogeneity analysis in the current body of work, both in terms of who is most affected by fluoride and when fluoride exposure is most harmful.

Our results reveal robust evidence that prenatal exposure to fluoridated water leads

6. For more details, see
https://www.eweb.org/yourpublicutility/news/wateruseinsummermorethantwiceashighaswinter
https://www.fsec.ucf.edu/en/Publications/pdf/FSEC-PF-464-15.pdf.

7. The optimal fluoride concentration in drinking water is inversely proportional to the average temperature. See https://www.cdc.gov/fluoridation/timelineforcommunitywaterfluoridation.

8. Even within the epidemiological literature, which focuses more explicitly on prenatal exposure, there is little evidence on how fluoride's effects differ across population subgroups (e.g., Goin et al., 2024; Malin et al., 2024; Hall et al., 2023; Abduweli Uyghurturk et al., 2020; Goyal et al., 2020)

to adverse birth outcomes. Firstly, fluoridation is associated with a modest but statistically significant reduction in birth weight by around 11.248 grams (0.34%). While notable, this pales in comparison to the effect on central nervous system anomalies (CNSA), which rise by nearly 50%, indicating a potentially severe risk to fetal neurodevelopment. We also observe smaller but significant increases in the onsets of very low birth weight (9.09%) and prematurity (3.05%). These findings are especially pronounced when exposure occurs in the first trimester, highlighting early gestation as a critical window.

Secondly, our results reveal that these effects are not evenly distributed across subgroups. Contrary to prior literature on environmental stressors (such as Fan and He, 2023; Hill and Ma, 2022; Currie et al., 2013), which often finds larger impacts among marginalized populations (see e.g., Guimbeau et al., 2024; Wang, Chen, and Li, 2022; Armand and Kim Taveras, 2021; Beland and Oloomi, 2019), we find stronger adverse effects among White, married, and more educated mothers, i.e., the less socioeconomically disadvantaged. These patterns may reflect greater fluoride intake through a higher water intake (Brooks et al., 2017; Rosinger and Herrick, 2016; Popkin, Barclay, and Nielsen, 2005) or trust (Rosinger et al., 2018; Onufrak et al., 2014; Goodman et al., 2013) in tap water among these groups in the US (or a combination of both). We further document sex- and birth-type-specific vulnerabilities, conditional on survival: male infants experience greater weight loss, while female infants are more prone to CNS anomalies.

Thirdly, we find that the effects associated with (low-dose) cumulative versus (high-dose) acute fluoride exposure are heterogeneous. Specifically, birth weight effects are more pronounced when early pregnancy coincides with the summer months, a period when overall water consumption is higher, suggesting that cumulative exposure to low fluoride concentrations matters for birth weight. In contrast, congenital anomalies, including CNSA and MSKLA, are more strongly associated with early pregnancy, overlapping with the winter season, possibly indicating that more acute and highly concentrated fluoride intake during early pregnancy could trigger congenital malformations. These seasonal patterns align with established distinctions in the environmental health literature between chronic and acute exposures, which often elicit different biological responses and developmental outcomes (e.g., Iban-Arias et al., 2025; Neto, Ferraro, and Vieira, 2023; Sarigiannis and Hansen, 2012).

Importantly, our back-of-the-envelope cost analysis shows that the combined hospital charges at birth and lifetime earnings losses or treatment associated with these adverse outcomes can outweigh the CDC's oft-cited $20 average return per dollar spent on water fluoridation (CDC, 2024). This potentially flips the intervention's net social benefit from positive to negative for the average exposed fetus.

The rest of the paper is organized as follows. Section 3.2 reviews the literature. Section 3.3 presents the data, while section 3.4 presents the empirical strategy. Section 3.5 discusses the results, while section 3.6 presents the robustness checks. Section 3.7 presents the cost of water fluoridation associated with these adverse birth outcomes. Lastly, section 3.8 concludes.

3.2 Related Literature

The literature on water quality and its effects on fetal and infant health outcomes is extensive (e.g., Fan and He, 2023; Flynn and Marcus, 2023; Currie et al., 2013), with numerous studies identifying harmful consequences of exposure to various contaminants (e.g., Marcus, 2025; Bhalotra et al., 2021; Sorensen et al., 2019; Keskin, Shastry, and Willis, 2017; Pitt, Rosenzweig, and Hassan, 2015; Rau, Urzúa, and Reyes, 2015; Aldy, 2014; Zhang, 2012), particularly during pregnancy. While many sources of water quality concerns arise unintentionally from industrial, agricultural, or environmental processes, some intentional public policies — such as water fluoridation — involve the deliberate introduction of chemical agents, which may also have unintended adverse effects.

3.2.1 Water Pollution

The sources of water pollution can be broadly classified into five categories: industrial, agricultural, ecological, climate change-induced, and infrastructure-related sources.

One well-documented case of industrial-induced water pollution with health consequences is the Deepwater Horizon oil spill, examined by Beland and Oloomi (2019). Using a differences in differences framework, the authors find significant increases in the rates of low birth weight and preterm births by 0.95 and 0.94 percentage points, respectively. These effects were particularly pronounced among Black and Hispanic mothers, younger women, and those with lower education or unmarried status. Within the literature focusing on industrial-induced water pollution (e.g., Ebenstein, 2012), subsequent studies have investigated the health impacts of other industrial activities that may contaminate water sources— most notably shale gas development, or fracking (e.g., Bartik et al., 2019; Hill, 2018; Hill and Ma, 2017). For instance, Hill and Ma (2022) using a difference-in-differences design, find that fetal exposure to fracking activity increases the incidence of preterm birth and low birth weight by approximately 11 to 13% in affected areas of the United States.[9]

Adverse health outcomes may also be triggered by prenatal exposure to water quality degradation resulting from agricultural activities (e.g., Mettetal, 2019; Lai, 2017). In the context of India, Brainerd and Menon (2014) demonstrate that the probability of neonatal mortality increases following exposure to agricultural toxins in water during the month of conception, especially among uneducated women in rural areas. While this study highlights the immediate impacts of agricultural water contamination, Fletcher and Noghanibehambari (2024) extend this line of inquiry by examining the long-term consequences of early-life pesticide exposure. Using a difference-in-differences framework, they find that in-utero exposure to pesticides reduces male life-expectancy by average 2.2 months—and by nearly a full year in areas with higher exposure levels.[10]

While industrial and agricultural activities are major anthropogenic drivers of water pollution, climate change introduces another layer of complexity to water quality and its impact on health (see Fletcher and Noghanibehambari, 2024; Da Mata et al., 2023). For example, Guimbeau et al. (2024) examine the effects of rising ocean salinity, driven by climate change, on children's anthropometric outcomes in Bangladesh. Their results indicate that increased salinity exposure in utero leads to significant declines in height-for-age, weight-for-height, and weight-for-age by 0.11, 0.13, and 0.15 standard deviation (SD), respectively.

Within the broader category of infrastructure failures, lead contamination in public water systems has been a major concern. Wang, Chen, and Li (2022) evaluate the short-term health consequences of the US Flint water crisis. Applying synthetic control methods, their analysis reveals an increase in low birth weight rates by 15.5 percent, particularly among Black mothers. The Flint water crisis was precipitated by a change in the water source and a complete failure to implement any anti-corrosion measures

9. Recent attention has also turned to microplastic pollution, where the in-utero exposure during the third trimester increases risks of low birth weight by 0.37 per 1,000 births (Du, Zhang, and Zou, 2024).

10. The impact of agriculture on water quality extends beyond direct chemical contamination. It can also stimulate blue-green algal blooms that, in turn, produce potent toxins such as microcystin, which have documented adverse health effects (Jones, 2019).

with the new supply. However, lead poisoning can also occur due to the shifts in water chemistry, which potentially reduce the effectiveness of existing anti-corrosion measures. This is the case of the US Newark water crisis.[11] Dave and Yang (2022) implement a generalized DiD approach and show that the in-utero exposure of lead raised the probability of low birth weight and preterm birth by 18 and 19 percent, respectively, specifically during the initial period following the crisis. This effect has dissipated over time due to information spillovers and avoidance behaviors.[12] Specifically, information spillovers—such as media reports and heightened public awareness of lead contamination—induced stress responses even among nearby populations not directly affected by the crisis, while avoidance behaviors included increased prenatal care and other protective health actions.

3.2.2 Water Fluoridation

The common thread across the aforementioned studies is that they focus on unintentional and well-recognized water pollutants such as oil, chemicals, heavy metals, and agricultural runoff. These externalities typically arise as byproducts of industrial, agricultural, or municipal processes. Nonetheless, public health interventions can implicitly result in similar "externalities" and widespread unintended health implications, as is the case with water fluoridation.

In the early 20th century, fluoride was discovered as a mineral essential for protecting teeth against decay and cavities. To this end, it has been used by dentists and, more broadly, added to public water supplies, within a specified limit set by the WHO, to promote public dental wellbeing. Its first use in public water supplies began in Grand Rapids, Michigan in 1945, followed by a rapid expansion across the US and the globe.[13] Despite its public health intentions, several fluoride overexposure incidents have raised safety concerns, including notable cases in New Mexico (Hoffman et al., 1980), Vermont (Vogt et al., 1982), Maryland (Anderson, Beard, and Sorley, 1980), and Alaska (Gessner et al., 1994).

A growing body of literature, specifically from the medical and epidemiological

11. https://en.wikipedia.org/wiki/Newark_water_crisis
12. See the work of Aizer and Currie (2019) and Aizer et al. (2018) for lead poisoning implications beyond birth outcomes.
13. https://en.wikipedia.org/wiki/Water_fluoridation_in_the_United_States

fields, has begun to examine the broader health effects of fluoride exposure. However, these studies have largely focused on establishing correlations rather than identifying causal effects. For instance, fluoride exposure has been associated with dental and skeletal fluorosis (Fawell et al., 2006), and more recently, with neurodevelopmental concerns such as reduced IQ and neurotoxicity in children (NTP, 2024; Grandjean et al., 2024; Green et al., 2019). Some studies suggest that these adverse effects can occur even at fluoride levels deemed safe by international standards (Taylor et al., 2025). Moreover, exposure to fluoride, particularly during pregnancy, has been linked to poorer nervous system and cardiac autonomic functioning in infants (Krzeczkowski et al., 2024). Arun, Rustveld, and Sunny (2022) report an increased risk of low and very low birth weight, especially among minorities (e.g., Hispanic). However, this study lacks precise prenatal exposure timing—they infer birth outcomes based on maternal fluoride exposure at the time of the survey, which might not even coincide with the pregnancy period. Moreover, the study does not exploit any quasi-experimental variation in exposure to identify the causal effect of prenatal fluoride exposure on birth outcomes; instead, it relies on a simple linear regression. Their relatively small sample size further raises concerns about external validity and the robustness of subgroup analyses. To address some of these gaps, a few studies have considered an explicitly defined prenatal exposure window, showing that such adverse birth outcomes are more noticeable when exposure happens during the first trimester (Ortíz-García et al., 2022). However, most of the aforementioned limitations persist and have contributed to the emergence of conflicting findings, with some studies demonstrating no correlation between prenatal fluoride exposure and birth outcomes (Goodman et al., 2022).

Despite the extensive epidemiological and observational research on fluoride exposure, most existing studies remain limited in their efforts to identify causal effects (see, e.g., Choi et al., 2012; Ibarluzea et al., 2022). In response, a new strand of the literature has emerged within the health economics field to shed light on the causal effects of fluoride exposure on economic and health outcomes. A prominent contribution is the work of Glied and Neidell (2010), who investigate the implications of fluoride exposure on labor market outcomes rather than health directly. They use exogenous variation in fluoride access to show that fluoridation increases women's earnings only in the US by 4 percent, especially among those from low socioeconomic backgrounds. Aggeborn and Öhman (2021) take a more direct approach by examining natural variation in fluoride levels in Sweden on dental health, cognitive ability, and labor income. The authors find positive effects of fluoride on dental health and labor income, but no clear channel on cognitive outcomes. The observed positive effects were also more evident in lower socioeconomic groups. However, their study has been criticized mainly for (1) fluoride exposure mismeasurement and for (2) using lifetime averages of exposure from birth to age 18-20. These methodological issues likely obscured any potential neurological effects of fluoride exposure during early childhood or prenatal periods. Additionally, they do not provide sufficient information on how these effects could have happened, i.e., the underlying mechanisms, which does not speak to prenatal fluoride exposure.

In an attempt to overcome these limitations, Roberts (2024) provides more robust evidence by focusing specifically on fluoride exposure from birth to age five in the US. This study utilizes the 1992 Fluoride Census and implements a stacked DiD framework. On one hand, this dataset permits examining a much larger sample than that of Aggeborn and Öhman (2021). On the other hand, this design improves upon earlier work by better accounting for geographic and temporal heterogeneity in exposure, while establishing clean treated and control groups. By employing this design, the study effectively mitigated potential confounders and birth cohort effects. Roberts (2024) finds that early-life fluoride exposure negatively affects long-term outcomes such as economic self-sufficiency (1.9 SD), physical health (1.2 SD), and high school graduation (1.5 percentage points) and military service rates (1 percentage points). These effects are more pronounced for men.

While this recent strand of the literature has made notable contributions, important gaps remain. First, prior work has largely focused on postnatal rather than prenatal exposure. Second, the outcomes studied have been long-term, such as economic self-sufficiency and labor income, leaving short-term indicators, like birth outcomes, which are predictors of long-term outcomes, relatively unexplored. Third, the literature on the causal impact of fluoride exposure on birth outcomes is still nascent, with a clear gap isolating the impact of in-utero exposure. Our study addresses these gaps by examining the short-term effects of prenatal fluoride exposure on birth outcomes using a DiD approach. In doing so, we reassess the effect of deliberately added fluoride—long regarded as safe by public health authorities—as a potential water pollutant with implications for fetal development, thus contributing to both the environmental health literature and the broader body of work on the fetal origins hypothesis (Barker and Osmond, 1986).

3.3 Data

In this study, we combine data from two key sources: the US National Vital Statistics System (1982-1993) and the 1992 Fluoride Census.[14] The Vital Statistics provide detailed birth-level information, while the Fluoride Census offers community-level data on the timing and nature of water fluoridation across the US. We use the Fluoride Census and construct comprehensive county-level fluoridation data and merge it with the county-level Vital Statistics data. This allows us to examine the relationship between water fluoridation and infant health outcomes at the county level.

3.3.1 Birth Outcomes and Covariates

The Vital Statistics provides data on our main outcomes of interest, which capture key dimensions of infant health: birth weight, premature births, and congenital anomalies. For birth weight, we use a continuous variable measuring the infant's weight at birth (in grams) and construct two binary indicators: very low birth weight (< 1500 g) and low birth weight (< 2500 g). For premature births, we construct a binary indicator for whether the gestational length is less than 38 weeks. In the case of congenital anomalies, we focus on the most commonly reported malformations, which we group into four categories based on the International Classification of Diseases system. Accordingly, we create four dummy variables indicating the presence of the following conditions at birth: (i) Central Nervous System Anomalies (CNSA), (ii) Neural Tube Defects (NTD), (iii) Musculoskeletal Anomalies (MSKLA), and (iv) Non-Central Nervous System Anomalies (NCNSA).[15]

In addition to the aforementioned health outcome variables, this dataset provides detailed information on maternal, pregnancy, and birth characteristics, which we include as control variables. These include maternal age, race, marital status, and education; pregnancy-related factors such as the number of prenatal visits, indicators of

14. https://www.cdc.gov/nchs/nvss/index.htm
https://www.cdc.gov/fluoridation/index.html

15. International Classification of Diseases (ICD): https://archive.cdc.gov/www_cdc_gov/ncbddd/birthdefects/surveillancemanual/chapters/chapter-6/chapter6-2.html
While NTD are technically a subset of nervous system disorders in the ICD classification, we treat them as a separate category due to their distinct developmental timing (closure of the neural tube during weeks 3–4 of gestation), their substantial public health burden, and their separate reporting in birth defect surveillance and prevention policies (e.g., folic acid fortification). This approach aligns with CDC surveillance guidelines and targeted prevention efforts (Wald, Morris, and Blakemore, 2018; Williams et al., 2015).

high-risk pregnancy, and birth characteristics including the infant's sex, birth order (parity), and birth type.

3.3.2 Fluoridation Data and Treatment Definition

To measure exposure to fluoride, our main explanatory variable, we need county-level data on water fluoridation expansion. To construct this dataset, we rely on data from the 1992 Fluoride Census conducted by the US Department of Health and Human Services (CDC). This census provides detailed information on the timing of water fluoridation expansion across communities. Water fluoridation began in 1945 and expanded rapidly until the early 1990s. By 1989, over 55% of the US population had access to water with a dentally significant fluoride concentration. Using these records, we construct a county-level dataset that documents the introduction of fluoridated water systems. A county is defined as "treated" if at least one community within it implemented fluoridation. In cases where multiple communities within a county adopted fluoridation at different times, we assign the earliest implementation date as the county-level treatment date. By doing so, we follow a conservative approach, where some births labelled 'treated' will in fact still be unexposed. This non-differential exposure misclassification attenuates our estimate towards zero, producing a lower-bound (i.e., conservative) estimate of the effect size.

It is important to note that the Fluoride Census 1992 distinguishes between two water systems: *adjusted* and *natural*. *Adjusted* systems are those where fluoride was deliberately added to supplies with initially deficient levels in order to reach an optimal concentration. This optimal level varies according to the community's average maximum daily air temperature.[16] In contrast, *natural* systems contain fluoride due to naturally occurring geological sources, without policy intervention, but levels can naturally exceed the recommended threshold. Since our analysis focuses specifically on the effect of policy-induced water fluoridation, we restrict the main analysis to adjusted systems and exclude communities with naturally occurring fluoride from our sample. In total, we analyze 1,949 counties with artificial fluoridation and exclude 379 counties where fluoride occurs naturally.

We focus on (adjusted) water fluoridation for both conceptual and empirical reasons. First, adjusted systems represent deliberate policy interventions, implemented by local governments to achieve specific public health objectives. Studying these cases

16. The optimal fluoride level is inversely proportional to the average maximum daily air temperature.

allows us to evaluate the effects of intentional fluoridation programs, which is critical for informing future policy decisions. In contrast, naturally occurring fluoride levels are a function of local geology and hydrology and may not be representative of the conditions or motivations underlying treated systems. Second, since water fluoridation is implemented as a deliberate policy decision, it offers a clearer identification of exposure timing and intensity. Specifically, in adjusted systems, the timing and dose of fluoridation are well documented and plausibly exogenous to short-term changes in birth outcomes. In contrast, natural fluoride exposure, driven by geological factors, may have cumulative or lagging effects that are harder to isolate or attribute precisely to any policy decision. This distinction underscores why our focus on *adjusted* systems is particularly suited for assessing the causal impact of intentional fluoridation programs. As a robustness check, we consider whether similar patterns hold in naturally fluoridated areas, acknowledging that the mechanisms at play may differ. These results are included in the appendix for completeness but are not central to our identification strategy.

3.3.3 Sample Construction and Study Period

Based on data availability and to ensure population representativeness, we imposed the following restrictions. We start by excluding counties that had implemented fluoridation but for which the exact implementation date was not reported. We also drop counties with populations under 100,000, as FIPS codes for these smaller counties are not available in the Vital Statistics. At the individual level, we remove all births associated with missing gestational age. Furthermore, we exclude extreme values, such as cases with more than 48 prenatal visits or with maternal ages below 18 or above 45. Additionally, we exclude counties with fewer than 100 reported births over the study period to ensure a sufficient sample size for reliable estimation. We further drop observations in which the county of the mother's residence differs from the surveyed county of birth, to ensure accurate fluoride exposure. This restriction helps avoid contamination from cross-county movement and supports the assumption that exposed mothers were residing in, and therefore likely consuming water from, the fluoridated system attributed to their county of residence.

Our final dataset links county-level fluoridation exposure to birth-level outcomes for the years 1982–1993, the period during which county FIPS codes are consistently available in the Vital Statistics. We focus on this period for two reasons. First, The Vital Statistics are publicly available from 1968 onward, but FIPS codes are not included for the years 1968–1981. While the alphabetical order of counties is available for that period, retrieving FIPS codes for this earlier period is feasible. However, expanding our dataset to include that period would substantially increase the data dimension, making the analysis computationally expensive. Second, the year 1992 marks the end of this major fluoridation expansion wave. In total, 2,042 counties had fluoridated their water systems by 1992. By that year, 2,042 counties had implemented water fluoridation. Of these, 267 counties initiated fluoridation during our study period, representing approximately 14% of the total. The remaining 86% had already been treated and are included in our dataset as early-treated counties. This constructed dataset allows us to exploit variation in the timing of fluoridation across counties using a difference-in-differences design described in Section 3.4. A full description of all variables is provided in Table 3.1.

3.4 Empirical Strategy

To estimate the effect of in-utero exposure to fluoride, we rely on a difference-in-differences (DiD) framework, which exploits variation in the timing of water fluoridation across counties. To do so, we compare birth outcomes of pregnant women in counties that adopt fluoridation to those in counties that have not yet done so, before and after implementation, we estimate the causal effect of in-utero fluoride exposure. Specifically, we estimate the following two-way fixed effects (TWFE) model using ordinary least squares (OLS):

$$Y_{icmt} = \alpha_0 + \alpha_1 Treatment_{icmt} + \alpha_3 X_{icmt} + \mu_c + \omega_m + \delta_t + \epsilon_{icmt} \tag{3.1}$$

where Y_{icmt} denotes the birth outcome for infant i, born in county c, in month m and year t. The outcomes we examine include continuous birth weight (in grams), indicators for very low ($<$ 1500g) and low birth weight ($<$ 2500g), prematurity (gestational length $<$ 38 weeks), and congenital anomalies: CNSA, NTD, MSKLA, and NCNSA.

The main treatment variable, $Treatment_{icmt}$, equals one if the mother of infant i (born in month m and year t) is exposed to fluoridated water in her county of residence c during her pregnancy. The vector X_{icmt} includes maternal characteristics (age, race, education, marital status), pregnancy characteristics (number of prenatal visits), and birth characteristics (infant sex, birth order, and birth type).[17] County fixed-effects (μ_c) account for time-invariant county-level confounders; month (ω_m) and year (δ_t) fixed-effects control for time-variant factors. The error term is presented by ϵ_{icmt}. Standard errors are clustered at the county level.[18]

Table 3.2 reports summary statistics for covariates across treated ($Treatment = 1$) and untreated ($Treatment = 0$) women. Although these comparisons mix pre- and post-treatment periods, the covariates appear broadly similar across groups.[19] In Table 3.3, we compare the means of covariates associated with those who were never exposed to fluoride during pregnancy (never treated) to those who were ever exposed, between the pre-treatment (i.e., not-yet-treated) and post-treatment months. Across all comparisons, the treated and control groups are, on average, comparable.

The validity of our empirical strategy relies on the assumption that, in the absence of water fluoridation, birth outcomes in treated and control groups would have followed similar trends, i.e., the parallel trends assumption. Figures 3.1 and 3.2 display the dynamic effects of fluoridation on each birth outcome using an event study framework.[20] In general, for all outcomes, there is no evidence of significant differences in the pre-treatment period, which supports the assumption of parallel trends. That is, prior to water fluoridation, birth outcomes evolved similarly across the treated and control groups. After fluoridation begins, we observe statistically significant changes in outcomes, particularly a decrease in birth weight and increases in the rates of premature births and congenital anomalies. These effects appear to grow over time, especially for pregnancies with longer exposure windows. This pattern suggests a cumulative effect of fluoride exposure, with earlier and more prolonged prenatal exposure yielding more

17. In this specification, we include indicators for mother being White or Black, respectively (reference group: race other than Black or White). All results, including all the conducted heterogeneity analyses and robustness checks, remain robust if we include only an indicator for White (reference group: non-White).

18. Since the treatment (water fluoridation) varies across counties, we cluster standard errors at the county level to account for potential error correlation within counties. This approach also follows the literature (e.g., Roberts, 2024). Our results remain robust when clustering at the state level instead.

19. We formally assess balance using standardized differences, which are all below 0.1—a conventional threshold indicating good balance. We favor standardized differences over statistical significance tests such as the Wald test, as the latter can be misleading in large samples with unequal group sizes.

20. We conduct the event study by regressing the outcome variables on the $D_{(j)}$: month indicators relative to water fluoridation date, $j \in \{-9, .., 0, ..9\}$ and X, μ, ω, δ, as in Eq. 3.1. $D_{(0)}$ corresponds to the month of water fluoridation; this is the reference group. $D_{(-j)}$ and $D_{(j)}$ are j month before and after fluoridation, respectively, for $j \in \{-8, .., 0, ..8\}$. $D_{(-9)}$ and $D_{(9)}$ is at least 9 months before and after fluoridation, respectively. We cluster standard errors at the county level. Exposure during the first trimester corresponds to $D_{(9)}$, $D_{(8)}$, and $D_{(7)}$, second trimester corresponds to $D_{(6)}$, $D_{(5)}$, and $D_{(4)}$, and third trimester corresponds to $D_{(3)}$, $D_{(2)}$, and $D_{(1)}$.

pronounced impacts on birth outcomes. Section 3.5 provides a more detailed examination of these findings.

3.5 Results

Given that we lack individual-level data on water consumption and therefore cannot definitively identify who consumed fluoridated water, the effects presented should be interpreted as the average Intent-to-Treat (ITT) effect. This represents the average impact of living in a county that implemented water fluoridation during pregnancy, rather than the average effect on those who actually consumed fluoridated water. Consequently, the true effect of water fluoridation might be larger, suggesting that our estimates could be considered conservative or potential lower bounds.

3.5.1 Main Analysis

Tables 3.4 and 3.6 report the estimated effects of prenatal exposure to water fluoridation on birth weight, prematurity, and congenital anomalies. As shown in Table 3.4, water fluoridation is associated with a statistically significant reduction in birth weight. Specifically, birth weight decreases by 11.248 grams, representing a 0.34% decline relative to the mean birth weight in the control group. To assess whether the observed effect is primarily driven by infants with normal birth weight (i.e., $\geq$ 2500 g), we re-estimate Eq. 3.1 for the births classified as low and very low birth weight. The results, reported in Table 3.5, show no statistically significant effect of water fluoridation among these subgroups. Nonetheless, water fluoridation increases the likelihood of very low birth weight by 0.01 percentage points (9.09%) and the probability of preterm birth by 0.5 percentage points (3.05%). Both effects are statistically significant at the 5% level. Turning to Table 3.6, the only statistically significant change in congenital anomalies is found for central nervous system malformations. Following fluoridation, the incidence of CNSA increases by 0.02 percentage points, corresponding to a 50% increase relative to the mean. This effect is statistically significant at the 1% level.

To put these magnitudes in perspective, the estimated percentage increase in adverse birth outcomes, such as the incidence of very low birth weight and preterm births, could parallel the effects stemming from well-known and heavily regulated pollutants (see e.g., Wang, Chen, and Li, 2022; Beland and Oloomi, 2019). Additionally, the increase in CNSA represents a substantially larger effect than those on birth weight and prematurity, suggesting a potentially more specific or pronounced vulnerability of the developing nervous system to fluoride exposure.

3.5.2 Heterogeneity Analysis

We now explore whether the effects of prenatal fluoride exposure vary across different periods of pregnancy and maternal, birth, and seasonal characteristics. In this subsection, we first examine heterogeneity by trimester of exposure.

3.5.2.1 Trimester of Exposure

Figures 3.1 and 3.2 suggest that early prenatal exposure may play a particularly important role in driving adverse outcomes. Specifically, the estimated effects become significant among women with longer durations of exposure to fluoridated water, highlighting the importance of exposure timing and accumulation. To formally assess which stage of gestation is most sensitive to fluoride exposure, we estimate a model that interacts our treatment indicator with indicators for each trimester of pregnancy:

$$Y_{icmt} = \beta_0 + \beta_1 First_{icmt} + \beta_2 Second_{icmt} + \beta_3 Third_{icmt} + \beta_4 X_{icmt} + \mu_c + \omega_m + \delta_t + \epsilon_{icmt} \tag{3.2}$$

In Eq. 3.2, *First*, *Second*, and *Third* denote exposure to water fluoridation during the first, second, and third trimester, respectively. The results, reported in Table 3.7, reveal that exposure during the first trimester only, but not the subsequent trimesters, is associated with a 17.664 g decline in birth weight, corresponding to a 0.53% reduction relative to the mean. Additionally, the probability of very low birth weight increases by 0.1 percentage points, which is about 9.1%, among women exposed in their first trimester only, but not the other trimesters. Both low birth weight and preterm birth incidents appear to be influenced by exposure during not only the first but also the second trimester. Specifically, exposure in the first trimester increases the probability of low birth weight and prematurity by 0.3 (5.3%) and 1.4 percentage points (8.5%), respectively. Similar effects are associated with second-trimester exposure. In summary, the results indicate that the first trimester is the primary window of heightened susceptibility to fluoride exposure, with some effects also associated with exposure during starting second trimester. In contrast, no adverse effects are detected in the third trimester. This underscores the importance of early exposure timing and cumulative fluoride intake in shaping birth outcomes.[21]

3.5.2.2 Mother Characteristics

The existing literature suggests that the adverse effects of water pollutants often fall disproportionately on socioeconomically disadvantaged and racially marginalized groups (see Section 3.2). Motivated by this pattern, we explore whether a similar distributional inequality arises from exposure to water fluoridation, a well-intentioned public health intervention. Specifically, we examine heterogeneity in treatment effects by maternal race, marital status, and educational level.

Race: The results of prenatal exposure to water fluoride by maternal race are presented in Tables 3.8 and 3.9. As shown in Table 3.8, birth weight declines for infants of both White and Black mothers following fluoridation, with comparable magnitudes: 11.03 grams (0.33 percent) and 13.63 grams (0.44 percent), respectively. However, the effects on adverse birth outcomes differ by group. Among White mothers, water fluoridation is linked to significant increases in the rates of very low birth weight (0.1 percentage points or 14.29%) and premature birth (0.610 percentage points or 4.3%). In contrast, these effects are not statistically significant for Black mothers. Similarly, the incidence of CNSA increases significantly only among newborns of White mothers compared to black mothers. For White mothers, the increase is 0.02 percentage points, or approximately 50%, whereas for Black mothers, it is 0.03 percentage points, or 10%.

Marital Status: A similar pattern emerges when we stratify the sample by marital status. As reported in Table 3.10, birth weight declines, almost similarly, following fluoridation for both married (11.46 g or 0.34%) and single mothers (13.929 g or 0.44%). Nonetheless, only infants born to married mothers show increased risks of very low birth weight (0.1 percentage points or 14.28%) and premature birth (0.6 percentage points or 4.3%). Turning to congenital anomalies, Table 3.11 reveals that fluoridation increases the likelihood of CNSA by 0.06 percentage points (150%) among infants of

21. We do not report results for congenital anomalies by trimester of exposure due to their rarity. The number of observed cases within each trimester is too small to yield reliable estimates.

married mothers only. Interestingly, among infants of single mothers, the probability of NTD rises by 0.07 percentage points, equivalent to a 175% increase.[22]

Education: Finally, we explore differences by maternal education level. Tables 3.12 and 3.13 compare the effects of water fluoridation for mothers with at least a college degree versus those with only elementary or no formal education. In Table 3.12, birth weight declines in both groups, but more substantially among higher-educated mothers: 14.75 g (0.43%) vs. 9.458 g (0.29%). However, the probabilities of very low birth weight and premature birth increase significantly only among infants of less-educated mothers, by 0.1 percentage points (9%) and 0.5 percentage points (2.6%), respectively. For congenital malformations (Table 3.13), we observe a comparable increase in CNSA across maternal education levels, with a slightly higher effect among infants born to more educated mothers. Specifically, the likelihood of CNSA rises by approximately 0.03 percentage points (57%) among infants of higher-educated mothers, compared to 0.02 percentage points (50%) among infants of less-educated mothers.

To summarize, water fluoridation seems to influence the least socioeconomically disadvantaged population the most. These findings may initially appear counterintuitive, as prior literature often links greater environmental harm to socioeconomically disadvantaged or racially marginalized populations. However, several plausible explanations exist. The pronounced effects on very low birth weight, prematurity, and CNSA among the more advantaged groups may reflect differences in water consumption behavior. In particular, greater trust in and use of public water supplies, especially post-water fluoridation due to its dental benefits, among White, married, and higher-educated mothers could lead to higher actual exposure to fluoridated water during pregnancy, increasing the likelihood of these adverse birth outcomes (See e.g., Javidi and Pierce, 2018; AWWA, 2024). This contrasts with many well-documented environmental pollutants, where regulatory and industrial siting decisions have disproportionately affected marginalized communities. As for NTD, the finding that these increase only among infants born to single mothers may point to nutritional or behavioral vulnerabilities. Single mothers may have reduced access to folate-rich diets or prenatal supplements, both of which are critical for preventing neural tube defects (see e.g., De-Regil et al., 2015; Molloy et al., 2009). In addition, differences in the timing, quality, or

22. When examining percentage changes (relative change), outcomes that occur infrequently (i.e., rare events) are more likely to show large relative increases or decreases, even if the actual number of cases (absolute change) changes only slightly (Noordzij et al., 2017). Also, see http://stats.mom.gov.sg/SL/Pages/Absolute-vs-Relative-Change-Pitfalls.aspx

frequency of prenatal care could also play a role. In our data, for example, the average number of prenatal visits among married mothers is 22.4% higher than for single mothers, potentially contributing to earlier detection or mitigation in that group.[23]

3.5.2.3 Birth Characteristics

Next, we examine whether the effect of water fluoridation is influenced by birth characteristics, such as infant sex and birth order (parity).

Infant Sex: The impact of water fluoridation varies by the infant's sex. Specifically, male infants appear more susceptible to reductions in birth weight and premature births, whereas female infants are more likely to experience congenital anomalies. This pattern is consistent with evidence that male fetuses are generally more vulnerable to adverse prenatal conditions and face higher risks of miscarriage or stillbirth (e.g., Barrett and Lessing, 2021; Al-Qaraghouli and Fang, 2017; Mondal et al., 2014). In contrast, female fetuses are more resilient and more likely to survive, which may explain why congenital anomalies are more commonly observed among females. As shown in Table 3.14, birth weight decreases significantly among males by 15.092 grams, representing a 0.445% reduction. Additionally, the probability of very low birth weight and preterm birth increases only for male infants, by 0.1 and 0.7 percentage points, respectively, corresponding to relative increases of 9.09% and 4.12%. These effects are not statistically significant for female infants. In contrast, Table 3.15 shows that both male and female newborns experience higher rates of CNSA following fluoridation, with comparable increases of 0.02 percentage points, which is about 40% for males and 50% for females. However, only female infants exhibit increased risks for other types of anomalies. Specifically, the probabilities of NTD and NCNSA each rise by 0.1 percentage points, representing increases of over 100% for NTD and around 50% for NCNSA.

Birth Order (Parity): The effect of water fluoridation also appears to vary by birth order. As shown in Table 3.16, the reduction in birth weight is nearly identical for both first births (11.49 grams or 0.35%) and non-first births (10.93 grams or 0.32%), with both effects statistically significant. However, only non-first births exhibit significantly increased risks of adverse outcomes. Specifically, very low birth weight rises by 0.1 percentage points (9.1 %), and the probability of preterm birth increases by 0.6 percentage

23. The average number of prenatal visits is 11.437 among married mothers and 9.346 among single mothers.

points (3.53%). Turning to congenital anomalies, the effects of fluoridation on CNSA are consistent across both groups. As shown in Table 3.17, the probability of CNSA increases by approximately 0.02 percentage points for both first and non-first births, translating to a relative increase of around 40% in the former case and 50% in the latter case. These patterns may reflect differences in maternal physiological reserves, health behavior, and prenatal care between first and non-first pregnancies. For instance, summary statistics indicate that women having their first birth reported more prenatal visits on average (11.315), on average, than those with previous births (10.614).

Taken together, these results indicate that the impact of water fluoridation is not evenly distributed across infant subgroups. Male infants are particularly vulnerable to reductions in birth weight and increased risk of prematurity, consistent with prior research highlighting the heightened biological sensitivity of male fetuses to environmental stressors (see e.g., Fletcher and Noghanibehambari, 2024; Sanders and Stoecker, 2015; Almond and Currie, 2011; Banerjee et al., 2010). In contrast, female infants show greater susceptibility to congenital anomalies, particularly NTD and CNSA, which may reflect sex-specific developmental pathways, warranting further research into underlying biological mechanisms. Lastly, while birth weight reduction and the risk of CNSA are consistent across first and later-born infants, non-first births are more likely to exhibit very low birth weight and prematurity, suggesting that cumulative maternal strain or differential prenatal behaviors may mediate this vulnerability.

3.5.2.4 Seasonal Characteristics

Finally, we examine whether the effect of fluoride exposure varies by the season in which conception occurs. Indeed, environmental stressors are known to produce heterogeneous physiological effects depending on the pattern of exposure: acute versus chronic (see Katoto et al., 2021; Neto, Ferraro, and Vieira, 2023). By focusing on the season of conception, we aim to differentiate between these two modes of fluoride exposure. Seasonal patterns in water consumption may shape exposure intensity. During the summer months, individuals tend to drink more water, potentially increasing their overall fluoride intake. In contrast, in winter, although total water consumption may decline, people are more likely to consume hot beverages. Boiling water, common in colder months, does not remove fluoride; rather, it can increase its concentration due to evaporation, thereby raising the risk of acute fluoride exposure.[24] Accordingly, this analysis aims to distinguish between the effects of consistent, sustained, low-level exposure (summer) and abrupt, high-concentration exposure (winter). Since Section 3.5.2.1 showed that the first trimester is the most critical window for fluoride exposure, we limit our analysis to births where the first trimester occurred during either summer (June, July, August) or winter (December, January, February). We exclude spring and fall to avoid contamination from transitional weather patterns, where drinking behaviors may be inconsistent. The results are summarized in Tables 3.18 and 3.19. When the first trimester overlaps with the summer season, fluoride exposure is associated primarily with a significant decrease in birth weight by 18.05 grams (0.54%) and an increase in the likelihood of very low birth weight by 0.1 percentage points (9.1%). These effects exceed the overall effects seen in Section 3.5. In contrast, congenital anomalies appear more responsive to winter exposure. Specifically, the probabilities of CNSA, NTD, and MSKLA increase by 0.02, 0.03, and 0.1 percentage points, respectively, which is equivalent to relative increases of 50%, 75%, and 50%. These effects are also greater than the documented overall increases as in Section 3.5.

In conclusion, birth weight appears more sensitive to cumulative exposure over time, whereas congenital anomalies may be triggered by short-term, high-concentration exposure. This seasonal distinction offers new insight into the possible pathways through which fluoride affects fetal development. It is also consistent with a well-established body of research distinguishing between acute and chronic exposures to environmental stressors, which often result in different biological responses and health outcomes (see e.g., Iban-Arias et al., 2025; Neto, Ferraro, and Vieira, 2023; Sarigiannis and Hansen, 2012).

3.6 Robustness Checks

To assess the validity of our main results, we perform a series of robustness checks. These aim to rule out potential threats to our identification strategy (e.g., avoidance behavior or selective migration and fertility) as well as test the sensitivity of our findings

24. See https://www.livingwhole.com.au/what-happens-when-you-boil-fluoride

to sample restrictions (e.g., naturally occurring fluoride, high-risk pregnancies), alternative model specifications (e.g., exclusion of key controls), falsification strategies (e.g., placebo tests), and identification approaches (e.g., stacked DiD).

3.6.1 Avoidance Behavior or Selective Migration and Fertility

Our identification strategy could be compromised by avoidance behavior or selective migration and fertility. In this section, we attempt to rule out these threats. Regarding avoidance behavior, fluoridation decisions were made at the community (and thus county) level, rather than by individuals. Moreover, during our study period, water fluoridation was widely promoted as a public health success, primarily for its dental health benefits. Concerns about potential adverse health effects were not yet widespread. Public trust in municipal water quality was further reinforced by the Safe Drinking Water Act (SDWA), enacted in 1974, which established stringent federal standards for public water systems.[25] These regulations applied to all public supplies, including those subject to fluoridation, and contributed to high levels of consumer confidence in tap water throughout the 1980s and early 1990s. In contrast, bottled water is not covered by the SDWA, and bottled water producers are not federally required disclose fluoride content, making it potentially a less transparent alternative.[26] Taken together, these factors suggest that avoidance behavior is minimal.

As for selective migration, households could have relocated toward or away from fluoridated counties once fluoridation schemes were announced, or that parents might have altered fertility timing in response to perceived benefits or risks. To test for such sorting, we regress all the maternal, pregnancy, and infant characteristics on an indicator of water fluoridation implementation. The results are presented in Table 3.20. Across all specifications, the coefficients are generally insignificant for most of the observable characteristics, indicating no change in the composition of residents and births coinciding with the fluoridation rollout. This finding supports that our estimated treatment effects are unlikely to be driven by selective migration or fertility.

3.6.2 Naturally Occurring Fluoride

While our primary analysis focuses on artificially adjusted fluoride levels through deliberate public health interventions, fluoride can also naturally occur in groundwater,

25. https://www.epa.gov/laws-regulations/summary-safe-drinking-water-act
26. https://en.wikipedia.org/wiki/Water_fluoridation_in_the_United_States

often exceeding recommended public health limits. Analyzing counties with naturally occurring fluoride provides additional insight into whether our main findings reflect fluoride exposure more broadly or specifically arise from public fluoridation practices. Accordingly, we examine whether our results remain robust when exposure is defined by naturally occurring fluoride concentrations. To this end, we estimate the following regression:

$$Y_{icmt} = \alpha_0 + \alpha_1 Natural_{icmt} + \alpha_2 X_{icmt} + \mu_j + \omega_m + \delta_t + \epsilon_{icmt}, \tag{3.3}$$

where $Natural_{icmt}$ indicates whether the mother of infant i, born in month m of year t, resides in a county with naturally occurring (rather than adjusted) fluoride during her pregnancy. State fixed-effects (μ_j) replace county fixed-effects due to perfect collinearity between county characteristics and naturally occurring fluoride indicators. Standard errors are clustered at the county level.

The results are summarized in Tables 3.21 and 3.22. Naturally occurring fluoride exhibits no statistically significant effect on birth weight, very low birth weight, or prematurity. However, consistent with our main findings, we find a statistically significant increase in CNSA anomalies. Specifically, residing in counties with naturally occurring fluoride is associated with a 0.014 percentage point increase in the probability of CNSA—slightly lower than the 0.02 percentage point effect observed for policy-induced fluoride. This finding suggests that fluoride concentrations, regardless of whether introduced deliberately or naturally occurring, may disproportionately affect CNSA outcomes. Furthermore, it provides supportive evidence for our previous findings regarding seasonal heterogeneity, particularly that acute, high-concentration exposure (as may occur during winter) might play a critical role in triggering CNSA anomalies.

3.6.3 High-Risk Pregnancies

In our main analysis, we did not account for underlying maternal medical conditions that could contribute to adverse birth outcomes. As described in Section 3.3, we now include in Eq. 3.1 an indicator for whether any of the following risk factors are present: anemia, diabetes, hypertension, or eclampsia. We repeat this estimation only for congenital anomalies, as these risk factors are available over the full period for congenital anomalies only but not birth weight (i.e., starting in 1989).[27] The updated estimates are presented in Table 3.23. We find that the effect of water fluoridation remains robust to the inclusion of these controls. Specifically, the increase in central nervous system anomalies persists, with an estimated effect of 0.023 percentage points, nearly identical to the original estimated effect in Section 3.5.1.

3.6.4 Mother, Birth, and Pregnancy Characteristics

Next, we assess whether our results are sensitive to the inclusion/exclusion of maternal (age, race, education, marital status), birth (infant sex, birth order, birth type), and pregnancy-related (number of prenatal visits) characteristics, which were included as controls in Equation 3.1. To do so, we estimate three alternative specifications. In the first, we exclude all mother, birth, and pregnancy controls. In the second, we exclude only the birth and pregnancy controls, retaining maternal characteristics. In the third, we exclude only the maternal controls, while keeping birth and pregnancy characteristics. The results for each outcome variable are summarized in Tables 3.24 through 3.31. In general, the effects on birth weight, very low birth weight, prematurity, and CNSA remain robust across most specifications.

3.6.5 Placebo Test

To further examine the robustness of our findings, we conduct two placebo tests.

In the first placebo test, we randomly assign treatment status to observations. Specifically, we generate a pseudo-treatment variable to simulate random treatment exposure. We consider two variants of this placebo design. In the first, we perform a simple random assignment of treatment without any covariate adjustment. In the second, we enhance the credibility of the placebo by performing stratified random assignment, using all available maternal, pregnancy, and birth characteristics as stratifying variables. This ensures that treatment is balanced across subgroups while preserving the joint distribution of key covariates. In both specifications, we re-estimate Eq. 3.1 using the pseudo-treatment indicator and assess whether any significant effects emerge. The results are presented in Tables 3.32 through 3.35. As expected, the placebo results yield null effects across all outcomes, thereby reinforcing the causal interpretation of our findings.

27. We do not conduct the analysis for birth weight and premature births because it would result in the loss of more than 10 million observations.

In the second placebo test, we assess the extent to which observed covariates predict treatment status. Specifically, we regress the treatment indicator on all maternal, birth, and pregnancy-related controls. If exposure to water fluoridation during pregnancy is exogenously determined, we would expect these observable characteristics to be uncorrelated with treatment assignment. Of the ten covariates included, the results presented in Table 3.36 show that eight have statistically insignificant coefficients. Overall, these results suggest that prenatal exposure to water fluoridation is largely uncorrelated with observed characteristics, supporting the assumption that there is no selection based on the observables.

3.6.6 Stacked Difference-in-Differences

Finally, we assess the robustness of our results by adopting an alternative estimation strategy, the stacked difference-in-differences (stacked DiD) approach, which is particularly suited for settings with staggered policy adoption. The staggered DiD design, as presented the TWFE specification described in Eq. 3.1, rely on a crucial assumption: homogeneous treatment effects, meaning the treatment's impact does not vary across groups or over time. The rationale behind this assumption is that the TWFE model makes 2×2 comparisons, including those between earlier-treated units (serving temporarily as controls) and later-treated units. Such comparisons remain valid only if the effect of treatment remains stable across both periods and groups, irrespective of whether the treatment itself was randomly assigned (Baker, Larcker, and Wang, 2022). However, this assumption could be violated in our context. Roberts (2024) outlines several plausible reasons why the homogeneity assumption might not hold. For instance, other fluoride sources (e.g., dental treatments and dietary intake) and access to dental care may have varied significantly across birth cohorts included in their analysis. This is less problematic in our work since our analysis is explicitly concerned with the short-term effects of prenatal fluoride exposure, confined to a narrower timeframe than Roberts (2024), who examines the long-run consequences of early childhood fluoride exposure. Nonetheless, Roberts (2024) raises another point, which remains a legitimate concern for our study since we both construct the fluoridation data from the same underlying source. Specifically, heterogeneous treatment effects may arise given that water fluoridation occurs at the level of public water systems, where many counties comprise multiple systems, alongside residents who rely on private well water. Accordingly, the proportion of county residents consuming fluoridated water could vary substantially across counties.

To address potential biases arising from such heterogeneity, recent econometric literature proposes various alternative DiD approaches, which primarily differ based on how they define suitable control groups and integrate covariates into their analyses (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Cengiz et al., 2019). Specifically, the estimator developed by Callaway and Sant'Anna (2021) offers a flexible framework that computes group-time average treatment effects (ATTs) using only clean comparisons between treated units and appropriate control groups (e.g., never-treated or not-yet-treated). However, this approach is computationally intensive in settings like ours with hundreds of treatment cohorts, each requiring a separate estimation step. Likewise, Sun and Abraham (2021) propose an interaction-weighted estimator that addresses the shortcomings of TWFE by estimating event-time-specific treatment effects for each cohort through a fully interacted event study design. Although this method is computationally more tractable than Callaway and Sant'Anna (2021)'s, its coefficients do not correspond to a single ATT. Instead, they reflect weighted averages of cohort-specific dynamic effects, making their interpretation less straightforward, particularly when the focus is on summarizing the average impact of treatment rather than its temporal dynamics.

Given these limitations, and following recent empirical applications including Roberts (2024) and Cengiz et al. (2019), we adopt the stacked DiD approach as a robustness check. This design involves constructing event-specific datasets, where for each "event" (defined by the timing of fluoridation initiation), treated and "clean" control groups are identified, given a predetermined time frame. These event-specific datasets are subsequently combined or "stacked" relative to their event times. Thus, for any given treatment onset time, clean control observations include only those units that remain untreated throughout the specified evaluation window while excluding observations that were previously treated or soon-to-be treated. As a result, the control group shrinks over time, as treated counties progressively exit the control pool. Hence, even with heterogeneous treatment effects, comparisons are made with respect to clean controls.

To implement this approach, we first aggregate observations to the month-year-county level by computing their means, thereby transforming our repeated cross-sectional dataset into a panel structure.[28] Then, we explicitly incorporate indicators for each stacked event (i.e., stack level), denoted by j, in the collapsed version of Eq. 3.1, as such

28. This step reduces data dimensionality; otherwise, estimation would be computationally burdensome, given the dataset's substantial initial size.

$$Y_{cmtj} = \alpha_0 + \alpha_1 Treatment_{cmtj} + \alpha_3 X_{cmtj} + \mu_c + \omega_m + \delta_t + \epsilon_{cmtj}, \tag{3.4}$$

We estimate four specifications for equation 3.4: (i) without explicitly specifying a window size, thus defaulting to the entire sampling period; (ii) specifying a window of one year before and one year after treatment, closely aligning with the gestation period; and (iii) again without specifying a window but explicitly restricting comparisons to observations that are never treated.[29] In all these estimations, sample weights are applied.[30] Standard errors are clustered at the county level.

Across all specifications, the effects of water fluoridation on birth weight (in grams) and the probability of CNSA (in percentage points) remain robust. In the first case (i), where no window is specified, Tables 3.37 and 3.38 show that water fluoridation decreases birth weight by 19.608 g and increases the likelihood of CNSA by 0.042 percentage points. These effects are larger than those observed in the main analysis (11.248 g and 0.023 percentage points, respectively); however, the effects on very low birth weight and premature births lose statistical significance. A similar pattern holds when a one-year window is specified (i.e., case ii) shown in Tables 3.39 and 3.40: birth weight declines by 12.995 g, and the probability of CNSA increases by 0.054 percentage points following fluoridation. Finally, when no window is specified, and only the never-treated observations serve as the control group (i.e., case iii), only the effects on birth weight and CNSA remain statistically significant, with signs and magnitudes similar to those in case (i) when both the never and not-yet-treated comprise the control group (18.431 g and 0.039 percentage points, respectively), as summarized in Tables 3.41 and 3.42.

3.7 Cost of Water Fluoridation

To place our estimated health effects in economic perspective, we compute two sets of monetary costs based on our estimates: immediate hospital expenditures incurred at birth and long-term costs that accrue over the life course. All figures are expressed in 2024 dollars and summarized in Tables 3.43 and 3.44. It should be noted that we rely on the estimates from the main analysis (Section 3.5.1), which are the most conservative.

29. In cases (i) and (ii), the control group is comprised of the never and not-yet treated.

30. The basic unweighted stack might implicitly weight cohorts unevenly; thus, it's recommended to apply sample weights so that each cohort's contribution to the final ATT is proportional to its size (see Wing, Freedman, and Hollingsworth, 2024).

As shown in the heterogeneity analyses (Section 3.5.2), the effects can be more pronounced. Therefore, the actual economic burden of water fluoridation can potentially exceed the costs we calculate below.

3.7.1 Short-term Costs

For the immediate costs, we rely on the following studies that estimated the hospital cost at birth: per-case hospital charges for very-low–birth-weight infants and premature births are taken from Rogowski (1998) and Kowlessar, Jiang, and Steiner (2013), respectively; the corresponding cost for neonatal admissions with CNSA is drawn from Koschnitzky et al. (2022). We inflate the original dollar values to 2024 using the Consumer Price Index (CPI) published by the US Bureau of Labor Statistics. Table 3.43 reveals that the total short-term hospital costs of water fluoridation are estimated to be at least $4 million for 100 births.

3.7.2 Long-term Costs

For the long-term costs, we evaluate the lifetime economic losses along three channels. First, we compute the productivity loss of premature births using estimates from Waitzman, Jalali, and Grosse (2021) (channel 1). Second, we assign a lifetime cost based on the CDC's 2016 estimate (CDC, 2016) (channel 2). All long-run values are inflated to 2024 dollars. As shown in Table 3.43, these two channels account for at least $315 billion of long-run losses for 100 births.

Third, we quantify the effect of birth weight reductions on adult earnings (channel 3). Specifically, the 0.34% decrease in mean birth weight we observe corresponds to a reduction of approximately 0.019 SD.[31] We convert this change into an earnings penalty using the elasticity reported in Lambiris et al. (2022). Assuming a 40-year working life, the present value loss in median personal income ranges from $304 to $743 per person (Table 3.44. Aggregated across all treated individuals in our sample (22.6 million births), the total present value loss ranges from $6.9 to $16.8 billion.

Altogether, these short- and long-term costs indicate that the economic burden of water fluoridation may substantially offset, if not exceed, the potential labor market

31. The standard deviation in the control group ($T = 0$) is 589.386.
We consider the 2023 real median personal income from the Federal Reserve Bank of St. Louis (FRED) to compute the effect on earnings.

benefits documented in prior work (see e.g., Glied and Neidell, 2010; Aggeborn and Öhman, 2021).

3.8 Conclusion

This paper revisits a public-health intervention long considered unequivocally beneficial and cost-effective. Leveraging staggered roll-outs of community water fluoridation between 1982 and 1992, we find that prenatal exposure lowers mean birth-weight by about 11 grams (0.34%), raises the probability of very-low birth-weight and prematurity, and, most strikingly, nearly doubles the incidence of central nervous system anomalies. These results are generally robust across several robustness checks. Heterogeneity analyses reveal three main results. First, exposure during the first trimester is the most vulnerable window. Second, the adverse effects are largest for infants born to White, married, and better-educated mothers, that is, the groups that are typically more advantaged. Third, cumulative low-dose exposure during early pregnancy primarily harms birth weight, whereas acute high-dose exposure disproportionately triggers central nervous system malformations. Although large-scale community fluoridation is most prevalent in the US, the results uncovered in this study are likely to translate to other developed regions, where fluoride is added to water.

Placing these estimates in monetary terms shows that hospital costs at delivery plus discounted lifetime productivity losses can offset the dental-care savings traditionally used to justify fluoridation. Our back-of-the-envelope calculations imply a present-value income loss of roughly $304–$743 per exposed infant. Thus, every dollar spent on fluoridation might generate negative net returns once unintended health damages are internalized. Accordingly, the evidence presented here calls for a careful re-evaluation of community fluoridation programs, including a reassessment of target fluoride concentrations, seasonal dosing practices, and the continued necessity of universal exposure in light of widespread access to alternative fluoride sources.

Tables and Figures

FIGURE (3.1) Event Study for Birth Weight and Premature Births

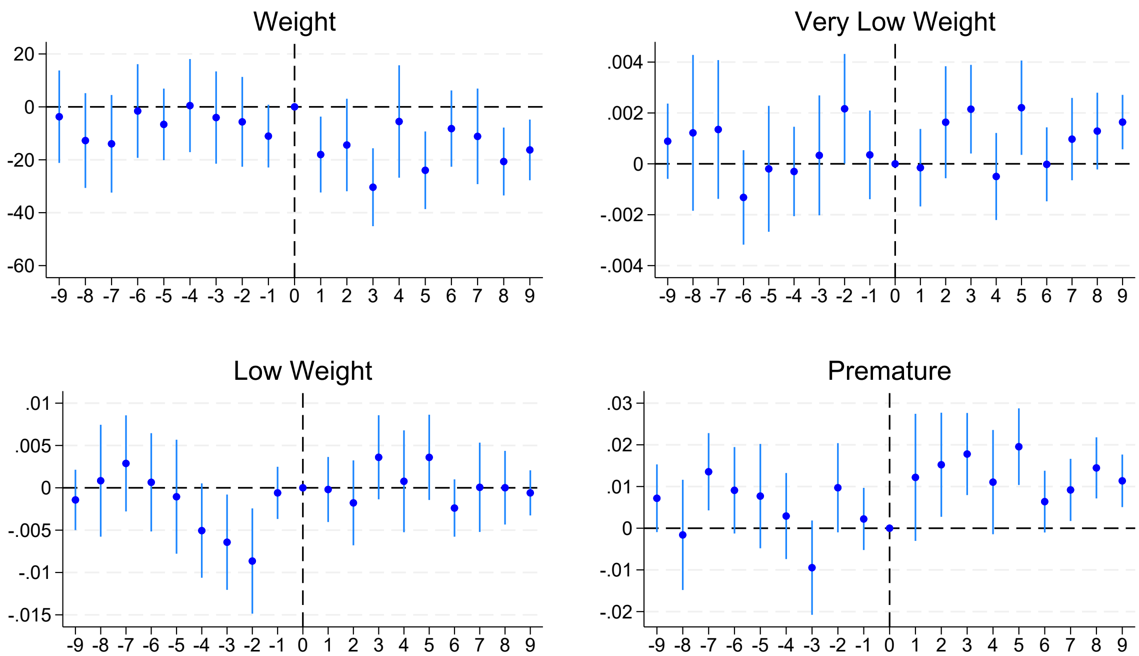

Note: The x-axis corresponds to the time (months) relative to water fluoridation data. Negative (positive) values are the number of months prior (post) the fluoridation. On the x-axis, 1,2, and 3 correspond to the third trimester; 4, 5, and 6 correspond to the second trimester; 7, 8, and 9 correspond to the first trimester. The y-axis measures the changes in the outcome variables at each month, with their 95% confidence interval. Standard errors are clustered at the county level. Outcome variables: birth weight (in g) and indicators of very low weight, low weight, and premature births. These dynamic plots are generated while accounting for all covariates (see footnote 18).

FIGURE (3.2) Event Study for Congenital Anomalies

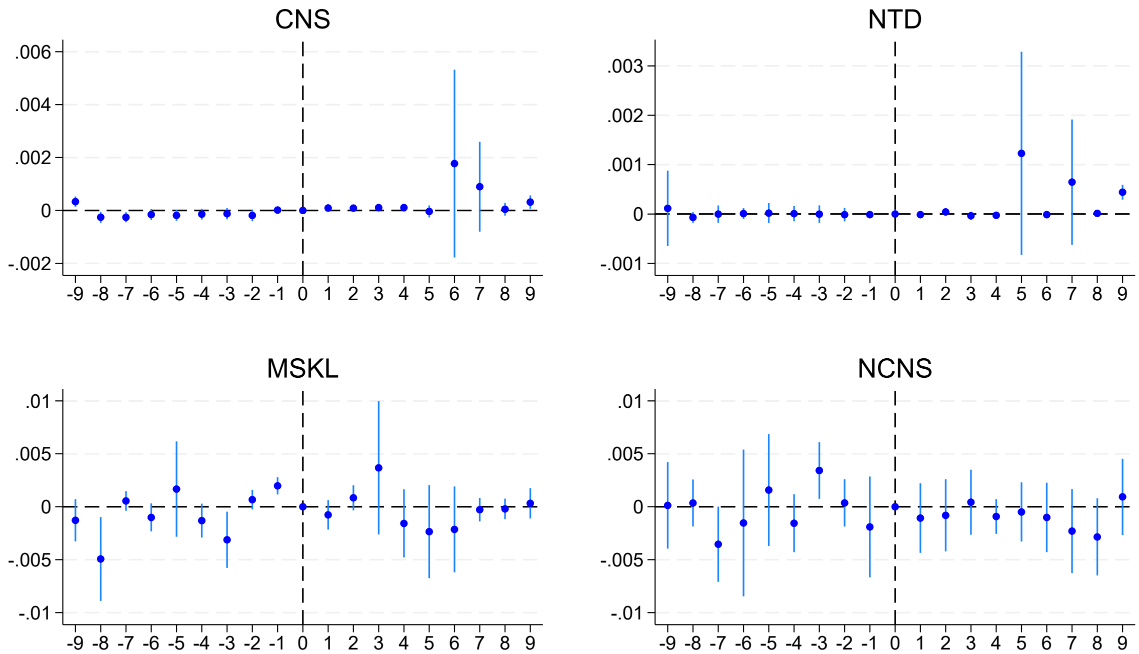

Note: The x-axis corresponds to the time (months) relative to water fluoridation data. Negative (positive) values are the number of months prior (post) the fluoridation. On the x-axis, 1,2, and 3 correspond to the third trimester; 4, 5, and 6 correspond to the second trimester; 7, 8, and 9 correspond to the first trimester. The y-axis measures the changes in the outcome variables at each month, with their 95% confidence interval. Standard errors are clustered at the county level. Outcome variables: indicators of central nervous system anomalies (CNS), neural tube defects (NTD), musculoskeletal (MSKL) anomalies, and non-CNS anomalies. These dynamic plots are generated while accounting for all covariates (see footnote 18).

TABLE (3.1) Definition of Variables

Variable		Definition
		Panel A: Birth Outcomes
weight	Birth Weight	Number of grams of birth weight
very low weight	Very Low Birth Wright	= 1 if birth weight < 1500g
low weight	Low Birth Wright	= 1 if birth weight < 2500g
premature	Premature Birth	= 1 if gestation length < 38 weeks
CNSA	Central Nervous System Anomalies	= 1 if the following is reported: Hydrocephalus, Microcephaly, Other CNS Disorders
NTD	Neural Tube Defects	= 1 if the following is reported: Anencephaly, Spina Bifida
MSKLA	Musculoskeletal Anomalies	= 1 if the following is reported: Other Musculoskeletal/Integumental Anomalies
NCNSA	Non-Central Nervous System Anomalies	= 1 if the following is reported: Omphalocele, Heart Malformation, Cleft Lip/Palate, Down's Syndrome
		Panel B: Mother Characteristics
age	Age	Age of mother in years (18–44 years)
race	Race	Dummy variables for race: White, Black, Other
status	Marital Status	Dummy variables for marital status: Married, Unmarried
educ	Education	Dummy variables for attaining: College Degree/Higher, High School, Elementary/No Formal Education
		Panel C: Pregnancy Characteristics
prenatal	Prenatal Visits	Number of prenatal visits (0–47 visits)
risky	Risky Pregnancy Factors	= 1 if medical risk factors are reported: anemia, diabetes, hypertension (chronic or pregnancy-related), eclampsia
		Panel D: Birth Characteristics
sex	Infant's sex	Dummy variables for infant's sex at birth: Male, Female
parity	Birth Parity	Dummy variables for birth parity: First Birth, Not the First Birth
type	Birth Type	Dummy variables for birth type: Singleton Births, Twins (or more) births

TABLE (3.3) Summary Statistics by Treatment Status

Variable	Never Treated	Not-Yet Treated	Treated
Age	26.442	25.538	26.42
White	0.740	0.846	0.775
Black	0.212	0.127	0.180
Other Race	0.048	0.028	0.044
Elementary or No Education	0.043	0.024	0.039
High School	0.514	0.462	0.476
College or Higher	0.311	0.247	0.329
Prenatal Visits	10.315	10.862	10.90
First Birth	0.399	0.408	0.401
Singleton	0.979	0.982	0.978
Male	0.512	0.512	0.512
Observations	4,125,310	273,722	20,680,696

Notes: This table summarizes the mean of the explanatory variables among the never treated (column 1) and treated (columns 2 and 3 in the pre-treatment and post-treatment months).

TABLE (3.2) Summary Statistics of Explanatory Variables

Variable	Treated					Control				
	Observations	**Mean**	**Std. Dev.**	**Min**	**Max**	**Observations**	**Mean**	**Std. Dev.**	**Min**	**Max**
Panel A: Mother Characteristics										
Age	22,569,123	26.419	5.434	17	45	4,877,108	26.391	5.471	17	45
White	22,569,123	0.782	0.413	0	1	4,877,108	0.746	0.435	0	1
Black	22,569,123	0.174	0.379	0	1	4,877,108	0.208	0.406	0	1
Other Race	22,569,123	0.044	0.206	0	1	4,877,108	0.046	0.210	0	1
Elementary or No Education	22,569,123	0.039	0.194	0	1	4,877,108	0.042	0.200	0	1
High School	22,569,123	0.468	0.499	0	1	4,877,108	0.511	0.45	0	1
College or Higher	22,569,123	0.332	0.471	0	1	4,877,108	0.307	0.461	0	1
Panel B: Pregnancy Characteristics										
Prenatal Visits	20,680,969	11.012	4.033	0	47	4,399,032	10.349	4.182	0	47
Panel C: Birth Characteristics										
First Birth	22,569,123	0.402	0.490	0	1	4,877,108	0.4	0.49	0	1
Singleton	22,569,123	0.978	0.146	0	1	4,877,108	0.979	0.144	0	1
Male	22,569,123	0.512	0.5	0	1	4,877,108	0.512	0.5	0	1

Notes: The left panel presents the descriptive statistics of the treated group ($Treatment = 1$), while the right panel presents the descriptive statistics of the control group ($Treatment = 0$). Column 1 corresponds to the number of observations. Columns 2, 3, 4, and 5 summarize the mean, standard deviation, minimum, and maximum, respectively. Panel A, B, and C list the mother, pregnancy, and birth characteristics.

TABLE (3.4) Effect of Water Fluoridation on Birth Weight and Premature Births

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
Treatment	-11.248***	0.001**	0.001	0.005 **
	(3.831)	(0.0003)	(0.001)	(0.002)
White	129.010***	0.002***	-0.008***	-0.023***
	(7.568)	(0.0004)	(0.001)	(0.001)
Black	-66.208***	0.013***	0.028***	0.054***
	(7.078)	(0.001)	(0.001)	(0.002)
Married	75.273***	-0.002***	-0.018***	-0.027***
	(2.944)	(0.0002)	(0.001)	(0.001)
Age	-0.021	0.0004***	0.001***	0.002***
	(0.173)	(0.00002)	(0.0001)	(0.0001)
High School	-17.774***	0.002***	0.006***	0.005***
	(4.652)	(0.001)	(0.001)	(0.001)
College or Higher	39.37***	0.001	-0.007***	-0.009 ***
	(6.210)	(0.001)	(0.001)	(0.001)
Prenatal Visits	23.009***	-0.003***	-0.005***	-0.0123***
	(0.506)	(0.0001)	(0.0001)	(0.0003)
First Birth	-95.292***	0.005***	0.015***	0.010***
	(1.608)	(0.0002)	(0.0003)	(0.001)
Singleton	973.291***	-0.085***	-0.363***	-0.448***
	(3.87)	(0.002)	(0.002)	(0.003)
Male	124.072***	-0.00002	-0.011***	0.014***
	(0.668)	(0.00005)	(0.0003)	(0.0002)
Constant	1984.395***	0.108***	0.445***	0.707***
	(8.164)	(0.002)	(0.003)	(0.005)
Observations	25,059,925	25,080,000	25,080,000	25,080,000
Mean of Control	3334.065	0.011	0.057	0.164

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($Treatment = 0$) is reported in the last row. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.5) Effect of Water Fluoridation on Birth Weight of Infants with Very Low and Low Weight

	Very Low Weight = 1	**Low Weight = 1**
Treatment	-4.879	-0.592
	(10.511)	(3.345)
White	-21.078***	-22.732***
	(4.313)	(1.506)
Black	-56.355***	-33.524***
	(4.748)	(1.8767)
Married	-11.663*	2.114**
	(6.843)	(0.824)
Age	-2.620***	-1.504***
	(0.181)	(0.065)
High School	25.254***	-3.911***
	(7.978)	(1.055)
College or Higher	0.651	-5.408***
	(8.244)	(1.204)
Prenatal Visits	15.901***	6.574***
	(0.450)	(0.179)
First Birth	-39.104***	-10.466***
	(1.787)	(0.607)
Singleton	16.009***	97.110***
	(2.275)	(1.102)
Male	1.723	-14.606***
	(1.419)	(0.528)
Constant	974.229***	2132.640***
	(12.308)	(4.726)
Observations	257,066	1,368,079

Notes: Each column corresponds to regression results of birth weight in two sub-samples. In column 1, the sample is births with very low birth weight ($<$ 1500 g). In column 2, the sample is the births with low birth weight ($<$ 2500 g). Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.6) Effect of Water Fluoridation on Congenital Anomalies

	CNSA	**NTD**	**MSKLA**	**NCNSA**
Treatment	0.0002***	0.0003	0.0008	0.0002
	(0.0001)	(0.0003)	(0.0006)	(0.0002)
White	0.00002	0.00002	-0.0001	0.00009
	(0.00004)	(0.00003)	(0.0001)	(0.0001)
Black	0.00003	-0.0001***	0.0002	-0.0006***
	(0.00006)	(0.00004)	(0.0002)	(0.0001)
Married	-0.0002***	-0.00003	-0.00003	-0.0003***
	(0.00004)	(0.00002)	(0.00005)	(0.00008)
Age	0.00001***	-0.000001	0.00002***	0.00007***
	(0.000002)	(0.000001)	(0.000005)	(0.000005)
High School	-0.0001*	-0.0002***	-0.0003	-0.0004**
	(0.00008)	(0.00006)	(0.0002)	(0.0002)
College or Higher	-0.0003***	-0.0002***	-0.0002	-0.0008***
	(0.00008)	(0.00007)	(0.0002)	(0.0002)
Prenatal Visits	-0.00001***	-0.00002***	-0.0000001	-0.00004***
	(0.000004)	(0.000004)	(0.000009)	(0.00001)
First Birth	0.00005**	-0.00001	0.0005***	0.0001**
	(0.00002)	(0.00001)	(0.00005)	(0.00004)
Singleton	-0.0005***	-0.0004***	0.0001	-0.0012***
	(0.0001)	(0.00007)	(0.00009)	(0.0002)
Male	0.00007***	-0.00003**	-0.0003***	0.0003***
	(0.00003)	(0.00001)	(0.00004)	(0.00005)
Constant	0.001***	0.001***	0.0006	0.0029***
	(0.0002)	(0.0003)	(0.0005)	(0.0003)
Observations	5,547,005	9,400,798	9,400,798	9,400,798
Mean of Control	0.0004	0.0003	0.002	0.003

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. Congenital anomalies are reported starting 1989, so they are not present in our entire sample (1982-1993). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.7) Effect of Water Fluoridation on Birth Weight and Premature Births by Trimester of Exposure

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
Trimester 1	-17.664***	0.001***	0.003***	0.014***
	(3.763)	(0.0004)	(0.001)	(0.003)
Trimester 2	-11.196	0.001	0.003**	0.017**
	(7.304)	(0.001)	(0.002)	(0.007)
Trimester 3	-8.821	0.0002	0.002	0.002
	(7.597)	(0.001)	(0.002)	(0.006)
Observations	24,928,049	24,944,299	24,944,299	24,944,299
Mean of Control	3334.065	0.011	0.057	0.164

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). The mean of the outcome variable in the control group ($T = 0$) is reported in the last row. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.8) Effect of Water Fluoridation on Birth Weight and Premature Births by Mother's Race

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
		Panel A: White Mothers		
Treatment	-11.026***	0.001***	0.001*	0.006**
	(3.837)	(0.0003)	(0.0007)	(0.003)
Observations	19,473,735	19,487,823	19,487,823	19,487,823
Mean of Control	3394.017	0.007	0.046	0.139
		Panel B: Black Mothers		
Treatment	-13.634 **	0.0001	0.001	-0.003
	(6.579)	(0.002)	(0.003)	(0.004)
Observations	4,598,609	4,603,693	4,603,693	4,603,693
Mean of Control	3130.645	0.024	0.097	0.2497

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.9) Effect of Water Fluoridation on Congenital Anomalies by Mother's Race

	CNSA	**NTD**	**MSKLA**	**NCNSA**
	Panel A: White Mothers			
Treatment	0.0002***	0.0002	0.0004	4.48×10^{-6}
	(0.00004)	(0.0003)	(0.0005)	(0.0002)
Observations	4,233,994	7,164,914	7,164,914	7,164,914
Mean of Control	0.0004	0.0003	0.002	0.003
	Panel B: Black Mothers			
Treatment	0.0003**	0.003	0.006	0.003
	(0.0001)	(0.003)	(0.004)	(0.002)
Observations	1,031,403	1,747,825	1,747,825	1,747,825
Mean of Control	0.0003	0.0003	0.001	0.002

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.10) Effect of Water Fluoridation on Birth Weight and Premature Births by Mother's Marital Status

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
	Panel A: Married Mothers			
Treatment	-11.46***	0.001***	0.001*	0.006**
	(4.112)	(0.0002)	(0.001)	(0.002)
Observations	18,569,588	18,582,925	18,582,925	18,582,925
Mean of Control	3394.585	0.007	0.044	0.14
	Panel B: Single Mothers			
Treatment	-13.929**	0.001	0.001	0.003
	(6.069)	(0.001)	(0.002)	(0.005)
Observations	6,490,333	6,497,071	6,497,071	6,497,071
Mean of Control	3182.111	0.019	0.087	0.224

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.11) Effect of Water Fluoridation on Congenital Anomalies by Mother's Marital Status

	CNSA	**NTD**	**MSKLA**	**NCNSA**
	Panel A: Married Mothers			
Treatment	0.0006***	0.0001	0.0004	0.0004
	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Observations	3,903,102	6,539,236	6,539,236	6,539,236
Mean of Control	0.0004	0.0003	0.002	0.004
	Panel B: Single Mothers			
Treatment	-0.0005	0.0007***	0.002	0.0001
	(0.0004)	(0.0001)	(0.002)	(0.0002)
Observations	1,643,902	2,861,562	2,861,562	2,861,562
Mean of Control	0.0004	0.0003	0.002	0.002

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.12) Effect of Water Fluoridation on Birth Weight and Premature Births by Mother's Education

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
	Panel A: Mothers with a College Degree			
Treatment	-14.75***	0.0005	0.002	0.006*
	(5.09)	(0.0004)	(0.001)	(0.003)
Observations	8,842,253	8,848,404	8,848,404	8,848,404
Mean of Control	3400.924	0.008	0.044	0.14
	Panel B: Mothers without a College Degree			
Treatment	-9.458**	0.001**	0.001	0.005**
	(3.67)	(0.0004)	(0.001)	(0.002)
Observations	16,217,672	16,231,596	16,231,596	16,231,596
Mean of Control	3282.604	0.011	0.064	0.196

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight (< 1500 g). Low Weight: indicator of low weight (< 2500 g). Premature: indicator of premature births (gestation length < 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.13) Effect of Water Fluoridation on Congenital Anomalies by Mother's Education

	CNSA	**NTD**	**MSKLA**	**NCNSA**
	Panel A: Mothers with a College Degree			
Treatment	0.0003**	-0.0004	0.001	-0.0001
	(0.0001)	(0.0005)	(0.002)	(0.0005)
Observations	2,157,394	3,772,255	3,772,255	3,772,255
Mean of Control	0.0003	0.0002	0.002	0.003
	Panel B: Mothers without a College Degree			
Treatment	0.0002***	0.001	-0.0001	0.001
	(0.0001)	(0.0003)	(0.0004)	(0.0003)
Observations	3,389,611	5,628,543	5,628,543	5,628,543
Mean of Control	0.0004	0.0004	0.001	0.002

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.14) Effect of Water Fluoridation on Birth Weight and Premature Births by Infant's Sex

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
	Panel A: Male Births			
Treatment	-15.092***	0.001***	0.0004	0.007***
	(3.81)	(0.0004)	(0.001)	(0.002)
Observations	12,828,426	12,838,916	12,838,916	12,838,916
Mean of Control	3394.791	0.011	0.051	0.17
	Panel B: Female Births			
Treatment	-7.305*	0.0003	0.002	0.003
	(4.20)	(0.0004)	(0.001)	(0.003)
Observations	12,231,496	12,241,081	12,241,081	12,241,081
Mean of Control	3270.385	0.011	0.063	0.157

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight (< 1500 g). Low Weight: indicator of low weight (< 2500 g). Premature: indicator of premature births (gestation length < 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.15) Effect of Water Fluoridation on Congenital Anomalies by Infant's Sex

	CNSA	**NTD**	**MSKLA**	**NCNSA**
		Panel A: Male Births		
Treatment	0.0002***	-0.00003	0.0004	-0.001*
	(0.0001)	(0.001)	(0.0005)	(0.0004)
Observations	2,837,544	4,811,745	4,811,745	4,811,745
Mean of Control	0.0005	0.0003	0.002	0.003
		Panel B: Female Births		
Treatment	0.0002***	0.001***	0.001*	0.001**
	(0.0001)	(0.0001)	(0.001)	(0.0004)
Observations	2,709,461	4,589,053	4,589,053	4,589,053
Mean of Control	0004	0.0003	0.002	0.002

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.16) Effect of Water Fluoridation on Birth Weight and Premature Births by Birth Order (Parity)

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
		Panel A: First Birth		
Treatment	-11.492***	0.0004	0.002	0.005*
	(3.801)	(0.0004)	(0.001)	(0.003)
Observations	10,060,421	10,068,147	10,068,147	10,068,147
Mean of Control	3302.094	0.011	0.057	0.154
		Panel B: Not First Birth		
Treatment	-10.927**	0.001**	0.001	0.006**
	(4.441)	(0.0004)	(0.001)	(0.002)
Observations	14,999,499	15,011,848	15,011,848	15,011,848
Mean of Control	3355.346	0.011	0.056	0.17

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight (< 1500 g). Low Weight: indicator of low weight (< 2500 g). Premature: indicator of premature births (gestation length < 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.17) Effect of Water Fluoridation on Congenital Anomalies by Birth Order (Parity)

	CNSA	**NTD**	**MSKLA**	**NCNSA**
		Panel A: First Birth		
Treatment	0.0002**	0.001	0.001	-0.001
	(0.0001)	(0.0005)	(0.001)	(0.001)
Observations	2,201,211	3,707,743	3,707,743	3,707,743
Mean of Control	0.0005	0.0003	0.002	0.003
		Panel B: Not First Birth		
Treatment	0.0002***	0.0001	0.001	0.001
	(0.00005)	(0.0002)	(0.001)	(0.001)
Observations	3,345,794	5,693,055	5,693,055	5,693,055
Mean of Control	0.0004	0.0003	0.002	0.003

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.18) Effect of Water Fluoridation on Birth Weight and Premature Births by Season of Conception

	Weight	**Very Low Weight**	**Low Weight**	**Premature**
		Panel A: Summer Conceptions		
Treatment	-18.047***	0.001***	0.002	0.005
	(5.573)	(0.0004)	(0.002)	(0.004)
Observations	6,251,850	6,256,694	6,256,694	6,256,694
Mean of Control	3343.58	0.011	0.056	0.166
		Panel B: Winter Conceptions		
Treatment	5.823	-0.0001	-0.002	0.005
	(6.028)	(0.001)	(0.002)	(0.003)
Observations	2,025,265	2,026,926	2,026,926	2,026,926
Mean of Control	3331.271	0.011	0.056	0.157

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight (< 1500 g). Low Weight: indicator of low weight (< 2500 g). Premature: indicator of premature births (gestation length < 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.19) Effect of Water Fluoridation on Congenital Anomalies by Season of Conception

	CNSA	**NTD**	**MSKLA**	**NCNSA**
	Panel A: Summer Conceptions			
Treatment	0.0003	0.001***	-0.0005	-0.001***
	(0.0002)	(0.0001)	(0.001)	(0.0002)
Observations	1,386,934	2,359,873	2,359,873	2,359,873
Mean of Control	0.0004	0.0002	0.002	0.003
	Panel B: Winter Conceptions			
Treatment	0.0002**	0.0003**	0.001***	0.001
	(0.0001)	(0.0001)	(0.0002)	(0.001)
Observations	451,034	766,490	766,490	766,490
Mean of Control	0.0004	0.0004	0.002	0.003

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. The mean of the outcome variable in the control group ($T = 0$) is reported in the last row of each panel. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.20) Effect of County's Water Fluoridation Adoption on Observable Characteristics

	Panel A: Mother Characteristics					
	White	**Black**	**Married**	**Age**	**High School**	**College**
Fluoridation	0.044	-0.046	0.036	0.32**	-0.056***	0.039***
	(0.028)	(0.028)	(0.025)	(0.153)	(0.021)	(0.014)
	Panel B: Pregnancy and Birth Characteristics					
	Prenatal Visit	**First Birth**	**Singleton**	**Male**		
Fluoridation	0.621***	0.008	-0.0005	0.0002		
	(0.229)	(0.007)	(0.0003)	(0.0004)		

Notes: Each column corresponds to regression results of the following outcome variables. White: indicator of White mother. Black: indicator of black mother. Married: indicators of married mother. Age: mother's age. High School: indicator of mother's highest education level High School. College: indicator of mother's highest educational level College or more. Prenatal Visits: number of prenatal visits. First Birth: indicator of whether it is first birth. Singleton: indicator of whether birth is singleton or multiple (twins or more). Male: indicator of a male infant. *Fluoridation* is an indicator of the implementation of water fluoridation. ALl regressions are done without covariates. The number of observations for each regression is 27,446,231. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and state fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.21) Effect of Natural Fluoride Birth Weight and Premature Births

	Birth Weight	**Very Low Weight**	**Low Weight**	**Premature**
Natural	2.282	-0.001	-0.001	0.001
	(7.651)	(0.001)	(0.001)	(0.002)
White	111.978***	0.001***	-0.009***	-0.028***
	(10.384)	(0.001)	(0.002)	(0.003)
Black	-60.752***	0.012***	0.023***	0.041***
	(14.970)	(0.001)	(0.002)	(0.004)
Married	76.358***	-0.003***	-0.018***	-0.028***
	(4.613)	(0.0004)	(0.002)	(0.003)
Age	1.087***	0.0004***	0.001***	0.001***
	(0.297)	(0.0000)	(0.0001)	(0.0001)
High School	-18.261***	0.002**	0.006***	0.006***
	(3.673)	(0.001)	(0.001)	(0.002)
College or Higher	28.171***	0.001*	-0.005***	-0.006***
	(6.070)	(0.001)	(0.002)	(0.002)
Prenatal Visits	21.066***	-0.003***	-0.004***	-0.011***
	(0.716)	(0.0001)	(0.0002)	(0.0004)
First Birth	-87.042***	0.005***	0.014***	0.006***
	(3.427)	(0.0003)	(0.001)	(0.001)
Singleton	958.367***	-0.086***	-0.366***	-0.422***
	(7.947)	(0.003)	(0.003)	(0.007)
Male	121.723***	-0.0002	-0.011***	0.014***
	(1.016)	(0.0001)	(0.001)	(0.001)
Constant	1991.467***	0.107***	0.453***	0.671***
	(12.020)	(0.003)	(0.004)	(0.008)
Observations	5,989,584	5,995,055	5,995,055	5,995,055

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Natural* is an indicator of natural fluoride presence. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and state fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.22) Effect of Natural Fluoride on Congenital Anomalies

	CNSA	NTD	MSKL	NCNSA
Natural	0.0001**	0.00003	0.00005	0.0004
	(0.0001)	(0.0001)	(0.001)	(0.0003)
White	0.00003	-0.00002	-0.0003	0.0003*
	(0.0001)	(0.0001)	(0.0003)	(0.0001)
Black	-0.0001	-0.0002***	-0.0003	-0.0001***
	(0.0001)	(0.0001)	(0.0003)	(0.0002)
Married	-0.00003	-0.0001**	-0.0001	-0.0003***
	(0.00004)	(0.00004)	(0.0001)	(0.0001)
Age	0.000004	-0.000004	0.0000003	0.00007***
	(0.00001)	(0.000003)	(0.000007)	(0.00001)
High School	-0.0002**	-0.0001**	-0.0003	-0.0002
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
College or Higher	-0.0003***	-0.0002***	-0.0004	-0.0004*
	(0.0001)	(0.0001)	(0.0003)	(0.0002)
Prenatal Visits	-0.00002**	-0.00002***	0.0002	-0.00002*
	(0.00001)	(0.000004)	(0.00001)	(0.00001)
First Birth	0.0001	-0.0001***	0.0002***	0.00005
	(0.0001)	(0.00003)	(0.0001)	(0.0001)
Singleton	-0.0003	-0.0004***	0.0004*	-0.002***
	(0.0002)	(0.0001)	(0.0002)	(0.0004)
Male	0.0001**	-0.0000003	-0.0003***	0.0003581***
	(0.000034)	(0.00003)	(0.0001)	(0.0001)
Constant	0.001***	0.001***	0.002**	0.003***
	(0.0003)	(0.0002)	(0.001)	(0.0004)
Observations	1,221,954	2,072,220	2,072,220	2,072,220

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Natural* is an indicator of natural fluoride presence. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.23) Effect of Water Fluoridation on Congenital Anomalies After Accounting for Pregnancy Risk Factors

	CNSA	**NTD**	**MSKLA**	**NCNS**
Treatment	0.0002***	0.0002	0.001	0.0001
	(0.00004)	(0.0003)	(0.001)	(0.0002)
White	0.00003	0.00002	-0.0001	0.0001
	(0.00004)	(0.00003)	(0.0001)	(0.0001)
Black	0.00002	-0.0001***	0.0001	-0.001***
	(0.0001)	(0.00004)	(0.0002)	(0.0001)
Married	-0.0002***	-0.00003	-0.00003	-0.0003***
	(0.00003)	(0.00002)	(0.0001)	(0.0001)
Age	0.00001***	-0.000001	0.00002***	0.0001***
	(0.000002)	(0.000001)	(0.000005)	(0.000005)
High School	-0.0001*	-0.0001***	-0.0001	-0.0003*
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
College or higher	-0.0002***	-0.0002***	-0.00004	-0.001***
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Prenatal	-0.00002***	-0.00002***	-0.000001	-0.00005***
	(0.000004)	(0.000004)	(0.00001)	(0.00001)
First Birth	0.00004**	-0.00001	0.0004***	0.0001
	(0.00002)	(0.00001)	(0.00005)	(0.00004)
Singleton	-0.0005***	-0.0004***	0.0002**	-0.001***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)
Male	0.0001**	-0.00003**	-0.0003***	0.0003***
	(0.00003)	(0.00001)	(0.00004)	(0.00004)
Risky	0.0003***	0.0001*	0.0008***	0.001***
	(0.0001)	(0.00003)	(0.0001)	(0.0001)
Constant	0.001***	0.001***	0.001	0.003***
	(0.0001)	(0.0003)	(0.0005)	(0.0003)
Observations	5,499,434	9,331,632	9,331,632	9,331,632

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator for prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.24) Effect of Water Fluoridation on Birth Weight

	(1)	**(2)**	**(3)**	**(4)**
Treatment	-8.816*	-5.720*	-13.315***	-11.248***
	(3.811)	(3.468)	(4.468)	(3.831)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	27,423,297	27,423,297	25,059,925	25,059,925

Notes: The outcome variable is birth weight (in grams). Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.25) Effect of Water Fluoridation on Very Low Birth Weight

	(1)	**(2)**	**(3)**	**(4)**
Treatment	0.0004	0.0002	0.001**	0.001**
	(0.0003)	(0.0002)	(0.0003)	(0.0003)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	27,446,231	27,446,231	25,080,000	25,080,000

Notes: The outcome variable is an indicator of very low birth weight (weight $<$ 1500 g). Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. . *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.26) Effect of Water Fluoridation on Low Birth Weight

	(1)	(2)	(3)	(4)
Treatment	0.001	0.0002	0.002*	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	27,446,231	27,446,231	25,080,000	25,080,000

Notes: The outcome variable is an indicator of low birth weight (weight $<$ 2500 g). Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.27) Effect of Water Fluoridation on Premature Births

	(1)	(2)	(3)	(4)
Treatment	0.004*	0.003	0.006**	0.005**
	(0.002)	(0.003)	(0.002)	(0.002)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	27,446,231	27,446,231	25,080,000	25,080,000

Notes: The outcome variable is an indicator of premature births (gestation length $<$ 38 weeks). Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.28) Effect of Water Fluoridation on Central Nervous System Anomalies

	(1)	(2)	(3)	(4)
Treatment	0.0002***	0.0002***	0.0002***	0.0002***
	(0.00005)	(0.0001)	(0.00005)	(0.0001)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	5,674,243	5,674,243	5,547,005	5,547,005

Notes: The outcome variable is an indicator of central nervous system anomalies. Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.29) Effect of Water Fluoridation on Neural Tube Defects

	(1)	(2)	(3)	(4)
Treatment	0.0002	0.0003	0.0002	0.0003
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	9,627,439	9,627,439	9,400,798	9,400,798

Notes: The outcome variable is an indicator of central neural tube defects. Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.30) Effect of Water Fluoridation on Musculoskeletal Anomalies

	(1)	**(2)**	**(3)**	**(4)**
Treatment	0.001 (0.00049)	0.001 (0.0005)	0.001 (0.0005098)	0.001 (0.0005727)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	9,627,439	9,627,439	9,400,798	9,400,798

Notes: The outcome variable is an indicator of musculoskeletal anomalies. Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.31) Effect of Water Fluoridation on Non-Central Nervous System Anomalies

	(1)	**(2)**	**(3)**	**(4)**
Treatment	0.00002 (0.0002)	0.0003 (0.0003)	-0.0001 (0.0002)	0.0002 (0.0002)
Mother Controls	✗	✓	✗	✓
Pregnancy & Birth Controls	✗	✗	✓	✓
Observations	9,627,439	9,627,439	9,400,798	9,400,798

Notes: The outcome variable is an indicator of non-central nervous system anomalies. Each column corresponds to regression results of Eq. 3.1. The columns differ in whether the mother, pregnancy, and birth controls are included. Mother controls: age, race, education, and marital status. Pregnancy and Birth controls: number of prenatal visits, infant sex, birth order (parity), and birth status. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.32) Effect of Randomized Treatment on Birth Weight and Premature Births: Without Stratification

	Birth Weight	**Very Low Weight**	**Low Weight**	**Premature**
Randomized Treatment	-0.159	0.0001	0.00004	0.0002
	(0.236)	(0.00004)	(0.0001)	(0.0001)
White	129.011***	0.0021***	-0.008***	-0.023***
	(7.567)	(0.0004)	(0.001)	(0.001)
Black	-66.203***	0.013***	0.0282***	0.054***
	(7.078)	(0.001)	(0.001)	(0.002)
Married	75.280***	-0.003***	-0.018***	-0.027***
	(2.944)	(0.0002)	(0.001)	(0.001)
Age	-0.021	0.0004***	0.001***	0.002***
	(0.173)	(0.00002)	(0.0001)	(0.0001)
High School	-17.789***	0.002***	0.006***	0.005***
	(4.649)	(0.001)	(0.001)	(0.001)
College of Higher	39.355***	0.001	-0.007***	-0.009***
	(6.208)	(0.001)	(0.001)	(0.001)
Prenatal Visits	23.008***	-0.003***	-0.005***	-0.013***
	(0.506)	(0.0001)	(0.0001)	(0.0003)
First Birth	-95.289***	0.005**	0.015***	0.010***
	(1.608)	(0.0002)	(0.0003)	(0.001)
Singleton	973.289***	-0.085***	-0.363***	-0.448***
	(3.87)	(0.002)	(0.002)	(0.003)
Male	124.071***	-0.00002	-0.011***	0.0139***
	(0.668)	(0.00005)	(0.0003)	(0.0002)
Constant	1975.209***	0.109***	0.446***	0.711***
	(7.545)	(0.002)	(0.003)	(0.005)
Observations	25,059,924	25,079,999	25,079,999	25,079,999

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Randomized Treatment* is a random indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.33) Effect of Randomized Treatment on Congenital Anomalies: Without Stratification

	CNSA	**NTD**	**MSKL**	**NCNSA**
Randomized Treatment	-0.00001	0.000003	0.00002	-0.0001***
	(0.00002)	(0.00001)	(0.00003)	(0.00004)
White	0.00003	0.00002	-0.0001	0.0001
	(0.00004)	(0.00003)	(0.0001)	(0.0001)
Black	0.00003	-0.0001***	0.0002	-0.001***
	(0.0001)	(0.00004)	(0.0002)	(0.0001)
Married	-0.0002***	-0.00003	-0.00003	-0.0003***
	(0.00004)	(0.00002)	(0.0001)	(0.0001)
Age	0.00001***	-0.000001	0.00002***	0.0001***
	(0.000002)	(0.000001)	(0.000005)	(0.00001)
High School	-0.0001	-0.0002***	-0.0003	-0.0004**
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
College or Higher	-0.0003***	-0.0002***	-0.0002	-0.001***
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Prenatal Visits	-0.00001***	-0.00002***	-0.0000001	-0.00004***
	(0.000004)	(0.000004)	(0.00001)	(0.00001)
First Birth	0.0001**	-0.00001	0.0005***	0.0001**
	(0.00002)	(0.00001)	(0.0001)	(0.00004)
Singleton	-0.0005***	-0.0004***	0.0001	-0.001***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)
Male	0.0001**	-0.00003**	-0.0003***	0.0003***
	(0.00003)	(0.00001)	(0.00004)	(0.00005)
Constant	0.001***	0.001***	0.001***	0.003***
	(0.0002)	(0.0001)	(0.0003)	(0.0003)
Observations	5,547,005	9,400,798	9,400,798	9,400,798

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Randomized Treatment* is a random indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.34) Effect of Randomized Treatment on Birth Weight and Premature Births: With Stratification

	Birth Weight	**Very Low Weight**	**Low Weight**	**Premature**
Randomized Treatment	-0.033	0.00002	0.0001	-0.0002
	(0.213)	(0.00004)	(0.0001)	(0.0002)
White	128.36***	0.002***	-0.007***	-0.022***
	(7.56)	(0.0004)	(0.001)	(0.001)
Black	-66.23***	0.013***	0.028***	0.054***
	(7.08)	(0.001)	(0.001)	(0.002)
Married	74.58***	-0.002***	-0.018***	-0.027***
	(2.96)	(0.0002)	(0.001)	(0.001)
Age	0.002	0.0004***	0.001***	0.002***
	(0.172)	(0.00002)	(0.0001)	(0.0001)
High School	-18.49***	0.002***	0.006***	0.005***
	(4.66)	(0.001)	(0.001)	(0.001)
College or Higher	38.49***	0.001	-0.007***	-0.009***
	(6.21)	(0.001)	(0.001)	(0.001)
Prenatal Visits	23.36***	-0.003***	-0.005***	-0.013***
	(0.51)	(0.0001)	(0.0001)	(0.0003)
First Birth	-95.27***	0.005***	0.015***	0.010***
	(1.61)	(0.0002)	(0.0001)	(0.001)
Singleton	971.33***	-0.084***	-0.363***	-0.448***
	(3.87)	(0.002)	(0.002)	(0.003)
Male	124.13***	-0.00002	-0.011***	0.014***
	(0.67)	(0.0001)	(0.0003)	(0.0002)
Constant	1974.61***	0.108***	0.446***	0.713***
	(7.60)	(0.003)	(0.003)	(0.005)
Observations	25,009,578	25,029,504	25,029,504	25,029,504

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight (< 1500 g). Low Weight: indicator of low weight (< 2500 g). Premature: indicator of premature births (gestation length < 38 weeks). *Randomized Treatment* is a random indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.35) Effect of Randomized Treatment on Congenital Anomalies: With Stratification

	CNSA	**NTD**	**MSKL**	**NCNS**
Randomized Treatment	0.00002	-0.00001	-0.00001	0.00001
	(0.00002)	(0.00002)	(0.00003)	(0.00003)
White	0.00004	0.00002	-0.0001	0.0001
	(0.00004)	(0.00003)	(0.0001)	(0.0001)
Black	0.00005	-0.0001***	0.0002	-0.001***
	(0.0001)	(0.00003)	(0.0002)	(0.0001)
Married	-0.0002***	-0.00003	-0.00004	-0.0003***
	(0.00003)	(0.00002)	(0.0001)	(0.0001)
Age	0.00001***	-0.000001	0.00003***	0.0001***
	(0.000002)	(0.000001)	(0.000005)	(0.00001)
High School	-0.0001*	-0.0002***	-0.0003	-0.0004**
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
College or Higher	-0.0003***	-0.0002***	-0.0002	-0.001***
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Prenatal Visits	-0.00002***	-0.00002***	0.0000004	-0.00004***
	(0.0000042	(0.000004)	(0.00001)	(0.00001)
First Birth	0.0001***	-0.00001	0.0005***	0.0001**
	(0.00002)	(0.00001)	(0.0001)	(0.00004)
Singleton	-0.0005***	-0.0004***	0.0001	-0.001***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)
Male	0.0001**	-0.00003**	-0.0003***	0.0003***
	(0.00003)	(0.00001)	(0.00004)	(0.00005)
Constant	0.001***	0.001***	0.001***	0.003***
	(0.0002)	(0.0001)	(0.0003)	(0.0003)
Observations	5,534,053	9,378,422	9,378,422	9,378,422

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Randomized Treatment* is a random indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.36) Effect of Selection of Observable Characteristics on Mother's Treatment Status

	Treatment
White	-0.0001
	(0.0002)
Black	-0.0004
	(0.0003)
Married	-0.001***
	(0.0002)
Age	-0.00001
	(0.00001)
High School	0.001
	(0.003)
College or Higher	0.001
	(0.003)
Prenatal Visits	0.00004
	(0.00004)
First Birth	-0.0002**
	(0.0001)
Singleton	0.0002
	(0.0002)
Male	-0.00002
	(0.00003)
Constant	0.824***
	(0.002)
Observations	25,080,000

Notes: *Treatment*: indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.37) Effect of Water Fluoridation on Birth Weight and Premature Births using Stacked DiD: No Specified Window

	Birth Weight	**Very Low Weight**	**Low Weight**	**Premature**
Treatment	-19.608***	-0.0001	0.003	0.005
	(5.739)	(0.001)	(0.002)	(0.005)
White	25.865	-0.001	0.001	-0.028
	(31.982)	(0.004)	(0.012)	(0.019)
Black	-176.678***	0.007	0.031**	0.064***
	(35.295)	(0.005)	(0.015)	(0.023)
Married	65.248***	-0.003	-0.018***	-0.032***
	(13.162)	(0.002)	(0.006)	(0.009)
Age	4.502***	-0.0002	0.001**	-0.001
	(0.994)	(0.0001)	(0.0004)	(0.001)
High School	-35.378*	-0.001	0.011*	0.021*
	(19.076)	(0.003)	(0.006)	(0.012)
College or Higher	15.110	-0.0004	-0.003	0.014
	(20.676)	(0.003)	(0.006)	(0.013)
Prenatal Visits	19.150***	-0.0013***	-0.004***	-0.014***
	(1.325)	(0.0002)	(0.0004)	(0.001)
First Birth	-72.561***	-0.001	0.016***	0.005
	(10.462)	(0.001)	(0.004)	(0.007)
Singleton	903.737***	-0.074***	-0.3585***	-0.414***
	(30.128)	(0.015)	(0.021)	(0.023)
Male	117.183***	-0.001	-0.007**	0.014**
	(8.623)	(0.001)	(0.003)	(0.005)
Constant	2117.029***	0.01***	0.425***	0.726***
	(51.982)	(0.018)	(0.027)	(0.039)
Observations	2,692,339	2,692,581	2,692,581	2,692,581

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. Sample weights are applied. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.38) Effect of Water Fluoridation on Congenital Anomalies using Stacked DiD: No Specified Window

	CNSA	**NTD**	**MSKLA**	**NCNSA**
Treatment	0.0004**	0.0004	0.008	0.011
	(0.0002)	(0.0005)	(0.006)	(0.008)
White	-0.008**	-0.003	-0.005	0.017
	(0.004)	(0.002)	(0.006)	(0.010)
Black	-0.003	-0.003	-0.013**	0.015
	(0.004)	(0.003)	(0.007)	(0.009)
Married	-0.002	-0.001*	-0.015***	-0.003
	(0.002)	(0.001)	(0.004)	(0.004)
Age	-0.0001	-0.0001	0.0004	0.001
	(0.0001)	(0.0001)	(0.0002)	(0.001)
High School	-0.018*	-0.002***	-0.018*	-0.014
	(0.009)	(0.001)	(0.011)	(0.010)
College or Higher	-0.015*	0.0003	-0.012	-0.034
	(0.009)	(0.001)	(0.010)	(0.023)
Prenatal Visits	0.00003	5.89e-06	0.0001	-0.001
	(0.0001)	(0.00004)	(0.0003)	(0.001)
First Birth	0.001	0.001	0.0001	0.009
	(0.001)	(0.001)	(0.002)	(0.005)
Singleton	0.006**	-0.001	0.005	-0.006
	(0.003)	(0.002)	(0.007)	(0.013)
Male	0.001	0.001	0.004*	-0.002
	(0.002)	(0.001)	(0.002)	(0.004)
Constant	0.02**	0.007**	0.013	-0.016
	(0.008)	(0.003)	(0.012)	(0.025)
Observations	102,769	172,828	172,828	172,828

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. Sample weights are applied. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.39) Effect of Water Fluoridation on Birth Weight and Premature Births using Stacked DiD: One Year Before and After the Treatment

	Birth Weight	**Very Low Weight**	**Low Weight**	**Premature**
Treatment	-12.995**	-0.0002	0.003	0.009*
	(6.362)	(0.001)	(0.003)	(0.004)
White	18.098	-0.001	0.004	-0.031
	(35.191)	(0.004)	(0.013)	(0.021)
Black	-167.716***	0.005	0.032*	0.047*
	(40.413)	(0.005)	(0.016)	(0.026)
Married	45.788***	-0.0002	-0.017***	-0.023**
	(13.131)	(0.002)	(0.006)	(0.010)
Age	4.312***	-0.0002	0.001	-0.001*
	(1.030)	(0.0001)	(0.0004)	(0.001)
High School	21.893	-0.006	-0.0002	-0.021
	(26.257)	(0.004)	(0.008)	(0.015)
College or Higher	74.334***	-0.007	-0.013	-0.02
	(27.710)	(0.004)	(0.009)	(0.015)
Prenatal Visits	22.402***	-0.001***	-0.005***	-0.016***
	(1.416)	(0.0002)	(0.001)	(0.001)
First Birth	-81.337***	0.001	0.016***	0.009
	(11.411)	(0.001)	(0.005)	(0.007)
Singleton	899.394***	-0.072***	-0.359***	-0.415***
	(31.291)	(0.016)	(0.021)	(0.024)
Male	121.575***	-0.0003	-0.009**	0.016***
	(9.498)	(0.001)	(0.004)	(0.005)
Constant	2068.577***	0.099***	0.442***	0.792***
	(57.657)	(0.018)	(0.029)	(0.041)
Observations	656,465	656,526	656,526	656,526

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight (< 1500 g). Low Weight: indicator of low weight (< 2500 g). Premature: indicator of premature births (gestation length < 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. Sample Weights are applied. *** p-value ≤ 0.01, ** p-value ≤ 0.05, * p-value ≤ 0.1.

TABLE (3.40) Effect of Water Fluoridation on Congenital Anomalies using Stacked DiD: One Year Before and After the Treatment

	CNSA	NTD	MSKLA	NCNSA
Treatment	0.001***	0.0005	-0.002*	-0.0002
	(0.0002)	(0.0003)	(0.001)	(0.0005)
White	-0.002	-0.0003	0.004	0.005
	(0.004)	(0.003)	(0.005)	(0.01)
Black	0.002	-0.002	-0.003	0.007
	(0.004)	(0.003)	(0.007)	(0.011)
Married	0.0005	-0.001	-0.005*	-0.001
	(0.002)	(0.001)	(0.002)	(0.004)
Age	-0.0001	-0.0001	-0.0003	-0.0001
	(0.0001)	(0.0001)	(0.0003)	(0.0003)
High School	0.0005	-0.002	0.002	-0.001
	(0.001)	(0.001)	(0.002)	(0.003)
College or Higher	0.001	-0.0001	0.004	-0.001
	(0.001)	(0.002)	(0.003)	(0.003)
Prenatal Visits	0.00004	-0.00002	0.0005	0.0001
	(0.0001)	(0.0001)	(0.0003)	(0.0004)
First Birth	-0.0001	0.0002	-0.001	0.003
	(0.001)	(0.001)	(0.001)	(0.002)
Singleton	0.004	-0.001	0.003	-0.027
	(0.003)	(0.002)	(0.007)	(0.015)
Male	-0.001	0.001	0.002	-0.003
	(0.002)	(0.001)	(0.002)	(0.004)
Constant	-0.001	0.007	-0.001	0.029
	(0.006)	(0.004)	(0.01)	(0.020)
Observations	2,860	3,886	3,886	3,886

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. Sample Weights are applied. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.41) Effect of Water Fluoridation on Birth Weight and Premature Births using Stacked DiD: No Specified Window. Control Group: Never Treated Only

	Birth Weight	**Very Low Weight**	**Low Weight**	**Premature**
Treatment	-18.431***	0.001	0.001	0.004
	(5.478)	(0.001)	(0.002)	(0.004)
White	26.927	-0.001	0.001	-0.030
	(32.219)	(0.004)	(0.012)	(0.019)
Black	-175.163***	0.007	0.030**	0.064***
	(35.549)	(0.005)	(0.015)	(0.023)
Married	64.475***	-0.003	-0.017***	-0.032***
	(13.140)	(0.002)	(0.006)	(0.009)
Age	4.457***	-0.0001	0.0008**	-0.0006
	(0.993)	(0.0001)	(0.0004)	(0.0006)
High School	-36.340*	-0.001	0.01	0.021*
	(19.03)	(0.003)	(0.006)	(0.012)
College or Higher	14.555	-0.0007	-0.003	0.013
	(20.543)	(0.003)	(0.006)	(0.013)
Prenatal Visits	18.934***	-0.001***	-0.004***	-0.014***
	(1.316)	(0.0002)	(0.0004)	(0.001)
First Birth	-72.851***	-0.001	0.016***	0.004
	(10.480)	(0.001)	(0.004)	(0.007)
Singleton	905.756***	-0.076***	-0.351***	-0.416***
	(29.688)	(0.015)	(0.02)	(0.022)
Male	118.305***	-0.001	-0.007**	0.015**
	(8.56)	(0.001)	(0.003)	(0.005)
Constant	2118.028***	0.101***	0.419***	0.729***
	(52.236)	(0.018)	(0.027)	(0.037)
Observations	2,450,666	2,450,908	2,450,908	2,450,908

Notes: Each column corresponds to regression results of the following outcome variables. Weight: birth weight in grams. Very Low Weight: indicator of very low weight ($<$ 1500 g). Low Weight: indicator of low weight ($<$ 2500 g). Premature: indicator of premature births (gestation length $<$ 38 weeks). *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. Sample weight are applied. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.42) Effect of Water Fluoridation on Congenital Anomalies using Stacked DiD: No Specified Window. Control Group: Never Treated Only

	CNSA	**NTD**	**MSKLA**	**NCNSA**
Treatment	0.0004**	0.0001	0.006	0.008
	(0.0002)	(0.0003)	(0.004)	(0.006)
White	-0.008**	-0.003	-0.006	0.016
	(0.004)	(0.002)	(0.006)	(0.010)
Black	-0.003	-0.004	-0.014**	0.014
	(0.004)	(0.003)	(0.007)	(0.01)
Married	-0.002	-0.001*	-0.014***	-0.003
	(0.002)	(0.001)	(0.004)	(0.004)
Age	-0.0001	-0.0001	0.0003	0.001*
	(0.0001)	(0.0001)	(0.0003)	(0.001)
High School	-0.018*	-0.002***	-0.019*	-0.016
	(0.009)	(0.001)	(0.011)	(0.010)
College or Higher	-0.015	0.0005	-0.013	-0.034
	(0.008)	(0.001)	(0.010)	(0.023)
Prenatal Visits	0.00004	-0.00000	0.0001	-0.001
	(0.0001)	(0.00004)	(0.0003)	(0.001)
First Birth	0.001	0.001	-0.0002	0.008*
	(0.001)	(0.001)	(0.002)	(0.005)
Singleton	0.006**	-0.001	0.005	-0.007
	(0.003)	(0.002)	(0.007)	(0.013)
Male	0.001	0.001	0.004*	-0.002
	(0.002)	(0.001)	(0.002)	(0.004)
Constant	0.019**	0.007**	0.017	-0.011
	(0.008)	(0.003)	(0.012)	(0.023)
Observations	99,465	168,993	168,993	168,993

Notes: Each column corresponds to regression results of the following outcome variables. CNSA: central nervous system anomalies. NTD: neural tube defects. MSKLA: musculoskeletal anomalies. NCNSA: non-central nervous system anomalies. *Treatment* is an indicator of prenatal fluoride exposure. Standard errors in parentheses (.) and are clustered at the county level. Month, year, and county fixed-effects are included. Sample weights are applied. *** p-value $\leq$ 0.01, ** p-value $\leq$ 0.05, * p-value $\leq$ 0.1.

TABLE (3.43) Economic Costs (in US $) of Water Fluoridation

Outcome	Cost (base year)	Cost (2024)	Fluoride Effect	Fluoride Cost
	Panel A: Short-term Costs			
Very low weight	93,800	259,767	10	2,597,670
Premature	21,500	29,982	50	1,499,100
CNSA	110,500	135,954	2	271,908
Total				4,368,678
	Panel B: Long-term Costs			
Premature	4.8×10^9	6.3×10^9	50	315×10^9
CNSA	10,000,000	13,102,874	2	26,205,748
Total				315,026,205,748

Notes: Column 1 corresponds to the outcome variable: Very Low Weight: indicator of very low weight (< 1500 g); Premature: indicator of premature births (gestation length < 38 weeks); CNSA: indicator for central nervous system anomalies. Column 2 presents the cost (per case) in the base year (as in reference). Column 3 presents the cost in 2024. Column 4 presents the effects of water fluoridation per 100 births. Column 5 corresponds to the cost associated with water fluoridation per 100 births.

TABLE (3.44) Lifetime Earnings Following Prenatal Exposure to Water Fluoridation (Birth Weight Channel)

r	Annuity Factor	PV Loss per Person	PV Loss for Treated
1%	33.72	$743	$16.8 billion
3%	23.12	$509	$11.5 billion
5%	17.16	$378	$8.5 billion
7%	13.80	$304	$6.9 billion

Notes: The effect on Lifetime earning is based on Lambiris et al., 2022 estimate: 2.75% reduction in earnings following a 1 SD decrease in birth weight. Column 1 corresponds to the discount factor. Column 2 is the annuity factor computed based on 40 years of work. Column 3 is the present value (PV) loss per person. This is computed based on our estimated reduction of birth weight (0.34% $\simeq$ 0.019 SD) and the real personal income in the US for the year 2023 ($42,220). Column 4 represents the total present value loss of real personal income for the entire treated sample (22,569,123).

Bibliography

Abdul Baki, Ghina, Louis-Philippe Beland, Myra Yazbeck, and Aline Zayat. 2025. "Reactance, rationalization, and women's rights for safe abortion: Evidence from Roe vs. Wade's overturn." *Journal of Economic Behavior & Organization* 233:106956.

Abdul Baki, Ghina, Hassan Benchekroun, and Walid Marrouch. 2024. "Emissions taxes under double-sided environmental externalities." *Economics Letters* 242:111884.

Abdul Baki, Ghina, and Walid Marrouch. 2022. "Environmental taxation in the bertrand differentiated duopoly: new insights." *Resource and Energy Economics* 70:101329.

Abduweli Uyghurturk, Dawud, Dana E Goin, Esperanza Angeles Martinez-Mier, Tracey J Woodruff, and Pamela K DenBesten. 2020. "Maternal and fetal exposures to fluoride during mid-gestation among pregnant women in northern California." *Environmental Health* 19:1–9.

Adämmer, Philipp, Jan Prüser, and Rainer A Schüssler. 2024. "Forecasting macroeconomic tail risk in real time: Do textual data add value?" *International Journal of Forecasting.*

Aggeborn, Linuz, and Mattias Öhman. 2021. "The effects of fluoride in drinking water." *Journal of Political Economy* 129 (2): 465–491.

Aizer, Anna, and Janet Currie. 2019. "Lead and juvenile delinquency: new evidence from linked birth, school, and juvenile detention records." *Review of Economics and Statistics* 101 (4): 575–587.

———. 2014. "The intergenerational transmission of inequality: maternal disadvantage and health at birth." *science* 344 (6186): 856–861.

Aizer, Anna, Janet Currie, Peter Simon, and Patrick Vivier. 2018. "Do low levels of blood lead reduce children's future test scores?" *American Economic Journal: Applied Economics* 10 (1): 307–341.

Al-Qaraghouli, Mohammed, and Yu Ming Victor Fang. 2017. "Effect of fetal sex on maternal and obstetric outcomes." *Frontiers in Pediatrics* 5:144.

Aldy, Joseph E. 2014. "The labor market impacts of the 2010 Deepwater Horizon oil spill and offshore oil drilling moratorium." *National Bureau of Economic Research* No. 20409.

Almond, Douglas, Kenneth Y Chay, and David S Lee. 2005. "The costs of low birth weight." *The Quarterly Journal of Economics* 120 (3): 1031–1083.

Almond, Douglas, and Janet Currie. 2011. "Killing me softly: The fetal origins hypothesis." *Journal of Economic Perspectives* 25 (3): 153–172.

Anderson, R, JH Beard, and D Sorley. 1980. "Fluoride intoxication in a dialysis unit—Maryland." *MMWR* 29 (12): 134–6.

Aouadi, Amal, Mohamed Arouri, and David Roubaud. 2018. "Information demand and stock market liquidity: International evidence." *Economic Modelling* 70:194–202.

Ardia, David, Keven Bluteau, Kris Boudt, and Koen Inghelbrecht. 2022. "Climate change concerns and the performance of green vs. brown stocks." *Management Science.*

Ardia, David, Keven Bluteau, and Maxime Rüede. 2019. "Regime changes in Bitcoin GARCH volatility dynamics." *Finance Research Letters* 29:266–271.

Armand, Alex, and Ivan Kim Taveras. 2021. "The Effect of Nature's Wealth on Human Development: Evidence from Renewable Resources." *CEPR Discussion Paper* No. DP15680.

Arun, Aaditya Krishna, Luis Rustveld, and Ajeesh Sunny. 2022. "Association between water fluoride levels and low birth weight: National Health and Nutrition Examination Survey (NHANES) 2013–2016." *International Journal of Environmental Research and Public Health* 19 (15): 8956.

Aruoba, S Borağan, Andrés Fernández, Daniel Guzmán, Ernesto Pastén, and Felipe Saffie. 2024. "Pricing under distress." *National Bureau of Economic Research.*

AWWA. 2024. "U.S. survey gauges public perceptions of tap water." *American Water Works Association* https://www.awwa.org/AWWA-Articles/u-s-survey-gauges-public-perceptions-of-tap-water/.

Baker, Andrew C, David F Larcker, and Charles CY Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics* 144 (2): 370–395.

Baker, Scott R, Nicholas Bloom, and Steven J Davis. 2016. "Measuring economic policy uncertainty." *The Quarterly Journal of Economics* 131 (4): 1593–1636.

Banerjee, Abhijit, Esther Duflo, Gilles Postel-Vinay, and Tim Watts. 2010. "Long-run health impacts of income shocks: Wine and phylloxera in nineteenth-century France." *The Review of Economics and Statistics* 92 (4): 714–728.

Barbaglia, Luca, Sergio Consoli, and Sebastiano Manzan. 2023. "Forecasting with economic news." *Journal of Business & Economic Statistics* 41 (3): 708–719.

Barker, David JP, and Clive Osmond. 1986. "Infant mortality, childhood nutrition, and ischaemic heart disease in England and Wales." *The Lancet* 327 (8489): 1077–1081.

Barnett, A.H. 1980. "The Pigouvian tax rule under monopoly." *American Economic Review* 70 (5): 1037–1041.

Barrett, Emily S, and Jeremy Lessing. 2021. "Sex-specific impacts of prenatal stress." *Prenatal Stress and Child Development*, 243–277.

Bartik, Alexander W, Janet Currie, Michael Greenstone, and Christopher R Knittel. 2019. "The local economic and welfare consequences of hydraulic fracturing." *American Economic Journal: Applied Economics* 11 (4): 105–155.

Bashash, Morteza, Deena Thomas, Howard Hu, E Angeles Martinez-Mier, Brisa N Sanchez, Niladri Basu, Karen E Peterson, Adrienne S Ettinger, Robert Wright, Zhenzhen Zhang, et al. 2017. "Prenatal fluoride exposure and cognitive outcomes in children at 4 and 6–12 years of age in Mexico." *Environmental health perspectives* 125 (9): 097017.

Bauer, Michael D, and Glenn D Rudebusch. 2021. "The rising cost of climate change: evidence from the bond market." *The Review of Economics and Statistics*, 1–45.

Beland, Louis-Philippe, and Sara Oloomi. 2019. "Environmental disaster, pollution and infant health: Evidence from the Deepwater Horizon oil spill." *Journal of Environmental Economics and Management* 98:102265.

Bell, Micheile L, Devra L Davis, and Tony Fletcher. 2008. "A retrospective assessment of mortality from the London smog episode of 1952: the role of influenza and pollution." *Urban Ecology: An International Perspective on the Interaction Between Humans and Nature,* 263–268.

Benchekroun, Hassan, and Ngo Van Long. 1998. "Efficiency inducing taxation for polluting oligopolists." *Journal of Public Economics* 70 (2): 325–342.

———. 2002. "On the multiplicity of efficiency-inducing tax rules." *Economics Letters* 76 (3): 331–336.

Berger, David, Ian Dew-Becker, and Stefano Giglio. 2020. "Uncertainty shocks as second-moment news shocks." *The Review of Economic Studies* 87 (1): 40–76.

Bernard, Jean-Thomas, Lynda Khalaf, Maral Kichian, and Clement Yelou. 2018. "Oil price forecasts for the long term: expert outlooks, models, or both?" *Macroeconomic Dynamics* 22 (3): 581–599.

Bertero, Alessia, Mario Chiari, Nicoletta Vitale, Mariagrazia Zanoni, Elena Faggionato, Alberto Biancardi, and Francesca Caloni. 2020. "Types of pesticides involved in domestic and wild animal poisoning in Italy." *Science of the Total Environment* 707:136129.

Bhalotra, Sonia R, Alberto Diaz-Cayeros, Grant Miller, Alfonso Miranda, and Atheendar S Venkataramani. 2021. "Urban water disinfection and mortality decline in lower-income countries." *American Economic Journal: Economic Policy* 13 (4): 490–520.

Bharadwaj, Prashant, Juan Pedro Eberhard, and Christopher A Neilson. 2018. "Health at birth, parental investments, and academic outcomes." *Journal of Labor Economics* 36 (2): 349–394.

Bharadwaj, Prashant, Petter Lundborg, and Dan-Olof Rooth. 2018. "Birth weight in the long run." *Journal of Human Resources* 53 (1): 189–231.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes. 2007. "From the cradle to the labor market? The effect of birth weight on adult outcomes." *The Quarterly Journal of Economics* 122 (1): 409–439.

Blackman, Allen, and Winston Harrington. 2018. "The use of economic incentives in developing countries: Lessons from international experience with industrial air pollution." *The Theory and Practice of Command and Control in Environmental Policy,* 199–238.

Bolton, Patrick, and Marcin Kacperczyk. 2021. "Do investors care about carbon risk?" *Journal of Financial Economics* 142 (2): 517–549.

Bouoiyour, Jamal, Refk Selmi, et al. 2016. "Bitcoin: A beginning of a new phase." *Economics Bulletin* 36 (3): 1430–1440.

Bovenberg, A Lans, and Ruud A De Mooij. 1997. "Environmental tax reform and endogenous growth." *Journal of Public Economics* 63 (2): 207–237.

Brainerd, Elizabeth, and Nidhiya Menon. 2014. "Seasonal effects of water quality: The hidden costs of the Green Revolution to infant and child health in India." *Journal of Development Economics* 107:49–64.

Broersma, Marcel, and Todd Graham. 2013. "Twitter as a news source: How Dutch and British newspapers used tweets in their news coverage, 2007–2011." *Journalism Practice* 7 (4): 446–464.

Brooks, Carolyn J, Steven L Gortmaker, Michael W Long, Angie L Cradock, and Erica L Kenney. 2017. "Racial/ethnic and socioeconomic disparities in hydration status among US adults and the role of tap water and other beverage intake." *American Journal of Public Health* 107 (9): 1387–1394.

Buchanan, J.M. 1969. "External diseconomies, corrective taxes, and market structure." *American Economic Review* 59 (1): 174–177.

Cachon, Gérard P. 2014. "Retail store density and the cost of greenhouse gas emissions." *Management Science* 60 (8): 1907–1925.

Cahill, Daniel, and Zhangxin Frank Liu. 2021. "Limitations of imitation: Lessons from another Bitcoin copycat." *Journal of Corporate Finance* 69:101992.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics* 225 (2): 200–230.

Canton, Joan, Antoine Soubeyran, and Hubert Stahn. 2008. "Environmental taxation and vertical cournot oligopolies: How eco-industries matter." *Environmental and Resource Economics* 40:369–382.

CDC. 2024. "Return on Investment: Optimally Fluoridated Water." *US Centers for Disease Control and Prevention.*

———. 2016. "Zika in the Western Hemisphere: Risks and Response." *https://archive.cdc.gov/www_cdc_*

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The effect of minimum wages on low-wage jobs." *The Quarterly Journal of Economics* 134 (3): 1405–1454.

Chakraborti, Lopamudra, and Jay P Shimshack. 2022. "Environmental disparities in urban Mexico: Evidence from toxic water pollution." *Resource and Energy Economics* 67:101281.

Chang, Yang-Ming, and Manaf Sellak. 2023. "Environmental Policy in Vertical Markets with Downstream Pollution: Taxes Versus Standards." *The BE Journal of Economic Analysis & Policy* 23 (4): 897–923.

Chiang, I-Hsuan Ethan, W Keener Hughen, and Jacob S Sagi. 2015. "Estimating oil risk factors using information from equity and derivatives markets." *The Journal of Finance* 70 (2): 769–804.

Choi, Anna L, Guifan Sun, Ying Zhang, and Philippe Grandjean. 2012. "Developmental fluoride neurotoxicity: a systematic review and meta-analysis." *Environmental Health Perspectives* 120 (10): 1362–1368.

Cochrane, John. 2009. *Asset pricing: Revised edition.* Princeton university press.

Consoli, Sergio, Luca Barbaglia, and Sebastiano Manzan. 2022. "Fine-grained, aspect-based sentiment analysis on economic and financial lexicon." *Knowledge-Based Systems* 247:108781.

Cremer, Helmuth, Firouz Gahvari, and Norbert Ladoux. 1998. "Externalities and optimal taxation." *Journal of Public Economics* 70 (3): 343–364.

Currie, Janet. 2009. "Healthy, wealthy, and wise: Socioeconomic status, poor health in childhood, and human capital development." *JoUrnal of economIc lIteratUre* 47 (1): 87–122.

Currie, Janet. 2011. "Inequality at birth: Some causes and consequences." *American economic review* 101 (3): 1–22.

Currie, Janet, Joshua Graff Zivin, Katherine Meckel, Matthew Neidell, and Wolfram Schlenker. 2013. "Something in the water: contaminated drinking water and infant health." *Canadian Journal of Economics/Revue canadienne d'économique* 46 (3): 791–810.

Currie, Janet, and Enrico Moretti. 2007. "Biology as destiny? Short-and long-run determinants of intergenerational transmission of birth weight." *Journal of Labor economics* 25 (2): 231–264.

Da Mata, Daniel, Lucas Emanuel, Vitor Pereira, and Breno Sampaio. 2023. "Climate adaptation policies and infant health: evidence from a water policy in Brazil." *Journal of Public Economics* 220:104835.

Dave, Dhaval M, and Muzhe Yang. 2022. "Lead in drinking water and birth outcomes: A tale of two water treatment plants." *Journal of Health Economics* 84:102644.

David, M., A.D. Nimubona, and B. Sinclair-Desgagné. 2011. "Emission taxes and the market for abatement goods and services." *Resource and Energy Economics* 33 (1): 179–191.

David, M., and B. Sinclair-Desgagné. 2005. "Environmental regulation and the eco-industry." *Journal of Regulatory Economics* 28 (2): 141–155.

De-Regil, Luz Maria, Juan Pablo Peña-Rosas, Ana C Fernández-Gaxiola, and Pura Rayco-Solon. 2015. "Effects and safety of periconceptional oral folate supplementation for preventing birth defects." *Cochrane Database of Systematic Reviews,* no. 12.

Delle Monache, Davide, Andrea De Polis, and Ivan Petrella. 2024. "Modeling and forecasting macroeconomic downside risk." *Journal of Business & Economic Statistics* 42 (3): 1010–1025.

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez. 2017. "The price of variance risk." *Journal of Financial Economics* 123 (2): 225–250.

Diebold, Francis X, and Robert S Mariano. 2002. "Comparing predictive accuracy." *Journal of Business & Economic statistics* 20 (1): 134–144.

Dittmar, Lars, and Aaron Praktiknjo. 2019. "Could Bitcoin emissions push global warming above 2° C?" *Nature Climate Change* 9 (9): 656–657.

Du, Xinming, Shan Zhang, and Eric Zou. 2024. "Marine Microplastics and Infant Health." No. w33094.

Dyhrberg, Anne Haubo. 2016. "Bitcoin, gold and the dollar–A GARCH volatility analysis." *Finance Research Letters* 16:85–92.

Easley, David, Maureen O'Hara, and Soumya Basu. 2019. "From mining to markets: The evolution of bitcoin transaction fees." *Journal of Financial Economics* 134 (1): 91–109.

Ebenstein, Avraham. 2012. "The consequences of industrialization: evidence from water pollution and digestive cancers in China." *Review of economics and statistics* 94 (1): 186–201.

Ebert, U. 1991. "Pigouvian tax and market structure: The case of oligopoly and different abatement technologies." *FinanzArchiv* 49:154–166.

Ebert, Udo, and Oskar Von Dem Hagen. 1998. "Pigouvian taxes under imperfect competition if consumption depends on emissions." *Environmental and Resource Economics* 12:507–513.

El Ouadghiri, Imane, Khaled Guesmi, Jonathan Peillex, and Andreas Ziegler. 2021. "Public attention to environmental issues and stock market returns." *Ecological Economics* 180:106836.

Ellingsen, Jon, Vegard H Larsen, and Leif Anders Thorsrud. 2022. "News media versus FRED-MD for macroeconomic forecasting." *Journal of Applied Econometrics* 37 (1): 63–81.

Elnaboulsi, Jihad C, Wassim Daher, and Yiğit Sağlam. 2023. "Environmental taxation, information precision, and information sharing." *Journal of Public Economic Theory* 25 (2): 301–341.

Engle, Robert F, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. 2020. "Hedging climate change news." *The Review of Financial Studies* 33 (3): 1184–1216.

Estay, Manuel, and John K Stranlund. 2022. "Entry, location, and optimal environmental policies." *Resource and Energy Economics* 70:101326.

Fama, Eugene, Lars Peter Hansen, and Robert Shiller. 2013. "Asset prices." *The Royal Swedish Academy of Sciences.*

Fama, Eugene F, and Kenneth R French. 2015. "A five-factor asset pricing model." *Journal of Financial Economics* 116 (1): 1–22.

Fan, Maoyong, and Guojun He. 2023. "Clean water and infant health: evidence from piped water provision in China." *Journal of the Association of Environmental and Resource Economists* 10 (1): 159–193.

Fawell, J, K Bailey, J Chilton, E Dahi, L Fewtrell, and Y Magara. 2006. "Human health effects. Fluoride in drinking water." *World Health Organization.*

Fletcher, Jason, and Hamid Noghanibehambari. 2024. "Journal of Environmental Economics and Management." *Journal of Environmental Economics and Management* 123:102903.

Flynn, Patrick, and Michelle Marcus. 2023. "A watershed moment: the Clean Water Act and birth weight." *Journal of Human Resources.*

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as data." *Journal of Economic Literature* 57 (3): 535–74.

Gessner, Bradford D, Michael Beller, John P Middaugh, and Gary M Whitford. 1994. "Acute fluoride poisoning from a public water system." *New England Journal of Medicine* 330 (2): 95–99.

Ghysels, Eric, and Massimiliano Marcellino. 2018. *Applied economic forecasting using time series methods.* Oxford University Press.

Giacomini, Raffaella, and Halbert White. 2006. "Tests of conditional predictive ability." *Econometrica* 74 (6): 1545–1578.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber. 2021. "Climate change and long-run discount rates: Evidence from real estate." *The Review of Financial Studies* 34 (8): 3527–3571.

Glazyrina, Irina, Vasiliy Glazyrin, and Sergey Vinnichenko. 2006. "The polluter pays principle and potential conflicts in society." *Ecological Economics* 59 (3): 324–330.

Glied, Sherry, and Matthew Neidell. 2010. "The economic value of teeth." *Journal of Human Resources* 45 (2): 468–496.

Goin, Dana E, Amy M Padula, Tracey J Woodruff, Allison Sherris, Kiley Charbonneau, and Rachel Morello-Frosch. 2024. "Water fluoridation and birth outcomes in California." *Environmental Health Perspectives* 132 (5): 057004.

Goodkind, Andrew L, Benjamin A Jones, and Robert P Berrens. 2020. "Cryptodamages: Monetary value estimates of the air pollution and human health impacts of cryptocurrency mining." *Energy Research & Social Science* 59:101281.

Goodman, Alyson B, Heidi M Blanck, Bettylou Sherry, Sohyun Park, Linda Nebeling, and Amy L Yaroch. 2013. "Behaviors and attitudes associated with low drinking water intake among US adults, Food Attitudes and Behaviors Survey, 2007." *Preventing chronic disease* 10:E51.

Goodman, Carly, Meaghan Hall, Rivka Green, Richard Hornung, Esperanza Angeles Martinez-Mier, Bruce Lanphear, and Christine Till. 2022. "Maternal fluoride exposure, fertility and birth outcomes: the MIREC cohort." *Environmental Advances* 7:100135.

Gopu, Banu Preethi, Liane B Azevedo, Ralph M Duckworth, Murali KP Subramanian, Sherley John, and Fatemeh Vida Zohoori. 2022. "The relationship between fluoride exposure and cognitive outcomes from gestation to adulthood—a systematic review." *International journal of environmental research and public health* 20 (1): 22.

Goyal, Amit. 2012. "Empirical cross-sectional asset pricing: a survey." *Financial Markets and Portfolio Management* 26:3–38.

Goyal, Lajya Devi, Dapinder Kaur Bakshi, Jatinder Kaur Arora, Ankita Manchanda, and Paramdeep Singh. 2020. "Assessment of fluoride levels during pregnancy and its association with early adverse pregnancy outcomes." *Journal of Family Medicine and Primary Care* 9 (6): 2693–2698.

Grandjean, Philippe, Alessandra Meddis, Flemming Nielsen, Iben H Beck, Niels Bilenberg, Carly V Goodman, Howard Hu, Christine Till, and Esben Budtz-Jørgensen. 2024. "Dose dependence of prenatal fluoride exposure associations with cognitive performance at school age in three prospective studies." *European Journal of Public Health* 34 (1): 143–149.

Green, Rivka, Bruce Lanphear, Richard Hornung, David Flora, E Angeles Martinez-Mier, Raichel Neufeld, Pierre Ayotte, Gina Muckle, and Christine Till. 2019. "Association between maternal fluoride exposure during pregnancy and IQ scores in offspring in Canada." *JAMA pediatrics* 173 (10): 940–948.

Gronwald, Marc. 2019. "Is Bitcoin a Commodity? On price jumps, demand shocks, and certainty of supply." *Journal of International Money and Finance* 97:86–92.

Guimbeau, Amanda, Xinde James Ji, Zi Long, and Nidhiya Menon. 2024. "Ocean salinity, early-life health, and adaptation." *Journal of Environmental Economics and Management* 125:102954.

Hall, Meaghan, Bruce Lanphear, Jonathan Chevrier, Rick Hornung, Rivka Green, Carly Goodman, Pierre Ayotte, Esperanza Angeles Martinez-Mier, R Thomas Zoeller, and Christine Till. 2023. "Fluoride exposure and hypothyroidism in a Canadian pregnancy cohort." *Science of the Total Environment* 869:161149.

Harvey, Campbell R, and Yan Liu. 2021. "Lucky factors." *Journal of Financial Economics* 141 (2): 413–435.

Hill, Elaine, and Lala Ma. 2017. "Shale gas development and drinking water quality." *American Economic Review* 107 (5): 522–525.

Hill, Elaine L. 2018. "Shale gas development and infant health: evidence from Pennsylvania." *Journal of Health Economics* 61:134–150.

Hill, Elaine L, and Lala Ma. 2022. "Drinking water, fracking, and infant health." *Journal of Health Economics* 82:102595.

Hillert, Alexander, Heiko Jacobs, and Sebastian Müller. 2014. "Media makes momentum." *The Review of Financial Studies* 27 (12): 3467–3501.

Hoel, Michael. 1998. "Emission taxes versus other environmental policies." *Scandinavian Journal of Economics* 100 (1): 79–104.

Hoffman, Richard, Jonathan Mann, James Calderone, Jon Trumbull, and Michael Burkhart. 1980. "Acute fluoride poisoning in a New Mexico elementary school." *Pediatrics* 65 (5): 897–900.

Hong, Harrison, Frank Weikai Li, and Jiangmin Xu. 2019. "Climate risks and market efficiency." *Journal of Econometrics* 208 (1): 265–281.

Hong, Yongmiao, Fuwei Jiang, Lingchao Meng, and Bowen Xue. 2024. "Forecasting Inflation Using Economic Narratives." *Journal of Business & Economic Statistics,* no. just-accepted, 1–27.

Huber, Claus, and Franz Wirl. 1998. "The polluter pays versus the pollutee pays principle under asymmetric information." *Journal of Environmental Economics and Management* 35 (1): 69–87.

Hung, Man, Eric S Hon, Amir Mohajeri, Hyma Moparthi, Teresa Vu, Jason Jeon, and Martin S Lipsky. 2023. "A national study exploring the association between fluoride levels and dental fluorosis." *JAMA Network Open* 6 (6): e2318406–e2318406.

Iban-Arias, Ruth, Ariana Soares Dias Portela, Sibilla Masieri, Aurelian Radu, Eun-Jeong Yang, Lung-Chi Chen, Terry Gordon, and Giulio Maria Pasinetti. 2025. "Role of acute exposure to environmental stressors in the gut-brain-periphery axis in the presence of cognitive resilience." *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease* 1871 (5): 167760.

Ibarluzea, Jesús, Mara Gallastegi, Loreto Santa-Marina, Ana Jiménez Zabala, Enrique Arranz, Amaia Molinuevo, Maria-Jose Lopez-Espinosa, Ferran Ballester, Cristina M Villanueva, Isolina Riano, et al. 2022. "Prenatal exposure to fluoride and neuropsychological development in early childhood: 1-to 4 years old children." *Environmental Research* 207:112181.

IEA. 2021. "Net zero by 2050: A roadmap for the global energy sector."

Ino, Hiroaki, and Toshihiro Matsumura. 2021. "Optimality of emission pricing policies based on emission intensity targets under imperfect competition." *Energy Economics* 98:105238.

Isaksson, Lena Höglund. 2005. "Abatement costs in response to the Swedish charge on nitrogen oxide emissions." *Journal of Environmental Economics and Management* 50 (1): 102–120.

Javidi, Ariana, and Gregory Pierce. 2018. "US households' perception of drinking water as unsafe and its consequences: Examining alternative choices to the tap." *Water Resources Research* 54 (9): 6100–6113.

Jegadeesh, Narasimhan, and Di Wu. 2013. "Word power: A new approach for content analysis." *Journal of financial economics* 110 (3): 712–729.

Jones, Benjamin A. 2019. "Infant health impacts of freshwater algal blooms: evidence from an invasive species natural experiment." *Journal of Environmental Economics and Management* 96:36–59.

Kalamara, Eleni, Arthur Turrell, Chris Redl, George Kapetanios, and Sujit Kapadia. 2022. "Making text count: economic forecasting using newspaper text." *Journal of Applied Econometrics* 37 (5): 896–919.

Katoto, Patrick DMC, Amanda S Brand, Buket Bakan, Paul Musa Obadia, Carsi Kuhangana, Tony Kayembe-Kitenge, Joseph Pyana Kitenge, Celestin Banza Lubaba Nkulu, Jeroen Vanoirbeek, Tim S Nawrot, et al. 2021. "Acute and chronic exposure to air pollution in relation with incidence, prevalence, severity and mortality of COVID-19: a rapid systematic review." *Environmental Health* 20 (1): 41.

Katsiampa, Paraskevi. 2017. "Volatility estimation for Bitcoin: A comparison of GARCH models." *Economics Letters* 158:3–6.

Keskin, Pinar, Gauri Kartini Shastry, and Helen Willis. 2017. "Water quality awareness and breastfeeding: Evidence of health behavior change in Bangladesh." *Review of Economics and Statistics* 99 (2): 265–280.

Kohli, Varun, Sombuddha Chakravarty, Vinay Chamola, Kuldip Singh Sangwan, and Sherali Zeadally. 2023. "An analysis of energy consumption and carbon footprints of cryptocurrencies and possible solutions." *Digital Communications and Networks* 9 (1): 79–89.

Koschnitzky, Jenna E, Erwin Yap, Yifan Zhang, Monica J Chau, Ketan Yerneni, Anna Lisa Somera, Mark Luciano, and Abhay Moghekar. 2022. "Inpatient healthcare burden and variables influencing hydrocephalus-related admissions across the lifespan." *Journal of Neurosurgery* 139 (2): 502–511.

Kowlessar, Niranjana M, H Joanna Jiang, and Claudia Steiner. 2013. "Hospital stays for newborns, 2011." *Europe PMC.*

Krueger, Philipp, Zacharias Sautner, and Laura T Starks. 2020. "The importance of climate risks for institutional investors." *The Review of Financial Studies* 33 (3): 1067–1111.

Krzeczkowski, John E, Meaghan Hall, Dave Saint-Amour, Youssef Oulhote, Taylor McGuckin, Carly V Goodman, Rivka Green, Gina Muckle, Bruce Lanphear, and Christine Till. 2024. “Prenatal fluoride exposure, offspring visual acuity and autonomic nervous system function in 6-month-old infants.” *Environment International* 183:108336.

Lai, Wangyang. 2017. “Pesticide use and health outcomes: Evidence from agricultural water pollution in China.” *Journal of Environmental Economics and Management* 86:93–120.

Lambiris, Mark J, Mia M Blakstad, Nandita Perumal, Goodarz Danaei, Lilia Bliznashka, Günther Fink, and Christopher R Sudfeld. 2022. “Birth weight and adult earnings: a systematic review and meta-analysis.” *Journal of Developmental Origins of Health and Disease* 13 (3): 284–291.

Lin, Yongwei, Na Zhang, Jianfen Zhang, Junbo Lu, Shufang Liu, and Guansheng Ma. 2024. “Seasonality Affects Fluid Intake Behaviors among Young Adults in Hebei, China.” *Nutrients* 16 (11): 1542.

Liu, Yukun, Aleh Tsyvinski, and Xi Wu. 2022. “Common risk factors in cryptocurrency.” *The Journal of Finance* 77 (2): 1133–1177.

Luken, Ralph A. 2009. “Equivocating on the polluter-pays principle: The consequences for Pakistan.” *Journal of Environmental Management* 90 (11): 3479–3484.

Ma, June, Joshua S Gans, and Rabee Tourky. 2018. *Market structure in bitcoin mining.* Technical report. National Bureau of Economic Research.

Malin, Ashley J, Sandrah P Eckel, Howard Hu, E Angeles Martinez-Mier, Ixel Hernandez-Castro, Tingyu Yang, Shohreh F Farzan, Rima Habre, Carrie V Breton, and Theresa M Bastain. 2024. “Maternal urinary fluoride and child neurobehavior at age 36 months.” *JAMA Network Open* 7 (5): e2411987–e2411987.

Malmqvist, Erik, Davide Fumagalli, Christian Munthe, and DG Joakim Larsson. 2023. “Pharmaceutical Pollution from Human Use and the Polluter Pays Principle.” *Public Health Ethics* 16 (2): 152–164.

Manela, Asaf, and Alan Moreira. 2017. “News implied volatility and disaster concerns.” *Journal of Financial Economics* 123 (1): 137–162.

Mansur, Erin T. 2011. "Upstream versus downstream implementation of climate policy." In *The design and implementation of US climate policy,* 179–193. University of Chicago Press.

Marcus, Michelle. 2025. "Testing above the limit: Drinking water contamination and test scores." *Journal of Public Economics* 243:105313.

Marrouch, W., and B. Sinclair-Desgagné. 2012. "Emission taxes when pollution depends on location." *Environment and Development Economics* 17 (4): 433–443.

Mettetal, Elizabeth. 2019. "Irrigation dams, water and infant mortality: Evidence from South Africa." *Journal of Development Economics* 138:17–40.

Michaud, Pierre-Carl, Dana Goldman, Darius Lakdawalla, Adam Gailey, and Yuhui Zheng. 2011. "Differences in health between Americans and Western Europeans: effects on longevity and public finance." *Social science & medicine* 73 (2): 254–263.

Millock, Katrin, Céline Nauges, and Thomas Sterner. 2004. "Environmental taxes: a comparison of French and Swedish experience from taxes on industrial air pollution." *CESifo DICE report* 2 (1): 30–34.

Molloy, Anne M, Peadar N Kirke, James F Troendle, Helen Burke, Marie Sutton, Lawrence C Brody, John M Scott, and James L Mills. 2009. "Maternal vitamin B12 status and risk of neural tube defects in a population with high neural tube defect prevalence and no folic acid fortification." *Pediatrics* 123 (3): 917–923.

Mondal, Debapriya, Tamara S Galloway, Trevor C Bailey, and Fiona Mathews. 2014. "Elevated risk of stillbirth in males: systematic review and meta-analysis of more than 30 million births." *BMC medicine* 12:1–11.

Muth, Mary K, and Michael K Wohlgenant. 1999. "Measuring the degree of oligopsony power in the beef packing industry in the absence of marketing input quantity data." *Journal of Agricultural and Resource Economics,* 299–312.

Neto, Amanda Barbosa, Alexandre A Ferraro, and Sandra E Vieira. 2023. "Acute and subchronic exposure to urban atmospheric pollutants aggravate acute respiratory failure in infants." *Scientific reports* 13 (1): 16888.

Nimubona, A.D., and B. Sinclair-Desgagné. 2013. "The Pigouvian tax rule in the presence of an eco-industry." *Economics Bulletin* 33 (1): 747–752.

Noordzij, Marlies, Merel van Diepen, Fergus C Caskey, and Kitty J Jager. 2017. "Relative risk versus absolute risk: one cannot be interpreted without the other." *Nephrology Dialysis Transplantation* 32 (suppl_2): ii13–ii18.

NTP. 2024. "NTP monograph on the state of the science concerning fluoride exposure and neurodevelopment and cognition: a systematic review." *National Toxicology Program Monograph,* no. 8, NTP–MGRAPH.

OECD. 2001. *Environmentally related taxes in OECD countries: Issues and strategies.* Organisation for Economic Co-operation / Development.

Okorie, David Iheke, and Boqiang Lin. 2020. "Crude oil price and cryptocurrencies: evidence of volatility connectedness and hedging strategy." *Energy Economics* 87:104703.

Onufrak, Stephen J, Sohyun Park, Joseph R Sharkey, and Bettylou Sherry. 2014. "The relationship of perceptions of tap water safety with intake of sugar-sweetened beverages and plain water among US adults." *Public Health Nutrition* 17 (1): 179–185.

Oreopoulos, Philip, Mark Stabile, Randy Walld, and Leslie L Roos. 2008. "Short-, medium-, and long-term consequences of poor infant health: An analysis using siblings and twins." *Journal of Human Resources* 43 (1): 88–138.

Ortíz-García, Sofía G., Luisa E. Torres-Sánchez, Teresa V. Muñoz Rocha, Adriana Mercado-García, Karen E. Peterson, Howard Hu, Citlalli Osorio-Yáñez, and Martha María Téllez-Rojo. 2022. "Maternal urinary fluoride during pregnancy and birth weight and length: Results from ELEMENT cohort study." *Science of The Total Environment* 838:156459.

O'Connell, Joan, Jennifer Rockell, Judith Ouellet, Scott L Tomar, and William Maas. 2016. "Costs and savings associated with community water fluoridation in the United States." *Health Affairs* 35 (12): 2224–2232.

Papp, Anna, Douglas Almond, and Shuang Zhang. 2023. "Bitcoin and Carbon Dioxide Emissions: Evidence from Daily Production Decisions."

Pástor, L'uboš, Robert F Stambaugh, and Lucian A Taylor. 2022. "Dissecting green returns." *Journal of Financial Economics* 146 (2): 403–424.

Pearce, David, and R Kerry Turner. 1992. "Packaging waste and the polluter pays principle: a taxation solution." *Journal of Environmental Planning and Management* 35 (1): 5–15.

Pedersen, Lasse Heje. 2022. "Game on: Social networks and markets." *Journal of Financial Economics* 146 (3): 1097–1119.

Pesaran, M Hashem, and Allan Timmermann. 2009. "Testing dependence among serially correlated multicategory variables." *Journal of the American Statistical Association* 104 (485): 325–337.

Pigou, A. 1920. *The economics of welfare.* London: Macmillan.

Pindyck, Robert S. 2014. "Risk and return in the design of environmental policy." *Journal of the Association of Environmental and Resource Economists* 1 (3): 395–418.

Pitt, Mark M, Mark R Rosenzweig, and Nazmul Hassan. 2015. "Identifying the cost of a public health success: Arsenic well water contamination and productivity in Bangladesh." *National Bureau of Economic Research* No. 21741.

Popkin, Barry M, Denis V Barclay, and Samara J Nielsen. 2005. "Water and food consumption patterns of US adults from 1999 to 2001." *Obesity research* 13 (12): 2146–2152.

Rahman, Akim M, and Clive A Edwards. 2004. "Electricity: taxes on emission liabilities. An examination of the economic effectiveness of Polluter Pays Principles." *Energy Policy* 32 (2): 221–235.

Rau, Tomás, Sergio Urzúa, and Loreto Reyes. 2015. "Early exposure to hazardous waste and academic achievement: evidence from a case of environmental negligence." *Journal of the Association of Environmental and Resource Economists* 2 (4): 527–563.

Reboredo, Juan C, and Andrea Ugolini. 2018. "The impact of twitter sentiment on renewable energy stocks." *Energy Economics* 76:153–169.

Ren, Jie, Xi Chen, and Jian Hu. 2020. "The effect of production-versus consumption-based emission tax under demand uncertainty." *International Journal of Production Economics* 219:82–98.

Requate, T. 2006. “Environmental policy under imperfect competition.” In *The International Yearbook of Environmental and Resource Economics 2006/2007,* edited by T. Tietenberg and H. Folmer, 120–207. Cheltenham, UK: Edward Elgar.

Roberts, Adam. 2024. “What’s in the water? Long-run effects of fluoridation on health and economic self-sufficiency.” *Journal of Health Economics* 98:102934.

Rogowski, Jeannette. 1998. “Cost-effectiveness of care for very low birth weight infants.” *Pediatrics* 102 (1): 35–43.

Rosinger, A, and K Herrick. 2016. “Daily Water Intake among US Men and Women, 2009–2012; NCHS data brief, no 242.” *National Center for Health Statistics: Hyattsville, MD, USA.*

Rosinger, Asher Y, Kirsten A Herrick, Amber Y Wutich, Jonathan S Yoder, and Cynthia L Ogden. 2018. “Disparities in plain, tap and bottled water consumption among US adults: National Health and Nutrition Examination Survey (NHANES) 2007–2014.” *Public health nutrition* 21 (8): 1455–1464.

Royer, Heather. 2009. “Separated at girth: US twin estimates of the effects of birth weight.” *American Economic Journal: Applied Economics* 1 (1): 49–85.

Sadia, Mamoona, Abid Mahmood, Muhammad Ibrahim, Muhammad Kashif Irshad, Abul Hassan Ali Quddusi, Awais Bokhari, Muhammad Mubashir, Lai Fatt Chuah, and Pau Loke Show. 2022. “Microplastics pollution from wastewater treatment plants: A critical review on challenges, detection, sustainable removal techniques and circular economy.” *Environmental Technology & Innovation,* 102946.

Sanders, Nicholas J, and Charles Stoecker. 2015. “Where have all the young men gone? Using sex ratios to measure fetal death rates.” *Journal of Health Economics* 41:30–45.

Sandmo, Agnar. 2011. “Atmospheric externalities and environmental taxation.” *Energy Economics* 33:S4–S12.

———. 1978. “Direct versus indirect Pigovian taxation.” *European Economic Review* 7 (4): 337–349.

———. 1975. “Optimal taxation in the presence of externalities.” *The Swedish Journal of Economics,* 86–98.

Sarigiannis, Denis A, and Ute Hansen. 2012. "Considering the cumulative risk of mixtures of chemicals–a challenge for policy makers." *Environmental Health* 11:1–12.

Sartzetakis, E.S., and P. Tsigaris. 2005. "Environmental externalities in the presence of network effects: Adoption of low emission technologies in the automobile market." *Journal of Regulatory Economics* 28 (3): 309–326.

Schoonbeek, L., and F.P de Vries. 2009. "Environmental taxes and industry monopolization." *Journal of Regulatory Economics* 36 (1): 94–106.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson. 2022. "Measuring news sentiment." *Journal of econometrics* 228 (2): 221–243.

Silge, Julia, and David Robinson. 2017. *Text mining with R: A tidy approach.* " O'Reilly Media, Inc."

Simpson, R.D. 1995. "Optimal pollution taxation in a Cournot duopoly." *Environmental and Resource Economics* 6 (4): 359–369.

Sorensen, Lucy C, Ashley M Fox, Heyjie Jung, and Erika G Martin. 2019. "Lead exposure and academic achievement: evidence from childhood lead poisoning prevention efforts." *Journal of Population Economics* 32:179–218.

Stoll, Christian, Lena Klaaßen, and Ulrich Gallersdörfer. 2019. "The carbon footprint of bitcoin." *Joule* 3 (7): 1647–1661.

Stone, Richard. 2002. *Counting the cost of London's killer smog.*

Sugeta, Hajime, and Shigeru Matsumoto. 2007. "Upstream and downstream pollution taxations in vertically related markets with imperfect competition." *Environmental and Resource Economics* 38:407–432.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics* 225 (2): 175–199.

Taylor, Kyla W, Sorina E Eftim, Christopher A Sibrizzi, Robyn B Blain, Kristen Magnuson, Pamela A Hartman, Andrew A Rooney, and John R Bucher. 2025. "Fluoride exposure and children's IQ scores: a systematic review and meta-analysis." *JAMA Pediatrics.*

Tetlock, Paul C. 2007. "Giving content to investor sentiment: The role of media in the stock market." *The Journal of Finance* 62 (3): 1139–1168.

Thorsrud, Leif Anders. 2020. "Words are the new numbers: A newsy coincident index of the business cycle." *Journal of Business & Economic Statistics* 38 (2): 393–409.

Till, Christine, Philippe Grandjean, E Angeles Martinez-Mier, Howard Hu, and Bruce Lanphear. 2025. "Health Risks and Benefits of Fluoride Exposure During Pregnancy and Infancy." *Annual Review of Public Health* 46 (1): 253–274.

Timmermann, Allan. 2008. "Elusive return predictability." *International Journal of Forecasting* 24 (1): 1–18.

Tiwari, Aviral Kumar, Satish Kumar, and Rajesh Pathak. 2019. "Modelling the dynamics of Bitcoin and Litecoin: GARCH versus stochastic volatility models." *Applied Economics* 51 (37): 4073–4082.

Tobey, James A, and Henri Smets. 1996. "The Polluter-Pays Principle in the Context of Agriculture and the Environment." *World Economy* 19 (1): 63–87.

Treich, Nicolas. 2021. "Cultured meat: Promises and challenges." *Environmental and Resource Economics* 79 (1): 33–61.

Veneri, Federica, Inga Iamandii, Marco Vinceti, Linda S Birnbaum, Luigi Generali, Ugo Consolo, and Tommaso Filippini. 2023a. "Fluoride exposure and skeletal fluorosis: a systematic review and dose-response meta-analysis." *Current Environmental Health Reports* 10 (4): 417–441.

Veneri, Federica, Marco Vinceti, Luigi Generali, Maria Edvige Giannone, Elena Mazzoleni, Linda S Birnbaum, Ugo Consolo, and Tommaso Filippini. 2023b. "Fluoride exposure and cognitive neurodevelopment: systematic review and dose-response meta-analysis." *Environmental Research* 221:115239.

Villena, Mauricio G, and María José Quinteros. 2024. "Corporate Social Responsibility, Environmental Emissions and Time-Consistent Taxation." *Environmental and Resource Economics* 87 (1): 219–255.

Vinceti, Silvio Roberto, Federica Veneri, and Tommaso Filippini. 2024. "Water fluoridation between public health and public law: an assessment of regulations across countries and their preventive medicine implications." *Annali di Igiene, Medicina Preventiva e di Comunità* 36 (3).

Vogt, Richard L, Linden Witherell, Denise Larue, and Douglas N Klaucke. 1982. "Acute fluoride poisoning associated with an on-site fluoridator in a Vermont elementary school." *American Journal of Public Health* 72 (10): 1168–1169.

Waitzman, Norman J, Ali Jalali, and Scott D Grosse. 2021. "Preterm birth lifetime costs in the United States in 2016: an update." *seminars in Perinatology* 45 (3): 151390.

Wald, Nicholas J, Joan K Morris, and Colin Blakemore. 2018. "Public health failure in the prevention of neural tube defects: time to abandon the tolerable upper intake level of folate." *Public Health Reviews* 39:1–11.

Wan, Guanghua, Tongjin Zhang, and Xiaoshan Hu. 2025. "Income effect of prenatal sunlight exposure: Empirical evidence from China." *Health Economics* 34 (1): 45–67.

Wang, Rui, Xi Chen, and Xun Li. 2022. "Something in the pipe: the Flint water crisis and health at birth." *Journal of Population Economics* 35 (4): 1723–1749.

WEF. 2021. "Net-zero challenge: The supply chain opportunity."

Williams, Jennifer, Cara T Mai, Joe Mulinare, Jennifer Isenburg, Timothy J Flood, Mary Ethen, Barbara Frohnert, Russell S Kirby, Centers for Disease Control, Prevention, et al. 2015. "Updated estimates of neural tube defects prevented by mandatory folic acid fortification—United States, 1995–2011." *CDC: Morbidity and Mortality Weekly Report (MMWR)* 64 (1): 1–5.

Wing, Coady, Seth M Freedman, and Alex Hollingsworth. 2024. "Stacked difference-in-differences." WP 32054.

Yin, X. 2003. "Corrective taxes under oligopoly with inter-firm externalities." *Environmental and Resource Economics* 26 (2): 269–277.

Yin, Xing, Xiaolin Chen, Xiaolin Xu, and Lianmin Zhang. 2020. "Tax or subsidy? Optimal carbon emission policy: a supply chain perspective." *Sustainability* 12 (4): 1548.

Yu, Jun. 2002. "Forecasting volatility in the New Zealand stock market." *Applied Financial Economics* 12 (3): 193–202.

Zhang, Jianfen, Na Zhang, Junbo Lu, Shufang Liu, Yongwei Lin, and Guansheng Ma. 2025. "Seasonal fluctuation of total water intake and hydration status among young men and women: a prospective cohort study." *Frontiers in Nutrition* 11:1463501.

Zhang, Jing. 2012. "The impact of water quality on health: Evidence from the drinking water infrastructure program in rural China." *Journal of Health Economics* 31 (1): 122–134.

Zheng, Tingguo, Xinyue Fan, Wei Jin, and Kuangnan Fang. 2024. "Words or numbers? Macroeconomic nowcasting with textual and macroeconomic data." *International Journal of Forecasting* 40 (2): 746–761.

Zhu, Chenbo, Juntian Yue, and Jing Chen. 2022. "Green product development and order strategies for retailers." *Sustainability* 14 (15): 9556.