

UNIVERSITÀ COMMERCIALE “LUIGI BOCCONI”
PHD SCHOOL

PhD program in Economics and Finance

Cycle: 36th

Disciplinary field: SECS-P/06

**Essays on Climate Adaptation,
Agriculture, and Health in Developing
Countries**

Advisor: **Joseph-Simon Görlach**

Co-advisor: **Jérôme Adda**

PhD Thesis by

Iván Jomin Kim Taveras

ID number: 3141117

Year 2026

Acknowledgements

I thank God for this journey, and my family for their unconditional love: my loving parents, Kim Pyoung Ku (김 평 구) and Aida Taveras de Kim; my sister, Mabel; and my two grandmothers, Kim Young Sil (김 영 실) for her constant encouragement, and Marta Varela, who helped my parents raise me and has been my greatest champion. To Katarina, thank you for the laughter, encouragement, patience, and trips that made our work lighter. Knowing you has been a gift.

I would like to express my gratitude to my advisors, Jérôme Adda and Joseph-Simon Görlach. Their guidance, support, and insightful feedback have been instrumental throughout my PhD journey.

A special thank you goes to my friend and mentor, Alex Armand. His wisdom, encouragement, and work ethic were a constant source of strength and inspiration.

I am also immensely grateful for the friendships I found at Bocconi. My sincere thanks to my close friends and colleagues, with a special mention to Alberto, Cristóbal, Enrico, Felix, Julian, Maik, and Matteo. I also send a heartfelt thank you to my friends in Spain and the Dominican Republic for their constant encouragement from afar.

I would like to extend my appreciation to Castelbarco, the residence that welcomed me during my first year. It was the perfect place to meet wonderful people, and I continue to be inspired by its mission.

My gratitude extends to the entire community at Bocconi. I am thankful to the professors who taught me during the first two years of the programme, and to all faculty members and fellow PhD students. I also wish to thank the university's administrative and support staff, especially Angela, Giulia, Marika, Patrizia, and Rebecca, whose work behind the scenes makes our work possible.

Finally, I wish to gratefully acknowledge the generous financial support provided by Università Commerciale “Luigi Bocconi” and the Fondazione Invernizzi.

Thesis abstract

This thesis investigates the human consequences of two critical, yet understudied, aspects of climate change in the developing world: shifts in the timing of the rainy season and ocean acidification. It provides new empirical evidence on whether and how vulnerable populations adapt, the limits of their adaptation strategies, and the resulting impacts on various outcomes.

Chapter 1 examines climate risk in Sub-Saharan Africa. Using micro-data from six countries, I show that a one-week delay in the rainy season onset reduces yields by 2% and consumption by 1%. Damages disproportionately harm vulnerable female-managed plots, whereas education and wealth build resilience. Farmers adapt by delaying planting, but this is insufficient due to two informational frictions. Specifically, *false onsets*—rare but misleading early rains followed by a dry spell—trigger damaging premature planting. Negative impacts are concentrated in locations experiencing long-term climatic shifts, indicating a persistent failure to adapt. Projections suggest cumulative discounted losses from 2025 to 2050 could reach 10% of 2024 real GDP. These findings establish shifts in seasonal timing as a first-order economic threat and highlight the value of short-range forecasts.

Chapters 2 and 3 shift focus to the world's oceans, establishing a causal link between ocean acidification and human health. Chapter 2 shows that increased ocean acidity adversely affects small-scale fisheries in 36 low- and middle-income countries. This shock propagates to households through a nutritional channel: reduced seafood availability leads to higher prices and diminished nutrient consumption by pregnant women, significantly increasing neonatal mortality risk. The analysis finds no evidence of contemporaneous adaptation.

Building on these findings, Chapter 3 investigates long-run adaptation. I find no evidence that historically exposed communities have developed resilience; conversely, they appear more vulnerable. After quantifying the substantial historical mortality burden, I project future impacts: under a high-emissions scenario (RCP8.5) without adaptation, ocean acidification could cause 77.2 million cumulative neonatal deaths by 2100. However, a low-emissions pathway (RCP2.6) could avert nearly 40 million of these deaths. Collectively, these essays demonstrate that subtle environmental changes have economically important consequences, underscoring the need for policies that enhance adaptive capacity and address informational barriers.

Contents

Acknowledgements	2
Thesis abstract	3
1 Shifting Seasons: Agricultural Adaptation and Resilience in Africa	6
1.1 Introduction	7
1.2 Linking the onset of the rainy season to crop yields	11
1.3 Data	13
1.3.1 Household, individual, and plot-level data.	13
1.3.2 Nutritional outcomes from Demographic and Health Surveys.	15
1.3.3 Onset of the rainy season and weather variables.	15
1.4 Empirical strategy	18
1.5 Results	20
1.5.1 The impact of rainy season onset on agricultural productivity	20
1.5.2 Impact on household welfare: Consumption and nutrition	22
1.5.3 Farmer adaptation strategies in response to onset variation	23
1.5.4 Heterogeneity of impacts: Vulnerability, resilience, and adaptation	26
1.5.5 Limits to adaptation: False onsets and farmer awareness	28
1.5.6 Projecting future damages	30
1.6 Conclusion	32
Appendix A1: Data and methodological procedures	46
A1.1 Variables and data sources	46
A1.2 Descriptive statistics	47
A1.3 Weight Construction	50
Appendix B1: Supplementary results	51
B1.1 Trends in rainy season onset, cessation, and length	51
B1.2 Issues related to identification	54
B1.3 Robustness of impact on productivity	57
B1.4 Stability of impacts on productivity over time	63
B1.5 Supplementary results on adaptation strategies	64
B1.6 Simulation methodology to correct for aggregation bias	68
B1.7 The role of adjusting planting dates	72
B1.8 Additional results for nutritional outcomes	73
B1.9 Stability of planting months	75
B1.10 False onsets: Robustness and alternative agronomical pathways	75
B1.11 Additional results on heterogeneity	77
Appendix C1: Methodology for Projecting Economic Damages	80
2 From Sea to Shore: The Impact of Ocean Acidification on Child Health	86

2.1	Introduction	87
2.2	Background	90
2.3	Data	91
2.4	Empirical strategy	94
2.5	Results	97
2.5.1	Defining the causal pathway of ocean acidity	98
2.5.2	The effect of early-life exposure	100
2.5.3	Heterogeneity by resource exploitation	104
2.6	Conclusions	105
	Appendix A2: Data and methodological procedures	125
A2.1	Variables, data sources, and the selection of DHS surveys	125
A2.2	Descriptive statistics	129
	Appendix B2: Supplementary results	131
B2.1	Coastal features and income processes	131
B2.2	Alternative definitions of coastal area	135
B2.3	Fertility	136
B2.4	Fishing and seafood dependency	139
B2.5	Selective migration	146
B2.6	Issues related to identification and measurement	149
B2.7	Aggregate shocks	151
B2.8	Responses in food consumption	153
B2.9	Robustness to assuming a shorter gestation period	153
B2.10	Early-life mortality rates	155
B2.11	Supplementary results on inference	155
B2.12	Parental investments and adaptation	156
B2.13	Physical development and adulthood outcomes	159
3	The Future Impact of Ocean Acidification on Child Health	163
3.1	Introduction	164
3.2	Data	166
3.3	Methodology	167
3.3.1	Testing for long-run adaptation	168
3.3.2	Estimating historical impacts and projecting future burden	170
3.4	Results	172
3.4.1	Evidence on long-run adaptation	172
3.4.2	The historical impact of ocean acidification	173
3.4.3	Projections of future mortality burden	174
3.5	Conclusion	175

Chapter 1

Shifting Seasons: Agricultural Adaptation and Resilience in Africa

Abstract

Climate risk threatens Sub-Saharan Africa's rainfed agriculture. Using microdata from six countries, I show that a one-week delay in the onset of the rainy season reduces yields by 2% and consumption by 1%. Damages disproportionately harm the most vulnerable, with the effects being most pronounced on female-managed plots, while education and wealth build resilience. Farmers adapt by delaying planting, but this is insufficient due to informational frictions. False onsets—misleading early rains followed by a dry spell—trigger premature planting, increasing damages. The negative impacts are concentrated in locations experiencing long-term climatic shifts, indicating a persistent failure to adapt. Projecting these damages forward reveals a substantial threat: cumulative discounted losses from 2025 to 2050 could reach up to 10% of 2024 real GDP. These findings establish shifts in seasonal timing as a first-order economic threat and highlight the value of short-range forecasts in mitigating this risk.

Keywords: Rainy season onset, Agricultural productivity, Adaptation, Africa, LSMS

JEL Codes: Q12, Q54, O13

1.1 Introduction

Sub-Saharan Africa's reliance on rainfed agriculture makes its population exceptionally vulnerable to weather risk. In a region where approximately 95% of cropland is rainfed and 55-62% of the workforce is employed in agriculture (IPCC, 2022), shifts in weather patterns pose a substantial threat to livelihoods. The *timing* of the rainy season is a particularly important feature of weather that is often overlooked. The onset of the rainy season has a direct biophysical impact on agricultural productivity by shifting crops' exposure to weather conditions throughout critical growth stages, and an indirect effect through an adaptation channel. Through this channel, weather is not merely an input but also an uncertain signal that informs farmers' decisions across several dimensions but most importantly their planting schedule (Kala, 2019). As climate change evolves, understanding farmers' adaptation to variation in timing is essential. This requires both identifying the strategies farmers employ and the frictions that limit the effectiveness of their responses or turn their heuristics into costly mistakes. First, it can inform the design of interventions to address the growing impacts of climate change; the strategies farmers already employ are, by revealed preference, those they deem most cost-effective among their feasible options (Hultgren et al., 2025). Second, it helps identify the most promising areas for intervention, especially where the adoption of agricultural technologies is low (Suri et al., 2024).

This paper addresses a central question for climate resilience: How do weather signals shape the adaptation to, and economic impacts of, shifting seasons? To answer this, the analysis proceeds in three parts. First, I quantify the causal impacts of short-run variation in the onset of the rainy season on agricultural yields and household welfare. Second, I examine the adaptation strategies farmers employ and the informational frictions that limit their response. Finally, I link these short-run semi-elasticities to long-term climatic trends to project the future economic damages. The analysis focuses on six Sub-Saharan African countries: Ethiopia, Malawi, Mali, Niger, Nigeria, and Tanzania, which collectively represent 35% of the region's population and 40% of its economy (The World Bank, 2025). Because the start of the rainy season cannot be observed directly, I use daily weather and soil data to identify the first date when conditions are suitable for planting.

Before the main empirical analysis, I establish the climatic context by documenting a significant long-term shift in the timing of the rainy season. The data reveal a stark trend for the rainy season onset: across all six countries, it has shifted significantly later between 1979 and 2024. The magnitude of this delay ranges from approximately 6 days in Tanzania to over 25 days in Nigeria. The trend for cessation dates is more varied; in some countries, such as Niger and Nigeria, the season's end has also become later, while in others like Mali and Tanzania, it now arrives earlier. The net result for most countries is a discernible shortening of the rainy season. For example, the growing period in Mali has shortened by approximately 26 days, while in

Niger, where the later onset was offset by a later cessation, the length of the rainy season has remained broadly unchanged.

My identification strategy is designed to first quantify the net economic impacts of onset timing, and then to disentangle the role of the adaptation channel in driving these impacts. The strategy exploits substantial year-to-year variability in onset timing around long-term trends, focusing on short-run, plausibly exogenous shocks. To isolate these shocks, I compare variation within locations over time, controlling for local trends and regional conditions using a multi-way fixed effects structure that includes high-resolution grid-cell fixed effects and climatic-area-by-country-by-year fixed effects. To control for the direct biophysical channel—such as changes in precipitation patterns resulting from a delayed rainy season—the model includes a rich set of weather controls. By accounting for these aggregate weather impacts, the strategy estimates the net effect of a pure timing shock, distinct from overall seasonal weather.

I find that a one-week delay lowers crop yields by approximately 2%. This effect is asymmetric, driven entirely by later-than-usual onsets, and is also short-lived, with no spillover to future agricultural seasons. These findings are robust across several specification checks, and the statistical inference holds when accounting for spatial correlation in onset dates.

A delayed onset also harms household welfare, lowering per capita consumption by 1-1.2% per week of delay and worsening women's nutritional status. It is therefore critical to understand if and how farmers respond. I find that they primarily adjust along two low-cost margins: they plant later and use more seeds, with limited evidence of other short-run changes in fertiliser use, crop choice, or labour allocation. Focusing on the most intuitive strategy—shifting the planting schedule—the analysis reveals that the adjustment is insufficient, as a one-week delay in the onset leads to a planting delay of just over half a day. To confirm that this muted adaptation response is a real phenomenon and not merely a statistical artifact driven by the coarse, monthly nature of the survey data, I design a simulation exercise to correct for this measurement error and find that the true daily-level planting adjustment, while larger than the initial estimate, remains modest. The key puzzle, therefore, is why farmers do not delay planting more substantially when the rainy season arrives late.

Part of the answer lies in exploring who is most affected and how they adapt. The damages from a delayed onset are not evenly distributed; the negative impacts are borne disproportionately by the most vulnerable, with plots managed by women suffering significantly larger yield losses. Conversely, resilience is systematically linked to several farmer characteristics which induce heterogeneous incentives for adaptation. Education and household wealth make farmers less constrained, which should, in theory, improve their ability to execute optimal adaptive strategies like adjusting planting dates. In contrast, technologies such as irrigation make them less vulnerable to the timing of the onset, weakening their incentive to adapt on this margin.

This heterogeneity raises an important question: do these resilient farmers achieve better out-

comes because they are better at adapting their planting schedules? The evidence suggests this is largely not the case. An analysis of planting decisions reveals no statistically significant difference in timing adjustments for farmers with more education or more assets. The findings for technology are mixed and align with their differing incentives: access to irrigation is associated with a tendency to delay planting *less*, likely because a reliable water source reduces farmers' dependence on the rainy season for optimal growing conditions. In contrast, farmers using improved seeds are the only group to show a significantly more responsive timing adjustment. The fact that most resilient groups do not exhibit superior adaptation on this key margin suggests their resilience stems from an ability to buffer shocks rather than from superior timing decisions.

The answer to both puzzles—why adaptation is insufficient overall, and why superior adaptation on timing is not observed even among less constrained farmers—points to an informational barrier. First, farmer awareness of the underlying shock appears low and non-linear. Households are not more likely to report having suffered a negative shock to their crops in response to a typical late onset; awareness is triggered only by extreme events when the onset is exceptionally late. Even then, farmers appear to misattribute the cause, reporting the shock as a drought rather than linking the damage to the shift in seasonal timing.

Second, farmers can be misled by incorrect signals. I introduce the concept of *false onsets*: early rains followed by a dry spell. Farmers often adapt by following a common heuristic: they wait for a significant rainy event, often verifying the soil moisture, and then plant immediately afterwards (see, for instance, [Marteau et al. 2011](#)). A false onset is so damaging because it turns this reasonable strategy into a costly mistake. During the dry spell following a false onset, crops are likely to fail. Indeed, I find that farmers exposed to a false onset during a later-than-usual season plant earlier, which then requires replanting when the original seeds fail to grow. Ultimately, this combination of a false start and a late true onset more than doubles the negative impact on productivity.

The findings in this paper, which estimate the cost of short-run weather variability, have direct implications for understanding the long-term economic consequences of climate change. The analysis shows that the negative impact of a delayed onset on productivity is driven entirely by locations already experiencing a secular trend towards later rainy seasons. This evidence of a persistent failure to adapt to gradual climatic shifts provides the justification for using the estimated short-run damages as a guide for future losses. To quantify this threat, I project the Net Present Value (NPV) of damages to the year 2050. The results are stark: under a *Business as Usual* scenario, cumulative discounted losses could reach approximately 10% of real GDP in 2024 for the most affected country, Mali. While this projection is grounded in observed behavior, assuming a complete absence of future adaptation may be unrealistic. I therefore supplement the main analysis by considering two optimistic counterfactuals: one where the yield impact of onset delays linearly decreases to zero by 2050, and another that accounts for

the mitigating effect of future irrigation expansion. Even under these more optimistic scenarios, the remaining economic costs are substantial. Furthermore, the analysis reveals a significant *Sustainability Benefit*: a less severe climate path could avert over \$20 billion in damages in Nigeria alone. These findings establish that subtle shifts in seasonal timing, a less-studied feature of climate change, represent a first-order economic threat to livelihoods in the region.

This paper contributes to three main strands of literature. First, it makes a novel contribution to the literature on adaptation to climate change in low-income settings. A large body of work, comprehensively reviewed by [Carleton et al. \(2024\)](#), documents how farmers in low-income settings adapt to weather shocks along various margins. Examples from this literature are geographically diverse, including studies on adjusting agricultural inputs in Kenya ([Jagnani et al., 2020](#)), investing in irrigation in India ([Taraz, 2017](#)), adopting soil and water conservation practices in Peru ([Tambet and Stopnitzky, 2021](#)), and changing land use in both India and Peru ([Aragón et al., 2021](#)). However, this literature has largely focused on responses to shocks in temperature and total precipitation, with less attention paid to the *timing* of the rainy season. This paper argues that this is an important omission because shocks to timing introduce unique informational frictions, adding a layer of complexity to a setting where learning is already complex and costly ([Laajaj and Macours, 2024](#)). My analysis shows these frictions have real consequences: farmers' adaptation is ultimately insufficient. By documenting their spontaneous ex-post responses, I find that while the most intuitive margins are adjusting planting dates and seed use, these adjustments are often too small or poorly timed to prevent losses, presenting a puzzle as to what limits a more effective response.

Second, this paper contributes to the extensive literature on information and knowledge constraints in agriculture. A large body of work documents how farmers may forgo profitable strategies, particularly technology adoption, because of constraints (for a review, see [Suri et al. 2024](#)). I show that the insufficiency of adaptation is driven by two key informational frictions. The first is a failure of attribution. Because year-to-year onset variation is not always a discrete, extreme event, its impact has low salience. My analysis shows that farmers struggle to distinguish it from normal productivity fluctuations; only when the delay is large enough do they report an adverse shock, and even then, they tend to misattribute it to drought rather than the underlying timing shift. This finding is consistent with [Patel \(2025\)](#), which, in the context of soil salinity and agricultural productivity in Bangladesh, provides a framework where such attribution failures arise naturally from an identification problem: farmers interpret ambiguous signals through the lens of their priors, leading to persistent misperceptions, especially when environmental shifts are subtle rather than salient. This paper provides evidence on this phenomenon in the context of seasonal timing in SSA. I introduce and test the importance of the *false onset*—early rains followed by a dry spell. As climate change systematically delays the onset, farmers' search for planting signals increasingly occurs at the wrong time of year, turning a traditionally reliable heuristic into a costly mistake that more than doubles the neg-

ative impact on productivity. These insights are useful for policy design, as addressing these informational constraints through tools like reliable short-range forecasts could enhance the strategies farmers already find intuitive, a potentially more effective approach than promoting costly and complex technologies that have historically suffered from low take-up (Suri et al., 2024). Together, these findings motivate the expansion of reliable weather forecasts to serve as an early-warning system, contributing to recent work on the value of both short-range (Carmacho and Conover, 2019; Yegbemey et al., 2023; Rudder and Viviano, 2024) and long-range forecasts (Giné et al., 2018; Rosenzweig and Udry, 2019; Burlig et al., 2024).

Third, this paper contributes to the literature quantifying the economic impacts of weather shocks by examining a critical, yet understudied, dimension of weather risk. An extensive literature provides robust evidence on the impacts of shocks to temperature and rainfall quantity on agricultural and economic outcomes (Deschênes et al., 2009; Deschênes and Moretti, 2009; Barreca et al., 2016; Aragón et al., 2021). This paper complements this body of work by analysing the economic consequences of a different feature of weather: the timing of the rainy season. In the rainfed agricultural systems that dominate Sub-Saharan Africa, a delay in the onset can shift critical crop growth stages into periods of greater heat stress or misalign harvest with labor cycles, even when total seasonal rainfall is normal. Despite its agronomic importance, the economic literature on onset timing is sparse, with the few existing studies focused on India (Amale et al., 2023; Burlig et al., 2024). I complement existing evidence by providing multi-country evidence from Sub-Saharan Africa, the region most reliant on rainfed agriculture. To do so, I introduce a robust, high-resolution agronomic measure of onset, trace its effects beyond agricultural production to welfare indicators, and uncover how these damages are borne disproportionately by the most vulnerable farmers (Udry, 1996). By first estimating the impact of onset timing and then identifying its mechanisms, my findings show that these damages stem largely from responses to uncertain signals, not just from inevitable biophysical constraints.

The rest of the paper is structured as follows. Section 1.2 discusses the causal pathways linking the rainy season onset to crop yields. Section 1.3 describes the data sources and construction of key variables. Section 1.3.3 presents the evidence on long-term trends in the rainy season calendar. Section 1.4 outlines the empirical strategy in detail. Section 1.5 presents the main results, and Section 1.6 discusses the findings and concludes.

1.2 Linking the onset of the rainy season to crop yields

The timing of the rainy season onset affects agricultural productivity through two distinct pathways, as illustrated in Figure 1.1: a direct **biophysical channel** that impacts crop growth, and an indirect **adaptation channel** that operates through farmer decision-making.

The biophysical channel: Direct impacts on crop growth. Shifts in the actual onset date of the rainy season directly change the biophysical conditions crops experience throughout their development. The timing of the onset largely dictates when the entire growing period begins. This shift means all subsequent crop life cycle stages occur at different points in the calendar year, potentially exposing them to less favourable weather and soil conditions regarding temperature, sunlight, and moisture, even if the total length of the season were unchanged. For instance, critical phases like germination, flowering, or grain and tuber filling might be pushed into periods of excessive heat or suboptimal solar radiation, thereby negatively impacting yields irrespective of season length (Jägermeyr and Frieler, 2018; Yang et al., 2024). Of course, a common direct consequence of a delayed onset is also a shortening of the effective growing season, a trend observed in the study regions (see Section 1.3.3). Such a reduction in available time can prevent crops from accumulating sufficient warmth or completing their development, leading to lower yields. Lastly, these onset-driven climatic shifts can indirectly lower yields by changing the exposure to pests and diseases, whose life cycles also respond to environmental changes (Yang et al., 2024).

The adaptation channel: Weather as a signal for farmer decisions. Beyond its direct effects, year-to-year variation in the onset is an important source of uncertainty and a critical **signal** for agricultural producers. In anticipation of, and in response to, this signal, farmers must make crucial management decisions that significantly influence eventual crop yields. These decisions span a range of practices, including the choice of crops to plant, the selection of specific crop varieties, the quantity and type of inputs (such as seeds, fertilisers, and pesticides), the allocation of labour, and, crucially, the timing of planting activities (Jägermeyr and Frieler, 2018; Marteau et al., 2011; Minoli et al., 2022). To interpret this uncertain signal, farmers often rely on local knowledge and heuristics, such as observing rainfall patterns and soil moisture, to guide their planting decisions and mitigate risks like early-season dry spells (Wolf et al., 2015; Marteau et al., 2011). The effectiveness and feasibility of these adjustments in the face of onset variability are central to understanding agricultural outcomes and represent the core of the adaptation channel. Persistent changes in seasonal timing can also make traditionally grown crop varieties less suitable over time, requiring farmers to adapt their crop choices in the long run.

The role of other timing features. While both the length and cessation of the rainy season are important, the relationship with productivity is not always straightforward. For example, an extended season is not inherently superior if those additional days expose the crop to end-of-season stresses like drought or increased pest pressure. Indeed, adaptive farming practices might involve selecting cultivars or adjusting planting schedules to achieve a strategically shorter but better-timed growing period to avoid such terminal stresses (Minoli et al., 2022). As the empirical evidence detailed in Section 1.5 will show, it is the year-to-year change in *onset timing* that is the primary aspect of seasonal timing that causally affects productivity in the

sampled countries, with cessation timing and overall length having less significant independent effects once the onset is controlled for.

1.3 Data

I combine detailed household and plot-level survey data with high-resolution climate information to analyse the impact of rainfall onset timing. Appendix A1.1 provides further details on variable construction and specific data sources.

1.3.1 Household, individual, and plot-level data.

The core household, individual, and plot-level information is sourced from a recently released harmonised panel dataset compiling the Living Standards Measurement Study - Integrated Surveys on Agriculture (LSMS-ISA) waves for seven Sub-Saharan African countries ([World Bank, 2024](#)). I use data for six of these countries: Ethiopia, Malawi, Mali, Niger, Nigeria, and Tanzania. While some central and northern regions of Ethiopia experience a secondary, sporadic rainy season (the *Belg*), the analysis focuses on the main agricultural cycle tied to the primary *Kiremt* rainy season. Uganda is excluded from the analysis because its predominantly bimodal rainfall pattern is unsuitable for my current methodology for identifying a single main rainy season onset, which relies on defining a consistent agroclimatic year. The included survey waves span the period from 2008 to 2022. Although the harmonised dataset is meant to be a panel, I treat the data as repeated cross-sections as agricultural plots cannot be reliably tracked across waves. Appendix A1.1 provides a detailed list of the specific survey waves included for each country (see Table A1.2).

The LSMS-ISA surveys are designed to be representative of the household and smallholder agriculture sectors within the sampled countries, though it is important to note that specific sampling and coverage might vary slightly by country and wave.¹ Figure A1.1 shows the broad geographical coverage of the sampled LSMS communities. Locations are spread widely across the geography of each country. Naturally, the sample excludes hyper-arid areas where agriculture is not practiced, which is most evident in the desert regions of northern Mali and Niger. The surveys collect comprehensive data relevant to this study. The household questionnaire records demographics, education, labour participation, household nonfarm activities, shocks experienced, and crucially provides an annualised per capita measure of consumption. The agricultural questionnaires gather detailed plot-level information, including inputs, crop choices, yields, and farming practices such as planting and harvest timing. For all included surveys, I focus on the main agricultural growing season. Importantly, the harmonised dataset

¹Refer to official LSMS-ISA documentation and [World Bank \(2024\)](#) for details.

provides plot area measurements based primarily on Global Positioning System (GPS) readings, addressing known inaccuracies in self-reported areas; imputation methods based on self-reports and administrative data are used when GPS measures are missing (see [World Bank, 2024](#), for details). The dataset also provides standardised groupings for crops, which I use to analyse crop choice decisions.² The harmonised dataset defines the main crop on a plot as the crop within these categories having the highest reported monetary value.

To construct the final estimating samples, I applied several restrictions. Survey waves without geocoordinates were excluded. I also excluded the 2019–2021 Tanzania Wave 5, as its multi-year survey period and lack of planting-year information made it impossible to reliably match plots to a specific agricultural season. Only plots reporting a complete harvest are included, as incorporating yields from incomplete harvests could introduce bias into the productivity estimates. This decision is supported by formal tests showing that the timing of the rainy season onset does not significantly predict the probability of observing a fully harvested plot; Table B1.1, Column (1), presents this result. Lastly, I retain only plots where planting occurred within the same calendar year as the survey interview, a restriction that drops approximately 5% of plots and primarily excludes perennial crops. These criteria define the final plot-level sample used for productivity analyses.

After applying these criteria, the final sample for plot-level analysis comprises approximately 130,000 plot-wave observations. For the household-level analyses of consumption, I restrict the sample to the approximately 29,000 household-wave observations from the agricultural module that report a complete harvest for all plots. This represents 75% of all household-wave observations. The individual-level sample includes all working-age individuals from these selected households.

This restriction on complete harvests for the consumption analysis is necessary because the surveys provide an annualised consumption measure without specifying its timing relative to the harvest. Including households yet to complete their harvest would mean comparing post-harvest households with those still in their lean season, the period between planting and harvesting, introducing significant unobserved heterogeneity. While this raises a valid concern about selection bias—for instance, a late onset could delay a harvest, making it more likely the survey occurs before completion—I explicitly test for this. Column (2) of Table B1.1 confirms that the onset timing does not significantly predict the probability of observing a complete harvest, suggesting the restriction does not systematically bias the results. Furthermore, as shown in Table 1.2, the main findings on consumption are robust to including all households, regardless of their main occupation sector. Appendix A1.1 details the construction of key outcome and control variables.

²The categories are barley, wheat, rice, sorghum, maize, millet, perennials such as fruit and tree crops, legumes, root crops, nuts, and others.

1.3.2 Nutritional outcomes from Demographic and Health Surveys.

To examine whether the impacts of rainfall onset variability extend to nutritional outcomes, I incorporate data from all 25 available DHS waves for the sampled countries that provide geocoded cluster information. A detailed list of the specific DHS survey years I included for each country is available in Table A1.2. DHS surveys are nationally representative, collecting comprehensive data on health and population with a particular focus on maternal and child health. Women aged 15–49 are the primary respondents, providing information on household characteristics, birth histories, wealth, and human capital, alongside objective anthropometric measurements (weight and height) for themselves and their children under five.

For this analysis, I primarily use samples of adult female respondents with available objective anthropometric data: a main sample from agricultural households, and a full sample of all women to test for spillovers. The key nutritional outcomes I examine for women are weight-for-height z-scores and indicators for whether she is classified as wasted or underweight. Analyses using data for children under five (weight-for-height z-scores, wasting, underweight) are presented in the appendix. Appendix A1.1 details the construction of all nutritional variables.

1.3.3 Onset of the rainy season and weather variables.

I rely on daily climate information sourced from a dataset of agrometeorological indicators derived from ERA5 reanalysis, provided by the Copernicus Climate Change Service (C3S) Climate Data Store (Boogaard et al., 2020). This dataset covers the period January 1979 to December 2024 at a $0.1^\circ \times 0.1^\circ$ spatial resolution, corresponding to approximately 11×11 kilometres at the equator. Throughout the text, I will use the terms $0.1^\circ \times 0.1^\circ$ grid cell and *location* interchangeably. I extract daily 2-meter minimum and maximum temperature, 2-meter dewpoint temperature, total precipitation, surface solar radiation, and 10-meter wind speed magnitude. These data provide the basis for calculating the rainy season onset and for constructing the weather variables used as controls in my empirical analysis (as detailed in Section 1.4).

To define the rainy season calendar, I calculate its onset, cessation, and overall length using a dynamic agronomic model that simulates the daily soil water balance (Takele and Dell’Acqua, 2023). The model requires a defined search window, comprising an earliest possible onset, a latest possible onset, and a latest possible cessation date. To set these parameters on a location-specific basis and account for regional differences, I first establish the usual start and end dates of the wet season by analysing long-term daily rainfall patterns, following Dunning et al. (2016).³ Using these location-specific dates, I define the search window as 90 days be-

³Appendix A1.1 provides details on how these typical start and end dates for each location are derived from historical rainfall data.

fore and after the usual start for the onset, and 90 days after the usual end for the cessation. Within this window, the model simulates day-to-day changes in soil moisture by combining daily weather data with information on local elevation (Danielson and Gesch, 2011) and the soil's water-holding capacity (Leenaars et al., 2018).

The *onset* date is then identified based on criteria ensuring sustained water availability suitable for crop growth: the ratio of actual-to-potential evapotranspiration (ETa/PET) must exceed 0.5 for seven consecutive days, followed by a 20-day period where simulated plant available water remains above the wilting point. The *cessation* date is defined as the first day when the ETa/PET ratio falls below 0.5 for seven consecutive days, followed by 12 consecutive non-growing days during which plant-available water remains below the wilting point over the root zone. The *length* of the rainy season is then calculated as the total number of days from the identified onset date to cessation date. Finally, the calculated onset and cessation dates for each location and year are converted into the corresponding week number of the year.

This agronomic definition, grounded in soil water availability, is employed for two primary reasons. First, by incorporating potential evapotranspiration (atmospheric water demand) and soil water holding capacity, it moves beyond rainfall thresholds to identify when sufficient moisture is likely available in the root zone to support germination and early plant growth. This provides a robust indicator of the start of the viable planting window, distinguishing true onsets from potential false starts caused by initial rains followed by dry, high-evaporation conditions. While the optimal planting date within this window varies depending on the maturation period of specific crops, this calculated onset marks the earliest point at which planting becomes reliably feasible for rainfed agriculture.

Second, this focus on soil moisture aligns closely with the heuristics farmers themselves report using when making planting decisions. Rather than relying solely on rainfall patterns, farmers commonly assess soil moisture directly to judge planting suitability. Evidence supporting the importance of soil moisture cues in farmer decision-making comes from diverse contexts, including Ghana (Antwi-Agyei et al., 2022), semi-arid India (Giné et al., 2018), and Niger (Marteau et al., 2011). Therefore, an onset definition based on simulated soil water balance provides a measure that is not only agronomically relevant but also resonates with practical farming strategies.

LSMS-ISA and DHS communities are matched to their corresponding $0.1^\circ \times 0.1^\circ$ ERA5 grid cell using their GPS coordinates to link survey data with the calculated onset week and the grid-level weather controls derived from the same ERA5-derived dataset. Finally, to account for broad agro-ecological characteristics, locations are classified into climatic areas based on an updated global Köppen-Geiger climate classification (Metzger et al., 2023). The detailed classes are aggregated into Tropical, Arid, and Temperate zones to ensure sufficient observations for fixed effect estimation. Appendix A1.1 provides further details on data processing.

The rainy season calendar over time. My analysis primarily focuses on the onset of the rainy season for two main reasons. First, while the entire rainy season calendar exhibits year-to-year variation, the timing of the onset is particularly critical. It is around this period that farmers make crucial planting decisions which have substantial implications for plot productivity (see Section 1.2). Second, as this section will show, onset timing across the sampled countries reveals a clear, consistent pattern of change over recent decades, a characteristic less evident in cessation dates.

To estimate the trends, I constructed a panel dataset at the $0.1^\circ \times 0.1^\circ$ grid cell level, covering the period 1979–2024. For each cell and year, I use the day of year (DOY) for the onset, cessation, and length of the rainy season, as detailed in this section. This sample is restricted to only those grid cells that spatially match the locations of households in the LSMS-ISA surveys. I then estimate country-specific trends using a fixed-effects model, with standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level (approximately 55×55 kilometres).

Figure 1.2 presents the country-specific trend analyses for the onset variable. The results consistently show that the onset of the rainy season is occurring later across all sampled countries. The total delay over the historical period is substantial, ranging from 5.8 days in Tanzania to 25.1 days in Nigeria.

Beyond the shift in average onset timing, I also explored whether the year-to-year variability is changing.⁴ To do this, I de-trended the onset data using the residuals from the superior-fitting quadratic models. Figure B1.2 plots the interquartile range (IQR) of these residuals within 5-year periods, revealing differing trends across the sample. In Ethiopia, Malawi, and Tanzania, there is an increasing trend in the IQR. This growing unpredictability has direct consequences for farmers: it makes traditional knowledge less reliable and increases the risk of costly planting errors, such as being misled by a false onset. Conversely, Mali, Niger, and Nigeria show a tendency towards decreasing IQR over time. For farmers in these regions, the rainy season, while still shifting, is becoming more predictable year-to-year, which can lower uncertainty and facilitate adaptation.

Regarding the cessation of the rainy season (Figure B1.3), trends are more heterogeneous. While Ethiopia, Niger and Nigeria exhibit a tendency towards later cessation dates, Malawi, Mali, and Tanzania show the season ending earlier. Despite these varied patterns, the net effect on the length of the rainy season (Figure B1.4) is a discernible decrease over time across most of the study area, with the exception of Niger, where the season length remains roughly unchanged. This reduction in the growing period is primarily driven by the consistent and significant delay I observe in the season's onset, reinforcing the importance of onset timing as the principal determinant of changes in the effective length of the agricultural season.

⁴Testing for a time trend in the interquartile range of the onset residuals is conceptually equivalent to testing for a form of heteroskedasticity where the error variance changes over time, e.g., $\sigma_{it}^2 = f(t)$. This speaks to whether the *predictability* of the onset is changing.

1.4 Empirical strategy

To estimate the causal effect of rainy season onset timing on agricultural, economic, and nutritional outcomes, I leverage temporal variation in onset dates within specific, fine-resolution geographic locations. Identification relies on comparing outcomes within the same $0.1^\circ \times 0.1^\circ$ grid cell across years, exploiting variations in onset timing relative to the cell's long-run average conditions and net of local trends. This approach matches survey units via their geocoordinates to grid cell ℓ , within climatic area a and country c , in year t . The analysis occurs at both the plot and household level for the LSMS-ISA data, and at the woman-level for the DHS data, as detailed in Section 1.3.

Given this identification strategy, the benchmark specification I estimate is:

$$Y_{ilt} = \beta W_{\ell t} + \mathbf{X}'_{ilt} \gamma + \mu_\ell + \delta_{cat} + \epsilon_{ilt} \quad (1.1)$$

Here, Y_{ilt} represents the outcome of interest for unit i (a plot, household, or individual). Key outcomes from the LSMS-ISA include agricultural productivity at the plot-level, household per capita consumption at the household-level, and farmer adaptation strategies. From the DHS, I examine women's nutritional outcomes at the individual-level. The key independent variable, $W_{\ell t}$, is the calendar week of the rainy season onset. The coefficient of interest, β , thus captures the change in the outcome associated with a one-week delay. As context for the magnitude of these year-to-year shocks, the pooled trend across all sampled cells reveals a significant long-term delay in onset dates of approximately 18 days between 1979 and 2024 (see Section 1.3.3). Conditional on the model's fixed effects, a one standard deviation shift in onset timing corresponds to approximately two weeks.

Identification of the causal effect β relies on the multi-way fixed effects structure, which isolates plausibly exogenous short-run deviations in $W_{\ell t}$. The model includes two key sets of fixed effects. First, μ_ℓ represents high-resolution grid-cell fixed effects ($0.1^\circ \times 0.1^\circ$). These *location fixed effects* account for all time-invariant characteristics specific to location ℓ , such as its average climate, soil properties, and market access. However, detecting the potentially subtle impact on nutritional outcomes is challenging, particularly with the more limited sample of agricultural households in the DHS of approximately 59,000 woman-year observations. To ensure the model retains sufficient statistical power for this specific analysis, the resolution for these fixed effects is therefore coarsened to the $0.5^\circ \times 0.5^\circ$ grid-cell level. Second, δ_{cat} represents a full set of fixed effects for each unique combination of climatic area, country, and year. These *spatially specific time effects* are crucial for identification for two primary reasons: they flexibly control for any underlying secular trends in outcomes, and they absorb any unobserved shocks common to a climatic region within a given country in a particular year. As detailed in Appendix A1.2, a grid cell ℓ is considerably smaller than any climatic area within a coun-

try (Figure A1.1), and each climatic area-by-country group contains a large number of unique grid cells with sampled households (Table A1.3). This nested variation allows the model to distinguish between broad, regional shocks and time-invariant local characteristics.

The vector $\mathbf{X}_{ilt}$ incorporates available time-varying controls to account for remaining observable differences. First, and most importantly for the conceptual framework outlined in Section 1.2, a comprehensive set of weather controls is included to separate the effect of onset timing from aggregate seasonal weather. By accounting for total rainfall and temperature patterns, these controls allow the model to estimate the net impact of a shift in seasonal timing, holding aggregate conditions constant. The analysis then proceeds to test the hypothesis that adaptation failures are a key driver of these net negative impacts. This distinction is crucial for policy. As the subsequent analysis will show, a large portion of these damages stems from identifiable adaptation failures—rather than inevitable biophysical consequences—which implies that these losses can be mitigated with targeted interventions, such as improved weather forecasts. These weather controls, measured at the grid-cell level over the calendar year, include: the number of days with precipitation below the 25th and above the 75th percentiles of the historical distribution, daily maximum relative humidity, average daily temperature, total annual precipitation, and harmful degree days (days with maximum temperature exceeding the 90th percentile). Appendix B1.3 presents estimates using these same controls aggregated over different temporal windows.

Second, specific controls are also included depending on the level of analysis. For plot-level outcomes, controls include household size, total farm area, and plot manager characteristics (age, sex, marital status, education). For household-level outcomes, controls include indicators for any household member’s education, non-farm enterprise ownership, number of plots, and an urban indicator. For individual-level nutritional outcomes from the DHS, controls include the woman’s years of education, her age and age squared, household head’s age, household size, and an urban indicator.

Finally, ϵ_{ilt} represents the idiosyncratic error term. Because outcomes and the onset of the rainy season exhibit spatial correlation, I cluster standard errors at a larger spatial unit: the $0.5^\circ \times 0.5^\circ$ grid-cell level. All descriptive statistics and regression analyses use the survey weights provided by LSMS-ISA and DHS. These weights are adjusted for the pooling of multiple survey waves per country to ensure representativeness. For plot-level analyses, household weights are further rescaled by the number of plots per household. Appendix A1.3 provides full details.

The demanding fixed-effects structure may raise concerns about whether sufficient variation remains for identification. For transparency, all regression tables report the number of observations used for identification and those dropped due to singleton groups. Furthermore, I provide two pieces of evidence in Appendix B1.2 to address these concerns directly. A variance de-

composition (Figure B1.5) confirms that substantial within-location variation remains in the key variables after accounting for fixed effects. A plot of the residualised onset of the rainy season (Figure B1.6) shows a symmetric, zero-centred distribution consistent with an idiosyncratic shock, providing supportive evidence for the exogeneity assumption.

1.5 Results

This section presents the main empirical findings. I begin by showing that a delayed rainy season onset reduces agricultural productivity and, in turn, harms household welfare, as measured by consumption and nutritional outcomes. I then investigate how farmers adapt to this shock, focusing on adjusting their planting schedules to offset losses. Finally, I explore a key challenge to this strategy: how a false onset, a misleading signal that prompts early planting, can lead to costly mistakes and worsen outcomes.

1.5.1 The impact of rainy season onset on agricultural productivity

The analysis begins by assessing the impact of rainy season onset timing on agricultural productivity. My findings indicate that a delayed start to the rainy season causally reduces agricultural yields.

Table 1.1 presents the estimated impact of a one-week delay in the rainy season onset on agricultural productivity, measured as the logarithm of yield in constant 2020 US dollars per hectare. Across all specifications, a later onset is associated with a statistically significant reduction in yields. In my preferred specification (Column 3), which includes a comprehensive set of location fixed effects ($0.1^\circ \times 0.1^\circ$ grid cell), climatic area by country by year fixed effects, weather controls, and plot-level demographic and farm characteristics, a one-week delay in the rainy season onset causes an approximate 2% decrease in agricultural yields. This effect is statistically significant at conventional levels.

The magnitude of this impact is economically important. To put this into perspective, a one standard deviation delay in the onset of the rainy season—which corresponds to one and a half weeks, conditional on the included fixed effects—leads to a reduction in agricultural productivity of about 3%. This estimate is remarkably stable: the impact of a one-week delay is a consistent 2% reduction in yields across all specifications, irrespective of the inclusion of weather or plot-level controls (Table 1.1, Columns 1-3).

To further explore the nature of this relationship, I check whether the impact of onset timing is uniform or asymmetric. For this, I did not use the continuous onset week variable. Instead, I binned each location-year's onset timing based on its historical distribution within its $0.5^\circ \times 0.5^\circ$ grid cell. Specifically, I created indicator variables for whether the onset fell into the

bottom, middle, or top tercile. I then re-estimated equation (1.1) using indicator variables for the earliest and latest terciles, with the middle tercile serving as the omitted category. The results, presented in Figure 1.3, reveal a stark asymmetry. An onset occurring in the earliest third of the historical distribution has an estimated effect on log yields of -0.3%, which is statistically indistinguishable from zero when compared to an onset in the middle tercile. In contrast, an onset occurring in the latest third of the historical distribution leads to a statistically significant reduction in log yields by approximately 12%, compared to an onset in the middle tercile. This confirms that the adverse productivity consequences are predominantly driven by significantly later-than-usual rainy seasons. This finding holds when controlling for grid-cell-specific linear trends (Figure B1.7), which I implement at the $0.5^\circ \times 0.5^\circ$ binning level. This check motivated by recent concerns in the literature about potential biases in binned climate analyses (Jones et al., 2025). See Appendix B1.2 for further details.

My benchmark empirical strategy, detailed in Section 1.4, imposes a stringent requirement on the data by leveraging only within-cell variation in onset timing, net of controls and local trends. Consequently, any remaining threat to identification would need to stem from an unobserved, cell-level time-varying factor that is correlated with the year-to-year deviations in onset timing and directly affects agricultural productivity through a channel other than the onset timing itself. Such a specific confounder is arguably less likely given that the timing of the rainy season, and its interannual variability, is substantially influenced by large-scale climatic patterns, such as the El Niño Southern Oscillation, the Indian Ocean Dipole, and Tropical Atlantic Variability, which affect precipitation regimes and climate extremes across Sub-Saharan Africa (Zita et al., 2025).

Nevertheless, I verify the robustness of this core finding through an extensive series of checks, detailed in Appendix B1.3. First, the negative and significant effect of onset timing remains stable even when employing different definitions of local trends by varying the geographic level at which such trends are defined; this is noteworthy as altering the local trend specification effectively changes the definition of the onset shock by modifying the variation used for identification (Table B1.3). Second, the impact is confirmed to be contemporaneous: a placebo test using future onset ($t+1$) shows a coefficient close to zero and statistically insignificant, while lagged onset ($t-1$) also has no significant effect on current yields, indicating the impact is not persistent across years (Figure B1.8). Third, when examining other features of the rainy season, I find that cessation timing or rainy season duration do not significantly affect productivity once onset is controlled for, reinforcing the primary role of onset (Table B1.4). Fourth, the onset effect is not merely capturing other weather-related shocks, as it remains robust to the inclusion of several weather controls defined at finer temporal frequencies (quarterly and monthly) that capture variations in temperature, precipitation patterns including droughts and floods, and relative humidity (Table B1.5). Fifth, the finding is robust to various alternative transformations of the yield variable (Table B1.6) and alternative measures of output (Table

B1.7). Notably, the magnitude of the yield reduction is larger when measured in kilograms per hectare than in US dollars, suggesting a potential local price increase in response to the negative supply shock. To address this potential attenuation bias, I re-estimate the model valuing output at constant national median prices per crop (see Section B1.3 for details); this yields a larger negative impact of approximately 2.7%, confirming that local price responses likely dampen the benchmark estimate (Table B1.8). Additionally, to ensure results are not driven by composition effects from aggregating diverse crop types, I restrict the sample to cereal crops only; the negative impact remains consistent in magnitude and significance (Table B1.9). The shock also affects the extensive margin, as a one standard deviation later onset increases the probability of crop failure by approximately 5 percent. Lastly, statistical inference is robust, holding under different clustering assumptions for standard errors (Table B1.10) and confirmed by permutation-based inference (Figure B1.9).

1.5.2 Impact on household welfare: Consumption and nutrition

The reduction in agricultural productivity documented above has direct consequences for household welfare, with economic impacts that affect the broader local economy.

First, the shock to agricultural output translates into lower household spending. For agricultural households, the group most directly affected, a later onset of the rainy season leads to a statistically significant decrease in per capita consumption. As shown in Table 1.2, Columns (1)–(3) indicate that a one-week delay reduces per capita consumption by approximately 1.0% to 1.2%. To put this magnitude in perspective, a one standard deviation shift in onset timing reduces consumption by approximately 1.5%.

Importantly, these negative economic consequences are not confined to agricultural households. As shown in Columns (4)–(6) of Table 1.2, the effect on per capita consumption remains statistically significant for the full sample. This indicates the presence of local spillover effects, as the initial shock to agricultural income propagates to the non-farm economy. A primary channel for this spillover is a contraction in local demand: as farmers' incomes fall, they reduce their spending on local goods and services, depressing economic activity for the entire community. This is consistent with the finding in Table B1.16 that a late onset is associated with a significant reduction in wage employment opportunities. Beyond this direct spillover, a broader general equilibrium effect could also be at play, operating through prices. A smaller harvest may increase local food prices, which would reduce the real purchasing power of all households and transmit the shock throughout the entire local economy.

Table 1.3 shows that a delay in the onset of the rainy season leads to a statistically significant deterioration in the short-term nutritional status of women in agricultural households. Specifically, a one-week delay is associated with a 0.007 standard deviation decrease in the weight-for-height z-score. While this effect is statistically significant, its magnitude is modest. The

standardised effect indicates that for a one-standard-deviation delay in the onset, the z-score decreases by approximately 2% of the sample mean.

Importantly, this nutritional decline is not severe enough to increase the likelihood of women being classified as clinically wasted or underweight, as shown by the statistically insignificant results in columns (2) and (3). The size of the impact on weight-for-height is consistent with the analysis of productivity and consumption decreases. Table 1.3 shows that a delay in the onset of the rainy season leads to a statistically significant deterioration in the short-term nutritional status of women in agricultural households.

This negative shock, however, does not appear to be transmitted to young children. Using the anthropometric data for children, I find no statistically significant impact of a delayed onset on weight-for-height z-scores or the probabilities of being wasted or considered underweight (Table B1.19). This divergence—a nutritional cost for mothers but not for children—could reflect incomplete risk sharing within households, where adults, particularly mothers, shield younger members from shocks, potentially at their own expense (Dercon and Krishnan, 2000).

The main analysis employs location fixed effects at the $0.5^{\circ} \times 0.5^{\circ}$ grid-cell level to maintain statistical power. Table B1.20 re-estimates the model with the finer $0.1^{\circ} \times 0.1^{\circ}$ fixed effects used in the agricultural analysis. In this more demanding specification, the point estimate remains negative and similar in magnitude, but the larger standard errors render it statistically insignificant.

Furthermore, the negative nutritional impact of a late onset extends beyond agricultural households. When the sample is expanded to include all women, a delayed onset still causes a statistically significant decrease in weight-for-height z-scores (Table B1.21), suggesting the presence of general equilibrium spillovers that affect the wider community.

1.5.3 Farmer adaptation strategies in response to onset variation

Given these significant adverse impacts, I next analyse how farmers respond to year-to-year variability in the start of the rainy season. I examine adjustments in key agricultural practices, focusing on four primary margins: the timing of planting and other on-farm activities, the intensity and type of input use, choices regarding crop types, and the allocation of household labour to off-farm sectors.

The most direct way farmers can respond to a varying onset is by adjusting their planting schedule. Table 1.4 (Column 1) shows that a later start of the rainy season leads to a delay in farmers' planting. Specifically, a one-week delay in the agronomically-defined onset week causes an average delay in the approximated first planting day of the year by 0.56 days (p-value < 0.05).

The estimated responsiveness from the main text—an average planting delay of just half a day

for a full week of onset delay—is difficult to interpret directly. The coarse, monthly nature of the planting data introduces two potential and opposing sources of measurement error, making the direction of the net bias theoretically unclear.

The first and more intuitive source is attenuation bias. As shown in Figure B1.13, the median planting *month* is remarkably stable over time, with a Pearson correlation of approximately 0.90 across survey waves. This stability, however, can mask significant daily-level adjustments. For the 90% of plots where planting is completed within a single calendar month, any shift in planting dates that does not cross a month-end boundary (e.g., a 10-day delay from June 5th to June 15th) is recorded as a zero-day change. This systematically understates the true responsiveness of farmers’ planting decisions.

Conversely, the data may also suffer from an opposing boundary-crossing bias that exaggerates responsiveness. If a small, multi-day delay in the onset causes a farmer’s planting date to shift from the end of one month to the beginning of the next (e.g., a 5-day delay from May 30th to June 4th), the monthly approximation would record this as a 30-day shift. In such cases, the regression would overstate the true sensitivity to the onset shock.

To resolve this ambiguity and estimate the net effect of these two opposing biases, I implement a simulation exercise to recover the unbiased daily-level coefficient. The full methodology is detailed in Appendix B1.6, but in brief, I use the observed monthly planting data to fit location-specific daily probability distributions, which I then use in an iterative search algorithm to find the true daily coefficient that would produce my observed monthly coefficient. I conduct this exercise using two different distributional assumptions: a parametric skew-normal and a non-parametric kernel density estimation (KDE).

The simulation results confirm that while both biases exist, the attenuation from unobserved within-month adjustments is the dominant effect. As shown in Figure B1.12, the corrected daily coefficient is larger in magnitude than the original estimate. However, the magnitude of this correction is modest. Using the KDE-based draws, the corrected coefficient is only 29% larger than the biased estimate. This finding is significant. While the simulation corrects for a statistical bias, the relatively small size of the correction implies that the true daily-level response of farmers is itself limited. In other words, the measurement error is not masking a large underlying behavioral response. This suggests that farmers are not adjusting their planting schedules to the onset of the rains as flexibly as might be expected, pointing towards the presence of other frictions or decision-making heuristics that Section 1.5.5 explores.

Beyond timing, a later onset week causes a statistically significant increase in seed quantity. As shown in Table 1.4 (Column 4), the coefficient of 0.032 indicates that a one-week delay in onset leads to an approximate 3.2% increase in seed quantity used per hectare. This could reflect a strategy to ensure adequate plant density under potentially less favourable or more uncertain early growing conditions. Alternatively, an increase in seed quantity might also indicate

re-planting efforts following early crop failure, a channel for which I will present suggestive evidence when discussing the impacts of false onsets in Section 1.5.5. This increase in quantity, however, does not translate into higher expenditure on seeds. Table B1.13 (Column 1) in Appendix B1.5 shows no significant change in seed value per hectare, suggesting that farmers likely use more of the same type of seeds they customarily use, such as saved seeds or common local varieties, rather than switching to more expensive, potentially improved, varieties as an immediate response to onset variation.

In contrast to these adjustments in planting timing and seed quantity, I find limited evidence that onset timing causes short-run changes in other major input applications. Table 1.4 (Column 5) shows no statistically significant change in the probability of using fertiliser, with the small coefficient indicating a negligible and insignificant effect on the likelihood of fertiliser application. Further analysis in Table B1.13 (Columns 2 and 3) indicates no significant causal effect of onset week on the quantity of inorganic fertiliser applied (kg/ha), nor on the value of inorganic fertiliser (USD/ha), where the point estimate for value suggests a decrease but is not statistically distinguishable from zero. Similarly, the probability of using pesticides (Table 1.4, Column 6) does not respond significantly to onset timing. I also find no significant adjustments in on-farm labour intensity. Table B1.13 (Columns 4 and 5) shows that onset week does not significantly change either total family labour days per hectare or hired labour days per hectare; the point estimate for hired labour days indicates a decrease, but this effect is not statistically significant.

When examining crop choice, I find no significant shifts in the main crop grown in response to the onset week of a particular year. Table 1.4 (Columns 7-9) shows no statistically significant change in the probability that the main crop on the plot is a cereal, a tuber, or a legume. However, a more detailed analysis of which crops are present on a plot reveals a specific adjustment margin. While a later onset does not affect the probability of a plot containing cereals or legumes, it does lead to a small but statistically significant decrease in the probability of it containing tubers (Table B1.14). This aligns with the finding on the composition of plot value, where a later onset also causes a minor decrease in the share of value attributed to tubers. Overall, these findings suggest that while major decisions regarding crop portfolios are largely inflexible in the face of year-to-year onset variability, farmers may make minor adjustments away from certain crops like tubers when the season starts late.

Furthermore, I explored whether onset timing influences broader land use decisions at the household level. As detailed in Table B1.15, I find no statistically significant impact of onset week on extensive margin decisions such as the number of fallow plots or the number of plots cultivated. Similarly, there is no significant effect on the overall scale of land operation, as measured by the logarithm of total cultivated area or the logarithm of total farm size. This suggests that households do not significantly alter these decisions in immediate response to year-to-year variations in onset timing.

Finally, I investigate whether households reallocate labour to off-farm sectors as an adaptive response. Contrary to a simple story of labour substitution, the evidence suggests that a delayed onset reduces local wage-employment opportunities. Using individual-level data from the LSMS-ISA, I find that a one-week delay in the onset of the rainy season causes a small but statistically significant decrease in the probability of individuals engaging in any wage work. Specifically, Panel A of Table B1.16 shows a 0.1 percentage point decrease for wage work over the past 12 months, representing a 2% reduction relative to the sample mean, and a 0.2 percentage point decrease for work in the past 7 days, a 3.7% reduction. Panel B provides evidence that this finding holds for the general working-age population and is not driven by other members of the household, as it remains significant when the sample is restricted to only household heads. Furthermore, there is no significant effect on working specifically in agriculture or services, nor on engagement in household business activities.

Taken together, these results suggest that rather than households readily shifting their labour to off-farm work, opportunities for local wage employment diminish when the agricultural season starts late, consistent with a negative shock to the local economy. An alternative explanation involving an increase of on-farm labour demand seems unlikely, as direct measures of on-farm labour days (both family and hired) do not show a significant increase; if anything, the point estimates for hired labour days, while not statistically significant, suggest a decrease (Table B1.13).

1.5.4 Heterogeneity of impacts: Vulnerability, resilience, and adaptation

While the average effect of a delayed onset on productivity is negative, the impact is not uniform. To identify the characteristics associated with vulnerability and resilience, I explore how the effect varies across key farmer, household, and technological dimensions. The results, presented in Table 1.5, reveal a significant and economically meaningful degree of heterogeneity.

Socioeconomic and demographic factors play a role in shaping a household's resilience. Column (1) shows that the negative impact of a delayed onset is exacerbated on plots managed by women. The interaction term is negative and significant (-0.006), indicating that the yield loss for female-managed plots is approximately a third larger than for male-managed plots in the face of the same one-week delay. Conversely, human capital appears to be a key mitigating factor. As shown in column (2), the interaction with a manager's formal education is positive and significant (0.006), implying that for educated managers, the yield penalty is reduced by over a quarter. Economic resources provide a similar buffering effect. Column (3) shows that households with above-median assets suffer a significantly smaller yield loss, suggesting that wealthier households are better able to deploy capital to smooth the shock. Finally, the analysis highlights the critical role of technology. For irrigated plots, the positive and significant interaction term (0.019) in Column (4) almost perfectly cancels out the baseline negative effect.

The use of improved seeds, shown in Column (5), also provides a statistically significant, albeit smaller, mitigating effect.

I explore additional dimensions of heterogeneity in Appendix Table B1.24. Notably, I find that market integration—measured by whether the household sold any produce (Column 5)—is associated with significantly greater resilience, consistent with the hypothesis that market access facilitates consumption smoothing or is correlated with higher adaptive capacity. Furthermore, while both female land management and poverty (lower assets) predict greater vulnerability, it is worth noting that coming from a poor household and having at least one female-managed plot are not strongly correlated in this sample; thus, while both factors drive effects in the expected direction, they likely capture distinct sources of disadvantage.

A crucial question that arises from these findings is whether these more resilient groups achieve better outcomes because they adapt their practices more effectively. I test this by examining if they adjust their planting schedules more responsively to a delayed onset. The results, shown in Table B1.25, suggest this is generally not the case. There is no statistically significant evidence that farmers with more education or more assets adjust their planting dates any differently than their less-resourced counterparts. The findings for technology are more nuanced. For irrigated plots, there is weak evidence that farmers delay planting *less*, perhaps because access to water reduces the urgency of optimal timing. In contrast, farmers using improved seeds are the only group to show a significantly more responsive adjustment, delaying their planting by an additional 0.34 days for every week of onset delay.

Taken together, these findings paint a clear picture. While all rainfed agriculture is exposed to shifting seasonal timing, economic damages are disproportionately borne by marginalised producers. Resilience is systematically linked to education, wealth, and access to adaptive technologies. However, the source of this resilience is not fully explained by a comprehensively superior management response to the shock. For what is arguably the most critical margin—planting time—the analysis reveals that most resilient farmers do not adapt their schedules more effectively. The response on other input use is more nuanced. While wealthier and more educated farmers show a statistically significant differential increase in the use of inputs like pesticides, the magnitude of this adjustment is economically small (see Table B1.26). The most meaningful adaptive response observed is in seed use, where more educated farmers significantly increase the quantity planted after a late onset. This selective and partial pattern of adaptation—absent on the crucial timing margin, marginal for some inputs, and robust only for seeds—suggests that the resilience of better-off farmers stems primarily from their greater capacity to buffer shocks, rather than from consistently superior real-time management. This implies that significant informational barriers to optimal adaptation remain, leaving room for improvement across the board. Furthermore, an examination of how these heterogeneous impacts have changed over the sample period reveals a concerning trend: the vulnerability gap for female-managed plots has significantly widened in more recent years, while the protec-

tive effect of formal education has strengthened (Table B1.12). This suggests that underlying inequalities may be amplifying the impacts of shifting seasonal timing over time.

1.5.5 Limits to adaptation: False onsets and farmer awareness

While strategically delaying planting can significantly mitigate productivity losses—a finding supported by an instrumental variable analysis presented in Appendix B1.7—the fact that not all farmers perfectly adjust their timing suggests that adaptation is not free of risk. In the sample, farmers appear to plant too early, and therefore miss the agronomically correct window, on approximately one-third of all plots. Specifically, 34% of plots are planted before the modelled onset date—recall that this is the earliest point at which soil moisture is sufficient to reliably sustain a crop. Planting before this window opens is a high-risk strategy, as it exposes seeds to early-season dry spells that can lead to crop failure. While one might hypothesize that this early planting is driven by off-farm labour commitments, this is unlikely to be the primary explanation, as only 3% of working-age individuals in the sample report having any wage employment. A more plausible explanation is that farmers’ decisions are constrained by misleading weather signals and imperfect awareness of onset delays. In this section, I provide suggestive evidence for both of these channels.

A primary source of confusing signals for farmers is a *false onset* of the rains. I develop a precise, data-driven definition to capture this specific weather pattern. The process first defines a relevant *search window* for each location, symmetrically set at 60 days around its long-term average onset date. This local baseline is calculated as the mean onset for each $0.5^\circ \times 0.5^\circ$ grid cell using annual data since 2000. This window helps to avoid identifying spurious weather events far outside the planting season. Within this plausible window, a false onset is identified by a two-part sequence. First, there must be a *wet spell* that could plausibly trigger planting: two consecutive days with at least 20mm of cumulative rainfall (Marteau et al., 2011). Second, to be considered *false*, this must be followed by a *damaging dry spell*: seven consecutive days with less than 0.1mm of rain each, which must begin within 20 days of the wet spell. For a location-year to be classified as a false onset, this entire wet-then-dry sequence must conclude 21 days before the date of the true agronomical onset, to avoid any overlap. This is a relatively rare event, affecting 4.5% of plot-years in the sample, and on average, it occurs 26 days before the true onset. To test its impact, I extend my benchmark regression to include a dummy variable for whether a false onset occurred and its interaction with the onset week, using a fully flexible specification where all fixed effects are also interacted with the false onset indicator.

A false onset should only harm productivity if it tricks farmers into planting at the wrong time. The results in Table 1.6 show this is precisely what happens. Column (1) reveals that when a false onset occurs, farmers react by planting significantly earlier than they otherwise would. The interaction term shows that for each week the true onset is delayed, the presence of a

false signal causes planting to occur approximately one day earlier. The consequences of this mistake are evident in Column (2). For each week of onset delay, a false onset is associated with an additional 10% increase in seed quantity, a substantial rise that strongly suggests a narrative of initial crop failure followed by costly replanting. Ultimately, this disruption has a dramatic impact on productivity. As shown in Column (3), the negative effect of a one-week delay in the rainy season is amplified more than three-fold when it is preceded by a false onset, with the total yield loss increasing from 2.3% to 8.9%. While the evidence on planting dates and seed quantities points to timing errors as the primary damage channel, other agronomic or economic pathways could also be at play. I explicitly test for these alternatives in Table B1.23. The results show no significant change in pesticide use, casting doubt on the hypothesis that false onsets primarily harm yields by increasing pest prevalence. Similarly, I find no effect on the use of hired labour, ruling out a labour market response as an alternative damage channel. The only significant effect is a sharp decrease in fertilizer use. This finding reinforces the main hypothesis: after a false onset induces a planting error and subsequent crop failure, farmers reduce their investment on the compromised plot. This core finding is robust to the specific parameters of the false onset definition (Table B1.22). Results are qualitatively unchanged when using alternative definitions. For instance, I test for a wider search window (70 days) and require a longer, more damaging dry spell (10 days) to ensure the signal is truly relevant. In these alternative specifications, a false onset continues to significantly worsen the impact of a delayed season on yields, with the total negative effect remaining large and statistically significant.

Beyond being misled by false onsets, a second informational barrier appears to be the low salience of the underlying shock itself. To explore whether farmers perceive and report the negative shock from a delayed rainy season onset—a potential barrier to adaptation—I use self-reported, plot-level data from the LSMS-ISA surveys asking if a plot experienced a negative shock. I examine two binary outcomes: suffering (1) any crop shock, and (2) a drought shock specifically. I estimate equation (1.1), also replacing the continuous onset variable with indicators for the top and bottom terciles of the historical distribution to test for non-linearities in perception. The results, presented in Table 1.7, indicate that farmers’ awareness is indeed non-linear. A typical delay in the rainy season does not significantly increase the probability that a farmer reports a negative shock, suggesting the effect may be difficult to distinguish from normal year-to-year variation. Awareness is triggered only by extreme events; farmers are significantly more likely to report a shock only when the onset is exceptionally late (in the top tercile). Even then, they tend to misattribute the shock primarily to drought rather than the shift in timing. These patterns of non-linear awareness and misattribution hold when controlling for grid-cell-specific linear trends. This robustness check, detailed in Appendix B1.2 (Table B1.2), addresses potential biases from unobserved local trends highlighted by Jones et al. (2025).

Given these apparent informational constraints, access to timely information should mitigate

losses by helping farmers better anticipate the true onset or react more swiftly (Burlig et al., 2024). To explore this, I examine whether household ownership of mobile phones, radios, and televisions—assets commonly associated with information access—moderates the impact of a delayed onset on productivity. Following the same approach as in Section 1.5.4, I find that owning any of these assets is associated with significantly smaller yield losses following a delayed onset. These plots experience yield reductions that are approximately 20% smaller than those in households without such assets (Table B1.27). To mitigate concerns about omitted variables, these specifications control for a household’s level of wealth (proxied by an indicator for being below the median asset index), which itself is not strongly correlated with ownership of any single information asset. While this evidence is correlational, it provides suggestive support that informational barriers are indeed a key constraint, and that improving farmers’ access to information could enhance resilience.

1.5.6 Projecting future damages

A key question is whether the estimated short-run effects of weather shocks can inform our understanding of the long-term consequences of climate change. A common approach to testing for long-run adaptation is to compare locations with different initial climatic conditions; for instance, locations with higher average temperatures might be better adapted to a heatwave (Dell et al., 2014). For the timing of the rainy season, however, the absolute day of the year is less important than its deviation from local expectations. I therefore adapt this approach to compare locations based on their long-term trends in onset timing. If farmers adapt to persistent environmental change, one might expect those in locations where the onset is secularly arriving later to be more resilient to a late-onset shock than farmers in areas with a more stable climate.

To test this, I classify locations as either *trending* or *non-trending* based on the long-run linear trend in onset timing over the last four decades, estimated at the $0.5^\circ \times 0.5^\circ$ grid-cell level. A location is defined as trending if it exhibits a statistically significant positive trend. I then expand equation (1.1) by interacting the onset week of the rainy season with an indicator for being in a trending location.

The results, presented in Table 1.8, provide strong evidence of a failure to adapt. The negative impact of a delayed onset on agricultural yields is driven entirely by locations already experiencing a long-term climate trend toward later rainy seasons. In non-trending locations, a later-than-usual onset has no statistically significant effect on productivity. In stark contrast, in trending locations, a one-week delay causes a significant yield loss. These vulnerable, trending locations constitute the vast majority of the sample, representing approximately 90% of plot-wave observations.

This lack of long-run adaptation is further corroborated by two pieces of evidence: a failure to learn even from recent, year-to-year experience, and an increase in the effect of a delayed onset

on productivity in recent years. First, if farmers were updating their strategies based on past weather, a delayed onset in the previous year ($t - 1$) should lead them to adjust their behavior in the current year (t). I test for this inter-annual learning across all major adaptation margins and find no evidence of a systematic response. As shown in Table B1.17, the previous year's onset has no statistically significant effect on planting dates, input use, or harvest timing. Second, I test whether the average impact of onset delays has diminished within the sample period itself, which could indicate adaptation over the medium term. Interacting the onset week with an indicator for later survey waves in countries with sufficient panel length (Ethiopia, Malawi, Nigeria) reveals that the negative impact has significantly increased in more recent years (Table B1.11).

Taken together, these three pieces of evidence suggest a persistent failure to adapt, which provides the justification for using short-run estimates to project future damages. To quantify this threat, I project the Net Present Value (NPV) of damages from 2025 to 2050, presented as a percentage of 2024 real GDP. The calculation combines my benchmark estimate of a 2% yield loss per week with country-specific climate shocks. Given the evidence that the negative impact of onset delays has intensified in recent years (Table B1.11), employing the average effect across the entire sample period likely yields a conservative estimate of future damages. These shocks are scenario-dependent: under the Business as Usual scenario, I use the full projected onset delay from Figure B1.1, while the Toward Sustainability scenario assumes a 50% less severe delay, reflecting climate mitigation efforts. The annual loss is calculated in absolute monetary terms by combining these shocks with FAO scenarios for future real agricultural shares and real GDP growth. The model also accounts for a data-driven expansion of irrigated land, which is forecast based on historical trends (see Figure C1.1); however, these projections suggest that only Ethiopia, Mali, and Tanzania will see meaningful increases in irrigation by 2050. Finally, I use a time-varying discount rate from the DICE-2023 model. The full methodology is detailed in Appendix C.

Figure 1.4 presents the main finding. Cumulative damages will be high under either climate path, but a more sustainable trajectory can offset a sizeable portion of the economic losses. Under the Business as Usual scenario (solid red line), cumulative discounted damages by 2050 are substantial, reaching approximately 10% of 2024 GDP in Ethiopia and Mali. The shaded area quantifies the Sustainability Benefit of a less severe climate shock combined with a more resilient economy, averting over \$4.3 billion in damages for Ethiopia and \$20.2 billion for Nigeria.

Of course, this projection is pessimistic as it does not allow for long-run adaptation beyond the expansion of irrigation. Figure 1.5 addresses this by comparing the baseline Business as Usual damages with an optimistic counterfactual where farmers adapt successfully, causing the impact of onset delays to decrease linearly to zero by 2050. While this Adaptation Benefit is significant, the remaining costs are still substantial.

In my model, I already include irrigation as a key adaptation technology. However, I found its projected effect is limited. To test this further, I simulated an aggressive (and costly) scenario that expanded irrigated farmland to 20% of the total by 2050. Comparing the results to the baseline reveals that the benefits of this expansion are modest for most countries (see Figure C1.2).

Finally, it is crucial to acknowledge the large uncertainty surrounding these estimates. The projections for GDP, agricultural shares, climate shocks, and adaptation capacity are all subject to uncertainty. While it is not possible to model this fully, Figure C1.3 and Table C1.1 present the Business as Usual projections using the 95% confidence interval of the benchmark yield loss coefficient. They reveal that even under the most optimistic lower-bound estimate, the cumulative economic damages remain large and economically significant for all countries.

1.6 Conclusion

This paper investigates the consequences of shifting seasonal timing for agricultural communities in Sub-Saharan Africa. Using decades of high-resolution climate data, I document a clear trend: the onset of the rainy season is occurring progressively later across six major African countries, leading to shorter growing seasons. The analysis shows that these delays have significant negative consequences. A one-week delay in the onset causally reduces agricultural yields by approximately 2%. When the onset is exceptionally late (in the latest tercile of the historical distribution), these losses deepen to roughly 12%. To put this magnitude in perspective, this impact is remarkably comparable to the damages caused by other major climate shocks, such as the 12% yield reduction from severe drought found by Corno et al. (2020) and the 11% decline from heat stress estimated by Aragón et al. (2021). This highlights that shifts in seasonal timing are a first-order threat comparable to traditional weather shocks.

These damages are not evenly distributed; they are borne disproportionately by the most vulnerable, particularly on plots managed by women, while education, wealth, and technology build resilience. In response, farmers adapt primarily along low-cost margins—delaying planting and increasing seed quantities—but these adjustments are insufficient to prevent losses to productivity and welfare.


A central puzzle emerging from the analysis is why farmers do not fully leverage planting date adjustments, a strategy this paper shows is both powerful and inexpensive. The evidence points strongly toward informational constraints as the primary barrier. This conclusion is reinforced by the finding that even farmers who are more resilient to the initial shock do not systematically adapt their planting timing more effectively. The problem is not merely a lack of resources but a widespread informational challenge, best illustrated by the *false onset*. This paper demonstrates that these misleading early rains trick farmers into planting prematurely,

a mistake that ultimately more than doubles the negative impact of a late onset of the rainy season.

The scale of these projected damages lends urgency to several clear policy implications. First, the demonstrated vulnerability of farmers to false onsets highlights a high potential return for investments in accessible and reliable short-range weather forecasts. Such services could create enormous value not by replacing farmer knowledge, but by helping to differentiate a true onset from a false one. Second, the heterogeneity of the impacts underscores that interventions must be targeted. Policy should not only deliver weather information but also work to reduce the underlying vulnerabilities that amplify damages, for instance, by improving women's access to resources and promoting education. Third, the limited adaptation along the crop choice margin suggests a need for continued investment in agricultural R&D to develop and disseminate seed varieties with more flexible growing periods. Finally, the evidence of negative spillovers to the entire local economy points to a need for social protection systems that are responsive enough to respond to shocks in seasonal timing, not just catastrophic events like major droughts.


The findings in this paper point to several promising avenues for future research. Building on this analysis, future work using true panel data could offer deeper insights into the dynamic, household-level adaptation strategies that unfold over multiple seasons. Further research could also explore a broader set of adaptation margins not explicitly covered here, including migration, the role of social networks, and investments in irrigation. Finally, linking these climate shocks to local market data would be a valuable extension, illuminating how market access and price signals mediate both the damages and the capacity for adaptation, providing a richer understanding of this critical climate vulnerability.

Figure 1.1: Channels of impact from shifting seasonal timing




Note. Two primary channels through which the timing of the onset of the rainy season affects economic outcomes. The *Biophysical channel* represents the direct impact of weather conditions (e.g., temperature, moisture) on crop growth and productivity, independent of farmer actions. The *Adaptation channel* represents how farmers’ responses to weather signals affect productivity. The empirical strategy first estimates the net impact of onset timing (controlling for the biophysical channel) and then uses subsequent analysis to identify the adaptation channel as the key mechanism, as detailed in Section 1.4.

Figure 1.2: Country-specific trends in rainy season onset




Note. Each panel in the figure displays the annual mean onset day of the year (DOY) for $0.1^\circ \times 0.1^\circ$ grid-cells matched to LSMS-ISA survey locations within the respective country from 1979–2024. The *Mean onset (DOY)* is the average across space for each year of the day of the year when the rainy season started conditions specified in Section 1.3 are met. The dashed red line represents the fitted linear trend, estimated using country-specific fixed-effects models. The total predicted change in the mean onset day of the year from 1979–2024, as implied by the model, is: Ethiopia (19.8 days), Malawi (17.8 days), Mali (11.9 days), Niger (7.6 days), Nigeria (25.1 days), and Tanzania (5.8 days).

Figure 1.3: Onset of the rainy season – asymmetric effects




Note. The figure displays OLS regression coefficients for indicator variables representing the bottom tercile and top tercile of the historical rainy season onset distribution within $0.5^\circ \times 0.5^\circ$ grid cells. The middle tercile is the omitted reference category. The dependent variable is log agricultural yields (2020 USD per hectare). The regression includes the benchmark set of fixed effects ($0.1^\circ \times 0.1^\circ$ grid cell and climatic area by country by year), weather controls, and plot-level controls as specified in Equation 1.1. Bars represent 95% confidence intervals based on standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level.

Figure 1.4: Projected damages by 2050



Note. Cumulative discounted damages from 2025 to 2050, expressed as a percentage of 2024 real GDP. Damages are calculated using the benchmark estimate of a 2% yield loss per week of onset delay. The solid line shows damages under the FAO's Business as Usual economic scenario and the full projected onset delay. The dashed line shows damages under the Toward Sustainability scenario, which assumes a 50% less severe onset delay. The shaded area represents the economic benefit of the sustainable path. All projections account for the mitigating effect of future irrigation expansion and use a time-varying discount rate from the DICE-2023 model. See Appendix C for full details.

Figure 1.5: Accounting for future adaptation



Note. Cumulative discounted damages from 2025 to 2050 under the Business as Usual scenario with an optimistic counterfactual, expressed as a percentage of 2024 real GDP. The solid line (No adaptation) shows the baseline projection using a constant 2% yield loss per week of onset delay. The dashed line (Decreasing effect) shows a scenario where this yield impact linearly declines to zero between 2025 and 2050. The shaded area represents the total economic benefit of this optimistic, catch-all adaptation scenario, quantified in millions of US dollars.

Table 1.1: Onset of the rainy season and agricultural productivity

Dependent variable:	Log yields (2020 USD per hectare)		
	(1)	(2)	(3)
Onset of the rainy season	-0.022 (0.007) [0.002]	-0.021 (0.008) [0.005]	-0.021 (0.007) [0.002]
Standardised effect	-0.036	-0.033	-0.033
Mean (dep. var.)	5.744	5.744	5.744
Identifying observations	130,556	130,556	129,919
Singleton observations	182	182	180
Countries	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022
Weather controls	No	Yes	Yes
Demographic controls	No	No	Yes

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. A full list of controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 1.2: Onset of the rainy season and per capita consumption

Dependent variable:	Log per capita consumption					
	Agricultural households			All households		
	(1)	(2)	(3)	(4)	(5)	(6)
Onset of the rainy season	-0.012 (0.004) [0.006]	-0.010 (0.005) [0.041]	-0.010 (0.005) [0.047]	-0.008 (0.002) [0.001]	-0.007 (0.003) [0.006]	-0.007 (0.003) [0.010]
Standardised effect	-0.020	-0.016	-0.015	-0.016	-0.014	-0.013
Mean (dep. var.)	5.970	5.970	5.970	6.227	6.227	6.225
Identifying observations	29,335	29,335	29,269	70,748	70,748	70,298
Singleton observations	338	338	337	457	457	457
Countries	6	6	6	6	6	6
Survey year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022
Weather controls	No	Yes	Yes	No	Yes	Yes
Demographic controls	No	No	Yes	No	No	Yes

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of per capita consumption, measured in 2020 US dollars. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. In Columns (1)–(3), the sample is restricted to households engaged in the agricultural sector and reporting a complete harvest. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 1.3: Onset of the rainy season and women’s nutrition

Dependent variable:	Weight-for-height	Wasted	Underweight
	(1)	(2)	(3)
Onset of the rainy season	-0.007 (0.004) [0.040]	0.001 (0.001) [0.261]	0.001 (0.001) [0.378]
Standardised effect	-0.016	0.001	0.001
Mean (dep. var.)	-0.850	0.110	0.142
Identifying observations	59,014	59,014	51,733
Singleton observations	16	16	15
Countries	6	6	6
Interview year range	1995–2018	1995–2018	1992–2022

Note. OLS regression estimates of Equation 1.1. The data are from the DHS (Croft et al., 2018) for women in agricultural households. The dependent variables are defined as follows. *Weight-for-height* is a z-score, representing the standard deviation from the median of the DHS reference population. *Wasted* is an indicator for a weight-for-height z-score below -2. *Underweight* is an indicator for a BMI below 18.5. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include weather controls, individual-level controls, climatic area by country by year fixed effects, and location fixed effects at the $0.5^\circ \times 0.5^\circ$ grid-cell level. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

Table 1.4: Onset of the rainy season and crop management

Dependent variable:	Timing			Inputs			Main crop is		
	Planting (DOY)	Planting duration	Harvest (DOY)	Seeds	Fertiliser	Pesticide	Cereals	Tubers	Legumes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Onset of the rainy season	0.557 (0.222) [0.012]	-0.008 (0.003) [0.016]	0.620 (0.942) [0.510]	0.032 (0.011) [0.006]	-0.001 (0.002) [0.664]	-0.000 (0.002) [0.771]	0.000 (0.002) [0.817]	-0.001 (0.001) [0.410]	0.000 (0.001) [0.910]
Standardised effect	0.842	-0.012	0.936	0.046	-0.001	-0.001	0.001	-0.001	0.000
Mean (dep. var.)	188.934	1.105	250.370	3.453	0.516	0.066	0.572	0.069	0.122
Identifying observations	125,609	125,609	106,970	120,669	135,081	135,306	136,855	136,855	136,855
Singleton observations	74	74	117	197	173	173	175	175	175
Countries	5	5	5	6	6	6	6	6	6
Planting year range	2009–2022	2009–2022	2009–2022	2009–2022	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variables are defined as follows. *Planting (DOY)* is the approximated first planting day of the year; this approximation assumes planting occurred on the 15th of the reported month. *Planting duration* is the number of distinct months planting occurred on the plot. *Harvest (DOY)* is the approximated first harvest day of the year, assuming harvest occurred on the 15th of the reported month. *Seed* is the log of seed quantity used on the plot, measured in kilograms per hectare. *Fertiliser* is an indicator variable taking the value 1 if any fertiliser, whether organic or inorganic, was used, and 0 otherwise. *Pesticide* is an indicator variable that equals 1 if any pesticide was used and 0 otherwise. *Cereals*, *Tubers*, and *Legumes* are indicator variables taking the value 1 if the main crop by production value on the plot belongs to the respective category, and 0 otherwise. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors, which are clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides further details on variable construction, selected surveys, and weighting procedures.

Table 1.5: Heterogeneous effects on agricultural productivity

Dependent variable:	Log yields (2020 USD per hectare)				
	(1)	(2)	(3)	(4)	(5)
Onset of the rainy season	-0.020 (0.007) [0.004]	-0.024 (0.007) [0.001]	-0.015 (0.007) [0.022]	-0.021 (0.007) [0.002]	-0.017 (0.010) [0.098]
× Manager is female	-0.006 (0.001) [0.000]				
× Manager has formal education		0.006 (0.002) [0.001]			
× Below median assets			-0.006 (0.002) [0.001]		
× Irrigated plot				0.019 (0.007) [0.009]	
× Improved seeds					0.005 (0.002) [0.033]
Mean (dep. var.)	5.743	5.743	5.769	5.748	5.705
Identifying observations	129,805	129,805	118,366	128,373	105,557
Singleton observations	180	180	179	177	215
Countries	6	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2011–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 1.6: The role of false onsets

Dependent variable:	Planting (DOY)	Seeds	Log yields (2020 USD per hectare)
	(1)	(2)	(3)
Onset of the rainy season	0.446 (0.226) [0.049]	0.033 (0.012) [0.009]	-0.024 (0.007) [0.001]
× False onset	-0.982 (0.422) [0.020]	0.099 (0.029) [0.001]	-0.062 (0.031) [0.043]
Mean (dep. var.)	188.916	3.453	5.744
Identifying observations	125,588	120,650	129,899
Singleton observations	95	216	200
Countries	5	6	6
Planting year range	2009–2022	2009–2022	2008–2022

Note. OLS regression estimates based on Equation (1.1). *Planting (DOY)* is the approximated first planting day of the year; this approximation assumes planting occurred on the 15th of the reported month. *Seeds* is the log of seed quantity used on the plot, measured in kilograms per hectare. *Log yields (2020 USD per hectare)* is the log of yields, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. The analysis examines the interaction between the rainy season onset week and a dummy variable indicating a *False onset*. Benchmark fixed effects are interacted with the false onset dummy. All specifications include a full set of weather and plot-level controls. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

Table 1.7: Farmers' awareness

Dependent variable:	Crops shock		Drought shock	
	(1)	(2)	(3)	(4)
Onset of the rainy season	0.003 (0.002) [0.301]		0.001 (0.003) [0.821]	
Earliest onset tercile		-0.005 (0.017) [0.769]		0.006 (0.016) [0.709]
Latest onset tercile		0.031 (0.012) [0.007]		0.034 (0.013) [0.008]
Standardised effect	0.004		0.001	
Mean (dep. var.)	0.411	0.411	0.202	0.202
Identifying observations	136,420	136,420	132,044	132,044
Singleton observations	173	173	180	180
Countries	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable in each column is a binary indicator for whether a farmer reported a negative shock on a given plot. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. *Earliest onset tercile* and *Latest onset tercile* are indicators for whether the onset of the rainy season for a location falls into the earliest or latest tercile of its historical distribution, with the middle tercile serving as the omitted category. These terciles are defined separately for each $0.5^\circ \times 0.5^\circ$ grid cell based on its long-term (1979–2020) onset history. For comparison, Panel A presents results using the continuous onset week variable. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, and a full set of weather and plot-level controls. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides detailed information on all variables.

Table 1.8: Heterogeneity of onset impacts by long-run climate trend

Dependent variable:	Log yields (2020 USD per hectare)		
	(1)	(2)	(3)
Onset of the rainy season	0.005 (0.009) [0.595]	0.004 (0.009) [0.656]	-0.000 (0.008) [0.981]
× Positive and significant trend	-0.033 (0.011) [0.003]	-0.032 (0.012) [0.009]	-0.026 (0.011) [0.017]
Mean (dep. var.)	5.744	5.744	5.744
Identifying observations	130,556	130,556	129,919
Singleton observations	182	182	180
Countries	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022
Weather controls	No	Yes	Yes
Demographic controls	No	No	Yes

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. A location is designated as *trending* if its $0.5^\circ \times 0.5^\circ$ grid cell has a statistically significant positive long-run trend in onset. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. A full list of controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Bibliography

- Amale, H. S., P. S. BIRTHAL, and D. S. Negi (2023). Delayed Monsoon, Irrigation and Crop Yields. *Agricultural Economics* 54(1), 77–94.
- Antwi-Agyei, P., A. J. Dougill, R. Doku, and R. C. Abaidoo (2022). The Role of Soil Moisture Information in Developing Robust Climate Services for Smallholder Farmers: Evidence from Ghana. *Agronomy* 12(2), 541.
- Aragón, F. M., F. Oteiza, and J. P. Rud (2021, February). Climate Change and Agriculture: Subsistence Farmers' Response to Extreme Heat. *American Economic Journal: Economic Policy* 13(1), 1–35.
- Barreca, A., K. Clay, O. Deschênes, M. Greenstone, and J. S. Shapiro (2016). Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century. *Journal of Political Economy* 124(1), 105–159.
- Boogaard, H., J. Schubert, A. De Wit, J. Lazebnik, R. Hutjes, and G. Van der Grijn (2020). Agrometeorological Indicators from 1979 to Present Derived from Reanalysis. Accessed on 01-03-2025.
- Burlig, F., A. Jina, E. M. Kelley, G. V. Lane, and H. Sahai (2024, 2). Long-Range Forecasts as Climate Adaptation: Experimental Evidence from Developing-Country Agriculture. Working Paper 32173, National Bureau of Economic Research.
- Camacho, A. and E. Conover (2019). The Impact of Receiving SMS Price and Weather Information on Small Scale Farmers in Colombia. *World Development* 123, 104596.
- Carleton, T., E. Duflo, B. K. Jack, and G. Zappalà (2024). Chapter 4 - Adaptation to climate change. In L. Barrage and S. Hsiang (Eds.), *Handbook of the Economics of Climate Change*, Volume 1, pp. 143–248. North-Holland.
- Corno, L., N. Hildebrandt, and A. Voena (2020). Age of Marriage, Weather Shocks, and the Direction of Marriage Payments. *Econometrica* 88(3), 879–915.
- Croft, T. N., A. M. J. Marshall, and C. K. Allen (2018). Guide to DHS Statistics. Demographic and Health Surveys Program.
- Danielson, J. J. and D. B. Gesch (2011). Global Multi-Resolution Terrain Elevation Data 2010 (GMTED2010). Open-File Report 2011-1073, U.S. Geological Survey.
- Dell, M., B. F. Jones, and B. A. Olken (2014). What Do We Learn from the Weather? the New Climate-Economy Literature. *Journal of Economic Literature* 52(3), 740–98.

- Dercon, S. and P. Krishnan (2000). In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia. *Journal of Political Economy* 108(4), 688–727.
- Deschênes, O., M. Greenstone, and J. Guryan (2009). Climate Change and Birth Weight. *American Economic Review* 99(2), 211–17.
- Deschênes, O. and E. Moretti (2009). Extreme Weather Events, Mortality, and Migration. *The Review of Economics and Statistics* 91(4), 659–681.
- Dunning, C. M., E. C. L. Black, and R. P. Allan (2016). The Onset and Cessation of Seasonal Rainfall over Africa. *Journal of Geophysical Research: Atmospheres* 121(19), 11405–11424.
- FAO (2018). The future of food and agriculture – Alternative pathways to 2050. Technical report, FAO, Rome.
- Giné, X., R. M. Townsend, and J. I. Vickery (2018). Forecasting When It Matters: Evidence from Semi-Arid India. *Journal of Development Economics* 131, 150–167.
- Hultgren, A., T. Carleton, M. Delgado, D. R. Gergel, M. Greenstone, T. Houser, S. Hsiang, A. Jina, R. E. Kopp, S. B. Malevich, K. E. McCusker, T. Mayer, I. Nath, J. Rising, A. Rode, and J. Yuan (2025, jun). Impacts of Climate Change on Global Agriculture Accounting for Adaptation. *Nature* 642(8068), 644–652.
- IPCC (2021). Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press. In press.
- IPCC (2022). Africa. In H.-O. Pörtner, D. C. Roberts, M. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, and B. Rama (Eds.), *Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pp. 1285–1455. Cambridge, UK and New York, NY, USA: Cambridge University Press.
- Jägermeyr, J. and K. Frieler (2018). Spatial Variations in Crop Growing Seasons Pivotal to Reproduce Global Fluctuations in Maize and Wheat Yields. *Science Advances* 4(11), eaat4517.
- Jagnani, M., C. B. Barrett, Y. Liu, and L. You (2020, 05). Within-Season Producer Response to Warmer Temperatures: Defensive Investments by Kenyan Farmers. *The Economic Journal* 131(633), 392–419.

- Jones, B. F., J. Moscona, B. A. Olken, and C. von Dessauer (2025, September). With or Without U? Binning Bias and the Causal Effects of Temperature Shocks.
- Kala, N. (2019). Learning, adaptation, and climate uncertainty: evidence from Indian agriculture. Technical report, MIT Sloan School of Management.
- Klein Goldewijk, K., A. Beusen, J. Doelman, and E. Stehfest (2017). Anthropogenic land use estimates for the Holocene – HYDE 3.2. *Earth System Science Data* 9(2), 927–953.
- Laajaj, R. and K. Macours (2024, April). The Complexity of Multidimensional Learning in Agriculture.
- Leenaars, J. G. B., L. Claessens, G. B. M. Heuvelink, T. Hengl, M. Ruiperez González, L. G. J. van Bussel, N. Guilpart, H. Yang, and K. G. Cassman (2018). Mapping Rootable Depth and Root Zone Plant-Available Water Holding Capacity of the Soil of Sub-Saharan Africa. *Geoderma* 324, 18–36.
- Marteau, R., B. Sultan, V. Moron, A. Alhassane, C. Baron, and S. B. Traoré (2011). The Onset of the Rainy Season and Farmers’ Sowing Strategy for Pearl Millet Cultivation in Southwest Niger. *Agricultural and Forest Meteorology* 151(10), 1356–1369.
- Metzger, M. J., J. Rombouts, L. Reino, M. Campos, A. Lomba, C. A. Guerra, I. Sattler, S. Ceaşu, D. Singing, S. Herrando, L. Brotons, F. Sayol, J. Santana, A. Delgado, M. J. Santos, J. P. Honrado, S. Schindler, D. Kübler, L. Mo, D. Poursanidis, A. Didem, P. Crawford, L. Halada, N. Abrantes, T. Domingos, and H. M. Pereira (2023). Integrative Climate Classification for the Global Land Surface. *Scientific Data* 10(1), 708.
- Minoli, S., J. Jägermeyr, S. Asseng, A. Urfels, and C. Müller (2022). Global Crop Yields Can Be Lifted by Timely Adaptation of Growing Periods to Climate Change. *Nature Communications* 13(1), 7079.
- Patel, D. (2025, March). Environmental Beliefs and Adaptation to Climate Change. Working Paper.
- Rosenzweig, M. R. and C. R. Udry (2019, 5). Assessing the Benefits of Long-Run Weather Forecasting for the Rural Poor: Farmer Investments and Worker Migration in a Dynamic Equilibrium Model. Working Paper 25894, National Bureau of Economic Research.
- Rudder, J. and D. Viviano (2024). Learning from Weather Forecasts and Short-Run Adaptation: Evidence from an At-Scale Experiment. Working Paper.
- Suri, T., C. Udry, J. C. Aker, C. B. Barrett, L. F. Bergquist, M. Carter, L. Casaburi, R. D. Osei, D. Gollin, V. Hoffmann, T. Jayne, N. Karachiwalla, H. Kazianga, J. Magruder, H. Michelson, M. Startz, and E. Tjernstrom (2024, 3). Agricultural Technology in Africa. *VoxDevLit* 5(2).

- Takele, R. and M. Dell’Acqua (2023). *AquaBEHER: Estimation and Prediction of Wet Season Calendar and Soil Water Balance for Agriculture*. R package version 1.4.0.
- Tambet, H. and Y. Stopnitzky (2021). Climate Adaptation and Conservation Agriculture among Peruvian Farmers. *American Journal of Agricultural Economics* 103(3), 900–922.
- Taraz, V. (2017). Adaptation to climate change: historical evidence from the Indian monsoon. *Environment and Development Economics* 22(5), 517–545.
- The World Bank (2025). World Development Indicators. Accessed on June 20, 2025.
- Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household. *Journal of Political Economy* 104(5), 1010–1046.
- Wolf, J., K. Ouattara, and I. Supit (2015). Sowing Rules for Estimating Rainfed Yield Potential of Sorghum and Maize in Burkina Faso. *Agricultural and Forest Meteorology* 214-215, 208–218.
- Wollburg, P., T. Bentze, Y. Lu, C. Udry, and D. Gollin (2024). Crop yields fail to rise in smallholder farming systems in sub-Saharan Africa. *Proceedings of the National Academy of Sciences* 121(21).
- World Bank (2024). A Longitudinal Cross-Country Dataset on Agricultural Productivity and Welfare in Sub-Saharan Africa. Policy Research Working Paper, World Bank Group, Washington, D.C.
- Yang, Y., D. Tilman, Z. Jin, P. Smith, C. B. Barrett, Y.-G. Zhu, J. Burney, P. D’Odorico, P. Fanktke, J. Fargione, J. C. Finlay, M. C. Rulli, L. Sloat, K. J. van Groenigen, P. C. West, L. Ziska, A. M. Michalak, the Clim-Ag Team, D. B. Lobell, M. Clark, J. Colquhoun, T. Garg, K. A. Garrett, C. Geels, R. R. Hernandez, M. Herrero, W. D. Hutchison, M. Jain, J. M. Jungers, B. Liu, N. D. Mueller, A. Ortiz-Bobea, J. Schewe, J. Song, J. Verheyen, P. Vitousek, Y. Wada, L. Xia, X. Zhang, and M. Zhuang (2024). Climate Change Exacerbates the Environmental Impacts of Agriculture. *Science* 385(6713).
- Yegbemey, R. N., G. Mensch, and C. Vance (2023). Weather information and agricultural outcomes: Evidence from a pilot field experiment in Benin. *World Development* 167(C).
- Zita, L. E., F. Justino, C. Gurjão, J. Adamu, and M. Talacuece (2025). Spatio-Temporal Characteristics of Climate Extremes in Sub-Saharan Africa and Potential Impact of Oceanic Teleconnections. *Atmosphere* 16(1), 86.

APPENDIX

Supplementary material to *Shifting Seasons: Agricultural Adaptation and Resilience in Africa*

Iván Kim Taveras

Appendix A: Data and methodological procedures

A1.1 Variables and data sources

Variable	Description
<i>Agricultural productivity</i>	Two primary measures of agricultural yields are used: output in constant 2020 US dollars per hectare and output in kilograms per hectare. These are calculated at the plot level.
<i>Climatic areas</i>	Household GPS coordinates from the LSMS-ISA surveys are spatially matched to define unique climatic areas based on the Köppen-Geiger climate classification (Metzger et al., 2023). Due to the requirements of fixed effect estimation needing sufficient observations within each group, these detailed classifications are aggregated into three broader zones: Tropical, Temperate, and Arid.
<i>Crop types and main crop</i>	Standardised crop groupings are utilised as defined in the LSMS-ISA harmonised dataset (World Bank, 2024). There are ten categories created to ensure comparability across countries: barley, wheat, rice, sorghum, maize, millet, perennials (e.g., fruit and tree crops), legumes, root crops, and nuts. Crops that do not fall into any of these categories are classified as <i>others</i> . The <i>main crop</i> on a given plot is identified as the crop within these categories that has the highest production value, following the methodology of World Bank (2024).
<i>False onset</i>	An area is considered exposed to a false onset if, within a 30-day window around its usual onset date (see <i>Usual rainy season dates</i>), it experiences a 2-day wet spell with at least 20mm of rainfall, which is then followed by a 7-day dry spell (where a day counts as dry if it receives less than 0.1 mm of rain) within the subsequent 20 days.
<i>Inputs</i>	Plot-level information is collected on the use of various agricultural inputs. <i>Seeds</i> refers to the quantity of seeds used, measured in kilograms per hectare for a given plot. <i>Seed value</i> is the total value of seeds used per hectare on a plot, expressed in constant 2020 US dollars (see <i>Unit values</i>). <i>Improved seeds</i> is an indicator variable taking the value 1 if the seeds planted are not traditional varieties (i.e., are improved or high-yielding), and 0 otherwise. <i>Fertiliser</i> is an indicator variable that equals 1 if either organic or inorganic fertilisers were used on the plot, and 0 otherwise. Similarly, <i>Pesticides</i> is an indicator variable taking the value 1 if pesticides were applied to the plot, and 0 otherwise.
<i>Irrigated</i>	An indicator variable at the plot level that takes the value 1 if any form of irrigation was used on the plot during the agricultural season, and 0 otherwise.
<i>Labour outcomes</i>	Individual-level data from LSMS-ISA are used to construct indicator variables for individuals of working age regarding their employment activities. <i>Any wage work (Last 12 months)</i> indicates if an individual reports working for a wage in the past twelve months. <i>Agriculture (Last 12 months)</i> indicates if an individual reports wage work specifically in the agricultural sector in the past twelve months. <i>Services (Last 12 months)</i> indicates if an individual reports wage work in the service sector in the past twelve months. For a shorter recall period, <i>Any wage work (Last 7 days)</i> indicates if an individual reports working for a wage in the past seven days. <i>Household business (Last 7 days)</i> indicates if they report working for a household business in the past seven days. Each of these is an indicator variable taking the value 1 if the condition is met, and 0 otherwise.
<i>Nutrition</i>	The DHS records objective measurements performed by its data collection team, utilizing the CDC Standard Deviation-derived Growth Reference Curves for standardized distributions (Croft et al., 2018). The indicators used are as follows: <i>w/h (weight-for-height)</i> represents the <i>z</i> -score from the reference curve, with <i>wasted</i> being an indicator variable equal to 1 if the <i>w/h</i> <i>z</i> -score is less than -2, and 0 otherwise. Similarly, <i>h/a (height-for-age)</i> is the <i>z</i> -score from the reference curve, where <i>stunted</i> is an indicator variable equal to 1 if the <i>h/a</i> <i>z</i> -score is less than -2, and 0 otherwise. <i>BMI</i> is defined as the ratio of weight in kilograms to the square of height in meters, excluding pregnant women. Finally, <i>Underweight</i> is an indicator variable set to 1 if a woman's BMI is below 18.5.
<i>Per capita consumption</i>	Household consumption is aggregated per capita and expressed in constant 2020 US dollars. This aggregate generally excludes expenditure on rent and durables. It is important to note that Tanzania Wave 5, which includes durables, is excluded from the main analyses as detailed in Section 1.3. Furthermore, comparable consumption aggregates are not available for Malawi Waves 3 and 4, and Mali Wave 2.
<i>Plot area</i>	The area of each agricultural plot is measured in hectares. This information is primarily sourced from GPS device measurements as provided in the LSMS-ISA harmonised dataset (World Bank, 2024). In instances where GPS measures are unavailable, imputation methods based on self-reports and administrative data are employed, as detailed by World Bank (2024).
<i>Proxies for access to information</i>	Individual indicators are created for specific assets: <i>Owns mobile phone</i> (equal to 1 if the household owns any mobile telephone, including basic, feature, regular, or smart phones where distinguished), <i>Owns Television</i> (equal to 1 for television set ownership), and <i>Owns Radio</i> (equal to 1 for radio ownership, including radio cassette or radio with flash drive). Additionally, for Nigeria only, an indicator <i>Internet Access</i> is included, sourced from a different survey module. Note that specific item labels used in the construction vary slightly across countries and waves. The construction aims for maximum consistency across surveys.

(continued on next page)

Variable	Description
<i>Rainy season calendar</i>	Variables summarising the rainy season calendar are computed for each $0.1^\circ \times 0.1^\circ$ grid cell and year using the AquaBEHER R package (Takele and Dell’Acqua, 2023). This package employs an agronomic definition based on soil water balance. The package requires the user to specify an earliest possible onset date, a latest possible onset date, and a latest possible cessation date for each location. These are determined flexibly: the earliest onset is set 90 days before the location’s usual onset date, the latest onset is 90 days after the usual onset date, and the latest cessation is 90 days after the usual cessation date (see <i>Usual rainy season dates</i>). The <i>Onset</i> is then defined as the first day when the actual-to-potential evapotranspiration ratio is greater than 0.5 for seven consecutive days, followed by a 20-day period in which plant-available water remains above the wilting point over the root zone of the soil layer. The <i>Cessation</i> is the first day when the actual-to-potential evapotranspiration ratio falls below 0.5 for seven consecutive days, followed by 12 consecutive non-growing days during which plant-available water remains below the wilting point over the root zone. The <i>Length</i> of the rainy season is the total number of days from the onset to the cessation of the season.
<i>Unit values</i>	To enable cross-country analysis, key agricultural input and output variables are valued in constant 2020 US dollars, following the methodology of the LSMS-ISA harmonised dataset (World Bank, 2024). This involves calculating median prices for specific categories of inputs and outputs at the enumeration area (EA) level. If fewer than 10 price observations exist at the EA level, medians are calculated using progressively larger administrative units up to the national level. These derived prices are then multiplied by standardised quantities. The resulting local currency unit values are converted to 2020 US dollars using annual exchange rates and CPI data from the World Bank Open Data Initiative.
<i>Usual rainy season dates</i>	The usual onset and cessation dates of the rainy season for each $0.1^\circ \times 0.1^\circ$ grid cell are determined following the methodology of Dunning et al. (2016). This method first computes the climatological mean rainfall for each day of the calendar year and the overall climatological daily mean rainfall for the location. From these, the climatological cumulative daily rainfall anomaly is calculated for each day of the year. The day of the minimum in this cumulative anomaly series marks the beginning of the climatological water season (usual onset), and the day of the maximum in the series marks its end (usual cessation). These usual dates provide the baseline for defining the search window in the AquaBEHER package (see <i>Rainy season calendar</i>).
<i>Weather controls</i>	Daily climate variables are sourced from the ERA5 reanalysis dataset ($0.1^\circ \times 0.1^\circ$ resolution), provided by the Copernicus Climate Change Service (C3S) Climate Data Store (Boogaard et al., 2020). For use as controls in regression analyses, these daily data are aggregated. Unless otherwise specified, variables are aggregated over the calendar year. As a robustness check, alternative aggregation periods such as quarterly and monthly aggregations are used. <i>Temperature</i> is the average of the daily minimum and maximum temperatures, expressed in $^\circ\text{C}$. <i>Total precipitation</i> is the sum of daily precipitation over the aggregation period, measured in mm. <i>Relative humidity</i> is a measure of the amount of water vapour in the air compared to the maximum amount of water vapour the air can hold at that temperature, expressed as a percentage. <i>Harmful degree days</i> quantifies the number of days where the maximum temperature exceeds the 90th percentile of a location’s historical temperature distribution. Additionally, the number of days with precipitation above the 50th, 75th, and 90th percentiles of a location’s historical distribution, and below the 5th, 15th, and 25th percentiles, are computed to capture variations in rainfall intensity and dry spells. For these percentile-based measures, the historical distributions are calculated at the $0.5^\circ \times 0.5^\circ$ grid-cell level to ensure robust local benchmarks.

Note. For time-varying variables, missing values are linearly interpolated.

Table A1.2: Sampled surveys


Country	LSMS-ISA		DHS
	Survey years	Waves	Survey years
Ethiopia	2011, 2013, 2015, 2018, 2022	1–5	2000, 2005, 2011, 2016, 2019
Malawi	2010, 2013, 2016, 2019	1–4	2000, 2004, 2010, 2015
Mali	2014, 2017	1–2	1995, 2001, 2006, 2012, 2018
Niger	2011, 2014	1–2	1992, 1998, 2012
Nigeria	2010, 2012, 2015, 2018	1–4	1990, 2003, 2008, 2013, 2018
Tanzania	2008, 2010, 2012	1–3	1999, 2010, 2015, 2022

Note. Year reported corresponds to the start year for surveys spanning multiple calendar years.

A1.2 Descriptive statistics

This subsection presents descriptive statistics for LSMS coordinates on sampled countries along with the geographic coverage of these coordinates within country and climatic area.

Figure A1.1: Area covered by the study



Note. The map displays the locations of sampled LSMS communities (pink dots) across the three primary climatic zones within the six study countries—Tropical, Arid, and Temperate—based on Metzger et al. (2023). These large colored regions correspond to the climatic area-by-country groups that define the spatially-specific time effects (δ_{cat}). The inset, showing the border between Nigeria and Niger, illustrates the resolution of the $0.1^\circ \times 0.1^\circ$ grid cells where location fixed effects (μ_ℓ) are applied.

Table A1.3: Geographical distribution of sampled $0.1^\circ \times 0.1^\circ$ grid cells within climatic area-by-country groups

Country	Climatic Area	No. of unique sampled grid cells ($0.1^\circ \times 0.1^\circ$)
Ethiopia	Arid	372
	Temperate	903
	Tropical	890
Malawi	Arid	1
	Temperate	339
	Tropical	171
Mali	Arid	876
	Tropical	654
Niger	Arid	246
Nigeria	Arid	467
	Tropical	2100
Tanzania	Arid	30
	Temperate	343
	Tropical	1461
Total		8,853

Note: Number of unique $0.1^\circ \times 0.1^\circ$ grid cells containing LSMS communities that are nested within each of the broader climatic area-by-country groups. These counts correspond to the number of location fixed effects (μ_ℓ) being absorbed by each spatially-specific time fixed effect (δ_{cat}) in the benchmark regression model.

Table A1.4: Descriptive statistics

	Mean (1)	Std. dev. (2)	Min (3)	Max (4)	Observations (5)
A. Plot					
Yield (2020 USD/ha)	1,574.65	3,263.98	0.00	19,227.22	135,726
Yield (kg/ha)	4,706.72	9,465.56	0.00	53,049.94	136,348
Crop failure (zero yield)	0.03	0.18	0.00	1.00	136,348
Plot suffered crop shock	0.35	0.48	0.00	1.00	137,507
False onset	0.05	0.22	0.00	1.00	137,948
Onset of the rainy season (week of year)	24.74	12.46	1.00	50.00	137,948
Planting month	5.90	2.27	1.00	12.00	126,440
Number of planting months	1.22	0.72	1.00	12.00	126,440
Seed quantity (kg/ha)	501.62	1,612.81	0.00	8,532.42	123,314
Seed value (2020 USD/ha)	235.61	688.89	0.00	3,908.35	123,514
Used fertiliser	0.51	0.50	0.00	1.00	136,099
Used pesticides	0.10	0.30	0.00	1.00	136,363
Main crop is a cereal	0.54	0.50	0.00	1.00	137,948
Main crop is a tuber	0.15	0.35	0.00	1.00	137,948
Main crop is a legume	0.11	0.32	0.00	1.00	137,948
Number of seasonal crops	1.66	0.97	0.00	15.00	137,820
Female plot manager	0.20	0.40	0.00	1.00	137,828
Plot manager age	47.05	15.01	0.00	100.00	137,646
Plot manager has formal education	0.50	0.50	0.00	1.00	137,605
Plot manager has wage work (past 12m)	0.10	0.30	0.00	1.00	135,852
Plot manager has wage work (past 7d)	0.10	0.30	0.00	1.00	135,226
B. Households					
Per capita consumption (2020 USD)	544.25	428.44	88.12	3,669.61	29,683
Household size	5.91	3.49	1.00	84.00	36,024
HH head age	47.63	15.11	8.00	100.00	35,985
Female HH head	0.17	0.37	0.00	1.00	35,999
HH head has any wage work (past 12m)	0.11	0.31	0.00	1.00	35,962
HH head has non-farm wage work (past 12m)	0.08	0.27	0.00	1.00	35,962
Has primary-age children	0.76	0.43	0.00	1.00	31,756
Farm size (ha)	1.80	3.09	0.02	29.78	35,486
Number of fallow plots	0.14	0.49	0.00	12.00	35,709
Number of cultivated plots	4.73	5.28	0.00	51.00	36,091
Owns non-farm enterprise	0.34	0.47	0.00	1.00	36,064
Household asset index	0.06	0.92	-1.22	10.00	36,065
Urban household	0.10	0.30	0.00	1.00	36,091
Onset of the rainy season (week of year)	24.17	12.19	1.00	50.00	36,081
C. Individuals					
Has any wage work (past 12m)	0.03	0.18	0.00	1.00	208,521
Has any wage work (past 7d)	0.03	0.18	0.00	1.00	205,742
Female	0.50	0.50	0.00	1.00	211,194
Age	22.33	18.91	0.00	100.00	210,547
Married	0.35	0.48	0.00	9.00	204,155
Has formal education	0.54	0.50	0.00	1.00	203,053
Onset of the rainy season (week of year)	24.75	11.27	1.00	50.00	213,312
D. Adult women (DHS)					
Has any wage work (past 12m)	0.88	0.32	0.00	1.00	92,737
Age	18.96	3.40	10.00	46.00	79,901
Married	0.82	0.39	0.00	1.00	92,737
Years of education	3.46	3.77	0.00	20.00	92,730
HH head age	42.43	13.43	13.00	97.00	92,642
HH head is female	0.18	0.39	0.00	1.00	92,737
Household size	4.46	0.92	1.00	5.00	92,737
BMI	21.88	3.54	12.02	59.72	61,558
Underweight (BMI < 18.5)	0.14	0.34	0.00	1.00	51,796
Wasted	0.10	0.31	0.00	1.00	59,073
Weight-for-height z-score	-0.83	1.00	-3.99	5.96	59,073
Stunted	0.19	0.39	0.00	1.00	59,712
Height-for-age z-score	-1.09	1.06	-5.99	5.70	59,712
Onset of the rainy season (week of year)	33.28	14.49	1.00	50.00	92,737
E. ERA5 cells					
Onset of the rainy season (week of year)	26.04	12.75	1.00	50.00	255,072

Continued on next page

Table A1.4 – continued from previous page

	Mean (1)	Std. dev. (2)	Min (3)	Max (4)	Observations (5)
Cessation of the rainy season (week of year)	34.37	11.31	1.00	52.00	170,676
Length of the rainy season	121.18	55.74	25.00	337.00	170,676
False onset	0.02	0.14	0.00	1.00	255,507
Days with precipitation below 25th percentile	31.13	45.03	0.00	269.00	255,438
Days with precipitation above 75th percentile	93.56	22.82	0.00	226.00	255,438
Maximum relative humidity	74.17	15.52	32.93	99.93	255,438
Average daily temperature (Celsius)	24.36	4.24	8.48	31.30	255,438
Total precipitation (mm)	1,261.35	826.63	75.92	11,772.02	255,438
Harmful degree days	34.45	50.61	0.00	1,275.74	255,438

Note. To account for outliers, the following continuous variables are winsorised at the 1st and 99th percentiles: Yield (2020 USD/ha), Yield (kg/ha), Seed quantity (kg/ha), Seed value (2020 USD/ha), Per capita consumption (2020 USD), and Farm size (ha). Appendix A1.1 provides detailed information on all variables, selected surveys, and weighting procedures.

A1.3 Weight Construction

I start from the LSMS household sampling weight w_i . To pool survey waves within a country while preserving representativeness, I apply a country–wave factor $f_{ct} = W_{ct}/W_c$, where $W_{ct} = \sum_{i \in (c,t)} w_i$ is the sum of LSMS weights in country c , wave t , and $W_c = \sum_{t'} W_{ct'}$ sums over all waves used for that country. The final household weight is $w_i^* = w_i f_{ct}$. For plot-level analyses, if household i reports P_i plots, each plot p receives $w_{ip}^* = w_i^*/P_i$, ensuring households contribute proportionally regardless of plot count. I use w_i^* or w_{ip}^* for all descriptive statistics and regressions, according to the unit of analysis.


Appendix B: Supplementary results

B1.1 Trends in rainy season onset, cessation, and length

This appendix provides supplementary figures regarding the long-term trends in rainy season onset, cessation, and length, complementing the analysis presented in Section 1.3.3. The trend estimation methodology, including the fixed-effects model and data construction, is detailed in the main text. All trends cover the period 1979–2024 for $0.1^\circ \times 0.1^\circ$ grid cells matched to LSMS-ISA survey locations, with standard errors clustered at the $0.5^\circ \times 0.5^\circ$ level.

Forecast to 2100. Figure B1.1 plots the forecast to 2100. By 2100, the linear trend model predicts an average delay of at least ten days in all countries, with extreme cases like Ethiopia and Nigeria of more than a month.

Figure B1.1: Long-range forecast of the onset of the rainy season to 2100




Country	Forecasted change in mean onset	
	By 2050	By 2100
Ethiopia	11.5	33.5
Malawi	6.9	20.1
Mali	10.3	30.0
Niger	4.4	12.8
Nigeria	14.5	42.4
Tanzania	3.4	9.8

Note. Comparison of the long-range forecast of rainy season onset to 2100 using a linear versus. Each panel shows the annual mean onset (in DOY) for a country, with blue dots representing the historical data (1979–2024). The dashed red line show the out-of-sample forecasts from the linear model. The model is estimated using only historical data, and the trends are then extrapolated. The vertical dashed line marks the year 2024, the beginning of the forecast period. The accompanying table summarizes the total change in onset days predicted by the linear model for 2050 and 2100.

Trends in onset variability. Beyond shifts in the average onset timing, I also examined changes in its year-to-year variability. To do this, I first de-trended the onset data for each country by calculating the residuals from the country-specific fixed-effects models. Then, for each country, I computed the interquartile range (IQR) of these de-trended onset residuals within 5-year rolling windows. Figure B1.2 plots the evolution of this IQR over time for each country.


Figure B1.2: Country-specific trends in rainy season onset variability



Note: Each panel shows the 5-year rolling interquartile range (IQR) of de-trended onset Day of Year (DOY) residuals for LSMS-ISA matched grid cells within the respective country. De-trending is based on country-specific fixed-effects models. The red line represents the linear trend fitted to these 5-year IQR values. This figure illustrates changes in the year-to-year predictability of onset timing. A negative trend (e.g., Mali, Niger) indicates a decrease in variability, meaning onset timing has become more predictable. A positive trend (e.g., Malawi, Nigeria) indicates an increase in variability, suggesting onset timing has become less predictable.

Cessation trends. As noted in Section 1.3.3, trends in the cessation of the rainy season are more varied than onset trends. Figure B1.3 presents the country-specific trend analyses for the cessation DOY, highlighting this heterogeneity. For instance, while Niger and Nigeria show a tendency towards later cessation, Malawi, Mali, and Tanzania exhibit trends towards earlier cessation. A pooled analysis for cessation (not shown) consequently reveals a relatively flat and statistically insignificant overall trend.


Figure B1.3: Country-specific trends in rainy season cessation day of year



Note. Each panel displays the annual mean cessation day of the year (DOY) for $0.1^\circ \times 0.1^\circ$ grid cells matched to LSMS-ISA survey locations within the respective country from 1979–2024. The *Cessation of the rainy season* is the day of the year when conditions specified in Section 1.3 are met. The dashed red line represents the fitted linear trend, estimated using country-specific fixed-effects models. The total predicted change in the mean cessation day of the year from 1979–2024, as implied by the model, is: Ethiopia (1.8 days), Malawi (-7.8 days), Mali (-4.7 days), Niger (8.3 days), Nigeria (7.3 days), and Tanzania (-13.1 days).

Season length trends. The interplay of later onsets and diverse cessation patterns generally leads to a decrease in the length of the rainy season across the study area. The consistent and significant delay in the season’s onset is the primary driver of this reduction. Figure B1.4 illustrates these country-specific trends in the length of the rainy season.

Figure B1.4: Country-specific trends in rainy season length



Note. Each panel displays the annual mean of the length of the rainy season for $0.1^\circ \times 0.1^\circ$ grid cells matched to LSMS-ISA survey locations within the respective country from 1979–2024. The dashed red line represents the fitted linear trend, estimated using country-specific fixed-effects models. The total predicted change in the mean length of the rainy season in days over the full sample period, as implied by the model, is: Ethiopia (-14.4 days), Malawi (-25.6 days), Mali (-18.7 days), Niger (0.7 days), Nigeria (-14.9 days), and Tanzania (-14.1 days).

B1.2 Issues related to identification

Selection bias. A primary concern for identification is the potential for selection bias stemming from the sample restrictions used in the main analyses. The plot-level productivity analysis is restricted to plots that report a complete harvest, while the household-level consumption analysis uses an even stricter criterion, including only households for which *all* reported plots have a complete harvest. This could bias the estimates if the timing of the rainy season onset systematically affects the probability of an observation being included in these samples. For instance, a delayed onset could push the harvest date past the survey’s fieldwork period, leading to an underrepresentation of late-onset plots and households.

To formally test for this at both levels, I estimate linear probability models where the dependent variable is an indicator for meeting the respective sample selection criterion. As shown in Table B1.1, the results provide confidence that selection bias is not a significant concern. At the plot level, the coefficient on the rainy season onset week is small and not statistically significant, indicating that a late onset does not predict whether an individual plot reports a complete harvest. This result holds at the household level as well; onset timing does not significantly predict the probability that a household reports a complete harvest for all of its plots.


Table B1.1: Onset of the rainy season and selection into analysis sample

Dependent variable:	Complete harvest (plot) (1)	Complete harvest (household) (2)
Onset of the rainy season	-0.002 (0.001) [0.151]	-0.002 (0.003) [0.543]
Standardised effect	-0.003	-0.003
Mean (dep. var.)	0.904	0.732
Identifying observations	148,854	43,330
Singleton observations	177	370
Countries	6	6
Survey year range	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. The main independent variable, *Onset of the rainy season*, is the calendar week when the agronomic conditions specified in Section 1.3 are met. All specifications include high-resolution ($0.1^\circ \times 0.1^\circ$) grid-cell fixed effects, a full set of climatic area-by-country-by-year fixed effects, and a comprehensive set of plot-level and weather controls detailed in Section 1.3. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses; *p*-values are in brackets. Appendix A1.1 provides further details on all variable definitions, survey data, and weighting procedures.


Variation and identification. A common concern with demanding fixed-effects models is whether sufficient variation remains in the data to precisely identify the parameters of interest. This section presents two figures to address this concern. Figure B1.5 presents a variance decomposition for the main outcome and explanatory variables to show the primary sources of variation in the data. Figure B1.6 then plots the kernel density of the residualised onset of the rainy season. This residual represents the plausibly exogenous variation used for identification, and the figure is used to visually inspect its distribution for any signs of systematic bias.

Figure B1.5: Between and within variation decomposition of onset and yields



Note. Between and within standard deviations for three standardised variables: planting month, onset of the rainy season, and log yield. The decomposition is shown for three different panel data structures: climatic area by country, year, and $0.1^\circ \times 0.1^\circ$ grid cell. The sample used for this decomposition is the plot-level LSMS-ISA sample as described in Section 1.3.

Figure B1.6: Residualised onset of the rainy season




Note. The figure plots the kernel density of the residualised onset of the rainy season for the three main samples used in the analysis: the plot-level agricultural sample, the household-level consumption sample, and the women's nutrition sample (DHS). For each sample, the residuals are obtained by regressing the onset week on the full set of fixed effects used in its corresponding benchmark specification. For visual clarity, the distributions are trimmed at -10 and 10. The onset of the rainy season is defined as the week of the year when the agronomic conditions specified in Section 1.3 are met.

Robustness to place-specific trends. Recent literature highlights that panel analyses of climate impacts can produce biased estimates if unobserved, place-specific trends correlate with weather variables, a concern particularly relevant when using binned regressors (Jones et al., 2025). While my main analysis relies on a continuous measure of rainy season onset, this issue is pertinent to the tercile indicators I use in two secondary analyses: testing for asymmetric effects on yields (Figure 1.3) and examining farmer awareness of shocks (Table 1.7).

To address this potential issue, I augment my benchmark specification by controlling for grid-cell-specific linear trends, which Jones et al. (2025) identify as the most suitable strategy for a short panel. Given that I bin the onset of the rainy season at the $0.5^\circ \times 0.5^\circ$ level, I implement these trends by interacting fixed effects for each grid cell with a linear year trend. Although a linear trend assumption warrants caution, yields have largely stagnated in my study countries (Wollburg et al., 2024), making it a reasonable first-order approximation for smooth, unobserved spatial changes. Jones et al. (2025) provide further support, finding that crop yield estimates are less susceptible to this form of bias.

As I show in Figure B1.7 and Table B1.2, my findings are robust to this more demanding specification. The inclusion of these trends does not alter the key results: the negative impact on agricultural yields remains driven entirely by late onsets, and farmers' awareness of a shock is still only triggered by exceptionally late onsets.

Figure B1.7: Asymmetric effects — Correcting for binning bias



Note. The figure displays OLS regression coefficients for indicator variables representing the bottom tercile and top tercile of the historical rainy season onset distribution within $0.5^\circ \times 0.5^\circ$ grid cells. The middle tercile is the omitted reference category. The dependent variable is log agricultural yields (2020 USD per hectare). The regression augments equation (1.1) by including interactions between the $0.5^\circ \times 0.5^\circ$ grid-cell fixed effects and a linear time trend. Bars represent 95% confidence intervals based on standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level.

Table B1.2: Farmers' awareness — Correcting for binning bias

Dependent variable:	Crops shock (1)	Drought shock (2)
Earliest onset tercile	-0.011 (0.018) [0.530]	0.019 (0.018) [0.284]
Latest onset tercile	0.029 (0.014) [0.044]	0.039 (0.014) [0.005]
Standardised effect		
Mean (dep. var.)	0.411	0.202
Identifying observations	136,420	132,044
Singleton observations	173	180
Countries	6	6
Planting year range	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable in each column is a binary indicator for whether a farmer reported a negative shock on a given plot. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. *Earliest onset tercile* and *Latest onset tercile* are indicators for whether the onset of the rainy season for a location falls into the earliest or latest tercile of its historical distribution, with the middle tercile serving as the omitted category. These terciles are defined separately for each $0.5^\circ \times 0.5^\circ$ grid cell based on its long-term (1979–2020) onset history. For comparison, Panel A presents results using the continuous onset week variable. All specifications augment equation (1.1) by including interactions between the $0.5^\circ \times 0.5^\circ$ grid-cell fixed effects and a linear time trend. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides detailed information on all variables.

B1.3 Robustness of impact on productivity

This appendix presents a series of robustness checks for the main findings on the impact of rainy season onset on agricultural productivity, as discussed in Section 1.5.1 and presented in Table 1.1.

Alternative local trends. The stability of the estimated onset coefficient was tested against alternative specifications that incorporate different ways of controlling for local time trends. Introducing local time trends effectively changes the definition of the onset shock, as it alters

the variation from which the effect is identified by removing different local temporal patterns. Table B1.3 shows that the negative impact of a delayed onset on agricultural productivity remains consistent even when employing these more demanding specifications for local trends, which vary the geographic level at which such trends are defined.


Table B1.3: Alternative local trends

Dependent variable:	Log yields (2020 USD per hectare)			
	(1)	(2)	(3)	(4)
Onset of the rainy season	-0.017 (0.007) [0.012]	-0.019 (0.007) [0.004]	-0.016 (0.007) [0.034]	-0.020 (0.008) [0.012]
Standardised effect	-0.027	-0.030	-0.022	-0.022
Mean (dep. var.)	5.744	5.744	5.744	5.744
Identifying observations	129,919	129,919	129,919	129,913
Singleton observations	180	180	180	186
Countries	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022
Level of local trend	Country	Climatic area	Admin level 1	Admin level 2

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects. The columns differ by the inclusion of location by planting year trends. A full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures. The standardised effect shows the impact of a one standard deviation change in the onset week.

Past and future onsets. To assess whether the estimated impact is genuinely contemporaneous, I introduced past (one-year lag) and future (one-year lead, a placebo) season onset dates into the benchmark specification (equation 1.1). As shown in Figure B1.8, neither past nor future onset dates significantly affect current season productivity.

Figure B1.8: Timing



Note. Marginal effects of rainy season onset on log yields measured in 2020 USD per hectare. Estimates are based on equation 1.1. On the x-axis, *t* refers to the onset week of the actual planting year, *t-1* refers to the onset week from the previous year, and *t+1* refers to the onset week from the subsequent year, serving as a placebo test. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

The roles of rainy season cessation and length. To better understand the distinct role of rainy season onset relative to other timing features, I examined specifications that also incorporate

cessation timing and total season length. Importantly, the agronomic conditions defining cessation are not met in all location-years, leading to missing values for cessation and, consequently, length for some observations. Therefore, in Table B1.4, I restrict the sample to observations with non-missing cessation and length data.

Table B1.4: Cessation and length of the rainy season

Dependent variable:	Log yields (2020 USD per hectare)				
	(1)	(2)	(3)	(4)	(5)
Cessation of the rainy season	-0.004 (0.004) [0.254]	-0.003 (0.004) [0.450]			
Onset of the rainy season		-0.016 (0.008) [0.045]		-0.018 (0.008) [0.025]	-0.017 (0.008) [0.029]
Length in weeks			-0.003 (0.004) [0.514]	-0.003 (0.004) [0.426]	
Mean (dep. var.)	5.616	5.616	5.616	5.616	5.616
Identifying observations	106,511	106,511	106,511	106,511	106,511
Singleton observations	97	97	97	97	0
Countries	6	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met; *Cessation of the rainy season* is the week of the year when conditions specified in Section 1.3 are met; and *Length* is the total number of weeks the rainy season lasted, calculated from the onset and cessation dates. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, weather controls, and plot-level controls (a full list of these controls is presented in Section 1.3). The sample for these regressions is restricted to observations with non-missing data for cessation and length. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are reported in brackets. Appendix A1.1 provides further details on variable construction.

Alternative specifications for weather controls. Table B1.5 presents estimates of equation 1.1 using weather variables aggregated at different frequencies.

Table B1.5: Alternative weather controls

Dependent variable:	Log yields (2020 USD per hectare)		
	(1)	(2)	(3)
Onset of the rainy season	-0.021 (0.007) [0.002]	-0.023 (0.007) [0.002]	-0.019 (0.007) [0.009]
Standardised effect	-0.033	-0.033	-0.027
Mean (dep. var.)	5.744	5.744	5.744
Identifying observations	129,919	129,919	129,919
Singleton observations	180	180	180
Countries	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022
Time aggregation	Yearly	Quarterly	Monthly

Note: OLS regression estimates of Equation 1.1. The dependent variables is the log of yield (measured in 2020 US dollars per hectare). *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, and plot-level controls (for yield regressions) or household-level controls (for consumption regressions) as listed in Section 1.3. The columns differ by the temporal aggregation of the weather controls: Columns (1) and (3) use weather controls aggregated at the quarterly level, while Columns (2) and (4) use weather controls aggregated at the monthly level. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are reported in brackets. Appendix A1.1 provides further details on variable construction.

Alternative transformations of yield. The negative impact of a delayed onset on agricultural productivity persists when using alternative transformations of the yield variable. Table B1.6

shows results using $\log(1+\text{yield})$, the inverse hyperbolic sine (IHS) transformation of yield, and yield levels (after trimming the 1st and 99th percentiles to mitigate outlier influence).

Table B1.6: Alternative transformations of yield

Dependent variable transformation	Yield (trimmed) (1)	Log(1+Yield) (2)	IHS(Yield) (3)
Onset of the rainy season	-23.513 (10.898) [0.031]	-0.028 (0.009) [0.001]	-0.030 (0.009) [0.001]
Standardised effect	-36.476	-0.044	-0.046
Mean (dep. var.)	801.653	5.547	6.203
Identifying observations	133,530	134,880	134,880
Singleton observations	182	181	181
Countries	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022

Note: OLS regression estimates of Equation 1.1. Each column uses a different transformation of yield (measured in 2020 USD per hectare) as the dependent variable. Column (1) uses yield levels (measured in 2020 US dollars per hectare), with the 1st and 99th percentiles trimmed. Column (2) uses the log of (1 + yield). Column (3) uses the inverse hyperbolic sine (IHS) transformation of yield. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, weather controls, and plot-level controls (a full list of these controls is presented in Section 1.3). Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are reported in brackets. Appendix A1.1 provides further details on variable construction.

Alternative measures of output. The negative impact of a delayed onset extends to other measures of farm output. Table B1.7 presents results for the impact of onset timing on: output in kilograms, output in constant 2020 US dollars, log of yields in kilograms per hectare, and an indicator for crop failure, defined as zero reported yield.

Table B1.7: Alternative output measures

Dependent Variable:	Log output (2020 USD) (1)	Log output (kg) (2)	Log yield (kg per hectare) (3)	Crop failure (4)
Onset of the rainy season	-0.019 (0.008) [0.015]	-0.028 (0.008) [0.000]	-0.031 (0.006) [0.000]	0.002 (0.001) [0.047]
Standardised effect	-0.030	-0.044	-0.049	0.003
Mean (dep. var.)	4.331	5.407	6.824	0.037
Identifying observations	131,185	131,756	130,475	135,504
Singleton observations	179	178	179	178
Countries	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Constant national prices. The benchmark yield measure is valued using median prices at the enumeration area level. If a delayed onset causes a local negative supply shock, local prices may rise, attenuating the estimated negative impact on revenue. To address this, I construct an alternative yield measure where output is valued using constant national median prices. I use a hierarchical approach: prices are first assigned based on the national median for the specific *crop*; if unavailable, I use the national median for the *crop type*; finally, I index any remaining

gaps using a country-year weighted price index. Table B1.8 reestimates Table 1.1.

Table B1.8: Onset of the rainy season and agricultural productivity: Constant national prices

Dependent variable:	Log yields (2020 USD per hectare)		
	(1)	(2)	(3)
Onset of the rainy season	-0.028 (0.008) [0.001]	-0.027 (0.009) [0.003]	-0.027 (0.009) [0.002]
Standardised effect	-0.046	-0.041	-0.042
Mean (dep. var.)	5.751	5.751	5.751
Identifying observations	130,777	130,777	130,142
Singleton observations	182	182	180
Countries	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022
Weather controls	No	Yes	Yes
Demographic controls	No	No	Yes

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. Local prices per crop are calculated at constant national levels using a hierarchical procedure detailed in Section B1.3. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. A full list of controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Cereal crops subsample. To ensure that the results are not driven by composition effects from aggregating diverse crop types with different physical properties (for instance, heavier tubers vs. lighter cereals), I compute plot-level yields for cereals. Table B1.9 reestimates Table 1.1.

Table B1.9: Onset of the rainy season and cereals productivity


Dependent variable:	Log yields (2020 USD per hectare)		
	(1)	(2)	(3)
Onset of the rainy season	-0.023 (0.011) [0.037]	-0.023 (0.012) [0.051]	-0.019 (0.010) [0.064]
Standardised effect	-0.034	-0.032	-0.027
Mean (dep. var.)	6.762	6.762	6.761
Identifying observations	73,946	73,946	73,614
Singleton observations	113	113	113
Countries	5	5	5
Planting year range	2009–2022	2009–2022	2009–2022
Weather controls	No	Yes	Yes
Demographic controls	No	No	Yes

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield of cereals grown in a plot, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. A full list of controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Permutation-based inference. To further evaluate the statistical robustness of the estimated effect of rainy season onset on agricultural productivity, I implemented a series of permutation tests. These tests involve re-estimating the impact of rainy season onset (equation 1.1, using specifications from Table 1.1) across 500 iterations for each of three distinct randomisation scenarios. Scenario A permutes planting years within local $0.5^\circ \times 0.5^\circ$ grid cells, effectively reassigning the timing of the onset shock experienced by a household to that of another household in a different wave but within the same local area. Scenario B applies a similar permutation of planting years but within the broader scope of each country. Scenario C directly tests

the importance of an observation's specific geographic location by randomising its geographic identifiers within its country. This process reassigns the experienced weather and onset data to the household while holding its original planting year constant.

Figure B1.9: The effect on yields: permutation-based inference



Note. Empirical distributions of estimated coefficients for *Onset of the rainy season* from 500 permutations for three randomisation scenarios. Panel A: Planting year permuted within $0.5^\circ \times 0.5^\circ$ grid cells. Panel B: Planting year permuted within countries. Panel C: Geographic location (and thus onset/weather) permuted across communities within countries, holding planting year constant. Each distribution is generated by re-estimating the three benchmark specifications from Table 1.1. The sub-table reports the empirical p-values, indicating the proportion of permuted coefficients that are as extreme or more extreme than the actual estimated coefficients. All models include the full set of fixed effects and controls as per Equation 1.1. Appendix A1.1 provides detailed information on variables.

Alternative clustering of standard errors. The benchmark analysis clusters standard errors at the $0.5^\circ \times 0.5^\circ$ grid-cell level to account for potential spatial correlation in errors. To assess the sensitivity of statistic inference to this choice, Table B1.10 presents the estimated coefficient for rainy season onset on log yields from the preferred specification (Table 1.1, Column 3) under several alternative clustering assumptions. These include clustering at finer ($0.1^\circ \times 0.1^\circ$ grid cell) and coarser ($1^\circ \times 1^\circ$ grid cell) spatial units, as well as two-way clustering by spatial unit and planting year.

Table B1.10: Alternative clustering of standard errors

Dependent variable: <i>Level of clustering</i>	Log yields (2020 USD per hectare)					
	None	$0.1^\circ \times 0.1^\circ$	$0.5^\circ \times 0.5^\circ$	$1^\circ \times 1^\circ$	$0.5^\circ \times 0.5^\circ$ $\times$ planting year	$1^\circ \times 1^\circ \times$ planting year
	(1)	(2)	(3)	(4)	(5)	(6)
Onset of the rainy season	-0.021 (0.004) [0.000]	-0.021 (0.007) [0.001]	-0.021 (0.007) [0.002]	-0.021 (0.007) [0.004]	-0.021 (0.008) [0.023]	-0.021 (0.008) [0.025]
Standardised effect	-0.033	-0.033	-0.033	-0.033	-0.033	-0.033
Mean (dep. var.)	5.744	5.744	5.744	5.744	5.744	5.744
Identifying observations	129,919	129,919	129,919	129,919	129,919	129,919
Singleton observations	180	180	180	180	180	180
Countries	6	6	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Each column uses a different clustering level for standard errors. Standard errors are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

B1.4 Stability of impacts on productivity over time

This section tests whether the impact of a delayed rainy season onset is stable over time. To do this, I leverage the temporal dimension of the data for the subset of countries with a sufficient number of survey waves to identify within-location changes between two distinct periods. Thus, identification of the interaction terms involving the *Later waves* variable relies solely on the variation within Ethiopia, Malawi, and Nigeria, which are the only countries in the sample with four or more waves allowing for the definition of distinct early and late periods.

I create an indicator variable, *Later waves*, which takes the value 1 for more recent survey periods. The definition of this variable is country-specific to reflect the different survey timings: for Ethiopia, *Later waves* = 1 for Waves 3-5 (covering years 2015-2022) and 0 for Waves 1-2 (2011-2013); for Malawi, it equals 1 for Waves 3-4 (2016-2019) and 0 for Waves 1-2 (2010-2013); and for Nigeria, it equals 1 for Waves 3-4 (2015-2018) and 0 for Waves 1-2 (2010-2012). For all other countries (Mali, Niger, Tanzania), this variable effectively remains constant within the available waves, and thus observations from these countries do not contribute to identifying the coefficients on terms interacted with *Later waves*.

Table B1.11 tests for a change in the average impact over time by interacting the onset week with this *Later waves* indicator. Table B1.12 expands this to a triple-interaction model to test whether the buffering or vulnerability effects of farmer characteristics have also changed over time.

Table B1.11: Dynamic effect of the onset on agricultural productivity

Dependent variable:	Log yields (2020 USD per hectare)		
	(1)	(2)	(3)
Onset of the rainy season	-0.018 (0.007) [0.005]	-0.018 (0.007) [0.011]	-0.019 (0.006) [0.004]
× <i>Later waves</i>	-0.014 (0.005) [0.005]	-0.013 (0.005) [0.008]	-0.011 (0.004) [0.011]
Mean (dep. var.)	5.744	5.744	5.744
Identifying observations	130,556	130,556	129,919
Singleton observations	182	182	180
Countries	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022
Weather controls	No	Yes	Yes
Demographic controls	No	No	Yes

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. The regressions are run on the full sample, but the identification of the interaction term × *Later waves* relies solely on variation within Ethiopia, Malawi, and Nigeria. *Later waves* is an indicator variable equal to 1 for the following periods: Ethiopia (2015-2022), Malawi (2016-2019), and Nigeria (2015-2018); and 0 otherwise. The main effect of *Later waves* is absorbed by the fixed effects. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. A full list of controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B1.12: Dynamic heterogeneous effects of the onset on agricultural productivity

Dependent variable: Indicator:	Log yields (2020 USD per hectare)				
	Female-managed plot (1)	Educated plot manager (2)	Below median assets (3)	Irrigated (4)	Improved seeds (5)
Onset of the rainy season	-0.018 (0.006) [0.005]	-0.018 (0.007) [0.005]	-0.014 (0.006) [0.021]	-0.019 (0.007) [0.004]	-0.007 (0.009) [0.453]
× Later waves	-0.008 (0.004) [0.054]	-0.018 (0.005) [0.000]	-0.008 (0.005) [0.081]	-0.010 (0.004) [0.014]	-0.030 (0.011) [0.006]
× Indicator	-0.001 (0.002) [0.741]	0.002 (0.002) [0.409]	-0.005 (0.002) [0.009]	0.011 (0.009) [0.234]	0.008 (0.006) [0.201]
× Indicator × Later waves	-0.007 (0.003) [0.014]	0.009 (0.003) [0.001]	-0.004 (0.003) [0.169]	0.012 (0.010) [0.233]	-0.003 (0.007) [0.646]
Mean (dep. var.)	5.744	5.744	5.743	5.749	5.705
Identifying observations	129,919	129,919	129,886	128,486	105,671
Singleton observations	180	180	180	177	215
Countries	6	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2011–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. The regressions are run on the full sample, but the identification of interaction terms involving *Later waves* relies solely on variation within Ethiopia, Malawi, and Nigeria. *Later waves* is an indicator variable equal to 1 for the following periods: Ethiopia (2015–2022), Malawi (2016–2019), and Nigeria (2015–2018); and 0 otherwise. The *Indicator* variables in each column are defined as follows: (1) *Female-managed plot*: an indicator equal to 1 if the plot manager is female; (2) *Educated plot manager*: an indicator equal to 1 if the plot manager has any formal education; (3) *Below median assets*: an indicator equal to 1 if the household's asset index is below the sample median; (4) *Irrigated*: an indicator equal to 1 if the plot is irrigated; (5) *Improved seeds*: an indicator equal to 1 if improved seeds were used. All specifications include the relevant lower-order terms (e.g., *Onset* × *Indicator*, *Indicator* × *Later waves*) and the main effect of the *Indicator*. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

B1.5 Supplementary results on adaptation strategies

This appendix provides additional results for adaptation margins, supplementing the summary findings presented in Section 1.5.3 and Table 1.4. The tables below present these supplementary findings: Table B1.13 details responses in the use of various on-farm inputs including seeds, fertilisers, pesticides, and on-farm labour; Table B1.14 covers crop choice and plot composition; Table B1.15 shows changes in land use along the extensive and intensive margin; and Table B1.16 examines sectoral labour allocation.

Table B1.13: Onset of the rainy season and input use

Dependent variable:	Seed value	Inorganic fertiliser quantity	Inorganic fertiliser value	Days from family	Days from hired workers
	(1)	(2)	(3)	(4)	(5)
Onset of the rainy season	0.005 (0.009) [0.591]	0.002 (0.013) [0.880]	-0.014 (0.009) [0.125]	2.204 (24.010) [0.927]	-5.854 (4.320) [0.176]
Standardised effect	0.008	0.003	-0.020	3.539	-9.388
Mean (dep. var.)	2.953	3.906	4.787	1070.531	102.819
Identifying observations	121,001	44,516	46,098	117,596	117,145
Singleton observations	198	192	195	143	143
Countries	6	6	6	5	5
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variables are: *Seed value*, the value of seeds used, measured in 2020 US dollars per hectare; *Inorganic fertiliser quantity*, the natural logarithm of the quantity of inorganic fertiliser applied, measured in kilograms per hectare; *Inorganic fertiliser value*, the natural logarithm of the value of inorganic fertiliser used, measured in 2020 US dollars per hectare; *Days from family*, the total number of days per hectare family members worked on the plot; and *Days from hired workers*, the total number of days per hectare hired workers worked on the plot. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, weather controls, and plot-level controls as listed in Section 1.3. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are in brackets. Appendix A1.1 provides further details on variable construction.

Table B1.14: Onset of the rainy season and crop choice

Dependent variable:	Share of plot value attributed to			Plot contains		
	Cereals	Tubers	Legumes	Cereals	Tubers	Legumes
	(1)	(2)	(3)	(4)	(5)	(6)
Onset of the rainy season	0.001 (0.002) [0.596]	-0.002 (0.001) [0.013]	0.000 (0.001) [0.690]	-0.003 (0.002) [0.098]	-0.002 (0.001) [0.044]	-0.001 (0.002) [0.623]
Standardised effect	0.001	-0.004	0.001	-0.004	-0.004	-0.001
Mean (dep. var.)	0.408	0.056	0.126	0.645	0.076	0.234
Identifying observations	136,855	136,855	136,855	136,855	136,855	136,855
Singleton observations	175	175	175	175	175	175
Countries	6	6	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variables are: *Share of plot value attributed to cereals/tubers/legumes*, the proportion of the total production value of the plot derived from the respective crop category. *Plot contains cereals/tubers/legumes* are indicator variables taking value 1 if the plot contains any crop from the respective category, and 0 otherwise. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, weather controls, and plot-level controls as listed in Section 1.3. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are in brackets. Appendix A1.1 provides further details on variable construction.

Table B1.15: Onset of the rainy season and land use

Dependent variable:	No. fallow plots (1)	No. cultivated plots (2)	Log total cultivated area (3)	Log total farm size (4)
Onset of the rainy season	0.004 (0.003) [0.128]	0.016 (0.032) [0.610]	-0.003 (0.007) [0.631]	0.002 (0.006) [0.729]
Standardised effect	0.007	0.026	-0.005	0.003
Mean (dep. var.)	0.153	4.451	-0.134	0.055
Identifying observations	30,881	30,881	30,613	30,696
Singleton observations	313	313	320	320
Countries	6	6	6	6
Survey year range	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variables are household-level measures: *No. fallow plots*, the number of plots left unused by the household; *No. cultivated plots*, the number of plots used in agriculture by the household; Column (3) *Log Total Cultivated Area*, the natural logarithm of total area in hectares actively cultivated by the household; and Column (4) *Log total farm size*, the natural logarithm of total area in hectares to which the household has planting rights. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, weather controls, and household-level controls (a full list of these controls is presented in Section 1.3). Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are reported in brackets. Appendix A1.1 provides further details on variable construction, selected surveys, and weighting procedures. Data are from the LSMS-ISA household-level sample.

Table B1.16: Onset of the rainy season and sector choice

Dependent variable:	Last 12 months			Last 7 days	
	Any wage work (1)	Agriculture (2)	Services (3)	Any wage work (4)	Household business (5)
A. Working-age individuals					
Onset of the rainy season	-0.001 (0.001) [0.029]	-0.001 (0.000) [0.254]	-0.000 (0.000) [0.131]	-0.002 (0.000) [0.000]	-0.001 (0.001) [0.158]
Standardised effect	-0.002	-0.001	-0.001	-0.003	-0.002
Mean (dep. var.)	0.050	0.018	0.025	0.054	0.127
Identifying observations	224,735	224,735	224,735	223,915	224,140
Singleton observations	8	8	8	8	9
Countries	6	6	6	6	6
Survey year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022
B. Household head					
Onset of the rainy season	-0.002 (0.001) [0.167]	-0.001 (0.001) [0.464]	-0.001 (0.001) [0.438]	-0.003 (0.001) [0.005]	-0.002 (0.001) [0.136]
Standardised effect	-0.003	-0.001	-0.001	-0.005	-0.004
Mean (dep. var.)	0.121	0.032	0.069	0.130	0.197
Identifying observations	47,653	47,653	47,653	47,402	47,582
Singleton observations	353	353	353	353	355
Countries	6	6	6	6	6
Survey year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variables are indicator variables (1 if the condition is met, 0 otherwise) for individuals of working age: *Any wage work*, reports working for a wage in the past twelve months; *Agriculture*, reports working for a wage in the agricultural sector in the past twelve months; *Services*, reports working for a wage in the service sector in the past twelve months; *Any wage work*, reports working for a wage in the past seven days; *Household business*, reports working for a household business in the past seven days. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, weather controls, and individual-level controls (sex, age, age squared, marital status, formal education, urban residence, household size) as specified in Section 1.4. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses; *p*-values are in brackets. Appendix A1.1 provides further details on variable construction.

Table B1.17: Lagged onset on adaptation

Dependent variable:	Timing			Inputs			Main crop is		
	Planting (DOY)	Planting duration	Harvest (DOY)	Seeds	Fertiliser	Pesticide	Cereals	Tubers	Legumes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Onset of the rainy season, $t - 1$	0.036 (0.208) [0.862]	0.001 (0.003) [0.769]	-0.075 (0.999) [0.940]	-0.004 (0.009) [0.639]	-0.003 (0.002) [0.100]	0.002 (0.001) [0.098]	-0.003 (0.001) [0.047]	0.001 (0.001) [0.175]	0.000 (0.001) [0.760]
Standardised effect	0.062	0.002	-0.126	-0.007	-0.005	0.004	-0.005	0.002	0.001
Mean (dep. var.)	188.934	1.105	250.370	3.453	0.516	0.066	0.572	0.069	0.122
Identifying observations	125,609	125,609	106,970	120,669	135,081	135,306	136,855	136,855	136,855
Singleton observations	74	74	117	197	173	173	175	175	175
Countries	5	5	5	6	6	6	6	6	6
Planting year range	2009–2022	2009–2022	2009–2022	2009–2022	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variables are defined as follows. *Planting (DOY)* is the approximated first planting day of the year; this approximation assumes planting occurred on the 15th of the reported month. *Planting duration* is the number of distinct months planting occurred on the plot. *Harvest (DOY)* is the approximated first harvest day of the year, assuming harvest occurred on the 15th of the reported month. *Seed* is the log of seed quantity used on the plot, measured in kilograms per hectare. *Fertiliser* is an indicator variable taking the value 1 if any fertiliser, whether organic or inorganic, was used, and 0 otherwise. *Pesticide* is an indicator variable that equals 1 if any pesticide was used and 0 otherwise. *Cereals*, *Tubers*, and *Legumes* are indicator variables taking the value 1 if the main crop by production value on the plot belongs to the respective category, and 0 otherwise. *Onset of the rainy season, $t - 1$* is the week of the year when conditions specified in Section 1.3 in the year before the current season. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors, which are clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides further details on variable construction, selected surveys, and weighting procedures.


B1.6 Simulation methodology to correct for aggregation bias

The use of monthly-level planting data may introduce aggregation bias when I estimate the effect of daily-level weather shocks. To address this, I implement a simulation exercise designed to recover the true daily-level coefficient from my coarse monthly data. The simulation involves two main stages: first, generating daily planting probability distributions for each location, and second, using these distributions in an iterative search algorithm to find the unbiased daily coefficient.

Generating daily planting probability distributions. To disaggregate the observed monthly planting dates into daily probabilities, I generate a daily probability mass function (PMF) for each unique ERA5 location in my sample. This PMF represents the likelihood of planting occurring on any given day of the year. I employ two distinct methods to construct these distributions as a robustness check. For each method, I generate 20 unique *drawing pools*, one for each iteration of the main simulation. Each drawing pool contains 100 simulated planting dates for each plot observation, providing a rich set of potential outcomes for the simulation. To reduce noise from outlier responses, I only consider planting months that, for a given location, account for at least 1% of the total observations for that location.

The first method is a parametric approach that fits a three-parameter skew-normal distribution to the sequence of observed planting *months*. This captures the timing and potential asymmetry of the planting season. To correctly model seasons that cross over the new year (e.g., November–February), I transform the month data onto a continuous axis where January becomes month 13 and February becomes month 14, allowing the distribution to be fitted over an uninterrupted domain. The resulting continuous probability density function is then discretised to generate daily probabilities. Figure B1.10 provides a visual representation of this method for a random sample of locations.


Figure B1.10: Implied daily planting probabilities from kernel density estimation



Note. This figure displays a random sample of nine ERA5 locations to illustrate the fit of the skew-normal distribution. The blue bars represent the density of observed monthly planting dates, with the day-of-year (DOY) approximated as the 15th of the reported month. The red line is the implied daily planting probability, derived by fitting a three-parameter skew-normal distribution to the sequence of observed months for each location. To reduce noise from outlier responses, only planting months that account for at least 1% of the total observations for a given location are considered. The model accounts for year-crossing seasons by transforming the month variable onto a continuous axis (e.g., January becomes month 13) before fitting the distribution.

As an alternative, I use a non-parametric Kernel Density Estimation (KDE). This method is more flexible and fits a Gaussian KDE directly to the day-of-year (DOY) of each planting observation, approximated as the 15th of the reported month. To handle the circular nature of annual data, I employ a *mirroring* technique. Planting dates from early in the year (e.g., day 15) are duplicated and shifted forward by a year (to day 380), while points from late in the year (e.g., day 350) are duplicated and shifted backward (to day -15). This ensures the KDE treats the year as a continuous loop, correctly modeling the density across the year-end boundary. Figure B1.11 illustrates the fit of the KDE method.

Figure B1.11: Implied daily planting probabilities from kernel density estimation



Note. This figure displays a random sample of nine ERA5 locations to illustrate the fit of the Kernel Density Estimator (KDE). The blue bars represent the density of observed monthly planting dates, with the day-of-year (DOY) approximated as the 15th of the reported month. The red line is the implied daily planting probability, derived by fitting a Gaussian KDE to the DOY data for each location. To reduce noise from outlier responses, only planting months that account for at least 1% of the total observations for a given location are considered. The model accounts for the circular nature of annual data using a *mirroring* technique, where data points near the start and end of the year are duplicated and shifted to ensure a continuous density across the year-end boundary.


Simulation design. The simulation exercise uses an iterative search algorithm to find the true daily-level coefficient (β_{daily}) that, when subjected to the same temporal aggregation as my observed data, reproduces the coefficient estimated from the original monthly data. The process is as follows:

1. **Establish target coefficient.** First, I establish a benchmark by running my main regression specification. I convert the observed planting months to a day-of-year (approximated as the 15th of the month) and regress this on the onset of the rainy season. The resulting coefficient serves as the target for the simulation to match, denoted as β_{target} .
2. **Generate stable residuals.** Before the main simulation, I run a preliminary 10-iteration loop to generate a stable counterfactual residual for my treatment variable. This step purges the residual of any confounding effects that might arise from shifts in the year-level fixed effects, which can occur when the treatment effect is applied to the daily data.
3. **Iterative search for β_{daily} .** The core of the simulation is a `while` loop that searches for the true daily coefficient. The search begins with an initial guess where $\hat{\beta}_{daily} = \beta_{target}$.

In each iteration, I simulate a daily planting date for each observation by taking a random draw from the appropriate daily probability distribution and adding the treatment effect, which is calculated as the stable residual multiplied by the current guess for $\hat{\beta}_{daily}$. This new dataset of simulated daily planting dates is then aggregated back to the monthly level, mimicking the coarseness of the original survey data. I re-run the benchmark regression on this simulated monthly data to obtain a new coefficient, $\beta_{simulated}$. The algorithm compares $\beta_{simulated}$ to the target β_{target} , adjusts the guess for $\hat{\beta}_{daily}$ based on the difference, and the loop repeats. The process continues until $\beta_{simulated}$ converges to β_{target} (i.e., the difference is within a pre-defined tolerance).

The final, converged value of $\hat{\beta}_{daily}$ is my estimate of the true, unbiased daily-level coefficient. This entire procedure is performed independently for both the Skew-Normal and KDE-based drawing pools.

Figure B1.12: Correcting for aggregation bias in planting dates



Note. Comparison of regression coefficients from observed, monthly-reported planting data with two simulated point estimates of the underlying true daily coefficient. The *Benchmark* is the coefficient from a regression of planting day on the week of rainy season onset using the coarse monthly data; the vertical line represents its 95% confidence interval. This benchmark coefficient serves as the target for the simulation. The *Skew Normal* and *Kernel Density Estimation* points are the results of an iterative search for the true daily-level coefficient that, after being coarsened to a monthly frequency, reproduces the benchmark estimate. The y-axis represents the number of days planting is delayed for a one-week delay in the rainy season onset. All regressions are based on equation (1.1). See Appendix B1.6 for further details on the simulation.

B1.7 The role of adjusting planting dates

As established in the main text, farmers adapt to a delayed rainy season by shifting their planting dates. While this low-cost adaptation strategy is effective—a finding supported by the instrumental variable (IV) analysis presented in this appendix—quantifying the magnitude of the effect is challenging. The primary difficulty is the endogeneity of a farmer’s planting decision, which biases simple OLS estimates. For instance, the OLS estimate in Column 1 of Table B1.18 suggests only a minimal mitigating effect.

To obtain a causal estimate, I therefore employ an instrumental variable (IV) strategy that leverages local planting norms. The core intuition is twofold. First, for the instrument to be relevant, it must predict a farmer’s own planting date. I argue that the planting decisions of farmers in a broad surrounding area, driven by shared agro-climatic conditions and cultural norms, will be highly correlated with an individual farmer’s timing. Second, for the instrument to be valid, it must satisfy the exclusion restriction—meaning it should not directly affect a plot’s yield other than through the farmer’s own planting choice. A potential threat to this restriction would be an unobserved, spatially correlated shock like a local weather event that affects both neighbours’ planting and the farmer’s own yields. To address this, the model conditions on weather outcomes measured at the farmer’s own high-resolution $0.1^\circ \times 0.1^\circ$ cell. Conditional on these local weather controls, the planting dates of neighbors should not have a direct impact on a farmer’s own productivity. Based on this logic, I construct a leave-one-out mean instrument. I first create an adjusted planting date for each plot by taking the residual based on equation (1.1). This regression includes location fixed effects at the $0.5^\circ \times 0.5^\circ$ grid-cell level and the benchmark local trend, which effectively de-means and de-trends the planting date. The instrument is then the average of these residuals for all other plots within a broad $1^\circ \times 1^\circ$ region, excluding those from the plot’s own immediate $0.1^\circ \times 0.1^\circ$ cell.

The results of this IV strategy are presented in Columns (2) and (3) of Table B1.18. The reduced-form estimate in Column (2) is positive and statistically significant, confirming that during a delayed onset, crop yields are higher in regions where neighbouring farmers also delayed their own planting. Column (3) presents the main two-stage least squares (2SLS) estimate. This estimate must be interpreted with caution, as the instrument is weak by conventional standards (significant only at the 10% level). With this caveat in mind, the coefficient provides our best available, though suggestive, estimate of the causal effect. The magnitude of the coefficients implies that a strategic delay in planting of approximately two weeks is sufficient to offset half of the negative productivity shock from a one-week delay in the rainy season. This greater than one-to-one tradeoff can be understood through two complementary factors. First, the estimated benefit of delaying planting is likely an underestimate due to measurement error (Section 1.6.3); the monthly reporting of planting dates attenuates the coefficient, which in turn inflates the calculated number of weeks needed for the trade-off. Second, the IV strategy

identifies the benefit of delaying planting relative to the local average. In a context where the rainy season is systematically shifting later, this local average itself may be sub-optimal. If the entire local planting calendar is lagging behind the shifting climate, the measured coefficient only captures the benefit of a marginal adjustment, not a shift to a truly optimal schedule.

Table B1.18: Shifting dates as a strategy

Dependent variable: <i>Estimator</i>	Log yields (2020 USD per hectare)		
	OLS (1)	OLS (reduced form) (2)	2SLS (3)
Onset of the rainy season	-0.026 (0.009) [0.003]	-0.027 (0.009) [0.003]	-0.026 (0.010) [0.009]
× Planting date (week of the year)	0.000 (0.000) [0.028]		0.006 (0.003) [0.055]
× Planting date in neigh. cells (week of the year)		0.002 (0.001) [0.031]	
Mean (dep. var.)	5.781	5.782	5.782
Identifying observations	118,933	116,058	116,058
Singleton observations	73	75	71
Countries	5	5	5
Interview year range	2009–2022	2009–2022	2009–2022

Note. OLS and 2SLS regression estimates based on Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. This table presents three specifications. Column (1) shows the OLS estimate of the interaction between onset week and the plot's own planting date. Column (2) presents the reduced-form OLS estimate, using the interaction with the instrument. Column (3) presents the two-stage least squares (2SLS) instrumental variable estimate. The instrument, *Planting date in neigh. cells*, is the leave-one-out mean of adjusted planting dates of other plots in the surrounding $1^\circ \times 1^\circ$ grid cell. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects, climatic area by country by year fixed effects, and a full set of weather and plot-level controls. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

B1.8 Additional results for nutritional outcomes

This section presents supplementary analyses for the nutritional outcomes discussed in Section 1.5.2. First, it presents the results for children. Second, it tests the sensitivity of the main findings for women to the specification of the location fixed effects. Finally, it examines the results for women on a broader sample including those outside agricultural households.

Table B1.19: Onset of the rainy season and children's nutrition

Dependent variable:	Weight-for-height (1)	Wasted (2)	Underweight (3)
Onset of the rainy season	0.002 (0.002) [0.444]	0.000 (0.000) [0.671]	-0.000 (0.001) [0.857]
Standardised effect	0.004	0.000	-0.000
Mean (dep. var.)	-0.424	0.079	0.291
Identifying observations	53,047	53,047	52,628
Singleton observations	14	14	14
Countries	6	6	6
Interview year range	1990–2018	1990–2018	1990–2018

Note. OLS regression estimates of Equation 1.1. The data are from the DHS (Croft et al., 2018) for children in agricultural households. The dependent variables are defined as follows. *Weight-for-height* is a z-score, representing the standard deviation from the median of the DHS reference population. *Wasted* is an indicator for a weight-for-height z-score below -2. *Underweight* is an indicator for a BMI below 18.5. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include weather controls, individual-level controls, climatic area by country by year fixed effects, and location fixed effects at the $0.5^\circ \times 0.5^\circ$ grid-cell level. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

Table B1.20: Table 1.3 – Finer location fixed effects

Dependent variable:	Weight-for-height (1)	Wasted (2)	Underweight (3)
Onset of the rainy season	-0.005 (0.006) [0.350]	0.000 (0.002) [0.924]	0.001 (0.002) [0.660]
Standardised effect	-0.005	0.000	0.001
Mean (dep. var.)	-0.850	0.110	0.142
Identifying observations	58,670	58,670	51,322
Singleton observations	360	360	426
Countries	6	6	6
Interview year range	1995–2018	1995–2018	1992–2022

Note. OLS regression estimates of Equation 1.1. The data are from the DHS (Croft et al., 2018) for women in agricultural households. The dependent variables are defined as follows. *Weight-for-height* is a z-score, representing the standard deviation from the median of the DHS reference population. *Wasted* is an indicator for a weight-for-height z-score below -2. *Underweight* is an indicator for a BMI below 18.5. *Onset of the rainy season* is the week of the year of the onset. All specifications include weather controls, individual-level controls, climatic area by country by year fixed effects, and location fixed effects at the $0.1^\circ \times 0.1^\circ$ grid-cell level. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

Table B1.21: Table 1.3 – All women


Dependent variable:	Weight-for-height (1)	Wasted (2)	Underweight (3)
Onset of the rainy season	-0.003 (0.001) [0.016]	0.000 (0.000) [0.528]	-0.000 (0.000) [0.902]
Standardised effect	-0.007	0.000	-0.000
Mean (dep. var.)	-0.666	0.098	0.151
Identifying observations	255,422	255,422	229,479
Singleton observations	0	0	0
Countries	6	6	6
Interview year range	1995–2018	1995–2018	1992–2022

Note. OLS regression estimates of Equation 1.1. The data are from the DHS (Croft et al., 2018). The dependent variables are defined as follows. *Weight-for-height* is a z-score, representing the standard deviation from the median of the DHS reference population. *Wasted* is an indicator for a weight-for-height z-score below -2. *Underweight* is an indicator for a BMI below 18.5. All specifications include weather controls, individual-level controls, climatic area by country by year fixed effects, and location fixed effects at the $0.1^\circ \times 0.1^\circ$ grid-cell level. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

B1.9 Stability of planting months

Figure B1.13 illustrates the stability of the median planting month over time at the $0.5^\circ \times 0.5^\circ$ level. Panel A plots the median planting month from the most recent survey wave against the median from the prior wave, which are separated by approximately three years on average. Panel B presents the same comparison but uses the earliest available survey wave, showing stability over an average period of approximately six years. In both panels, the points cluster tightly around the 45-degree line, demonstrating a high correlation between the median planting months across survey waves.

Figure B1.13: Stability of the median planting month



Note. Median planting month over time at the $0.5^\circ \times 0.5^\circ$ level. For locations observed in at least two survey waves, Panel A plots the median planting month from the most recent wave against the median from the prior wave. Panel B plots the most recent wave against the earliest available wave to show stability over a longer period. Points are jittered to show density. The 45-degree dashed line represents perfect stability, where the median planting month does not change between waves. The Pearson's correlation coefficient is reported in the top-left.

B1.10 False onsets: Robustness and alternative agronomical pathways

This section tests whether the main findings are sensitive to the parameters used to define a false onset. Table 1.6 re-estimates the analysis using alternative definitions that vary the length of the search window around the usual onset date and set a more stringent threshold for the dry spell.

Table B1.22: Table 1.6 – Different definitions of false onsets

Dependent variable:	Planting (DOY)	Seeds	Log yields (2020 USD per hectare)
	(1)	(2)	(3)
A. 70-day search window and 7-day dry spell			
Onset of the rainy season	0.430 (0.228) [0.060]	0.031 (0.012) [0.012]	-0.025 (0.007) [0.001]
× False onset	-0.989 (0.377) [0.009]	0.086 (0.054) [0.108]	-0.066 (0.028) [0.018]
Mean (dep. var.)	188.916	3.453	5.744
Identifying observations	125,588	120,650	129,897
Singleton observations	95	216	202
Countries	5	6	6
Planting year range	2009–2022	2009–2022	2008–2022
B. 60-day search window and 10-day dry spell			
Onset of the rainy season	0.450 (0.230) [0.051]	0.030 (0.012) [0.014]	-0.023 (0.007) [0.001]
× False onset	-0.714 (0.316) [0.024]	0.077 (0.018) [0.000]	-0.036 (0.011) [0.001]
Mean (dep. var.)	188.923	3.453	5.744
Identifying observations	125,598	120,658	129,907
Singleton observations	85	208	192
Countries	5	6	6
Planting year range	2009–2022	2009–2022	2008–2022
C. 70-day search window and 10-day dry spell			
Onset of the rainy season	0.495 (0.223) [0.027]	0.030 (0.012) [0.012]	-0.023 (0.007) [0.001]
× False onset	-0.702 (0.314) [0.026]	0.078 (0.018) [0.000]	-0.036 (0.011) [0.001]
Mean (dep. var.)	188.923	3.453	5.744
Identifying observations	125,598	120,658	129,905
Singleton observations	85	208	194
Countries	5	6	6
Planting year range	2009–2022	2009–2022	2008–2022

Note. OLS regression estimates. The dependent variable for each column is listed at the top. The analysis examines the interaction between the rainy season onset week and a dummy variable indicating a *False onset*. Benchmark fixed effects are interacted with the false onset dummy. All specifications include a full set of weather and plot-level controls. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

Table B1.23: False onsets: Alternative pathways

Dependent variable:	Fertiliser (1)	Pesticide (2)	Days from hired workers (3)
Onset of the rainy season	-0.000 (0.002) [0.862]	-0.001 (0.002) [0.573]	-6.801 (4.498) [0.131]
× False onset	-0.071 (0.034) [0.040]	-0.015 (0.028) [0.580]	8.740 (5.504) [0.113]
Mean (dep. var.)	0.516	0.066	102.816
Identifying observations	135,059	135,284	117,139
Singleton observations	195	195	149
Countries	6	6	5
Planting year range	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates based on Equation (1.1). *Fertiliser* is an indicator variable taking the value 1 if any fertiliser, whether organic or inorganic, was used, and 0 otherwise. *Pesticide* is an indicator variable that equals 1 if any pesticide was used and 0 otherwise. *Days from hired workers*, the total number of days per hectare hired workers worked on the plot. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. The analysis examines the interaction between the rainy season onset week and a dummy variable indicating a *False onset*. All specifications include the main effect of the false onset dummy (absorbed) and fully flexible fixed effects (all benchmark fixed effects are interacted with the false onset dummy). All specifications include a full set of weather and plot-level controls. Standard errors, clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level, are in parentheses. *p*-values are in brackets. Appendix A1.1 provides further details.

B1.11 Additional results on heterogeneity

Table B1.24 presents additional results on heterogeneous impacts on productivity, extending the analysis in Section 1.5.4. Table B1.25, instead, presents results on heterogeneous impacts on planting dates. Table B1.26 presents results on heterogeneous impacts on other adaptation measures, such as the use of seeds and fertilisers.

Table B1.24: Additional results on heterogeneous impacts on productivity

Dependent variable:	Log yields (2020 USD per hectare)				
	(1)	(2)	(3)	(4)	(5)
Onset of the rainy season	-0.021 (0.007) [0.002]	-0.021 (0.007) [0.002]	-0.020 (0.007) [0.004]	-0.024 (0.007) [0.001]	-0.029 (0.007) [0.000]
× Manager has worked (past 12 m)	0.003 (0.002) [0.203]				
× Household head has worked (past 12m)		0.002 (0.002) [0.231]			
× Household head is female			-0.006 (0.001) [0.000]		
× Household head is educated				0.006 (0.002) [0.002]	
× Sold produce					0.015 (0.001) [0.000]
Mean (dep. var.)	5.749	5.742	5.743	5.743	5.743
Identifying observations	128,186	129,611	129,670	129,612	129,612
Singleton observations	177	179	180	179	179
Countries	6	6	6	6	6
Planting year range	2008–2022	2008–2022	2008–2022	2008–2022	2008–2022

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B1.25: Heterogeneous effects on planting dates

Dependent variable:	Planting DOY				
	(1)	(2)	(3)	(4)	(5)
Onset of the rainy season	0.542 (0.223) [0.015]	0.574 (0.224) [0.011]	0.555 (0.227) [0.015]	0.600 (0.215) [0.005]	0.736 (0.276) [0.008]
× Manager is female	0.093 (0.083) [0.264]				
× Manager has formal education		-0.034 (0.060) [0.570]			
× Below median assets			0.008 (0.069) [0.902]		
× Irrigated plot				-0.832 (0.430) [0.054]	
× Improved seeds					0.361 (0.122) [0.003]
Mean (dep. var.)	188.934	188.934	188.968	188.753	179.593
Identifying observations	125,609	125,609	125,533	124,531	108,328
Singleton observations	74	74	74	74	99
Countries	5	5	5	5	5
Planting year range	2009–2022	2009–2022	2009–2022	2009–2022	2011–2022

Note. OLS regression estimates of Equation 1.1. *Planting (DOY)* is the approximated first planting day of the year. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B1.26: Heterogeneous impacts on other adaptation measures

Dependent variable:	Seeds		Fertiliser		Pesticide	
	(1)	(2)	(3)	(4)	(5)	(6)
Onset of the rainy season	0.029 (0.011) [0.011]	0.027 (0.011) [0.015]	-0.001 (0.002) [0.605]	-0.000 (0.002) [0.960]	-0.001 (0.002) [0.462]	0.000 (0.002) [0.895]
× Manager has formal education	0.007 (0.003) [0.011]		0.000 (0.001) [0.553]		0.001 (0.000) [0.000]	
× Below median assets		0.004 (0.003) [0.140]		-0.001 (0.000) [0.002]		-0.001 (0.000) [0.000]
Mean (dep. var.)	3.453	3.453	0.516	0.516	0.066	0.066
Identifying observations	120,558	120,511	135,081	135,007	135,306	135,231
Singleton observations	197	197	173	173	173	173
Countries	6	6	6	6	6	6
Planting year range	2009– 2022	2009– 2022	2008– 2022	2008– 2022	2008– 2022	2008– 2022

Note. OLS regression estimates of Equation 1.1. *Seeds* is the log of seed quantity used on the plot, measured in kilograms per hectare. *Fertiliser* is an indicator variable taking the value 1 if any fertiliser, whether organic or inorganic, was used, and 0 otherwise. *Pesticide* is an indicator variable that equals 1 if any pesticide was used and 0 otherwise. *Days from hired workers* is the total number of days per hectare hired workers worked on the plot. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B1.27: Heterogeneous impacts by information access

Dependent variable:		Log yields (2020 USD per hectare)		
Indicator:	Owns mobile phone	Owns radio	Owns television	
	(1)	(2)	(3)	
Onset of the rainy season	-0.021 (0.007) [0.002]	-0.020 (0.007) [0.004]	-0.021 (0.007) [0.004]	
× Indicator	0.004 (0.002) [0.039]	0.003 (0.002) [0.104]	0.000 (0.004) [0.921]	
Mean (dep. var.)	5.769	5.750	5.710	
Identifying observations	118,329	114,985	110,279	
Singleton observations	179	181	182	
Countries	6	6	6	
Planting year range	2008–2022	2008–2022	2008–2022	

Note. OLS regression estimates of Equation 1.1. The dependent variable is the log of yield, measured in 2020 US dollars per hectare. *Onset of the rainy season* is the week of the year when conditions specified in Section 1.3 are met. The regressions are run on the full sample, but the identification of interaction terms involving *Later waves* relies solely on variation within Ethiopia, Malawi, and Nigeria. The *Indicator* variables in each column are binary: (1) *Owns mobile phone* equals 1 if the household owns at least one mobile phone; (2) *Owns radio* equals 1 if the household owns at least one radio; (3) *Owns television* equals 1 if the household owns at least one television set. Each specification interacts the onset week with the respective indicator. All specifications include the main effect of the *Indicator* and control for an indicator variable equal to 1 if the household's asset index is below the sample median. All specifications include $0.1^\circ \times 0.1^\circ$ grid-cell fixed effects and climatic area by country by year fixed effects. All specifications also include weather and plot-level controls; a full list of these controls is presented in Section 1.3. Standard errors clustered at the $0.5^\circ \times 0.5^\circ$ grid-cell level are reported in parentheses. *p*-values are reported in brackets. Appendix A1.1 provides detailed information on variables, selected surveys, and weighting procedures.

Appendix C: Methodology for Projecting Economic Damages

This appendix details the methodology used to calculate the Net Present Value (NPV) of future economic damages from delayed rainy season onsets, as presented in the main text. The calculation combines empirical estimates from this paper with country-specific projections of climate and economic variables.


Annual GDP loss. The annual loss in GDP for a given country c in year t is calculated in absolute monetary terms (constant 2015 USD). It is the product of the total agricultural yield loss and the projected real value of agricultural GDP. The full calculation for the annual loss is:

$$\text{Annual Loss}_{ct} = \underbrace{(D_{ct} \times \beta_{\text{yield}})}_{\text{Gross Yield Loss}} \times \underbrace{(1 - I_{ct})}_{\text{Irrigation Adjustment}} \times \underbrace{(A_{ct} \times G_{ct})}_{\text{Real Agri. Value Added}} \quad (1.2)$$

In this equation, D_t represents the cumulative delay in the rainy season onset in weeks, and β_{yield} is the semi-elasticity of crop yield to this delay (a 2.11% yield loss per week). This gross loss is adjusted by I_t , the projected share of irrigated agricultural land. To forecast I_t through 2050, several functional forms (linear, quadratic, exponential, logistic, Gompertz) are fitted to the historical trend from the HYDE 3.2 dataset (Klein Goldewijk et al., 2017). The model with the best fit for each country, determined by the lowest Bayesian Information Criterion (BIC), is used to extrapolate the trend.

The final two components, A_t and G_t , represent the projected real share of agriculture in the economy and the projected real GDP (in constant 2015 USD) for year t . The historical portion of these series (up to 2024) is constructed using World Bank indicators. To derive the historical real agricultural share (A_t), an implicit GDP deflator is calculated by dividing nominal GDP by real GDP. The real agricultural share is then the nominal agricultural value added divided by this deflator, all divided by real GDP. To address missing data for Ethiopia's nominal agricultural value for 2023 and 2024, values are imputed by applying the country's 10-year average growth rate from 2013-2022. Finally, to project both A_t and G_t forward from 2025 to 2050, the annual growth rates implied by the FAO's long-term agricultural and macroeconomic scenarios are applied to the final historical values.

Figure C1.1: Historical and projected share of irrigated cropland



Note. The figure shows the historical and projected share of total cropland that is irrigated for each country. The historical data are from the HYDE 3.2 dataset (Klein Goldewijk et al., 2017). The projection from 2018 to 2050 is a forecast based on the best-fitting time series model (selected via BIC) for each country's historical trend.

Scenario framework. To quantify the economic benefits of climate mitigation, the analysis compares two distinct future pathways from 2024 to 2050. These are adapted from the Food and Agriculture Organization's (FAO) long-term global scenarios (FAO, 2018), which are built upon the Shared Socio-economic Pathways (SSPs) framework, formulated by the Intergovernmental Panel on Climate Change Sixth Assessment Report (IPCC, 2021). The scenarios differ in their assumptions for future economic development and for the magnitude of the climate shock itself.

1. **Business as Usual (BAS):** This scenario is adapted from the IPCC's middle of the road pathway (corresponding to SSP2), which assumes that socio-economic and technological trends continue along historical patterns with uneven progress toward sustainability goals. For this scenario, I assume the delay in the onset of the rainy season is based on the projected delay shown in Figure B1.1.
2. **Toward Sustainability (TS):** This scenario is adapted from the IPCC's sustainability pathway (corresponding to SSP1), which assumes a global shift toward more equitable development and respect for environmental boundaries. To reflect the success of these mitigation efforts, the expected delay in the rainy season onset is assumed to be halved relative to the BAS scenario.

In Equation 1.2, these scenarios determine the future values for the cumulative delay (D_t), the agricultural share (A_t), and real GDP (G_t), allowing for a direct comparison of their economic consequences.

Social discount rate and NPV. To calculate the Net Present Value (NPV) of this stream of future damages, annual losses must be discounted to the present day. The social discount rate, r_t , is not assumed to be constant but is derived from the DICE-2023 model's Ramsey-rule framework, which accounts for time-varying consumption growth and climate-related capital risks. The formula for the discount rate in year t is:


$$r_t = \rho + \phi g_t - \frac{1}{2} \phi^2 \sigma_c^2 t + \beta_{clim} \pi_{capital} \quad (1.3)$$

The specific parameter values from the DICE-2023 framework are as follows: the rate of pure time preference $\rho = 0.001$; the elasticity of the marginal utility of consumption $\phi = 0.95$; the variance of the logarithm of consumption growth $\sigma_c^2 = (0.01)^2$; a parameter for the systematic risk of climate damages $\beta_{clim} = 0.5$; and the risk premium on capital $\pi_{capital} = 0.05$. The model also incorporates a time-varying growth rate of per capita consumption, g_t , which is assumed to decline linearly from 1.9% to 1.7% over the projection period.

The NPV of damages is then the sum of all discounted annual GDP losses from the start year ($t_0 = 2024$) to the end year ($T = 2050$):

$$NPV = \sum_{t=t_0}^T \frac{\text{Annual Loss}_t}{\prod_{i=t_0}^t (1 + r_i)} \quad (1.4)$$


Figure C1.2: The mitigating effect of irrigation expansion



Note. Benefit of future irrigation expansion under the Business as Usual scenario. The dashed line shows an optimistic counterfactual where the share of irrigated land increases linearly to 20% of agricultural land in all sampled countries. The solid line shows the main projection, which includes a data-driven forecast of irrigation expansion based on historical trends. The shaded area between the lines represents the irrigation benefit—the total value of damages averted due to this specific form of adaptation, quantified in millions of US dollars.

Uncertainty analysis. To provide a sense of the uncertainty surrounding these projections, two sensitivity analyses are conducted. First, uncertainty in the damage coefficient is assessed by recalculating the NPV of damages using the lower and upper bounds of the 95% confidence interval from the benchmark regression. Second, to model uncertainty in future adaptation, an optimistic scenario is run where the damage coefficient, β_{yield} , is assumed to decrease linearly to zero between 2024 and 2050.

Figure C1.3: Sensitivity of future damages to coefficient uncertainty



Note. Statistical uncertainty of the damage projection under the Business as Usual scenario. The solid line is the central projection, calculated using the point estimate of a 2% yield loss per week of onset delay. The shaded area represents the 95% confidence interval, with the upper and lower bounds calculated by re-running the projection using the corresponding bounds of the yield loss coefficient from the benchmark regression in Table 1.1.

Table C1.1: Uncertainty around estimates

Country	Lower Bound (95% CI)	Point Estimate	Upper Bound (95% CI)
Ethiopia	2.53%	6.83%	11.12%
Malawi	1.99%	5.37%	8.75%
Mali	3.70%	9.96%	16.23%
Niger	1.42%	3.81%	6.21%
Nigeria	2.62%	7.06%	11.50%
Tanzania	0.75%	2.02%	3.29%

Note. The table shows the projected NPV of damages by 2050 as a percentage of 2024 GDP under the Business as Usual scenario. Each column uses a different value for the yield loss per week of onset delay, corresponding to the point estimate and the 95% confidence interval from the benchmark regression.

A manuscript based on the following two chapters is conditionally accepted
for publication at the *Journal of the European Economic Association*.

Chapter 2

From Sea to Shore: The Impact of Ocean Acidification on Child Health

This chapter is co-authored with Alex Armand (NOVA SBE)

Abstract

Since the Industrial Revolution, ocean water acidity has risen by 26% due to anthropogenic emissions—a process known as *ocean acidification*—posing a risk for marine life and the communities depending on it. This paper examines the consequences of ocean acidification for child health, using data from coastal regions in 36 low- and middle-income countries from 1972 to 2018, encompassing 41% of the world’s coastal population. Leveraging short-term exogenous shifts in ocean acidity near human settlements for identification, we find that prenatal exposure to higher water acidity significantly raises the risk of death in the first months of life and impacts early childhood development. We show evidence consistent with these effects being associated with maternal malnutrition, as increased acidity reduces catches for small-scale fisheries, increasing seafood prices and reducing consumption of crucial nutrients. Our findings indicate limited adaptation to these impacts.

Keywords: Climate Change; Ocean; Acidification; Health; Mortality; Development

JEL Codes: I15, Q20, Q54, O10

2.1 Introduction

Since the Industrial Revolution, ocean water acidity has risen by 26%—a phenomenon known as *ocean acidification* (Doney et al., 2020). This outcome originates from the ocean’s key role in regulating the climate by absorbing carbon dioxide from the atmosphere. Due to increased emissions from human activities, the amount absorbed has surged over the past two centuries, disrupting the ocean’s chemical balance. The resulting increased acidity affects marine life and is likely endangering human communities whose economic development depends on marine resources (Dalgaard et al., 2020; IPCC, 2022b). Nevertheless, empirical evidence on these consequences remains limited, with most findings derived from economic or ecological models projecting the hypothetical effects of ocean acidification (see, e.g., Brander et al., 2012; Colt and Knapp, 2016).

This paper investigates the consequences of ocean acidification for child health, focusing on the region where the impact is expected to be most pronounced—in low- and middle-income countries (L&MICs). These countries are home to most of the over 3 billion people worldwide who depend on ocean-harvested resources to survive (FAO, 2022). In their coastal regions, seafood consumption forms a significant part of nutritional intake, predominantly supported by small-scale fisheries operating near human settlements (World Bank, 2012). This form of dependency is particularly vulnerable to local disruptions to marine resources, such as those caused by acidification, with children being particularly at risk (see Section 2.2 for a detailed discussion of these processes).

We exploit local dependency on marine resources to analyse the effects of ocean acidification in the coastal areas of 36 L&MICs, spanning large regions of Africa, Asia, and Latin America from 1972 to 2018. In 2020, these countries accounted for 39% of the world’s population, rising to 41% when focusing on coastal areas (United Nations, 2024). Our approach compares individuals and communities over time by matching their locations to temporal variation in the acidity of the nearest waters, as measured by pH—a logarithmic scale where lower values indicate the higher acidity of a solution. For identification, we focus on short-term exogenous changes in pH at these locations. In the short run, and similarly to weather patterns, pH in a specific point in the ocean deviates exogenously around a long-run (decreasing) trend, with waters being relatively more (or less) acidic.¹ Exploiting this property, we capture deviations using a linear framework accounting for unobserved heterogeneity through multi-way fixed effects (FEs). We support this approach with several checks described in Sections 2.4 and 2.5.

Following this approach, we first establish a link between ocean acidity and child health. Comparing marine catches in coastal areas over time, we show that higher acidity reduces the quan-

¹Ocean pH is affected by wind, temperature, sea ice, precipitation, runoff, and ocean circulation (Feely et al., 2008). As such, its local variation is analogous to weather patterns (Appendix B2.1), whose short-run deviations have been used in the literature to identify climatic shocks (Dell et al., 2014).

tity and value of seafood caught by small-scale fisheries, suggesting a net negative effect of water acidity on species harvested in proximity to coastal settlements. At the same time, higher acidity does not impact the catch of commercial fisheries or the overall economic activity in the coastal area. Our findings indicate that the negative shocks to small-scale fisheries increase local prices for seafood, translating into a diminished probability of seafood consumption and increased malnutrition for local populations, particularly affecting pregnant women. Maternal malnutrition is a critical risk factor for children's health, and can lead to developmental deficits and even to death ([Victora et al., 2021](#)).

Motivated by this link, we study the consequences of early-life exposure to varying degrees of ocean acidity. We use data from up to 1.5 million live births from 1972 to 2018, leveraging individuals' geolocation and month and year of birth to compute exposure. Comparing children (or siblings) born in the same location, but at different points in time, we reveal that experiencing higher acidity in nearby waters while *in utero* leads to increased mortality in the first months of life. A decrease of 0.01 pH units (equivalent to a 2.3% increase in acidity) raises neonatal mortality (death during the first month of life) by approximately 1.5 deaths per 1,000 live births in coastal communities. This effect represents a 5% increase relative to the sample average of neonatal mortality in the period of study. The largest impacts are observed in areas with a higher dependence on seafood and in regions where forms of fishing that deplete marine biodiversity are prevalent, in line with evidence on the cascading effects of over-exploitation (see, e.g., [Frank and Wilcove, 2019](#)). Experiencing higher acidity while *in utero* has instead no effect on mortality beyond the first months of life.

The increase in neonatal mortality following ocean acidification confirms a mechanism rooted primarily in changes in maternal health. However, investments in maternal and child health are unaffected, suggesting the lack of adaptation. Our results indicate that differential access to medical care and nutrient supplementation, as well as behavioural changes that could occur after parents observe their child's health, are unlikely to be at play.² The results also confirm the absence of significant income changes at the household level, as such changes would typically prompt adjustments in investment (see, e.g., [Baird et al., 2011](#)).

Finally, exposure to ocean acidity *in utero* not only impacts mortality rates, but also influences infants' development. Objective anthropometric measurements indicate that mortality due to prenatal exposure to ocean acidity is more common among children who would have otherwise had worse health outcomes. On average, children who survive past their first month of life have slightly better health indicators and thus represent a positively selected group. However, when examining gendered effects, we observe a significant increase in stunting among female children, which outweighs this positive selection. We present evidence consistent with these

²We lack direct measures of maternal stress. However, evidence indicates that maternal stress plays a significant role in response to traumatic events (see, e.g., [Aizer et al., 2016](#); [Persson and Rossin-Slater, 2018](#); [Menclova and Stillman, 2020](#); [Berthelon et al., 2021](#)), as opposed to the relatively mild changes in maternal health that we examine in later sections.

negative effects lasting into women's adulthood, accompanied by worse economic outcomes.

Our findings contribute to different strands of the literature. First, we further our understanding of the current effects of climate change. By providing evidence on the consequences for coastal communities of altering ocean acidity, our results contribute to the emerging literature on the impacts of ocean acidification on human communities (see, e.g., [Colt and Knapp, 2016](#)), an essential dimension to integrate into the general literature measuring the costs of climate change (see, e.g., [Auffhammer, 2018](#)). While many studies have explored the impact of climate change on human behaviour—focusing on issues like rising temperatures and shifting rainfall patterns ([Deschênes et al., 2009](#); [Deschênes and Moretti, 2009](#); [Barreca et al., 2016](#))—the ocean has received far less attention. In addition, due to the open-access nature of ocean resources, changes in their productivity are not comparable to changes in land and agricultural productivity (see, e.g., [Collier, 2010](#)), which have been more extensively studied in relation to climate change.

Second, we provide new evidence regarding the early stages of children's development. Numerous studies have examined shocks that are either directly observable or have direct effects on health, thus leading to adaptation or avoidance behaviour ([Almond et al., 2018](#) provides a review of this literature).³ Ocean pH is not directly observed or felt by individuals (e.g., it is not reported in weather forecasts or newspapers), it has no direct effect on health, and public awareness about its changing nature is highly limited.⁴ Our results suggest that shocks with these characteristics—being largely imperceptible, lacking direct health impacts, and receiving minimal public attention—generate a limited behavioural response, contrasting with evidence on adaptation following nutritional shocks such as famines and prolonged fasting ([Razzaque et al., 1990](#); [Almond and Mazumder, 2011](#); [Majid, 2015](#)), or nutrient supplementation ([Adhvaryu and Nyshadham, 2016](#)). In our setting, parents are likely either unaware of the consequences of altering their diets, or their health is not sufficiently affected to induce behavioural changes. These findings align with evidence on the limited knowledge of dietary consequences in poorer settings (see, e.g., [Hirvonen et al., 2017](#)) and with the high prevalence of micronutrient deficiency in L&MICs ([Lowe, 2021](#)). Such deficiency can occur without deficits in caloric intake, making detection difficult for non-experts.

Finally, this study provides new evidence on the importance of wildlife for human and economic development ([Michalopoulos and Papaioannou, 2013](#); [Bowles and Choi, 2019](#); [Dalgaard et al., 2020](#); [Mayshar et al., 2022](#)). Our findings complement the recent literature on biodiversity and poverty (see, e.g., [Dasgupta, 2021](#)) and on the role of overly exploitative practices

³Studies related to our setting cover atmospheric events ([Maccini and Yang, 2009](#); [Heft-Neal et al., 2018](#); [Geruso and Spears, 2018a](#); [Adhvaryu et al., 2024](#)) and environmental contamination or degradation ([Chay and Greenstone, 2003](#); [Arceo et al., 2016](#); [Isen et al., 2017](#); [Geruso and Spears, 2018b](#); [Black et al., 2019](#); [Berazneva and Byker, 2022](#)). On avoidance behaviour, see [Moretti and Neidell \(2011\)](#) on air pollution.

⁴The tangible effects of ocean pH could be discerned from seafood catch. However, awareness among the general public of a link between water acidity and fisheries is severely limited. Surveys conducted in richer countries suggest that only a tiny segment of the population is aware of this influence ([Gelcich et al., 2014](#); [Lotze et al., 2018](#)).

such as deforestation ([Burgess et al., 2012](#); [Jayachandran, 2013](#)), overfishing ([Stavins, 2011](#)), and poaching ([Kremer and Morcom, 2000](#)). Closely related to our study is the work of [Frank and Sudarshan \(2023\)](#), which highlights the role of the functional extinction of vultures in India on human mortality, and that of [Feir et al. \(2023\)](#), which shows how the loss of the bison in North America led to persistent negative consequences for Native Americans.

2.2 Background

This section summarises how ocean acidity can impact marine life and human behaviour, while we test these channels in Section [2.5.1](#). Concerning marine life, the scientific literature highlights two channels.

First, water acidity directly affects the physiology of marine species: at lower pH levels, many organisms must invest additional energy to maintain their metabolic processes and biological functions, with consequences for their survival, growth, development, and reproduction ([Gattuso and Hansson, 2011](#)). Although these effects are heterogeneous across and within species, biological effects are generally large and negative ([Kroeker et al., 2010](#); [Alter et al., 2024](#)). Laboratory experiments indicate negative responses to acidity in approximately 50% of species tested ([Wittmann and Pörtner, 2013](#)).

Second, water acidity degrades key marine habitats such as coral reefs and macroalgal forests. These areas serve as crucial feeding grounds for fish, making them important catchment areas for subsistence and artisanal fisheries ([Doney et al., 2020](#)). Degradation of these ecosystems disrupts marine food chains, not only reinforcing the direct effects of pH on the physiology of marine species, but also altering competition for food across species ([Sunday et al., 2017](#)). Scientific evidence indicates that commonly-consumed species, which typically boast better nutritional content, are more vulnerable to these effects, particularly in the presence of overfishing ([Jones and Cheung, 2018](#); [Hicks et al., 2019](#); [Maire et al., 2021](#)).

The overall net effect of these two channels on marine life is uncertain ([IPCC, 2022b](#)). Some species might benefit from the consequences of acidic conditions, resulting in compositional changes (i.e., an increase in the quantity or quality of some species and a reduction in others) rather than reduced seafood stock.⁵ Nevertheless, any change in species occurrence can impact harvest composition and, consequently, human nutrition and health. The reduced availability of commonly harvested species could result in income deterioration ([Colt and Knapp, 2016](#)), potentially reducing investments in health and nutrition among those relying on fisheries as a source of income, and in dietary shifts driven by price adjustments or changes in seafood quality.⁶

⁵Appendix [B2.1](#) discusses variables with a more direct effect on quantity, such as rising global sea surface temperatures ([Keeling et al., 2010](#)).

⁶The biological changes to marine life induced by water acidity are expected to result in reduced protein intake

The resulting effects on nutrition are expected to be more pronounced in vulnerable coastal and island communities in L&MICs. In these countries, seafood is a crucial source of nutrients, providing 26% of all consumed animal proteins, as compared to the global average of 17% (FAO, 2022). Countries such as Bangladesh, Cambodia, the Gambia, Ghana, Indonesia, Sierra Leone, and Sri Lanka reach peaks of at least 50% reliance on seafood for protein. In addition, most of this nutrient intake is supported by small-scale fisheries, definitionally more vulnerable to local shocks as compared to larger fisheries (FAO, 2023). L&MICs host 97% of all workers employed in marine capture, and more than 90% of them are engaged in small-scale fisheries supporting local consumption (see, e.g., Simmance et al., 2022). The ability to balance changes in local supply with imports is limited because L&MICs tend to export higher-quality seafood caught in their waters and supplement local demand with imports of lower-quality fish (McCauley et al., 2018).

Children, particularly, are at risk. Because nutritional alternatives and access to micronutrient supplementation are limited, seafood is recognised as an important source of macronutrients, such as proteins, and micronutrients, such as iron, iodine, zinc, vitamin A, vitamin D, vitamin B12, calcium, and essential fatty acids (United Nations, 2021). These nutrients are essential for maternal health and for fetal and child development.⁷ Early-life deficiencies in these nutrients can lead to severe health consequences, even including death (Victora et al., 2021).

2.3 Data

We collate myriad data sources. Appendices A2.1–A2.2 further detail the variables and data sources. Figure 2.1 shows the geographical coverage of the study.

Mortality, human capital, and adaptive behaviour. We collate and homogenise 92 household surveys from 36 countries collected by the Demographic and Health Surveys (DHS) from 1990 to 2018. Individual surveys provide nationally representative data on health and population in L&MICs, with a particular focus on maternal and child health, and they have been widely used to calculate mortality rates among children. The dataset is supplemented with objective measurements of human development and nutrition, such as height, weight, and haemoglobin concentration in blood samples. The programme surveys women aged 15–49 and includes information about their demographics, including wealth and human capital accumulation. Each surveyed woman’s birth history is recorded and includes information on their children’s year and month of birth, sex, birth order, whether they are twins, and the date of death when applicable. We assume measurement error related to mortality is minimal, as the timing of a child’s

and compromised seafood quality (Falkenberg et al., 2020).

⁷*Iron* and *iodine* support brain development and help prevent stillbirth. *Zinc* and *vitamin A* promote childhood survival and growth. *Calcium* and *vitamin D* reduce the risk of preterm delivery, while *vitamin B12* is vital for a healthy pregnancy and the development of the nervous system and brain in children. *Essential fatty acids* help prevent preeclampsia, preterm delivery, and low birth weight, and support children’s cognitive development.

death, being a tragic event, is unlikely to be forgotten.⁸

The primary sampling unit is a community (or cluster), representing a village or neighbourhood. Our dataset includes all available surveys with geographical coordinates and considers only countries with direct access to the ocean. We use all available surveys and re-weight observations to correct for oversampling of countries surveyed multiple times. Appendix A2.1 provides the full list of countries and surveys. Results are robust to different selection criteria. For questions omitted in specific survey rounds, we re-compute weights to account for this selection.

We restrict the sample to coastal areas. Using geolocation for communities, we compute the minimum straight distance to the shoreline, and following United Nations (2003), we define a *coastal area* as the buffer extending landward from the ocean’s shore up to a distance of 100 km. Individual characteristics tend to be comparable in magnitude between communities in the coastal and inland areas, but households in proximity to the ocean are slightly richer and exhibit lower mortality rates (Appendix A2.2).

Ocean acidity. We capture this chemical feature of the ocean using water pH—a logarithmic scale that indicates the acidity (basicity) of an aqueous solution at lower (higher) values—measured at the surface. Measurement of ocean pH using satellites for our time-space dimension is currently unavailable (Land et al., 2019). We therefore obtain the chemical features of the ocean from the Hadley Global Environment Model 2 – Earth System or HadGEM2–ES, developed by Collins et al. (2011) and Jones et al. (2011), which combines historical climate data, physical equations, and simulations to reconstruct past climate conditions, extrapolating from observations where observed data are incomplete.⁹ Section 2.5 discusses estimates controlling for confounders that could influence the pH measurement, such as anthropogenic waste.

Data are monthly global raster data at the $1^\circ \times 1^\circ$ resolution for the period 1972–2018.¹⁰ We match raster points to communities or coastal areas using a proximity criterion (see Section 2.4 for details about the matching procedure). In Appendix B2.1, summary statistics for matched points highlight that pH varies locally both within and across years around its long-run level, showing a high similarity to weather systems. Variation in pH originates from both the time

⁸The history of terminated pregnancies (i.e., pregnancies that did not result in a live birth regardless of the cause) is not recorded. Appendix B2.3 shows that ocean acidity in nearby waters does not affect the probability of experiencing a terminated pregnancy and that results are robust to accounting for recall bias in the reporting of a child’s death.

⁹Although the series matches well the available information from observational data of ocean features (see, e.g., Totterdell, 2019), our estimates might suffer from measurement error due to extrapolation. The possibility that such measurement error is correlated with unobservable determinants of local development or health is minimised by the climatology-based framework of the data (i.e., by construction, data cannot capture pH at the coast, where it may be influenced by local human activity). While there may be unobserved factors that influence both the outcome variables analysed in the paper and pH, we interpret our results assuming measurement error is uncorrelated with both. Regarding the use of climatology-based data in economics, see, e.g., Carleton et al. (2022); Adhvaryu et al. (2024); Matranga (2024).

¹⁰Data in this format was provided by the European Space Agency (ESA) Pathfinders–OA project.

and geographic dimensions, with comparable contributions from its between and within components.

We supplement our data with other variables that could jointly determine ocean acidity and outcomes in the coastal areas. First, we gather information about additional chemical features of the ocean with the same source used for pH. Second, we draw on the ERA5 database to supplement the data with other meteorological features in the same location in the ocean where pH is measured, including temperature and wind speed. Third, to control for weather characteristics inland, we include yearly rainfall and temperature data at the community level from the PRIO-GRID database.

Ocean exploitation. Data about seafood catch at the temporal and geographical resolution of the DHS data is unavailable, forcing us to work with data aggregated at the country level or for a restricted temporal horizon.

First, we gather data about quantity, price, and landed value of catches within the exclusive economic zones (EEZs) of each country in our sample from the Sea Around Us initiative (Pauly et al., 2020). An EEZ is a sea zone prescribed by the United Nations Convention on the Law of the Sea (UNCLOS) extending up to 200 nautical miles from the coast. Within this zone, the coastal state has special rights over the exploration and use of marine resources. Figure 2.1 shows the geographical coverage of these zones.

Data are disaggregated at the levels of fishing sector (industrial or small-scale), seafood group (defined by 11 commercial categories), and destination of use (direct human consumption or other uses). Quantities are reported in kilotons (kt), while landed values are computed using nominal ex-vessel prices in local currency and converted into 2010 US\$ real equivalents. Using ex-vessel species-level prices, we build yearly median prices for each commercial group.

We include catches whose destination of use is direct human consumption and distinguish between two sectors of activities: the *industrial sector* is large-scale and commercial, and it includes catch from large motorised vessels that is overwhelmingly sold commercially; the *small-scale fishing sector* comprises artisanal and subsistence fishing and includes vessels that primarily supply local consumption.¹¹ The resulting dataset is a yearly time series at the seafood group–country level covering the full period under analysis (1972–2018). In our selected group of countries, both quantities and landed values have been steadily increasing for both small-scale and industrial fishing (Appendix B2.4).

Second, for heterogeneity analyses, we supplement these data with more geographically granular data about exploitation intensity and type, but restricted to a shorter time period due to data availability. We consider a form of *extractive fishing* by focusing on industrial fishing. The

¹¹Catch from small-scale fisheries is notoriously under-reported in national statistics (Pauly and Zeller, 2016). The Sea Around Us explicitly accounts for this limitation, and it applies under-reporting adjustments to better reflect actual catches. Any measurement error related to under-reporting is not expected to correlate with local changes in ocean acidity.

Global Fishing Watch dataset provides us data on the hours industrial fishing vessels spend at specific geolocations. Because data are available only for 2012–2016, we build a global grid at the $1^\circ \times 1^\circ$ resolution summing fishing hours within each cell over the available period.

In addition, we define *night-time fishing* using the Automatic Boat Identification System for VIIRS Low Light Imaging Data (Elvidge et al., 2015). This system provides the time and geolocation of boats using nightlight as measured from satellite imaging. Because only 16% of fishing detected with this algorithm is also captured by industrial fishing (Kroodsma et al., 2018), night-time fishing tends to capture boats operating on a smaller and local scale, thus potentially contributing to the local economy. As with the measure of extractive fishing, we build a global grid at the $1^\circ \times 1^\circ$ resolution with the sum of all detected boats for the period in which data are available (2017–2019).

We normalise intensity from both activities to be between 0 (no presence) and 1 (high intensity). These measures aim to capture longer-run fishing patterns by averaging daily information over the full period in which data are available. However, if patterns of fishing were very different in the past, we would be capturing heterogeneity specific to exploitation in the period for which we have data availability. At least for extractive fishing, which tends to have low sensitivity to economic and environmental variation (Kroodsma et al., 2018), time-invariant heterogeneity is likely capturing suitability for industrial fishing. Appendix B2.4 shows that fishing patterns are primarily driven by differences in geography, while individual characteristics are comparable in areas with high versus low intensities of both types of fishing. The intensity of night-time fishing is comparable in areas with varying intensities of extractive fishing. Dependency on fish for nutrition is also highly stable over time.

Economic activity. We complement the data with the average night-time light emission from the calibrated DMSP-OLS Night-time Lights Time Series 4. Night-time luminosity measured by satellite images is a widely used proxy for economic activity and human development (see, e.g., Bruederle and Hodler, 2018). Yearly data are available for the period of 1992–2012. We normalise luminosity by population in the grid cell using the PRIO-GRID database, performing the analysis using night-time luminosity in a gridded dataset at the $0.5^\circ \times 0.5^\circ$ resolution, selecting only grid cells in the coastal area of sampled countries.

2.4 Empirical strategy

To estimate the impact of ocean acidification, we exploit temporal and geographical variation in ocean pH to compare communities as they face varying degrees of ocean acidity near their locations. We match communities with ocean pH using the nearest data point in the ocean. This point is likely the fishing ground of small-scale fisheries based in or near the community. Available evidence highlights that 84% of all small-scale fisheries operate within 20 km from

the shoreline (FAO, 2023).

We denote as $R_{c,mt}$ the ocean pH (multiplied by 100 to focus on changes of 0.01 units) matched to community or country c and measured at time mt , where m indicates the month and t the year. To match geographical and temporal variation in the unit of analysis with $R_{c,mt}$, we follow two approaches. First, contemporaneous exposure is computed by matching $R_{c,mt}$ either with the location and time of observation of the outcome variable. Second, early-life exposure is computed by matching $R_{c,mt}$ with an individual's location, month, and year of birth. As is standard in the literature, we assume that the survey location corresponds to the location of birth, also supported by the evidence suggesting the absence of selective migration in our setting (Section 2.5.2). When exposure is computed over multiple months, we average pH over that period. For instance, exposure *in utero* is the average $R_{c,mt}$ during the nine months preceding the date of birth.

Because pH is a logarithmic scale, we can interpret a decrease of 0.01 units as approximately a 2.3% increase in acidity. In our sample, 0.01 units correspond to the median within-year variation that a specific location experiences (i.e., the difference between the minimum and maximum pH in a specific year is on average 0.01), and to one-third of a standard deviation in *in-utero* exposure to pH.¹² To better understand the magnitude of the estimated effects, we also quantify the historical change in pH in the sampled area. The average reduction in pH from 1972 to 2018 was 0.075 units (0.016 units per decade). Therefore, the unit we analyse corresponds to the average reduction in pH experienced over approximately 6.25 years. Considering future projections of ocean acidity, the Intergovernmental Panel on Climate Change (IPCC, 2022a) predicts, under a business-as-usual emissions scenario, an average reduction of approximately 0.31 units of pH between 2018 and 2100 (0.037 units per decade). Under this scenario, a 0.01 decrease in pH would occur in just 2.67 years.

For identification, we exploit short-run exogenous deviations in pH levels from the spatially specific long-run trend (correcting for seasonality if the unit of analysis varies within year). Deviations are computed by capturing unobserved heterogeneity in the estimating equation using a set of FEs, which allows the isolation of deviations from the raw variation in pH. These FEs capture the time-invariant characteristics of the location of birth or observation (*location effects*); the common characteristics at the time of birth or observation (*non-spatial time effects*); and the trends and seasonality components of the ocean pH and the outcome variable of interest that are specific to a geographical region (*spatially specific time effects*). The latter are particularly important for identification because ocean acidification is spatially heterogeneous, with some regions experiencing faster or slower acidification, and more amplified or compressed within-year variation than others. The nature of abnormal deviation in the pH of our main independent variable is reinforced by the evolution over time of the sample aver-

¹²The standard deviation of pH is amplified by the large geographical area we cover. Conditional on FEs, a change of 0.01 in pH is roughly three standard deviations of residual variation (Appendix B2.1).

age short-run deviation (Appendix Figure B2.1). Exogeneity is supported by the balance of observable characteristics in areas affected by different deviations (Appendix B2.6).

We adapt this approach based on the geographical and temporal variation of the unit of analysis. When analysing data about seafood catch, the unit of analysis is at the seafood group–country–year level. We therefore estimate the effect of ocean pH in the nearest area of the ocean using the following specification:

$$y_{ic,t} = \beta R_{c,t} + \mathbf{X}_{c,t}\gamma + \Omega_{c,t} + \epsilon_{ic,t} \quad (2.1)$$

where $y_{ic,t}$ is the catch in the seafood group i fished in the EEZ of country c in year t . Because the unit of analysis covers an area of ocean larger than the resolution at which pH is observed, to compute $R_{c,t}$, we average all data points within the EEZ of country c in a year t . We label this variable *pH in proximity to the coast*. The specification includes a vector of weather control variables, $\mathbf{X}_{c,t}$, and the set of FEs that defines the identifying variation in terms of deviations from long-run patterns, $\Omega_{c,t}$.¹³ Location effects are captured by seafood group by country FEs; non-spatial effects by year of observation FEs; and spatially specific time effects by fishing area by year FEs. Fishing areas are geographical regions used for fisheries management and reporting, grouping multiple countries together (FAO, 2020). Finally, the idiosyncratic errors, $\epsilon_{ic,t}$, are assumed to be clustered at the EEZ level.

When analysing data about children and women, the unit of analysis is at the individual level. We estimate the effect of ocean pH in the nearest area of the ocean using the following specification:

$$y_{ic,mt} = \beta R_{c,mt} + \mathbf{X}_{ic,mt}\gamma + \Omega_{c,mt} + \epsilon_{ic,mt} \quad (2.2)$$

where $y_{ic,mt}$ is the outcome of interest for individual i at time mt in community c .

Because the geographical area of a community is smaller than the resolution at which pH is observed, $R_{c,mt}$ is the pH corresponding to time mt at the closest data point in the ocean, matched using the shortest straight-line distance. We label this variable as *pH in the nearest waters*. The specification includes a vector of demographic and weather control variables, $\mathbf{X}_{ic,mt}$, and the set of FEs that defines the identifying variation in terms of deviations from long-run patterns, $\Omega_{c,mt}$.¹⁴ Non-spatial time effects are captured by controlling for (interview or birth) month by year FEs. Spatially specific time effects are captured by macro-region by (interview or birth) year FEs, capturing local trends and, when within-year variation is observed, by macro-

¹³Controls include the average temperature and rainfall (and their interaction) of coastal areas, and the average oxygen concentration in the EEZ, another chemical feature of the ocean that is strongly correlated with ocean temperature (Free et al., 2019).

¹⁴When the outcome variable refers to children, *demographic controls* include the child's gender and birth order, the number of twins born with the child, mother's age at birth and at the time of the interview (including their square terms), mother's years of education, the household head's gender and age, and household size. When the outcome variable refers to adult women, these controls are limited to mother and household head characteristics. *Weather controls* are the same as in equation (2.1). In Section 2.5, we discuss the sensitivity to estimates of adding additional controls.

region by (interview or birth) month FEs, capturing local seasonality. A macro-region is a geographical area including multiple communities. We consider alternative definitions, such as administrative units like the district or the country of the community, and grid cells of different resolutions.¹⁵ Location effects depend instead on whether we are focusing on contemporaneous or early-life exposure. For contemporaneous impacts, we cannot exploit within-community variation because almost every individual in the community is interviewed in the same month. In this case, the *benchmark* specification includes location FEs, grouping multiple communities using grid cells. For early-life exposure, the *benchmark* specification includes community FEs, leveraging within-community temporal variation originating from birth dates. In this scenario, we can further exploit within-family variation by adding mother-specific FEs, controlling for mothers' time-invariant characteristics (*within-sibling* specification). Finally, the idiosyncratic errors, $\epsilon_{ic,t}$, are assumed to be clustered at the ocean raster data point (see Section 2.3).

We support the validity of the identifying assumptions in equations (2.1) and (2.2) with a variety of tests discussed in Section 2.5. In particular, we address issues related to non-random selection driven by FEs, which occurs from the loss of groups with only one observation and can lead estimates to differ from the population-wise average effect (Cameron et al., 2011). For example, the within-sibling identifying assumptions restrict the sample to mothers with at least two live births, who are generally older, have fewer years of education, were younger at the time of their first birth, and live in poorer households and communities (Appendix A2.2). Threats from this form of selection are limited by our measure of shocks being not only continuous, but also exhibiting a high degree of variation (the within-community variance in the identifying sample used by the benchmark specification is always positive). Nevertheless, in all estimation tables, we report the number of observations used in the estimation (*identifying observations*), and the number of observations dropped due to the identifying restrictions (*singleton observations*).

2.5 Results

We apply the methodology presented in Section 2.4 to estimate the impact of ocean pH on child health. Section 2.5.1 begins by discussing the causal pathway between ocean acidity and child health, focusing on the impact of contemporaneous exposure on fishing and on human nutrition. Section 2.5.2 presents results on the effect of early-life exposure on child mortality and development, and on health investments. Section 2.5.3 analyses how these effects vary according to the prevalent method of marine resource exploitation near the community.

¹⁵Grids allay concerns about the potential endogeneity of administrative boundaries. To guarantee sufficient variation in the measurement of ocean pH, which varies at the $1^\circ \times 1^\circ$ resolution, we consider grids at $5^\circ \times 5^\circ$ and $10^\circ \times 10^\circ$ resolutions.

2.5.1 Defining the causal pathway of ocean acidity

We begin by looking at the impact of contemporaneous ocean pH on fisheries. Table 2.1 shows estimates of equation (2.1) of the impact of ocean pH in the EEZ on the quantity and value of seafood catch, on the median price of seafood, and on aggregate economic activity proxied by satellite-based night-time luminosity. Dependent variables are reported using an inverse hyperbolic sine transformation to account for zero values. Results are robust to alternative transformations (Appendix B2.4).

We begin by focusing on small-scale fishing in columns (1)–(2). This activity primarily serves local consumption, and impacts the nutrition of coastal communities, as confirmed by a positive correlation between seafood catch derived from this activity and better nutritional indicators among women (Appendix B2.4). We observe that a decrease of 0.01 in pH leads to a significant decline of 0.13 log-points in the quantity caught and 0.20 log-points in the landed value. These results highlight that, at least for species harvested by small-scale fisheries, ocean acidification has a net negative effect (see Section 2.2). Effects are larger for seafood with a lower price, whose primary nutrient is essential fatty acids, and with a lower resilience to ocean acidification, but not statistically different (Appendix B2.4).

In columns (3)–(4), we focus instead on industrial fishing. Neither the quantity of seafood caught nor its value is influenced by ocean pH in the EEZ. These results highlight the resilience of industrial fishing to the shock and is consistent with evidence showing this sector’s ability to absorb shocks by diversifying catch or relocating fishing activities outside EEZs (see, e.g., Anderson et al., 2017), a possibility that is more limited for small-scale fishing.

Column (5) provides estimates on the effect on the overall median price of seafood. We observe that a decrease of 0.01 in pH leads to a significant increase of 0.09 log-points in the median price of seafood. Overall, these results suggest that the effects on small-scale fishing are enough to influence the median price of seafood, thus potentially influencing consumption choices.

Because fishing is an important economic activity in coastal areas, we want to exclude any income changes that occur alongside price changes. In column (6), we test whether ocean pH induces a short-term deterioration in the overall economic activity of coastal areas. We look at the effect of ocean pH on the average satellite-based night-time luminosity in the coastal area of each selected country. For this analysis, and for comparison with estimates in columns (1)–(5), we average night-time luminosity according to the definition of the coastal area of a country (see Section 2.3) and estimate equation (2.1) at the country level.

We find no effect on night-time luminosity. While we cannot exclude the possibility that ocean acidification may eventually influence the overall economic activity, these results suggest that the consequences of short-run variations in ocean pH are not driven by changes in aggregate income. One possibility is that the effects on fishing, which are specific to small-scale fisheries,

are too small to influence the whole economy or are specific to regions where night-time luminosity is not very responsive to changes in economic activity, such as poorer areas. Another alternative is that night-time luminosity responds to these shocks only over the long term.

Appendix B2.7 shows that these results are not specific to coastal areas. We show that a drought in the coastal area, a shock to agricultural productivity known to generate income changes (see, e.g., Barrios et al., 2010), leads to significant reductions in night-time luminosity. In addition, the lack of changes in labour supply induced by ocean acidity further suggests the absence of short-term impacts on the aggregate economic activity.

In the absence of income changes, the effects on seafood prices should reflect on consumption choices. We examine whether ocean acidity induces responses in nutrition by estimating the contemporaneous effect on women's fish consumption. With a limited number of surveys and respondents, the DHS programme asked whether a mother consumed different kinds of food in the 24 hours prior to the interview. Columns (1)–(2) in Table 2.2 show that a decrease in ocean pH lowers the probability of seafood consumption by 2.6 percentage points (or 8.8% over the sample mean of 29.6%). This reduction is specific to the consumption of seafood, as we observe no significant effect on the probability of consuming other food items (Appendix B2.8). These results suggest that adults do not compensate by adapting their diets.¹⁶

Columns (3)–(5) in Table 2.2 focus instead on malnutrition among women. For women who are not pregnant, we measure malnutrition using an indicator variable for whether the respondent is underweight, defined as having a body mass index (BMI) below 18.5. We supplement this measure with micronutrient deficiency, a direct measure of malnutrition for all women and for pregnant women. We proxy deficiency using objective measurements of anaemia, performed by the DHS enumerators on a random subset of women in the sample. Anaemia is characterised by low levels of haemoglobin, a protein in red blood cells that carry oxygen in the blood, and is often caused by iron deficiency.

In line with the evidence discussed in Section 2.2, the results indicate a pattern in which ocean acidification leads to changes in fish harvesting that impacts nutrition. A 0.01 decrease in ocean pH in nearby waters increases the probability of nearby women being underweight by 0.4 percentage points (or 3.3% over the sample mean of 12.0%). In addition, it leads to a higher prevalence of anaemia, but only among pregnant women. The existence of an effect specific to this vulnerable population is unsurprising because, during pregnancy, the human body requires more iron to supply the growing fetus, and with limited nutritional alternatives, seafood is an important source of iron (Luke, 1991; FAO, 2023). A 0.01 decrease in pH at the time of the measurement leads to an increase in anaemia prevalence of 1.7 percentage points among pregnant women (or 3.7% over the sample mean of 45.4%).

¹⁶The DHS provides only information on whether the respondent consumed a food item, but not the quantity consumed. We cannot exclude the possibility that respondents adapt their diets by increasing/decreasing the quantities consumed.

2.5.2 The effect of early-life exposure

This section focuses on the effect of being exposed early in life to varying degrees of ocean acidity in nearby waters. We analyse relevant effects on mortality, parental adaptation, and child development.

Mortality. To investigate the effect on early-life mortality and to isolate a channel operating through maternal malnutrition, we begin by studying *in-utero* exposure to varying degrees of ocean acidity in the waters nearest to people's places of birth.¹⁷ We estimate impacts on the likelihood of mortality at age x (in months). For each age x ranging from 1 month to 60 months, we estimate equation (2.2), restricting the sample to children who, at the time of the interview, were born at least x months before (independently from being alive). We select the sample based on time from birth to avoid selecting children who are alive and younger than x .¹⁸ The dependent variables are indicator variables equal to 1 if the child has died by age x from birth, and 0 otherwise (multiplied by 1,000 to relate coefficients to changes in deaths per 1,000 live births).

Figure 2.2 plots the coefficients. We observe that experiencing higher degrees of ocean acidity while *in utero* has a substantial impact on mortality. The effect peaks in the first month of life (corresponding to neonatal mortality), and remains significant across the very first months of life. A smaller net effect is observed beyond the first months of life, while the effect is not statistically different from zero after the first year of life. Because the initial increase in mortality is offset by later decreases, the pattern is consistent with a displacement of mortality hastened by experiencing worse conditions *in utero*—a mechanism known in the literature as *death harvesting* (see, e.g., Heutel et al., 2021).

Given the results on mortality, we focus on neonatal mortality. Table 2.3 presents estimates of the effect on the neonatal mortality rate (NMR)—the number of deaths in the first month of life per 1,000 live births. Panel A uses the benchmark specification, while Panel B uses the within-sibling specification. Columns (1)–(3) remove seasonality at the country level, while columns (4)–(6) remove seasonality at the grid cell level. Columns (1) and (4) do not include any control variables, columns (2) and (5) add weather controls, and columns (3) and (6) further add demographic controls. Figure 2.3 shows estimates using alternative specifications, including alternative sets of control variables, different time FEs, and different definitions of macro-regions.

A 0.01 decrease in pH significantly increases NMR by 1.42–2.12 deaths per 1,000 live births in

¹⁷We approximate the actual gestation period assuming a gestation period of nine months. Estimates assuming a gestation period of eight months, which can be interpreted as a lower bound of the effect, remain negative and statistically significant in most specifications (Appendix B2.9).

¹⁸The heaping of deaths at 1 year of age is common, while mortality at ages 2, 3, 4 and 5 is hardly affected by heaping (Croft et al., 2018). We observe no effect on the estimates due to these potential issues. Appendix B2.10 presents estimates of the effect on mortality rates at standard times.

our benchmark specification (panel A). Estimates using the within-sibling specification are similar (panel B). In terms of standardised effects (conditional on FEs), a one-standard-deviation negative shock leads to an increase in NMR of 0.53–0.56 deaths per 1,000 live births in the benchmark specification and 0.53–0.67 deaths per 1,000 live births in the within-sibling specification (Appendix Table B2.1).¹⁹ Adding control variables has a limited impact on the estimates of the effect, providing further evidence in support of the exogeneity of short-run deviations in pH. Significant effects are also found when varying the definition of coastal area, with the most affected communities living within 40 km from the shore (Appendix B2.2).

Estimates in Table 2.3 are robust to a wide variety of checks. First, while changing the set of FEs alters our identifying assumptions and our definition of deviation, estimates are always negative and significantly different from zero at standard confidence levels (Figure 2.3). Second, in Appendices B2.1, B2.2, and B2.7, we show that estimates are robust to adding controls for (potentially endogenous) confounders in both the location of birth and the location where pH is measured, such as adverse weather events (see, e.g., Gröger and Zylberberg, 2016), the presence of human activity proxied by pollution in coastal water, and the presence of conflict (see, e.g., Axbard, 2016), or to excluding areas that are generally subjected to vast anthropogenic waste, such as estuaries (see, e.g., Kennish, 2017). Third, selective migration does not drive estimates, as restricting our sample using information on whether the mother was living in the same location of the interview before the gestation period does not affect our conclusions (Appendix B2.5). Fourth, results are not driven by selection into identification and are robust to potential sources of measurement error associated with distances from the shore (Appendix B2.6). Finally, statistical inference is robust to alternative clustering assumptions about standard errors in equation (2.2) and to permutation-based inference, which artificially varies the exposure to the shock in both time and space (Appendix B2.11). The latter allows a rejection of the null hypothesis of a nil effect at the 5% significance level for all estimates in Table 2.3.

Figure 2.4 presents a heterogeneity analysis of the effect on neonatal mortality, distinguishing by exposure type. We highlight two main features. First, evidence suggests that the effect is driven by exposure during gestation to lower levels of pH. Panel A presents a binned analysis rather than a continuous one, displaying estimates from equation (2.2), where ocean pH while *in utero* is replaced by the share of gestation time during which children were exposed to ocean pH levels within specific ranges. These ranges are computed using the quartiles of the 1972–2018 pH distribution in the nearest waters. The effect is significantly different from zero only for exposure to lower levels of pH, indicating that accelerated acidification has a stronger negative impact than the potential positive effect of slowed acidification. Panel B

¹⁹The magnitude of these point estimates is smaller compared to interventions providing medical services to pregnant women. Nyqvist et al. (2019) show that introducing community health promoters in Uganda reduced neonatal mortality by 28%. Lazuka (2023) shows instead that introducing maternity wards in Sweden reduced neonatal mortality by 56%. Note, however, that the historical reduction in pH in the study period is seven times larger than the deviation in pH described by the estimate.

shows the estimates of equation (2.2) by adding ocean pH in the nearest waters one month before conception (10 months before birth), the month of birth, and 1–4 months after birth (a placebo period posterior to the period considered for the death). Effects are specific to the gestation period, reinforcing the role of maternal malnutrition. Further, we find no evidence of short-term responses in fertility or in the probability of a mother experiencing a terminated pregnancy (Appendix B2.3).

Second, impacts are concentrated in communities relying more heavily on the ocean's resources. Panel C of Figure 2.4 shows estimates of the effect on NMR, allowing estimates to vary flexibly with distance from the ocean's shore, and from other water bodies, like lakes and rivers. The largest effect is observed at the shore, while the estimate is not statistically different from zero at higher distances. On the contrary, the effect is homogeneous with respect to distance from other water bodies. In line with these results, effects are larger where seafood represents a higher share of total animal proteins consumed, in countries with a positive trade balance for fish products, and where small-scale fisheries are central, such as in proximity to reefs (Appendix B2.4).²⁰

Adaptation in health investments. Section 2.5.1 highlighted limited adaptation to ocean acidity in terms of dietary choices. In Table 2.4, we examine whether a mother alters health investments (before and after a child's birth) in response to experiencing varying degrees of ocean acidity during gestation. Columns (1)–(2) examine birth-level information regarding investments in antenatal care (attendance to health visits during pregnancy and the presence of health professionals during these visits) and care at the time of delivery (presence of health professionals during delivery and whether delivery was performed in a health centre). Both variables range from 0 (no) to 2 (high investment). Columns (3)–(5) focus on investments after the birth: postnatal healthcare, the completion of the cycle of basic vaccinations, and whether the child has ever been breastfed.²¹ Estimates are based on equation (2.2).

For both antenatal and delivery investments, we do not observe any significant effect. In line with these results, the effect on neonatal mortality is homogeneous in the birth order and sex of the child—two predictors of differential parental investments (Baird et al., 2011)—and across a wide array of individual characteristics (Appendix B2.12). Given that antenatal care is closely linked with nutrient supplementation plans during pregnancy, we also exclude this pathway. The lack of observable effects on postnatal care suggests limited adjustments in response to their child's health. We observe no effect on morbidity and anaemia prevalence among children at the time of the measurement (Appendix B2.12), suggesting no dietary changes among liv-

²⁰In Appendix B2.4, we show that *in-utero* exposure to higher seafood prices in the local market significantly contributes to the probability of neonatal death. Due to data limitations, we limit the analysis to the sample of the Philippines, one of the most fish-dependent countries globally.

²¹Because information on parental investments is not recorded for all children, but only for a subset within the household (generally the youngest child), we cannot estimate the effect on adaptation using the within-sibling specification.

ing children, and no evidence of adaptation through post-delivery migration (Appendix B2.5). Appendix B2.12 provides further evidence for the absence of adaptation. The effects on health investments are homogeneous across household wealth, maternal marital status, and maternal education.

Child development. Table 2.5 shows the effects of *in-utero* exposure to varying degrees of ocean acidity on early-life physical development, as assessed through anthropometric measurements—an important form of human capital accumulation. Columns (1)–(2) focus on the effects on weight-for-height (w/h), which captures insufficient food intake or high incidence of infectious diseases in temporal proximity to the measurement, and on height-for-age (h/a), which captures the past or cumulative effects of under-nutrition and infectious diseases since conception. Estimates in columns (3)–(4) focus instead on indicator variables for abnormally low values of w/h (*wasting*) and of h/a (*stunting*). All measures rely on objective measurements performed by the DHS enumerators on a random subset of children alive at the time of the interview, and they therefore need to be interpreted in light of the results on mortality. Panel A estimates the overall effects using equation (2.2), while panel B looks at heterogeneity by sex, introducing in equation (2.2) an interaction term between the ocean pH and an indicator variable for whether the child is female.²²

Among all children (panel A), we do not highlight any significant effect on w/h, h/a, or the prevalence of stunting, but we do observe a significant effect on the prevalence of wasting. A 0.01 decrease in pH in nearby waters reduces the probability of being wasted by 0.6 percentage points, or 7.5% over the sample mean of 8%. This effect is also captured by examining the probability of being underweight in the first months of life, which potentially indicates differences in birth weight (Appendix B2.12). While only the effect on wasting is statistically significant, the coefficients in panel A suggest that mortality selection prevails over a scarring effect, as living children who experienced higher degrees of water acidity tend to have better, rather than worse, indicators.²³

Although the effect on neonatal mortality is not heterogeneous by sex, when looking at heterogeneity by sex in the effect on child development (panel B), we highlight that the effects in panel A are driven primarily by male children. Among female children, results suggest the prevalence of a scarring effect (i.e., living children who experienced higher degrees of water acidity tend to have worse as compared to those that experienced lower degrees), concentrated in measures associated with height. Among boys, a 0.01 decrease in pH increases h/a by 0.03 standard deviations and the probability of being stunted by 1.1 percentage points (with a *p*-value of 0.15). Among girls, these outcomes are statistically different from those of boys. A 0.01 decrease in pH decreases h/a by 0.06 standard deviations and increases the probability of

²²Appendix B2.13 provides estimates of equation (2.2) splitting the sample into male and female children and shows results trimming *z*-scores at the 1st and 99th percentiles.

²³Refer to Deaton (2007) for a discussion on adult height and childhood mortality in poorer countries.

being stunted by 2.3 percentage points, as compared to boys. The effect on stunting is not only statistically different between male and female children, but it is significantly negative for girls (Appendix B2.13).

To understand whether these effects persist into adulthood, in Appendix B2.13 we examine these indicators for adult women, building their *in-utero* exposure to ocean pH in nearby waters by exploiting their month and year of birth, and the location of the interview. Because the temporal distance between exposure and the date of measurements is much larger than that of Table 2.5, these estimates implicitly assume no migration. While female migration in poorer settings is expected to occur within limited geographical distances (see, e.g., Rosenzweig and Stark, 1989; Mbaye and Wagner, 2017; Corno et al., 2020), we cannot exclude this possibility.

The results suggest that the scarring effect on girls is persistent in the long run. A 0.01 decrease in pH significantly decreases h/a by 0.1 standard deviations and increases the probability of being stunted by 0.7 percentage points. Adaptation at later ages could also play a role as the magnitude of the effect is smaller among adults than children. We also consider the impacts on women's economic well-being. A 0.01 decrease in pH significantly decreases adulthood wealth by 0.5% relative to the sample mean. This impact is accompanied by statistically significant decreases in the number of births per woman and the probability of working of 0.01 children and 1.4 percentage points, respectively. We do not observe any effect on schooling and cognitive skills.

2.5.3 Heterogeneity by resource exploitation

Due to the centrality of ocean exploitation for nutrition in coastal areas, we turn our attention to the heterogeneity of the effects discussed in Section 2.5.2 with respect to the type and intensity of fishing activities (see Section 2.3 for definitions and limitations related to this measure). For comparability, we quantify the effect of a one-standard-deviation decrease in nearby waters' pH that is experienced while *in utero* (labelled as an *acidity shock*), and we report estimates in terms of a percentage change with respect to the sample mean. Figure 2.5 plots the estimated effects at different intensities of night-time fishing (panel A), capturing the activity of boats operating on a smaller and more local scale, and of extractive fishing (panel B), which captures the activity of industrial fleets. In terms of outcomes, we consider impacts on NMR and an index of physical development among children, built by averaging available z-scores for w/h and h/a to capture multiple anthropometric insufficiencies. By averaging z-scores, our approach is similar to the multiple-inference approach of Anderson (2008), using as a control group the reference population used by DHS to compute z-scores.

Extractive fishing significantly reduces the ability to counteract shocks, amplifying their impacts. The effects on both NMR and physical development are homogeneous along the intensity of night-time fishing. Conversely, we observe heterogeneous effects by intensity of extractive

fishing. Areas characterised by high intensity present a significantly larger effect on NMR compared to areas without extractive fishing. An acidity shock leads to a 1.4% increase in NMR in areas where extractive fishing is absent and a 5.0% increase in areas where extractive fishing is largest. The mortality selection induced by these effects is captured in the heterogeneity of the impact on physical development. An acidity shock leads to an improvement in physical development by 0.7% in areas where extractive fishing is absent and by 4.3% in areas where extractive fishing is largest. Formal tests of heterogeneous impacts confirm these results (Appendix Table B2.11). Appendix B2.13 shows a similar analysis for outcomes among adults by focusing on economic well-being and on the physical development index among adult women, highlighting a similar pattern for the persistence of the effects presented in Figure 2.5.

2.6 Conclusions


Small changes in the ocean's chemical composition can have significant impacts on coastal communities. Our research demonstrates that increased ocean acidity negatively affects local fishing, in turn compromising nutritional quality in these areas. Deteriorating conditions in people's early lives raise neonatal mortality rates and influence mortality selection. Accordingly, in the absence of emissions-reduction strategies, the IPCC (2022a) predicts further significant increases in ocean water acidity by 2100. We should be cautious about the substantial mortality effects of ocean acidification, even with improvements in mitigation efforts due to economic development.

Our findings highlight the need for future research in two key areas. First, studies on climate change impacts should probe various alternative channels that have not yet been thoroughly investigated, through which climate-related shifts influence human and economic development. For instance, evidence concerning the ocean's role in these dynamics remains limited. While our research focuses on water acidity and its effects on marine life, this represents only one dimension of how a changing ocean could impact communities, especially those heavily reliant on marine resources. Gaining a deeper understanding of these mechanisms would enhance the design and targeting of policies to support vulnerable communities as they cope with climatic risks.

Second, by illustrating that wildlife serves as a crucial buffer against negative shocks, we emphasise the importance of research not only on policies promoting wildlife conservation, but also on strategies to mitigate the effects of reduced biodiversity. For instance, prioritising regulations in the industrial fishing sector and establishing exclusive artisanal fishing zones can be vital. While recent studies have highlighted the potential for effective policies (Frank and Oremus, 2023; Oremus et al., 2023), challenges remain, especially in countries with weak governance of natural resources. In the absence of effective conservation incentives, further research is necessary to allocate resources efficiently to communities in need of mitigation support. Our


results suggest a rationale for investing in targeted nutritional interventions early in life to address the reduced nutrient availability caused by negative shocks to natural resources. Such interventions have proven effective in mitigating both the short- and long-term consequences of malnutrition (see, e.g., [Gertler et al., 2014](#)). In addition, in light of the centrality of parental investments for early childhood development ([Attanasio et al., 2020](#)), our findings underscore the importance of awareness and education in influencing parental decisions regarding nutrition in low-income settings.

Figure 2.1: Area covered by the study




Note. Geographical distribution of DHS communities. The darker shaded area represents all countries surveyed by the DHS with access to the ocean (the full list is reported in Appendix A2.1). *Communities (coastal area)* are villages and neighbourhoods within 100 km from the ocean's shore. Most estimates in the paper include only these observations. *Communities (inland)* are villages and neighborhoods further than 100 km from the ocean's shore. Appendix A2.2 details the procedure followed to compute the distance from the shore. *Selected EEZs* refer to the Exclusive Economic Zones of all ocean-access countries included in the DHS survey (see Section 2.3 for the definition). In line with Pauly et al. (2020), we apply current EEZ boundaries (as depicted in the figure) to the whole study period to maintain consistency across years.

Figure 2.2: Early-life exposure and mortality




Note. Marginal effect of *in-utero* exposure to ocean pH in the nearest waters on the probability of death at month x (indicated on the horizontal axis). The dependent variable is an indicator variable equal to one if the child is dead at month x from birth, and zero if the child is alive, multiplied by 1,000. Estimates are based on equation (2.2), including community FEs, birth month by birth year FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors are clustered at the ocean raster data point. The 90% confidence interval is indicated by dotted lines, beyond which the intervals are progressively shaded up to the 99% level. Within confidence bounds, color intensity reflects the relative density of observations across iterations. It is calculated by comparing the density in each iteration to a range between the lower bound (adjusted by 0.7) and the 99th percentile of densities across all iterations. These parameters were chosen to improve visibility. Appendix A2.1 provides further information on the variables and the list of surveys included in the study.

Figure 2.3: Early-life exposure and neonatal mortality – alternative specifications




Note. Marginal effect of *in-utero* exposure to ocean pH in the nearest waters on NMR under alternative sets of FEs in the benchmark specification (panel A), and in the within-sibling specification (panel B). The dependent variable is an indicator variable equal to 1 if the child died within the first month of life and 0 if the child survived, multiplied by 1,000. *pH in the nearest waters (in utero)* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the child's community during the 9 months before birth. Marginal effects are estimated using equation (2.2) with the set of controls reported in the bottom panel. *Main controls* are the weather and demographic controls (see Section 2.4). *Interactions* are interaction terms between the birth month and indicator variables for different oceans (matched using the shortest straight-line distance from the community). *Main specifications* highlight the estimates presented in Table 2.3. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors are clustered at the ocean raster data point. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Figure 2.4: Early-life exposure and neonatal mortality – heterogeneity



Note. Marginal effects of *in-utero* exposure to ocean pH in the nearest waters on NMR, by type of pH deviation (panel A), timing of exposure (panel B), and distance from water bodies (panel C). In panel A, estimates are based on equation (2.2) where $R_{c,t}$ is substituted by the share of time children were exposed *in utero* to different levels of ocean pH. We classify values in four bins using the quartiles of the 1972–2018 distribution of pH in the $5^\circ \times 5^\circ$ grid cell matched to the child's location of birth. The bars (linked to the right vertical axis) presents the average share of time children were exposed *in utero* to levels of ocean pH in the corresponding quartile. The average pH corresponding to each quartile is 8.01, 8.04, 8.06, 8.09, respectively. In panel B, estimates are based on equation (2.2), in which the *pH in the nearest waters* at different points in time is the pH (multiplied by a factor of 100) in the ocean grid cell closest to the individual's community in the corresponding period relative to birth; when the period refers to multiple months, the value is averaged. In panel C, estimates are based on equation (2.2) introducing interactions between $R_{c,t}$ and a cubic polynomial in distance. In all panels, estimates are based on the benchmark specification, including community FEs, birth month by birth year FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). In panel C, we further allow FEs to be specific to the areas that are more or less vulnerable, defined by being closer or further away than 40 km from the shore (see Appendix B2.2). The dependent variable is *NMR*, an indicator variable equal to 1 if the child died within the first month of life and 0 if the child survived, multiplied by 1,000. The sample is restricted to communities in the coastal area (see Section 2.3). In panels A and B, confidence intervals are computed at 90% level. In panel C, the 90% confidence interval is indicated by dotted lines, beyond which the intervals are progressively shaded up to the 99% level. Within confidence bounds, color intensity shows the relative density of observations by distance from shore. It is calculated by comparing the square root of the density at each point to the square root of the 90th percentile of the overall density. These parameters were chosen to improve visibility. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Figure 2.5: Ocean acidity and resource exploitation



Note. Estimated impacts (and 90% confidence intervals) of a one-standard-deviation increase in water acidity (experienced *in utero*) on neonatal mortality and on physical development as a function of intensity of fishing (0 = no presence / 1 = high intensity). Panel A (B) focuses on night-time (extractive) fishing (see Section 2.3 for the definitions). Estimates are based on equation (2.2) introducing interaction terms between pH in the nearest waters (*in utero*) and a quadratic polynomial in the corresponding intensity. *Neonatal mortality* is an indicator variable equal to 1 if the child died within the first month of life and 0 if the child survived, multiplied by 1,000. *Physical development* is the average z-score of available anthropometric measures. *pH in the nearest waters (in utero)* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the individual's community during the 9 months before birth. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors are clustered at the ocean raster data point. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures. We exclude surveys for Peru as information for the intensity of night-time fishing is not available.

Table 2.1: Ocean acidity and marine catches

Dependent variables:	Marine catch for human consumption				Economic activity	
	Small-scale fishing		Industrial fishing		Median price	Night-time luminosity
	Quantity (1)	Value (2)	Quantity (3)	Value (4)	(5)	(6)
pH in proximity to the coast	0.132 (0.068) [0.061]	0.199 (0.081) [0.019]	0.019 (0.088) [0.829]	0.033 (0.099) [0.742]	-0.091 (0.045) [0.049]	0.006 (0.044) [0.896]
Mean (dep.var.)	1.50	1.51	1.00	1.00	1.00	2.26
Identifying observations	19,129	19,129	19,129	19,129	13,603	777
Singleton observations	0	0	0	0	0	0
Countries	36	36	36	36	36	36
Year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1992–2012

Note. Estimates are based on equation (2.1). Dependent variables in column headers are transformed using an inverse hyperbolic sine transformation to account for zero values (Appendix B2.4 reports results using alternative transformations). In columns (1)–(5), each observation is the catch or landed value or the median price for a specific seafood group in the corresponding Exclusive Economic Zone (EEZ; see, Section 2.3 for a definition). In column (5), each observation is a country's yearly average night-time luminosity in its coastal area (see Section 2.3). *pH in proximity to the coast* is the yearly average pH in the corresponding EEZ (multiplied by a factor of 100). Specifications in columns (1)–(4) include country by seafood group FEs, and fishing area by year FEs. The specification in column (5) includes country FEs, and fishing area by year FEs. All specifications include weather controls (see Section 2.4). Standard errors clustered at the EEZ level are reported in parentheses, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 2.2: Ocean acidity and nutrition among women

Dependent variable:	Nutrition		Prevalence of anaemia	
	Consumed seafood	Underweight	Overall	In pregnancy
	(1)	(2)	(3)	(4)
pH in the nearest waters	0.025 (0.012) [0.029]	-0.004 (0.002) [0.073]	0.001 (0.005) [0.841]	-0.017 (0.007) [0.013]
Mean (dep.var.)	0.296	0.120	0.427	0.454
Identifying observations	49,045	407,699	272,688	14,672
Singleton observations	2	3	2	36
Communities	5,952	24,301	17,371	8,993
Countries	14	32	26	26
Interview year range	1956–2000	1944–2003	1950–2003	1951–2003

Note. Estimates are based on equation (2.2). Dependent variables are reported in column headers: *consumed seafood* is an indicator variable equal to 1 if the respondent consumed seafood in the 24 hours previous to the interview, and 0 otherwise (information is available for mothers in the sample and for a selected number of countries; see Appendix A2.1); *underweight* is an indicator variable equal to 1 if the respondent has a BMI below 18.5, and 0 otherwise (information is available for all women with anthropometric measurement); *prevalence of anaemia* is an indicator variable equal to 1 if the respondent has haemoglobin levels below 110 g/L, and 0 otherwise (information is available for all women with blood samples). *pH in the nearest waters* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the female respondent's community in the month of the interview. The sample is restricted to communities in the coastal area (see Section 2.3). All specifications include location FEs using grid cells at the $1^\circ \times 1^\circ$ resolution, interview month FEs, interview year FEs, country by interview month FEs, country by interview year FEs, and control variables (see Section 2.4; weather controls correspond to the year of interview). Standard errors clustered at the ocean raster data point are reported in parentheses, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 2.3: Early-life exposure to ocean acidity and neonatal mortality

Dependent variable:	Neonatal mortality rate (deaths per 1,000 births)					
	(1)	(2)	(3)	(4)	(5)	(6)
A. Benchmark specification						
pH in the nearest waters (<i>in utero</i>)	-1.421 (0.691) [0.040]	-1.423 (0.682) [0.038]	-1.493 (0.663) [0.025]	-2.120 (0.755) [0.005]	-2.100 (0.762) [0.006]	-2.086 (0.739) [0.005]
Mean (dep.var.)	30.473	30.473	30.474	30.474	30.474	30.475
Identifying observations	1,583,706	1,583,706	1,581,815	1,583,703	1,583,703	1,581,812
Singleton observations	25	25	25	28	28	28
Communities	31,380	31,380	31,380	31,380	31,380	31,380
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
B. Within-sibling specification						
pH in the nearest waters (<i>in utero</i>)	-2.077 (0.873) [0.018]	-2.139 (0.854) [0.013]	-2.246 (0.837) [0.008]	-2.476 (0.954) [0.010]	-2.526 (0.953) [0.008]	-2.638 (0.936) [0.005]
Mean (dep.var.)	31.476	31.476	31.476	31.476	31.476	31.476
Identifying observations	1,474,945	1,474,945	1,474,945	1,474,941	1,474,941	1,474,941
Singleton observations	108,786	108,786	108,786	108,790	108,790	108,790
Communities	31,356	31,356	31,356	31,356	31,356	31,356
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes	-	Yes	Yes
Demographic controls	-	-	Yes	-	-	Yes
Seasonality	Country	Country	Country	Cell	Cell	Cell

Note. Estimates are based on equation (2.2). The dependent variable is an indicator variable equal to 1 if the child died within the first month of life and 0 if the child survived, multiplied by 1,000. *pH in the nearest waters (in utero)* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the child's community during the 9 months before birth. The sample is restricted to communities in the coastal area (see Section 2.3). All specifications include community FEs, birth year by birth month FEs, country by birth year FEs. Seasonality is captured by either country by birth month FEs or $5^\circ \times 5^\circ$ cell by birth month FEs. In panel B, community FEs are replaced by mother FEs. The full list of controls is presented in Section 2.4. Standard errors clustered at the ocean raster data point are reported in parentheses, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 2.4: Early-life exposure to ocean acidity and health investments

Dependent variables:	Antenatal	Delivery	Postnatal		
			Healthcare	Breastfed	Vaccinated
	(1)	(2)	(3)	(4)	(5)
pH in the nearest waters (<i>in utero</i>)	-0.002 (0.007) [0.820]	-0.002 (0.004) [0.727]	0.004 (0.009) [0.629]	0.001 (0.003) [0.699]	-0.002 (0.006) [0.804]
Mean (dep.var.)	1.808	1.359	0.441	0.972	0.435
Identifying observations	155,980	168,460	101,075	206,350	208,765
Singleton observations	217	481	3,078	2,336	2,269
Communities	14,669	18,481	18,445	28,029	27,887
Countries	29	29	34	36	36
Birth year range	1985–2018	1985–2018	2002–2018	1987–2018	1987–2018

Note. Estimates are based on equation (2.2). The dependent variables are reported in column headers: *antenatal* and *delivery* aggregate different investment indicators (see Appendix B2.12), ranging from 0 (no investment) to 2 (larger investment); *healthcare* is an indicator variable equal to 1 if the mother or the child younger than 2 years old received postnatal care within 2 days of birth; *breastfed* is an indicator variable equal to 1 if the mother reports ever breastfeeding the child, and 0 otherwise; *vaccinated* is an indicator variable equal to 1 if the mother reports or the vaccination card shows the completion of the basic cycle of vaccinations according to the World Health Organization (WHO), and 0 otherwise. *pH in the nearest waters (in utero)* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the child's community during the 9 months before birth. The sample is restricted to communities in the coastal area (see Section 2.3). Column (3) excludes the surveys for Indonesia and Morocco because information is not available in the corresponding surveys. For cross-survey comparability, the sample in columns (1)–(3) is restricted to the last birth, independently from the child being alive at the time of the interview, while in columns (4)–(5) is restricted to living children under three years old and can therefore be affected by mortality selection. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Standard errors clustered at the ocean raster data point are reported in parentheses, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 2.5: Early-life exposure to ocean acidity and physical development

Dependent variables:	z-scores		Indicators	
	Weight-for-height (1)	Height-for-age (2)	Wasted (3)	Stunted (4)
A. Overall effect				
pH in the nearest waters (<i>in utero</i>)	-0.021 (0.016) [0.196]	-0.012 (0.015) [0.405]	0.006 (0.003) [0.090]	0.004 (0.004) [0.279]
Mean (dep.var.)	-0.309	-0.984	0.080	0.234
Identifying observations	232,339	232,575	232,339	232,575
Singleton observations	1,106	1,124	1,106	1,124
Communities	24,824	25,110	24,824	25,110
Countries	33	33	33	33
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018
B. Heterogeneity by sex				
pH in the nearest waters (<i>in utero</i>)	-0.001 (0.017) [0.969]	-0.033 (0.020) [0.100]	0.008 (0.006) [0.179]	0.011 (0.007) [0.147]
× female	-0.014 (0.020) [0.461]	0.057 (0.028) [0.047]	-0.011 (0.010) [0.303]	-0.023 (0.010) [0.022]
Mean (dep.var.)	-0.312	-0.993	0.080	0.236
Identifying observations	226,567	226,685	226,567	226,685
Singleton observations	6,878	7,014	6,878	7,014
Communities	23,979	24,248	23,979	24,248
Countries	33	33	33	33
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018

Note. Estimates are based on equation (2.2). Dependent variables are reported in column headers: *weight-for-height* (*w/h*) and *height-for-age* (*h/w*) are z-scores from a reference scale; *wasted* and *stunted* are indicator variables equal to 1 for an abnormally low weight-for-height and height-for-age, respectively, and 0 otherwise. *pH in the nearest waters (in utero)* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the individual's community during the 9 months before the birth of the child. The sample is restricted to communities in the coastal area (see Section 2.3). All panels exclude the surveys for Indonesia, Pakistan, and the Philippines because information is not available in the correspondent surveys. Specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables. In panel B, FEs are sex-specific. Standard errors clustered at the ocean raster data point are reported in parentheses, -values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Bibliography

- Adhvaryu, A., P. Bharadwaj, J. Fenske, A. Nyshadham, and R. Stanley (2024). Dust and Death: Evidence from the West African Harmattan. *The Economic Journal* 134(659), 885–912.
- Adhvaryu, A. and A. Nyshadham (2016). Endowments at Birth and Parents’ Investments in Children. *The Economic Journal* 126(593), 781–820.
- Aizer, A., L. Stroud, and S. Buka (2016). Maternal Stress and Child Outcomes: Evidence from Siblings. *Journal of Human Resources* 51(3), 523–555.
- Almond, D., J. Currie, and V. Duque (2018). Childhood Circumstances and Adult Outcomes: Act II. *Journal of Economic Literature* 56(4), 1360–1446.
- Almond, D. and B. Mazumder (2011). Health Capital and the Prenatal Environment: The Effect of Ramadan Observance during Pregnancy. *American Economic Journal: Applied Economics* 3(4), 56–85.
- Alter, K., J. Jacquemont, J. Claudet, M. E. Lattuca, M. E. Barrantes, S. Marras, P. H. Manríquez, C. P. González, D. A. Fernández, M. A. Peck, et al. (2024). Hidden Impacts of Ocean Warming and Acidification on Biological Responses of Marine Animals Revealed through Meta-Analysis. *Nature Communications* 15(1), 2885.
- Anderson, M. L. (2008). Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American statistical Association* 103(484), 1481–1495.
- Anderson, S. C., E. J. Ward, A. O. Shelton, M. D. Adkison, A. H. Beaudreau, R. E. Brenner, A. C. Haynie, J. C. Shriver, J. T. Watson, and B. C. Williams (2017). Benefits and Risks of Diversification for Individual Fishers. *Proceedings of the National Academy of Sciences* 114(40), 10797–10802.
- Arceo, E., R. Hanna, and P. Oliva (2016). Does the Effect of Pollution on Infant Mortality Differ between Developing and Developed Countries? Evidence from Mexico City. *The Economic Journal* 126(591), 257–280.
- Attanasio, O., S. Cattan, E. Fitzsimons, C. Meghir, and M. Rubio-Codina (2020). Estimating the Production Function for Human Capital: REsults from a Randomized Controlled Trial in Colombia. *American Economic Review* 110(1), 48–85.
- Auffhammer, M. (2018). Quantifying Economic Damages from Climate Change. *Journal of Economic Perspectives* 32(4), 33–52.
- Axbard, S. (2016). Income Opportunities and Sea Piracy in Indonesia: Evidence from Satellite Data. *American Economic Journal: Applied Economics* 8(2), 154–94.

- Baird, S., J. Friedman, and N. Schady (2011). Aggregate Income Shocks and Infant Mortality in the Developing World. *Review of Economics and Statistics* 93(3), 847–856.
- Barreca, A., K. Clay, O. Deschênes, M. Greenstone, and J. S. Shapiro (2016). Adapting to Climate Change: The Remarkable Decline in the Us Temperature-Mortality Relationship over the Twentieth Century. *Journal of Political Economy* 124(1), 105–159.
- Barrios, S., L. Bertinelli, and E. Strobl (2010, 05). Trends in Rainfall and Economic Growth in aFrica: A Neglected Cause of the aFrican Growth Tragedy. *The Review of Economics and Statistics* 92(2), 350–366.
- Berazneva, J. and T. S. Byker (2022, 09). Impacts of Environmental Degradation: FOrest, Loss, Malaria, and Child Outcomes in NIgeria. *The Review of Economics and Statistics*, 1–46.
- Berthelon, M., D. Kruger, and R. Sanchez (2021). Maternal Stress during Pregnancy and Early Childhood Development. *Economics & Human Biology* 43, 101047.
- Black, S. E., A. Bütikofer, P. J. Devereux, and K. G. Salvanes (2019). This Is Only a Test? Long-Run and Intergenerational Impacts of Prenatal Exposure to Radioactive Fallout. *Review of Economics and Statistics* 101(3), 531–546.
- Bowles, S. and J.-K. Choi (2019). The Neolithic Agricultural Revolution and the Origins of Private Property. *Journal of Political Economy* 127(5), 2186–2228.
- Brander, L. M., K. Rehdanz, R. S. Tol, and P. J. Van Beukering (2012). The Economic Impact of Ocean Acidification on Coral Reefs. *Climate Change Economics* 3(1).
- Bruederle, A. and R. Hodler (2018). Nighttime Lights as a Proxy for Human Development at the Local Level. *PloS one* 13(9), e0202231.
- Burgess, R., M. Hansen, B. A. Olken, P. Potapov, and S. Sieber (2012). The Political Economy of Deforestation in the Tropics. *The Quarterly Journal of Economics* 127(4), 1707–1754.
- Burke, M., E. Gong, and K. Jones (2015). Income Shocks and Hiv in aFrica. *The Economic Journal* 125(585), 1157–1189.
- C3S (2017). ERA5: Fifth Generation of Ecmwf Atmospheric Reanalyses of the Global Climate. Copernicus Climate Change Service Climate Data Store (CDS).
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). Robust Inference with Multiway Clustering. *Journal of Business & Economic Statistics* 29(2), 238–249.
- Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, et al. (2022). Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits. *The Quarterly Journal of Economics* 137(4), 2037–2105.

- Chay, K. Y. and M. Greenstone (2003). The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession. *The Quarterly Journal of Economics* 118(3), 1121–1167.
- Christian, R. R. and S. Mazzilli (2007). Defining the Coast and Sentinel Ecosystems for Coastal Observations of Global Change. *Hydrobiologia* 577(1), 55–70.
- CIESIN-CIAT (2005). Gridded Population of the World, Version 3 (GPWv3) - Population Count Grid. Center for International Earth Science Information Network and Centro Internacional de Agricultura Tropical. Palisades, NY.
- Collier, P. (2010). *The Plundered Planet: Why We Must—And How We Can—Manage Nature for Global Prosperity*. Oxford University Press.
- Collins, W., N. Bellouin, M. Doutriaux-Boucher, N. Gedney, P. Halloran, T. Hinton, J. Hughes, C. Jones, M. Joshi, S. Liddicoat, et al. (2011). Development and Evaluation of an Earth-System Model—HadGEM2. *Geosci. Model Dev. Discuss* 4(2), 997–1062.
- Colt, S. G. and G. P. Knapp (2016). Economic Effects of an Ocean Acidification Catastrophe. *American Economic Review* 106(5), 615–19.
- Corno, L., N. Hildebrandt, and A. Voena (2020). Age of Marriage, Weather Shocks, and the Direction of Marriage Payments. *Econometrica* 88(3), 879–915.
- Croft, T. N., A. M. J. Marshall, and C. K. Allen (2018). Guide to Dhs Statistics. Demographic and Health Surveys Program.
- Dalgaard, C.-J., A. S. Knudsen, and P. Selaya (2020). The Bounty of the Sea and Long-Run Development. *Journal of Economic Growth* 25, 259–295.
- d’Andon, O. F., A. Mangin, S. Lavender, D. Antoine, S. Maritorena, A. Morel, G. Barrot, J. Demaria, and S. Pinnock (2009). Globcolour - The European Service for Ocean Colour. In *Proceedings of the 2009 IEEE International Geoscience & Remote Sensing Symposium*.
- Dasgupta, P. (2021). The Economics of Biodiversity: The Dasgupta Review. Technical report, London: HM Treasury.
- Deaton, A. (2007). Height, Health, and Development. *Proceedings of the national academy of sciences* 104(33), 13232–13237.
- Dell, M., B. F. Jones, and B. A. Olken (2014). What Do We Learn from the Weather? the New Climate-Economy Literature. *Journal of Economic Literature* 52(3), 740–98.
- Deschênes, O., M. Greenstone, and J. Guryan (2009). Climate Change and Birth Weight. *American Economic Review* 99(2), 211–17.

- Deschênes, O. and E. Moretti (2009). Extreme Weather Events, Mortality, and Migration. *The Review of Economics and Statistics* 91(4), 659–681.
- Doney, S. C., D. S. Busch, S. R. Cooley, and K. J. Kroeker (2020). The Impacts of Ocean Acidification on Marine Ecosystems and Reliant Human Communities. *Annual Review of Environment and Resources* 45(1), 83–112.
- Elvidge, C., D. Feng-Chi Hsu, K. E. Baugh, and T. Ghosh (2014). National Trends in Satellite Observed Lighting: 1992-2012. Ed. Qihao Weng. CRC Press.
- Elvidge, C. D., M. Zhizhin, K. Baugh, and F.-C. Hsu (2015). Automatic Boat Identification System for Viirs Low Light Imaging Data. *Remote sensing* 7(3), 3020–3036.
- Falkenberg, L. J., R. G. Bellerby, S. D. Connell, L. E. Fleming, B. Maycock, B. D. Russell, F. J. Sullivan, and S. Dupont (2020). Ocean Acidification and Human Health. *International Journal of Environmental Research and Public Health* 17(12), 4563.
- FAO (2019). Faostat – Food Balance Sheets. Food and Agriculture Organization of the United Nations.
- FAO (2020). Fao Major Fishing Areas. [Accessed on 01/11/2020 from <https://www.fao.org/fishery>]. Food and Agriculture Organization.
- FAO (2022). *The State of World Fisheries and Aquaculture*. Food and Agriculture Organization of the United Nations. Fisheries Department.
- FAO (2023). *Illuminating Hidden Harvests – The Contributions of Small-Scale Fisheries to Sustainable Development*. Food and Agriculture Organization of the United Nations, Duke University, & WorldFish.
- Feely, R. A., C. L. Sabine, J. M. Hernandez-Ayon, D. Ianson, and B. Hales (2008). Evidence for Upwelling of Corrosive “Acidified” Water onto the Continental Shelf. *Science* 320(5882), 1490–1492.
- Feir, D. L., R. Gillezeau, and M. E. C. Jones (2023). The Slaughter of the Bison and Reversal of Fortunes on the Great Plains. *The Review of Economic Studies*.
- Flanders Marine Institute (2019). Maritime Boundaries Geodatabase: Maritime Boundaries and Exclusive Economic Zones (200NM), Version 11. Available online at <https://www.marineregions.org/>.
- Frank, E. and K. Oremus (2023). Regulating Biological Resources: Lessons from Marine Fisheries in the United States. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2023-63).

- Frank, E. and A. Sudarshan (2023). The Social Costs of Keystone Species Collapse: Evidence from the Decline of Vultures in India. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2022-165).
- Frank, E. G. and D. S. Wilcove (2019). Long Delays in Banning Trade in Threatened Species. *Science* 363(6428), 686–688.
- Free, C. M., J. T. Thorson, M. L. Pinsky, K. L. Oken, J. Wiedenmann, and O. P. Jensen (2019). Impacts of Historical Warming on Marine Fisheries Production. *Science* 363(6430), 979–983.
- Gattuso, J.-P. and L. Hansson (2011). *Ocean Acidification*. Oxford, UK: Oxford University Press.
- Gelcich, S., P. Buckley, J. K. Pinnegar, J. Chilvers, I. Lorenzoni, G. Terry, M. Guerrero, J. C. Castilla, A. Valdebenito, and C. M. Duarte (2014). Public Awareness, Concerns, and Priorities about Anthropogenic Impacts on Marine Environments. *Proceedings of the National Academy of Sciences* 111(42), 15042–15047.
- Gertler, P., J. Heckman, R. Pinto, A. Zanolini, C. Vermeersch, S. Walker, S. M. Chang, and S. Grantham-McGregor (2014). Labor Market Returns to an Early Childhood Stimulation Intervention in Jamaica. *Science* 344(6187), 998–1001.
- Geruso, M. and D. Spears (2018a). Heat, Humidity, and Infant Mortality in the Developing World. NBER working paper no. 24870, National Bureau of Economic Research.
- Geruso, M. and D. Spears (2018b). Neighborhood Sanitation and Infant Mortality. *American Economic Journal: Applied Economics* 10(2), 125–62.
- Gröger, A. and Y. Zylberberg (2016). Internal Labor Migration as a Shock Coping Strategy: Evidence from a Typhoon. *American Economic Journal: Applied Economics* 8(2), 123–153.
- Heft-Neal, S., J. Burney, E. Bendavid, and M. Burke (2018). Robust Relationship between Air Quality and Infant Mortality in Africa. *Nature* 559(7713), 254.
- Heutel, G., N. H. Miller, and D. Molitor (2021). Adaptation and the Mortality Effects of Temperature across Us Climate Regions. *The Review of Economics and Statistics* 103(4), 740–753.
- Hicks, C. C., P. J. Cohen, N. A. Graham, K. L. Nash, E. H. Allison, C. D’Lima, D. J. Mills, M. Roscher, S. H. Thilsted, A. L. Thorne-Lyman, et al. (2019). Harnessing Global Fisheries to Tackle Micronutrient Deficiencies. *Nature* 574(7776), 95–98.
- Hirvonen, K., J. Hoddinott, B. Minten, and D. Stifel (2017). Children’S Diets, Nutrition Knowledge, and Access to Markets. *World Development* 95, 303–315.

- ICF (2018). Demographic and Health Surveys Standard Recode Manual for Dhs 7. Technical report, DHS Program.
- ICF (2019). Demographic and Health Surveys 1991-2018 (Various Datasets). Calverton, Maryland: ICF International. <https://www.dhsprogram.com>.
- IIPS and ICF (2017). National Family Health Survey Nfhs-4 2015-16: India. Technical report, International Institute for Population Sciences.
- International Monetary Fund (2025). International Financial Statistics. Accessed on 29/01/2025 from <https://data.imf.org>.
- IPCC (2022a). AR6 Srocc Data for Figure Spm.1: Past and Future Changes in the Ocean and Cryosphere. MetadataWorks.
- IPCC (2022b). Climate Change 2022: Impacts, Adaptation and Vulnerability - Chapter 3: Oceans and Coastal Ecosystems and Their Services. Working group ii contribution to the sixth assessment report of the IPCC, Intergovernmental Panel on Climate Change.
- Isen, A., M. Rossin-Slater, and W. R. Walker (2017). Every Breath You Take—Every Dollar You’ll Make: The Long-Term Consequences of the Clean Air Act of 1970. *Journal of Political Economy* 125(3), 848–902.
- Jayachandran, S. (2013). Liquidity Constraints and Deforestation: The Limitations of Payments for Ecosystem Services. *American Economic Review* 103(3), 309–13.
- Jones, C., J. Hughes, N. Bellouin, S. Hardiman, G. Jones, J. Knight, S. Liddicoat, F. O’Connor, R. J. Andres, C. Bell, et al. (2011). The HadGEM2-Es Implementation of CMIP5 Centennial Simulations. *Geoscientific Model Development* 4(3), 543–570.
- Jones, M. C. and W. W. Cheung (2018). Using Fuzzy Logic to Determine the Vulnerability of Marine Species to Climate Change. *Global change biology* 24(2), e719–e731.
- Keeling, R. F., A. Körtzinger, and N. Gruber (2010). Ocean Deoxygenation in a Warming World. *Annual Review of Marine Science* 2(1), 199–229. PMID: 21141663.
- Kennish, M. J. (2017). *Practical Handbook of Estuarine and Marine Pollution*. CRC press.
- Kremer, M. and C. Morcom (2000, March). Elephants. *American Economic Review* 90(1), 212–234.
- Kroeker, K. J., R. L. Kordas, R. N. Crim, and G. G. Singh (2010). Meta-Analysis Reveals Negative yet Variable Effects of Ocean Acidification on Marine Organisms. *Ecology letters* 13(11), 1419–1434.

- Kroodsma, D. A., J. Mayorga, T. Hochberg, N. A. Miller, K. Boerder, F. Ferretti, A. Wilson, B. Bergman, T. D. White, B. A. Block, et al. (2018). Tracking the Global Footprint of Fisheries. *Science* 359(6378), 904–908.
- Land, P. E., H. S. Findlay, J. D. Shutler, I. G. Ashton, T. Holding, A. Grouazel, F. Girard-Ardhuin, N. Reul, J.-F. Piolle, B. Chapron, Y. Quilfen, R. G. Bellerby, P. Bhadury, J. Salisbury, D. Vandemark, and R. Sabia (2019). Optimum Satellite Remote Sensing of the Marine Carbonate System Using Empirical Algorithms in the Global Ocean, the Greater Caribbean, the Amazon Plume and the Bay of Bengal. *Remote Sensing of Environment* 235, 111469.
- Lazuka, V. (2023). It’S a Long Walk: Lasting Effects of Maternity Ward Openings on Labor Market Performance. *Review of Economics and Statistics* 105(6), 1411–1425.
- Lotze, H. K., H. Guest, J. O’Leary, A. Tuda, and D. Wallace (2018). Public Perceptions of Marine Threats and Protection from around the World. *Ocean & Coastal Management* 152, 14–22.
- Lowe, N. M. (2021). The Global Challenge of Hidden Hunger: Perspectives from the Field. *Proceedings of the Nutrition Society* 80(3), 283–289.
- Luke, B. (1991). Nutrition during Pregnancy: Part I, Weight Gain; Part II, Nutrient Supplements. *JAMA* 265(2), 281–282.
- Maccini, S. and D. Yang (2009). Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall. *American Economic Review* 99(3), 1006–26.
- Maire, E., N. A. Graham, M. A. MacNeil, V. W. Lam, J. P. Robinson, W. W. Cheung, and C. C. Hicks (2021). Micronutrient Supply from Global Marine Fisheries under Climate Change and Overfishing. *Current Biology* 31(18), 4132–4138.
- Majid, M. F. (2015). The Persistent Effects of in Utero Nutrition Shocks over the Life Cycle: Evidence from Ramadan Fasting. *Journal of Development Economics* 117, 48–57.
- Matranga, A. (2024). The Ant and the Grasshopper: Seasonality and the Invention of Agriculture. *The Quarterly Journal of Economics*, 1467–1504.
- Mayshar, J., O. Moav, and L. Pascali (2022). The Origin of the State: Land Productivity or Appropriability? *Journal of Political Economy* 130(4), 1091–1144.
- Mbaye, L. M. and N. Wagner (2017). Bride Price and Fertility Decisions: Evidence from Rural Senegal. *The Journal of Development Studies* 53(6), 891–910.
- McCauley, D. J., C. Jablonicky, E. H. Allison, C. D. Golden, F. H. Joyce, J. Mayorga, and D. Kroodsma (2018). Wealthy Countries Dominate Industrial Fishing. *Science Advances* 4(8), eaau2161.

- Meiyappan, P. and A. K. Jain (2012). Three Distinct Global Estimates of Historical Land-Cover Change and Land-Use Conversions for over 200 Years. *Frontiers of Earth Science* 6(2), 122–139.
- Menclova, A. K. and S. Stillman (2020). Maternal Stress and Birth Outcomes: Evidence from an Unexpected Earthquake Swarm. *Health Economics* 29(12), 1705–1720.
- Michalopoulos, S. and E. Papaioannou (2013). Pre-Colonial Ethnic Institutions and Contemporary African Development. *Econometrica* 81(1), 113–152.
- Miller, D. L., N. Shenhav, and M. Z. Grosz (2021). Selection into Identification in Fixed Effects Models, with Application to Head Start. *Journal of Human Resources* forthcoming.
- Moretti, E. and M. Neidell (2011). Pollution, Health, and Avoidance Behavior: Evidence from the Ports of Los Angeles. *Journal of human Resources* 46(1), 154–175.
- Natural Earth (2018). Physical Vector Data Themes - Version 4.1.0. Accessed on 21/05/2018 from <https://www.naturalearthdata.com>.
- Nyqvist, M. B., A. Guariso, J. Svensson, and D. Yanagizawa-Drott (2019). Reducing Child Mortality in the Last Mile: Experimental Evidence on Community Health Promoters in Uganda. *American Economic Journal: Applied Economics* 11(3), 155–192.
- Oremus, K. L., E. G. Frank, J. J. Adelman, S. Cruz, J. Herndon, B. Sewell, and L. Suatoni (2023). Underfished or Unwanted? *Science* 380(6645), 585–588.
- Pauly, D. and D. Zeller (2016). Catch Reconstructions Reveal That Global Marine Fisheries Catches Are Higher Than Reported and Declining. *Nature Communications* 7(1), 10244.
- Pauly, D., D. Zeller, and M. Palomares (2020). *Sea around Us Concepts, Design and Data*. Data download on October 14, 2023 from www.seaaroundus.org.
- Persson, P. and M. Rossin-Slater (2018). Family Ruptures, Stress, and the Mental Health of the Next Generation. *American economic review* 108(4-5), 1214–1252.
- Philippine Statistics Authority (2020). Fish: Retail Prices of Agricultural Commodities. Accessed on 08/02/2020 from openstat.psa.gov.ph.
- Razzaque, A., N. Alam, L. Wai, and A. Foster (1990). Sustained Effects of the 1974–5 Famine on Infant and Child Mortality in a Rural Area of Bangladesh. *Population Studies* 44(1), 145–154. PMID: 11612523.
- Rosenzweig, M. R. and O. Stark (1989). Consumption Smoothing, Migration, and Marriage: Evidence from Rural India. *Journal of political Economy* 97(4), 905–926.

- Simmanee, F. A., G. Nico, S. Funge-Smith, X. Basurto, N. Franz, S. J. Teoh, K. A. Byrd, J. Kolding, M. Ahern, P. J. Cohen, et al. (2022). Proximity to Small-Scale Inland and Coastal Fisheries Is Associated with Improved Income and Food Security. *Commun earth & environ* 3(1), 174.
- Stavins, R. N. (2011). The Problem of the Commons: Still Unsettled after 100 Years. *American Economic Review* 101(1), 81–108.
- Sunday, J. M., K. E. Fabricius, K. J. Kroeker, K. M. Anderson, N. E. Brown, J. P. Barry, S. D. Connell, S. Dupont, B. Gaylord, J. M. Hall-Spencer, et al. (2017). Ocean Acidification Can Mediate Biodiversity Shifts by Changing Biogenic Habitat. *Nature Climate Change* 7(1), 81–85.
- Sundberg, R. and E. Melander (2013). Introducing the Ucdp Georeferenced Event Dataset. *Journal of Peace Research* 50(4), 523–532.
- Tollefsen, A. F., H. Strand, and H. Buhaug (2012). Prio-Grid: A Unified Spatial Data Structure. *Journal of Peace Research* 49(2), 363–374.
- Totterdell, I. (2019). Description and Evaluation of the Diat-Hadocc Model v1.0: The Ocean Biogeochemical Component of HadGEM2-Es. *Geoscientific Model Development* 12, 4497–4549.
- UDEL (2018). University of Delaware Terrestrial Precipitation.
- UNEP-WCMC (2018). Global Distribution of Coral Reefs. UNEP World Conservation Monitoring Centre and the WorldFish Centre.
- United Nations (2003). Ecosystems and Human Well-Being: A Framework for Assessment. United Nations, Island Press, Washington DC.
- United Nations (2021). The Role of Aquatic Foods in Sustainable Healthy Diets. UN Nutrition Discussion Paper.
- United Nations (2024). World Population Prospects 2024. Department of Economic and Social Affairs, Population Division.
- Victora, C. G., P. Christian, L. P. Vdaletti, G. Gatica-Domínguez, P. Menon, and R. E. Black (2021). Revisiting Maternal and Child Undernutrition in Low-Income and Middle-Income Countries: Variable Progress Towards an Unfinished Agenda. *The Lancet* 397(10282), 1388–1399.
- Wessel, P. and W. H. Smith (1996). A Global, Self-Consistent, Hierarchical, High-Resolution Shoreline Database. *Journal of Geophysical Research: Solid Earth* 101(B4), 8741–8743.

Wittmann, A. C. and H.-O. Pörtner (2013). Sensitivities of Extant Animal Taxa to Ocean Acidification. *Nature climate change* 3(11), 995–1001.

World Bank (2012). Hidden Harvest: The Global Contribution of Capture Fisheries. Report number 66469-glb, The World Bank, FAO, World Fish and Agriculture and Rural Development.

APPENDIX

Supplementary material to *From Sea to Shore: The Impact of Ocean Acidification on Child Health*

Appendix A: Data and methodological procedures

A2.1 Variables, data sources, and the selection of DHS surveys

Variable	Description
<i>Adaptation</i>	Information is based on parental health investments obtained from the DHS Programme (ICF, 2019). We homogenise information across surveys and make use of the following variables: <i>Antenatal investment</i> is equal to 0 if no antenatal visit is completed, 1 if at least one visit is completed but without a health professional, and 2 if at least one visit is completed with a health professional. In Appendix B2.12, this indicator is split into individual variables. <i>Any visit</i> is an indicator variable equal to 1 if the mother attended any visit during pregnancy for antenatal care, and 0 otherwise. <i>Number of antenatal care visits</i> is the number of visits attended during pregnancy for antenatal care (reported in logarithms, adding one unit to allow for zero values). <i>With health professional</i> is an indicator variable equal to 1 if the mother was attended by a health professional (doctor, nurse or other professional) during pregnancy, and 0 otherwise. <i>Delivery investment</i> is equal to 0 if delivery is performed outside a health centre without a health professional, 1 if performed outside a health centre with a health professional, and 2 if delivery is performed in a health centre with a health professional. In Appendix B2.12, this indicator is split into individual variables. <i>In health centre</i> is an indicator variable equal to 1 if the mother gave birth in a health centre, and 0 otherwise. <i>With health professional</i> is an indicator variable equal to 1 if delivery was attended by a health professional (doctor, nurse or other professional), and 0 otherwise. For <i>postnatal investment</i>, <i>healthcare</i> is an indicator variable equal to 1 if the mother or the child younger than 2 years old received postnatal care within 2 days of birth. <i>Breastfed</i> is an indicator variable equal to 1 if the mother reports ever breastfeeding the child, and 0 if the mother reports to have never breastfed the child. For cross-survey comparability, the sample is restricted to children who live with their mother and are alive, and are less than 3 years old. <i>Vaccinated</i> is an indicator variable equal to 1 if the mother reports or shows a vaccination card for the following doses: BCG, 3 doses of DPT-containing vaccines, 3 doses of polio vaccine (excluding polio vaccine given at birth), and 1 dose of MCV. It is 0 otherwise. The sample is restricted to children under 3 years old for comparability (Croft et al., 2018).
<i>Altitude</i>	Communities' elevation in meters from the SRTM–Digital Elevation Model for the specified coordinate location. The variable is available in the DHS surveys (ICF, 2019).
<i>Agricultural land</i>	It measures the percentage area of a cell in 1970 that is used for agricultural purposes as defined by the ISAM-HYDE historical landuse dataset (Meiyappan and Jain, 2012). The data are downloaded from the PRIO-GRID version 2.0 database (Tollefsen et al., 2012). It is a vector grid network with a resolution of 0.5°×0.5° covering all terrestrial areas of the world that is spatially merged to DHS clusters using their geolocation.
<i>Basemaps</i>	Basemaps were created using ArcGIS® software by Esri®. Basemaps are used in line with the Esri Master License Agreement, specifically for the inclusion of screen captures in academic publications. We use the <i>World Topographic Map</i> .
<i>Child mortality</i>	Information is based on the DHS Programme surveys (ICF, 2019). DHS surveys collect respondents' full birth history and includes information on all children's year and month of birth, sex, birth order, whether they are twins, and the date of death when it applies. Note that only live births are recorded. This information is also used to create <i>age at first delivery</i>, and <i>fertility</i> (the number of live births at the time of the interview). We build mortality rates by multiplying the following indicators by 1,000 (the variables are set to missing if the date of the interview is before the end of the period considered for defining mortality): <i>neonatal (NMR)</i>: indicator equal to 1 if the child died before their first month of life, and 0 otherwise; <i>post-neonatal (PMR)</i>: indicator equal to 1 if the child died between the ages of 1–11 months, and 0 otherwise; <i>child (CMR)</i>: indicator equal to 1 if the child died between the ages of 12–59 months, and 0 otherwise; <i>infant (IMR)</i>: indicator equal to 1 if the child died between the ages of 0–11 months, and 0 otherwise; <i>under-5 (U5MR)</i>: indicator to 1 if the child died between the ages of 0–59 months, and 0 otherwise. Note that the DHS Programme reports two ages of death. The first is self-reported, while the second gives a calculated age from reported information. When dates of birth are not disclosed, these are imputed by the DHS Programme (Croft et al., 2018). We also use 67 special cases of self-reported age of death (198 and 199, which indicate that age at death was reported as a number of days and that the exact number is unknown), but results are robust to dropping these cases.
<i>Chlorophyll</i>	Chlorophyll concentration in coastal waters is measured in mg/m ³ , using GlobColour GSM-weighted estimates. We use data from the GlobColour project (d'Andon et al., 2009), which provides monthly global rasters for the period September 1997–2018 at a 0.25°×0.25° resolution by merging satellite imagery from five different sources made available by the European Space Agency and NASA.
<i>Conflict</i>	Number of violent events (and fatalities) in each cell for a specific year. The data are obtained from the Uppsala Conflict Data Programme (UCDP) (Sundberg and Melander, 2013).

(continued on next page)

Variable	Description
<i>Distances</i>	For shorelines, straight-line distances between the DHS cluster and the closest shoreline. Water bodies are identified from the GSHHG database (Wessel and Smith, 1996). We use the following two bodies. For the <i>ocean's shoreline</i> , we consider level 1 (continental land masses and ocean islands, except Antarctica). For <i>other water bodies</i> , we consider levels 2, 3 and 4 (lakes, islands in lakes, and ponds in islands within lakes and all levels included in the river database). See Appendix A2.2 for details about the procedure. For <i>coral reefs</i> , distance (in a straight line) between the DHS cluster and the closest coral reef vertex from UNEP-WCMC (2018).
<i>Economic indicators</i>	In Section B2.4, we compute real prices in US dollars for the Philippines using the consumer price index and exchange rates from International Monetary Fund (2025).
<i>Economic well-being</i>	DHS asset-based wealth index ranging from 1 (poorest) to 5 (richest).
<i>EEZ</i>	Geographical boundaries of Exclusive Economic Areas (EEZ) provided by the Flanders Marine Institute (2019).
<i>Extractive fishing</i>	Total number of hours from industrial fishing activities in the cell built using data from the Global Fishing Watch (Kroodsma et al., 2018), which tracks more than 70,000 industrial fishing vessels from 2012 to 2016. Because variation is available only for the period 2012–2016, we first compute total fishing hours in a global grid at $1^\circ \times 1^\circ$ resolution and then average each cell over the available period.
<i>Fishing area</i>	Geographical regions used for fisheries management and reporting, grouping multiple countries together (FAO, 2020). Because countries might have access to multiple areas, we use the primary fishing area, defined as the one covering the largest part of a country's EEZ. For country-level regressions, we merge fishing areas assigned to a single country to avoid collinearity. This is the case of Peru-Guatemala, and of Egypt-East Africa.
<i>Food intake</i>	DHS surveys collect respondents' food consumption for a variety of items (ICF, 2019). This information is available only for a restricted number of surveys: Cambodia (2005), Dominican Republic (2007), Egypt (2008), Ghana (2008), Guatemala (2015), Guyana (2009), Haiti (2006), Liberia (2007), Madagascar (2008), Namibia (2006), Nigeria (2008), Philippines (2008), Sierra Leone (2008), and Timor-Leste (2009 and 2016). Beyond fish consumption, we consider the following variables: <i>other proteins</i> includes all types of animal proteins excluding seafood, eggs, legumes, and beans which are protein-rich; <i>carbohydrates</i> includes simple carbohydrates such as bread, noodles and other grains and excludes starches (complex carbohydrates); <i>fats</i> includes any type of oil, animal fat, and butter; <i>other iron-rich food</i> includes any poultry, red meat, liver, beans, legumes, nuts and dark leafy greens.
<i>Human capital</i>	We make use of <i>schooling</i> , i.e., the number of completed years of education based on the respondent's self-reported highest level of education (comparable across countries), and of <i>cognitive skills</i> , i.e., an indicator variable of whether the respondent is able to read a whole sentence in her native language (as observed by enumerators) or has, at least, completed secondary schooling.
<i>Labour supply</i>	Indicator variable equal to 1 if the respondent is working, and 0 otherwise. DHS surveys record the employment status of respondents at the time of the interview.
<i>Marriage</i>	DHS surveys collect respondents' civil status, date of birth and, when available, their partner's age in years. We make use of the following variables. <i>Married</i> is an indicator variable equal to 1 if the respondent is currently married or living in an union, and 0 otherwise.
<i>Migration</i>	We build migration indicators using DHS question <i>VI04</i> , described as the "number of years the respondent has lived in the village, town, or city where she was interviewed," was not included in DHS-VI surveys (ICF, 2018). Information is unavailable in the following surveys: Bangladesh 2011 and 2014, Benin 2012, Cambodia 2010 and 2014, Cameroon 2011, Comoros 2012, Congo Democratic Republic 2013, Côte d'Ivoire 1998 and 2012, Dominican Republic 2013, Egypt 2014, Gabon 2012, Ghana 2014, Guatemala 2015, Guinea 1999 and 2012, Haiti 2012, Honduras 2011, Indonesia 2003, Liberia 2013, Mozambique 2011, Myanmar 2016, Namibia 2013, Nigeria 2013, Pakistan 2006, Senegal 2010–2016, Sierra-Leone 2013, Tanzania 2010, and Togo 2013.
<i>Night-time luminosity</i>	Average night-time light emission from the $0.5^\circ \times 0.5^\circ$ DMSP-OLS Night-time Lights Time Series Version 4 calibrated (Elvidge et al., 2014). Values range between 0 (lowest) and 1 (highest observed value). The time series are available from 1992–2012 from Tollefsen et al. (2012). Data are spatially merged to DHS clusters using their geolocation. To compute values in the coastal area of country, we consider only values with the coastal area (see Section 2.3 for the definition).
<i>Night-time fishing</i>	We use Automatic Boat Identification System for VIIRS Low Light Imaging Data (Elvidge et al., 2015) to identify detections. Using individual daily detections (which include geolocation), we build a $1^\circ \times 1^\circ$ global grid with the sum of detections for the period 2017–2019. We classify as boats only the strongest detections (quality flag rating equal to 1). Data are not available over the South Atlantic Anomaly (DHS surveys for Peru are the only ones affected).
<i>Nutrition</i>	The DHS records objective measurements performed by the DHS data collection team. Standardised distributions are the CDC Standard Deviation-derived Growth Reference Curves (Croft et al., 2018). The following indicators are used: <i>anaemia</i> is an indicator variable equal to 1 if the woman has haemoglobin levels below 110 g/L, and 0 otherwise; <i>underweight</i> is, for children, an indicator variable equal to 1 if the weight-for-age <i>z</i> -score is smaller than 2 or, for adults, if the BMI is lower than 18.5, and 0 otherwise; <i>w/h</i> (<i>weight-for-height</i>) is the <i>z</i> -score from the reference curve, while <i>wasted</i> is an indicator variable equal to 1 if the weight-for-height <i>z</i> -score is smaller than 2, and 0 otherwise; <i>h/a</i> (<i>height-for-age</i>) is the <i>z</i> -score from the reference curve, while <i>stunted</i> is an indicator variable equal to 1 if the height-for-age <i>z</i> -score is smaller than 2, and 0 otherwise; <i>Physical development</i> is the average between height-for-age and weight-for-height <i>z</i> -scores from the reference curves.
<i>Ocean chemistry</i>	Data are obtained from the Hadley Global Environment Model 2 - Earth System or HadGEM2–ES, originally developed by Collins et al. (2011); Jones et al. (2011). Data are provided as monthly global rasters at the $1^\circ \times 1^\circ$ resolution for a series of chemical features of the ocean. We use two variables: pH at surface and dissolved O_2 concentration.

(continued on next page)

Variable	Description
<i>Ocean identifiers</i>	Ocean identifiers are obtained from Natural Earth (2018) . Each DHS cluster is matched using the shortest straight-line distance.
<i>Population</i>	Population estimates used in Figure B2.11 are taken from the Gridded Population of the World, version 3 (CIESIN-CIAT, 2005), as provided in PRIO-GRID, covering the years 1990, 1995, 2000, and 2005.
<i>Seafood catch/prices</i>	Quantity (in kilotons) and landed value in US\$ of catches within each Exclusive Economic Zones (EEZ). Data source is Sea Around Us initiative (Pauly et al., 2020). EEZ boundaries are obtained from Flanders Marine Institute (2019) . We select EEZs of the countries in our sample; we drop the only contested area in our sample (i.e., Western Sahara region) because DHS did not sample households in this region. Data on catches is at the level of seafood commercial categories, including: anchovies; cod-like; crustaceans; flatfishes; herring-like; molluscs; perch-like; salmon and smelts; scorpionfishes; sharks and rays; tuna and billfishes; and other fishes and invertebrates. We merge the category salmon and smelts with tuna and billfishes as the fishing of the former category is almost non-existent in the study area. We include catches direct human consumption from industrial and small-scale fisheries; we exclude catches from the recreational sector. <i>Median prices</i> are computed as the median ex-vessel species-level within each commercial group-year-EEZ. Abnormally small and large prices are set to missing using the 0.5 th and the 99.5 th percentiles of the distribution of prices within the corresponding commercial group in the full timeframe analysed. Section B2.4 uses monthly province-level retail seafood prices for the Philippines from 1990 to 2018, as provided by Philippine Statistics Authority (2020) .
<i>Seafood dependency</i>	Average fish protein supply as proportion of all animal protein supply (FAO, 2019). Data are missing for Comoros.
<i>Weather</i>	Yearly total amount of precipitation (in millimeters) in the cell is based on monthly meteorological statistics from the GPCP v.2.2 Combined Precipitation Data Set, available for the years 1979–2014. Yearly mean temperature (°C) in the cell is based on monthly meteorological statistics from GHCN/CAMS, which is available for the period 1948–2014. Data are downloaded from the PRIO-GRID v2.0 database (Tollefsen et al., 2012), a vector grid network with a resolution of 0.5°×0.5°, and spatially merged to DHS clusters using their geolocation. We supplement this information with weather data from the ERA5 dataset (C3S, 2017), which provides a 0.25°×0.25° monthly gridded dataset for a variety of weather and climatic variables. We obtain sea surface temperature (SST), measured at the same point as ocean pH, and wind speed, total precipitations and air (2-meter) temperature, measured at closest location to a community. For all variables, we average daily values to monthly data. For our comparison with droughts in Section B2.7 , we require an extended monthly precipitation series, as defining a shock depends on the long-run precipitation distribution. Therefore, we use the University of Delaware Terrestrial Precipitation dataset (UDEL, 2018), provided by the National Oceanic and Atmospheric Administration (NOAA), which includes monthly values from 1900 to 2012 at a resolution of 0.5°×0.5°.

Note. For time-varying variables, missing values are linearly interpolated.

Table [A2.2](#) presents the Demographic and Health Surveys (DHS) included in the analysis. The availability of multiple surveys for some countries can lead to issues related to survey selection. Table [A2.3](#) presents estimates of equation (2.2) assuming different rules for the selection of surveys. When including multiple surveys for the same country, each observation is weighted by the product of the DHS sampling weight with a re-weighting factor, i.e., the ratio between the sum of the DHS sampling weights at the country-survey level and the sum of the DHS sampling weights at the country level. For adult-level estimates, we re-weight observations following the same procedure, repeating the computation of weights for different variables because the inclusion in each survey is variable-dependent. For adult outcomes relative to schooling and work, we include only observations that completed both the education and work module. This selection affects only the India 2015–2016 survey, for which we select only the women that completed the *state module*), and we use the corresponding weights ([IIPS and ICF, 2017](#)).

Table A2.2: Sampled countries

Country	DHS surveys available	Birth years matched
Angola	2015	1980-2016
Bangladesh	2000, 2004, 2007, 2011, 2014	1972-2014
Benin	1996, 2001, 2012	1972-2012
Cambodia	2000, 2005, 2010, 2014	1972-2014
Cameroon	1991, 2004, 2011	1972-2011
Colombia	2010	1974-2010
Comoros	2012	1975-2012
DR Congo	2007, 2013	1974-2013
Côte d'Ivoire	1994, 1998, 2012	1972-2012
Dominican Republic	2007, 2013	1972-2013
Egypt	1992, 1995, 2000, 2005, 2008, 2014	1972-2014
Gabon	2012	1975-2012
Ghana	1993, 1998, 2003, 2008, 2014	1972-2014
Guatemala	2015	1978-2015
Guinea	1999, 2005, 2012, 2018	1972-2018
Guyana	2009	1974-2009
Haiti	2000, 2006, 2012, 2016	1972-2017
Honduras	2011	1975-2012
India	2015	1976-2016
Indonesia	2003	1972-2003
Kenya	2003, 2008, 2014	1972-2014
Liberia	2007, 2013	1972-2013
Madagascar	1997, 2008	1972-2009
Morocco	2003	1972-2004
Mozambique	2011	1975-2011
Myanmar	2016	1980-2016
Namibia	2000, 2006, 2013	1972-2013
Nigeria	1990, 2003, 2008, 2013, 2018	1972-2018
Pakistan	2006	1973-2007
Peru	2000, 2004, 2009	1972-2009
Philippines	2003, 2008, 2017	1972-2017
Senegal	1993, 1997, 2005, 2010, 2012, 2014, 2015, 2016	1972-2016
Sierra Leone	2008, 2013	1973-2013
Tanzania	1999, 2010, 2015	1972-2016
Timor-Leste	2009, 2016	1974-2016
Togo	1998, 2013	1972-2014

Note. From all DHS surveys available on May 2020, we include only surveys for countries with direct access to the ocean and surveys with available geocoding of primary sampling units. *Birth years matched* refers to child-level information and includes all observations in the birth histories (*DHS birth recode*) that are matched with data on ocean pH and are assigned to communities within 100 km from the shore.

Table A2.3: Robustness to selection of surveys

Dependent variable: <i>DHS surveys:</i>	NMR (deaths per 1,000 births)			
	<i>All</i> (1)	<i>Latest</i> (2)	<i>Largest</i> (3)	<i>Random</i> (4)
pH in the nearest waters (<i>in utero</i>)	-1.493 (0.663) [0.025]	-1.418 (0.701) [0.044]	-1.806 (0.658) [0.006]	-1.608 (0.676) [0.018]
Mean (dep.var.)	30.474	26.601	27.774	29.036
Identifying observations	1,581,815	794,713	866,814	757,132
Singleton observations	25	32	33	30
Communities	31,380	17,389	18,361	16,416
Countries	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and controls (see Section 2.4). In column (1), observations are re-weighted to correct for oversampling of countries surveyed multiple times (see Appendix A2.1). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. *Latest* indicates that only the latest survey is selected. *Largest* indicates that the survey with the largest number of observations is selected. *Random* indicates that one random survey is selected among the available ones. Appendix A2.1 provides further information on the variables and the list of surveys included in the study.

A2.2 Descriptive statistics

For each household, distance is the minimum straight distance from the DHS geocoded coordinates to the coast and closest alternative water source, computed using *v.distance* function in GRASS. Table A2.4 presents descriptive statistics for households living within and beyond 100 km from the shore. Figure A2.1 presents an example of the procedure. Table A2.5 shows the observable differences between mothers with a single child (excluded in the within-sibling specification) and mothers with multiple children.

Table A2.4: Descriptive statistics for coastal and inland areas

	Coastal area		Inland area		Observations
	Mean (1)	Std. dev. (2)	Mean (3)	Std. dev. (4)	
A. Children					
Child is alive	0.92	0.27	0.91	0.29	4555492
Child is female	0.48	0.50	0.48	0.50	4555492
Birth order	2.54	1.80	2.66	1.84	4555492
Birth order in a multiple birth	0.03	0.23	0.03	0.22	4555492
Years since birth	12.29	7.87	12.10	7.77	4555492
Mother's age at birth of child	24.43	5.76	24.16	5.54	4555492
pH in the nearest waters (<i>in utero</i>)	8.05	0.03	8.06	0.03	4555492
B. Adult women					
Age at first delivery	20.88	4.24	20.45	3.83	1385467
Current age	30.65	9.81	29.97	9.76	1951250
Years of schooling	7.51	4.89	6.22	5.12	1950621
pH in the nearest waters (<i>in utero</i>)	8.07	0.03	8.08	0.03	1474866
At least primary schooling	0.69	0.46	0.60	0.49	1951201
Married	0.67	0.47	0.70	0.46	1950104
Working	0.48	0.50	0.46	0.50	1355025
Household head is female	0.22	0.41	0.17	0.37	1951247
Household head's age	46.10	13.11	46.37	13.17	1949918
Household members	5.62	3.03	6.05	3.11	1951250
Household wealth	3.40	1.34	2.94	1.41	1755397
Living in urban area	0.53	0.50	0.34	0.47	1951250
Distance from shore	31.26	30.21	462.18	289.55	1951250
Distance from another water body	47.26	101.98	24.87	23.99	1951250
Altitude	189.80	408.19	489.32	612.81	1951250
Temperature (° C)	26.31	3.29	25.25	3.77	1951250
Precipitations (mm)	1560.29	674.43	1303.71	671.73	1951250
Intensity of extractive fishing	0.12	0.38	0.01	0.14	1951250
Intensity of night-time fishing	0.13	0.32	0.03	0.14	1951250
C. Mortality rates					
Neonatal	27.39	163.22	37.17	189.18	4545390
Postneonatal	23.51	151.51	24.18	153.62	4200570
Child	21.50	145.04	27.55	163.68	3265547
Infant	50.39	218.74	60.63	238.65	4355601
Under-five	73.76	261.38	89.29	285.16	3504461


Note. Descriptive statistics by proximity to the ocean for all communities in selected countries with access to ocean. Coastal area includes all communities within 100 km from the ocean's shore (see Section 2.3). Inland area includes all communities that are farther away than 100 km from the ocean's shore. Means are reported in columns (1) and (3), standard deviations are reported in columns (2) and (4). Column (5) presents the total number of observations. *Years since birth* is measured at the time of the interview and is independent from the child being alive. *Mortality rates* are relative to 1,000 live births. *pH in the nearest waters (in utero)* is defined in Section 2.4; it refers to the date of birth of the child in panel A and to the date of birth of the woman in panel B. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table A2.5: Comparison of mothers with a single child versus multiple children

	One child		Multiple children		Observations
	Mean (1)	SD (2)	Mean (3)	SD (4)	
A. Children					
Child is alive	0.97	0.16	0.92	0.27	1587285
Child is female	0.47	0.50	0.49	0.50	1587285
Birth order	1.00	0.00	2.68	1.82	1587285
Birth order in a multiple birth	0.00	0.00	0.04	0.24	1587285
Years since birth	6.03	6.54	12.86	7.73	1587285
Mother's age at birth of child	22.51	4.71	24.60	5.82	1587285
B. Adult women					
Age at first delivery	22.51	4.71	20.37	3.95	495310
Current age	28.55	7.99	36.19	7.66	495310
Years of schooling	8.75	4.68	6.07	4.81	495104
pH in the nearest waters (<i>in utero</i>)	8.07	0.03	8.08	0.02	302404
At least primary schooling	0.79	0.41	0.57	0.49	495286
Married	0.81	0.40	0.89	0.31	495309
Working	0.47	0.50	0.54	0.50	438928
Household head is female	0.23	0.42	0.19	0.39	495310
Household head's age	45.04	15.18	44.61	11.96	494936
Household members	5.13	3.08	5.71	2.89	495310
Household wealth	3.53	1.32	3.25	1.35	427608
Living in urban area	0.57	0.49	0.49	0.50	495310
Distance from shore	31.14	30.01	32.49	30.24	495310
Distance from another water body	39.09	81.17	46.49	100.25	495310
Altitude	178.93	396.59	187.05	400.57	495310
Temperature (° C)	26.40	3.21	26.41	3.14	495310
Precipitations (mm)	1611.70	660.16	1551.50	683.80	495310
Intensity of extractive fishing	0.11	0.34	0.10	0.33	495310
Intensity of night-time fishing	0.14	0.33	0.14	0.33	495310
C. Mortality rates					
Neonatal	12.77	112.30	28.74	167.07	1583731
Postneonatal	7.16	84.34	24.85	155.68	1470093
Child	6.38	79.62	22.20	147.33	1141371
Infant	17.95	132.77	53.01	224.05	1516640
Under-five	23.63	151.88	75.99	264.98	1217000

Note. Descriptive statistics by the number of children of the mother (reported in column headers). Means are reported in columns (1) and (3), standard deviations in columns (2) and (4). Column (5) presents the total number of observations. *Years since birth* is measured at the time of the interview and is independent from the child being alive. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Figure A2.1: Distance to ocean and other water sources: an example



Note. Geolocation of DHS communities (panel A) and closest points to the ocean's shore (panel B). Lines represent straight distance from a community to the closest point on the coast's shoreline or on the shoreline of another water body. Basemap source: Esri. See Appendix A2.1 for data sources and attributions.

Appendix B: Supplementary results

B2.1 Coastal features and income processes

Figure B2.1 shows descriptive statistics of average pH at surface (panel A), the evolution of the average deviation in pH over time (panel B), and the between and within decomposition of the overall variation of ocean pH while *in utero* as compared to NMR (panel C). Table B2.1 shows descriptive statistics of the measure of shock under the different specifications presented in Table 2.3, and the correspondent standardised effect.

Table B2.1: Standardised effects in Table 2.3

	Benchmark specification				Within-sibling specification			
	Mean (1)	SD (2)	Effect (3)	Std. effect (4)	Mean (5)	SD (6)	Effect (7)	Std. effect (8)
Shock (specification 1)	0.00	0.38	-1.42	-0.54	-0.00	0.30	-2.08	-0.63
Shock (specification 2)	0.00	0.37	-1.42	-0.53	-0.00	0.30	-2.14	-0.64
Shock (specification 3)	0.00	0.37	-1.49	-0.56	-0.00	0.30	-2.25	-0.67
Shock (specification 4)	0.00	0.26	-2.12	-0.55	-0.00	0.22	-2.48	-0.54
Shock (specification 5)	0.00	0.25	-2.10	-0.53	-0.00	0.21	-2.53	-0.53
Shock (specification 6)	-0.00	0.25	-2.09	-0.53	-0.00	0.21	-2.64	-0.55


Note. Descriptive statistics of shocks in ocean pH in the nearest waters (*in utero*) under the benchmark and the within-sibling specifications. Columns (3) and (7) report to the point estimates in Table 2.3. The standardised effect is rescaling point estimates in terms of standard deviations in the residual variation of ocean pH in the nearest waters (*in utero*). Residual variation is obtained from the residuals of a linear regression using ocean pH experienced *in utero* as dependent variable and the set of FEs used in equation (2.2) as independent variables.

We focus next on other features in the ocean and in coastal areas that could influence income processes in sampled communities. In terms of **other ocean's characteristics**, Table B2.2 presents estimates of the effect of ocean acidity in the nearest waters (*in utero*) on NMR using equation (2.2) and controlling for various ocean characteristics and inland weather conditions obtained from the ERA5 dataset. Panels A–D in Figure B2.2 show the time series and the seasonality component for these variables.


In terms of **pollution and other chemical features of the ocean**, Columns (3)–(4) in Table B2.2 present estimates controlling for pollution in coastal waters, proxied by satellite-based algae abundance (chlorophyll concentration) from the GlobColour project from 1997–2018. The presence of pollution also impacts the availability of another input to marine life that is more closely related to fish survival: oxygen. At low levels of concentration (hypoxic conditions), marine wildlife changes behaviour to reach areas with higher oxygen levels, while at extremely low levels (dead-zones), mortality prevails. Oxygen concentration is also affected by climate change because higher temperatures lead to reduced oxygen concentration (Free et al., 2019). In column (7), we also control for this variable obtained from the HadGEM2-ES model. Because pH and oxygen concentration are chemical properties determined by common factors, we always include as control the residual variation in oxygen concentration, rather than its levels. Residual variation is computed as residuals of a linear regression of oxygen concentration in grid cell i at time t on the contemporaneous pH in the same grid cell. Controlling for other

Figure B2.1: Variation in ocean acidity for communities in the coastal area


A. Yearly and monthly variation of pH in raster points matched with DHS



B. Average shock in ocean acidity (*in utero*)




C. Between and within variation decomposition of ocean pH (*in utero*) and NMR



Note. In panel A, yearly average pH at surface in the period 1972–2018 (left figure), and monthly comparison between mean pH for each year in the left axis, and median pH for the whole period in the right axis (right figure). Variation is restricted to cells matched to the sample's communities. In panel A, the solid red line shows the quadratic trend in the series. In panel B, evolution over time of the average deviation in acidity levels from spatially specific (and seasonally-adjusted) long-run trends. *pH in the nearest waters (in utero)* is defined in Section 2.4 and is computed using the benchmark specification. Variation is restricted to cells matched to the sample. The solid red line shows the quadratic trend over the period. In panel C, decomposition of the sample standard deviation of ocean pH experienced *in utero*, as compared to NMR. The sample is restricted to communities in the coastal area (see Section 2.3). Geographical and time variables for which the decomposition is computed are reported at the bottom of each figure. Appendix A2.1 provides further information on the variables and the list of surveys included in the study.

chemical features does not affect these estimates. Panels E–F in Figure B2.2 depict the time series and the seasonality component for these variables.

Figure B2.2: Additional weather characteristics in the ocean's matched areas



Note. Descriptive statistics of weather characteristics are measured either at the same point as ocean pH (panels A, E, F) or at the location of the closest community (panels B, C, D). The figures on the left display yearly averages, with a solid red line indicating the quadratic trend in the series. The figures on the right illustrate monthly averages for each year in the sample, with the darker line representing the median over the entire period. Variation is limited to cells matched to the sample's communities. Appendix A2.1 provides further information on the variables.

Table B2.2: Neonatal mortality and shocks to income processes


Dependent variable:	NMR (deaths per 1,000 births)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Closest ocean's point									
pH in the nearest waters (<i>in utero</i>)	-1.704 (0.723) [0.019]		-2.691 (1.301) [0.039]		-1.597 (0.692) [0.021]		-1.903 (0.803) [0.018]		-1.695 (0.694) [0.015]
<i>In-utero</i> sea surface temperature	-0.443 (0.784) [0.572]	0.433 (0.797) [0.587]							-0.412 (0.771) [0.594]
<i>In-utero</i> chlorophyll concentration			-0.278 (0.928) [0.765]	-0.560 (0.900) [0.534]					
<i>In-utero</i> oxygen concentration									-0.048 (0.335) [0.887]
Location of birth									
<i>In-utero</i> wind speed					0.564 (1.516) [0.710]	0.202 (1.623) [0.901]			0.289 (1.546) [0.852]
<i>In-utero</i> total precipitations, ERA5							0.001 (0.008) [0.845]	0.003 (0.008) [0.684]	
<i>In-utero</i> 2-meter temperature, ERA5							-0.986 (0.875) [0.261]	-0.088 (0.781) [0.910]	
<i>In-utero</i> temperature									-0.108 (0.425) [0.799]
<i>In-utero</i> total precipitations									-0.002 (0.002) [0.172]
Mean (dep.var.)	29.644	29.644	24.939	24.939	29.644	29.644	29.644	29.644	29.644
Identifying observations	1,518,360	1,518,360	451,226	451,226	1,518,360	1,518,360	1,518,360	1,518,360	1,518,360
Singleton observations	20	20	233	233	20	20	20	20	20
Communities	31,380	31,380	16,409	16,409	31,380	31,380	31,380	31,380	31,380
Countries	36	36	36	36	36	36	36	36	36
Birth year range	1979–2018	1979–2018	1998–2018	1998–2018	1979–2018	1979–2018	1979–2018	1979–2018	1979–2018

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. *In utero* indicates that the variable is the average value in the ocean grid cell closest to the child's community during the 9 months before birth. *Year of birth* indicates that the variable is the average value in the child's community's grid cell in the year of birth. The sample is restricted to communities in the coastal area (see Section 2.3). In columns (3)–(4), the sample is further restricted to births between 1997–2018 due to data availability (observations are reweighted to account for dropped surveys), and to areas away from estuaries to alleviate endogeneity concerns. Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, 5°×5° grid cell by birth month FEs, and demographic controls (see Section 2.4). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.2 Alternative definitions of coastal area


In terms of **proximity**, we vary the distance from the ocean's shore, which affects the total number of live births considered and the estimate of the effect on NMR (Figure B2.3). In terms of **altitude**, we combine proximity with elevation requirements (see, e.g., [Christian and Mazzilli, 2007](#); Figure B2.4 shows this selection process in our sample). Finally, we can exclude areas with higher contamination, such as **estuaries**. Table B2.3 shows how estimates of the effect of ocean pH in the nearest waters (*in utero*) on NMR varies under different definitions of a coastal area.

Figure B2.3: Sample selection by distance from shore



Note. Number of live births included in the sample (panel A), and marginal effects of ocean pH in the nearest waters (*in utero*) on NMR (panel B) by sample selection according to the distance from the shore. *pH in the nearest waters (*in utero*)* is defined in Section 2.4. Estimates are based on equation (2.2) when the sample is selected according to distance (reported in the horizontal axis). Each specification includes community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). The 90% confidence interval is indicated by dotted lines, beyond which the intervals are progressively shaded up to the 99% level. Within confidence bounds, color intensity shows the relative density of observations by distance from shore. It is calculated by comparing the square root of the density at each point to the square root of the 90th percentile of the overall density. These parameters were chosen to improve visibility. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Figure B2.4: Sample selection using proximity and altitude criteria



Note. Communities in coastal areas distinguished by altitude (panel A), and an example (panel B). The full list of countries and surveys included in the study is reported in Appendix A2.1. See Section 2.3 for a definition of coastal area.

Table B2.3: The effect on neonatal mortality: varying sample selection criteria

Dependent variable:	NMR (deaths per 1,000 births)					
	<i>Altitude criteria:</i>	<i>Distance restriction:</i>	<i>Exclusion of estuaries:</i>	<i>Altitude criteria:</i>	<i>Distance restriction:</i>	<i>Exclusion of estuaries:</i>
	$\leq 100m$	$\leq 100m$	$\leq 40km$	$\leq 40km$	$\leq 100m$	$\leq 100m$
	-	-	-	-	-	-
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	-1.632 (0.776) [0.036]	-1.601 (0.758) [0.035]	-2.929 (0.797) [0.000]	-3.080 (0.943) [0.001]	-2.945 (0.836) [0.000]	-3.075 (0.994) [0.002]
Mean (dep. var.)	31.116	31.431	29.489	29.631	29.938	30.113
Identifying observations	1,137,356	978,016	1,061,342	893,056	845,155	685,815
Singleton observations	19	15	25	21	22	18
Communities	22,612	18,801	21,682	17,616	17,600	13,789
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to coastal areas (see Section 2.3) and according to the criteria reported in column headers. *Estuaries* are defined as communities that are at a distance of 10 km or less from the ocean's shore and at a distance of 10 km or less from another water source. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (the full list of controls in Section 2.4). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.3 Fertility

We examine the timing of a woman's first birth and the occurrence of her last terminated pregnancy (defined as any pregnancy not resulting in a live birth). DHS surveys only provide information on age at first birth and the year of the last terminated pregnancy. We approximate a duration model by building a panel that tracks each woman from the age of 12 until she either experiences one of these events, or until she turns 33 years old. In our sample, 99% of women experience their first birth by age 33. If a woman who experienced any of these two events is missing information on her age at first birth or the year of her last terminated pregnancy, she is dropped from the panel. In addition, we drop observations that do not have information about migration, which allows considering only women that lived in the same location for the whole period in the panel (see Appendix B2.5). When considering terminated pregnancies, if a woman had a successful pregnancy prior to her most recent terminated pregnancy, we track her from the time of her last successful pregnancy instead of the age of 12. We estimate the effect of the average ocean pH in the nearest waters in the year $t - 1$ on the probability of a woman experiencing her first birth or last terminated birth in year t using the following specification:

$$y_{ivc,tab} = \beta R_{vc,t-1} + \mathbf{X}_{vc,t-1}\gamma + \Omega_{ivc,tab} + \epsilon_{ivc,tab} \quad (2.3)$$

where $y_{ivc,tab}$ is the outcome of interest for a woman i born in year b with age a located in community v of country c in year t , which takes the value 1 in the year the woman has her first successful birth, and 0 otherwise. $\mathbf{X}_{vc,t-1}$ is a vector of weather control variables, $\Omega_{ivc,tab}$ is a set of FEs, and $\epsilon_{ivc,tab}$ are idiosyncratic errors assumed to be clustered at the ocean raster data point. We include individual FEs to remove time-invariant characteristics at the individual, household, and location levels; country by age FEs to flexibly allow for the probability of first

birth by age to vary across countries; country by year of birth FEs to account for cohort-specific effects; and country by year FEs to account for local trends. Table B2.4 presents results for first birth and last terminated pregnancy, respectively.

Table B2.4: Exposure, first birth and terminated pregnancies

Dependent variable:	First birth		Last terminated pregnancy	
	(1)	(2)	(3)	(4)
pH in the nearest waters (previous year)	0.000 (0.003) [0.918]	-0.000 (0.003) [0.883]	-0.000 (0.001) [0.521]	-0.000 (0.001) [0.680]
Mean (dep.var.)	0.111	0.111	0.013	0.013
Identifying observations	1,096,011	1,096,011	1,204,419	1,204,419
Singleton observations	1,082	1,082	47	47
Communities	19,153	19,153	15,630	15,630
Countries	28	28	26	26
Birth year range	1960–2003	1960–2003	1961–2002	1961–2002
Weather controls	-	Yes	-	Yes

Note. Estimates are based on equation (2.3). The dependent variable is an indicator variable equal to 1 when the woman experienced the event described in the header, and 0 otherwise. *pH in the nearest waters (previous year)* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the child's community in the year before. The sample is restricted to communities in the coastal area (see Section 2.3). Additional restrictions are described in Section B2.3. Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include individual FEs, country by age FEs, country by year of birth FEs, and country by year FEs. Additionally, even-numbered columns include weather controls. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

In addition, Table B2.5 estimates the effect of exposure to changes in ocean pH at the time of the interview on the probability of being pregnant and allows for heterogeneous effects by socioeconomic status and education. We found no significant differential effect of the ocean pH on current fertility across different socioeconomic statuses or education levels. Table B2.6 replicates Panel A in Table 2.3 excluding women who report ever experiencing a terminated pregnancy. Finally, Table B2.7 replicates Panel A in Table 2.3 by restricting the sample to recent births (at most 10 years prior to the interview).

Table B2.5: Ocean acidity and current fertility

Dependent variable:	Currently pregnant		
	(1)	(2)	(3)
pH in the nearest waters (<i>in utero</i>)	0.000 (0.001) [0.936]	0.000 (0.001) [0.942]	-0.000 (0.001) [0.929]
× Richer household		0.000 (0.001) [0.426]	
Richer household		-0.006 (0.002) [0.000]	
× At least primary schooling			0.000 (0.001) [0.692]
At least primary schooling			-0.007 (0.003) [0.010]
Mean (dep.var.)	0.062	0.059	0.062
Identifying observations	693,353	601,757	693,353
Singleton observations	3	2	3
Communities	30,932	26,589	30,932
Countries	36	36	36
Birth year range	1940–2002	1952–2002	1940–2002

Note. Estimates are based on equation (2.2). The dependent variable is an indicator variable equal to 1 if the woman reports being pregnant at the time of the interview, and 0, otherwise. *pH in the nearest waters* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include location FEs using grid cells at the 1°×1° resolution, interview month FEs, interview year FEs, country by interview month FEs, country by interview year FEs, and control variables (see Section 2.4; weather controls correspond to the year of interview). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B2.6: Table 2.3 – exclude women reporting at least one terminated pregnancy

Dependent variable:	NMR (deaths per 1,000 births)					
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	-1.864 (0.841) [0.027]	-1.935 (0.828) [0.020]	-1.956 (0.806) [0.016]	-1.896 (0.854) [0.027]	-1.950 (0.861) [0.024]	-1.925 (0.834) [0.021]
Mean (dep.var.)	29.125	29.125	29.128	29.125	29.125	29.127
Identifying observations	1,145,187	1,145,187	1,143,788	1,145,180	1,145,180	1,143,781
Singleton observations	66	66	66	73	73	73
Communities	28,764	28,764	28,764	28,764	28,764	28,764
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes	-	Yes	Yes
Demographic controls	-	-	Yes	-	-	Yes
Seasonality	Country	Country	Country	Cell	Cell	Cell

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs. Seasonality is captured by either country by birth month FEs or 5°×5° cell by birth month FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B2.7: Table 2.3 – restricting the sample to recent births


Dependent variable:	NMR (deaths per 1,000 births)					
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	-2.535 (1.316) [0.055]	-2.402 (1.331) [0.072]	-2.441 (1.308) [0.063]	-2.024 (1.141) [0.077]	-2.020 (1.147) [0.079]	-2.106 (1.131) [0.063]
Mean (dep.var.)	26.914	26.914	26.917	26.914	26.914	26.918
Identifying observations	746,982	746,982	745,962	746,960	746,960	745,940
Singleton observations	142	142	142	164	164	164
Communities	31,183	31,183	31,183	31,182	31,182	31,182
Countries	36	36	36	36	36	36
Birth year range	1980–2018	1980–2018	1980–2018	1980–2018	1980–2018	1980–2018
Weather controls	-	Yes	Yes	-	Yes	Yes
Demographic controls	-	-	Yes	-	-	Yes
Seasonality	Country	Country	Country	Cell	Cell	Cell

Note. Estimates are based on equation (2.2) restricting the sample to births within 10 years of the interview. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, and control variables (see Section 2.4). Controls for local seasonality are either country by birth month FEs or 5°×5° cell by birth month FEs. Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.4 Fishing and seafood dependency

Sea Around Us. Figure B2.5 plots the time series of fish catch (in quantity and value) for the countries in the study sample, differentiated by sector of activity during the period 1972–2018 (panel A), and the total catch over the same period by geographical region, again restricting the sample to the countries being part of the study (panel B).

Figure B2.5: Descriptive statistics about fish catch in the study area



Note. Evolution over time of the quantity and landed value of fish catch in the study area (panel A), and total quantity and landed value by geographical region in the period 1972–2018 (panel B). Source is Pauly et al. (2020). See Section 2.3 for further details.

Table B2.8 shows estimates of the effect of pH on landed value, distinguishing by the price of

the commercial group, by the main nutrient of fish in the commercial group, and by resilience to acidification. Lower resilience is among crustaceans and molluscs ([Alter et al., 2024](#)). Table [B2.9](#) replicates Table [2.1](#) in the main text, but using alternative transformations of the outcome variables. Table [B2.10](#) estimates (conditional) correlations of measures of nutrition among women and country-level landed value at the time of the interview. Because variation is available only yearly and at the level of the country, we control for FAO fishing zone-specific FEs, trends and seasonality.

Table B2.8: Ocean acidity and marine catches – by type

Dependent variables:	Landed value, by category of seafood caught					
	Price		Main nutrient		Resilience	
	Low (1)	High (2)	Fatty fish (3)	Lean fish (4)	Lower (5)	Higher (6)
A. Landed value (small-scale)						
pH in proximity to the coast	0.196 (0.113) [0.092]	0.198 (0.079) [0.018]	0.229 (0.098) [0.025]	0.188 (0.083) [0.030]	0.293 (0.111) [0.012]	0.179 (0.084) [0.041]
Mean (dep.var.)	1.80	1.35	1.44	1.54	2.07	1.39
Identifying observations	6,863	12,234	5,217	13,912	3,478	15,651
Singleton observations	25	7	0	0	0	0
Countries	36	36	36	36	36	36
Year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
B. Landed value (industrial)						
pH in proximity to the coast	0.029 (0.111) [0.794]	0.031 (0.095) [0.748]	0.065 (0.132) [0.626]	0.021 (0.102) [0.841]	-0.031 (0.144) [0.830]	0.047 (0.102) [0.647]
Mean (dep.var.)	1.21	0.89	0.73	1.11	1.70	0.85
Identifying observations	6,863	12,234	5,217	13,912	3,478	15,651
Singleton observations	25	7	0	0	0	0
Countries	36	36	36	36	36	36
Year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
C. Median price						
pH in proximity to the coast	-0.034 (0.025) [0.174]	-0.077 (0.045) [0.095]	-0.087 (0.046) [0.066]	-0.091 (0.053) [0.096]	-0.130 (0.077) [0.102]	-0.077 (0.041) [0.069]
Mean (dep.var.)	0.49	1.53	0.78	1.10	1.19	0.95
Identifying observations	6,863	6,684	3,885	9,716	3,238	10,340
Singleton observations	25	31	2	0	25	0
Countries	36	36	36	36	36	36
Year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018

Note. Estimates are based on equation (2.1). Dependent variables in column headers are the landed value, by type of catch. Landed values are reported using an inverse hyperbolic sine transformation to account for zero values. In columns (1)–(2), price heterogeneity is defined by computing the average unit price by seafood group within each country for the whole period 1972–2018. *Low (high)* indicates values that are smaller or equal (larger) than the median value. In columns (3)–(4), *fatty fish* indicates seafood groups whose primary nutrient content are essential fatty acids, while *lean fish* indicates seafood groups whose primary nutrient content are essential proteins. In columns (5)–(6), *lower (higher)* indicates seafood groups whose resilience to acidification is lower (higher). *pH in proximity to the coast* is defined in Section 2.4. Specifications include country by seafood group FEs, and fishing area by year FEs. Standard errors (in parentheses) are clustered at the EEZ level, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B2.9: Ocean pH and marine catches – robustness checks

Dependent variables:	Marine catch for human consumption				Activity	
	Small-scale fishing Quantity (1)	Value (2)	Industrial fishing Quantity (3)	Value (4)	Median price (5)	Night-time luminosity (6)
A. Levels						
pH in proximity to the coast	7.160 (2.723) [0.012]	9.555 (3.508) [0.010]	3.654 (5.218) [0.488]	4.287 (5.247) [0.419]	-0.228 (0.111) [0.046]	-0.212 (0.147) [0.157]
Mean (dep.var.)	16.10	18.52	14.00	13.37	1.52	5.48
Identifying observations	19,129	19,129	19,129	19,129	13,603	777
Singleton observations	0	0	0	0	0	0
Countries	36	36	36	36	36	36
Year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1992–2012
B. Log-levels						
pH in proximity to the coast	0.123 (0.060) [0.048]	0.176 (0.070) [0.017]	0.023 (0.079) [0.774]	0.037 (0.088) [0.674]	-0.070 (0.034) [0.048]	-0.001 (0.033) [0.965]
Mean (dep.var.)	1.25	1.26	0.83	0.84	0.79	1.76
Identifying observations	19,129	19,129	19,129	19,129	13,603	777
Singleton observations	0	0	0	0	0	0
Countries	36	36	36	36	36	36
Year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1992–2012

Note. Estimates are based on equation (2.1). Dependent variables are reported in column headers with different transformations depending on the panel. In panel A, we use the variables in levels, but top-coding them at the 99th percentile of the distribution to limit the effect of outliers on the estimates. In panel B, we compute the logarithm of the variables, add unity before computing logarithms to avoid creating missing values. In columns (1)–(4), each observation is the catch or landed value for a specific commercial group in the correspondent exclusive economic zone (EEZ). In columns (5)–(6), each observation is the country correspondent to the EEZ. *pH in proximity to the coast* is defined in Section 2.4. Specifications in columns (1)–(4) include country by commercial group FEs, and fishing area by year FEs. The specification in column (5) includes country FEs, and fishing area by year FEs. Standard errors (in parentheses) are clustered at the EEZ level, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table B2.10: Fisheries and nutrition among women

Dependent variable:	Nutrition		Prevalence of anaemia	
	Consumed seafood (1)	Underweight (2)	Overall (3)	In pregnancy (4)
Landed value (small-scale)	0.474 (0.234) [0.043]	0.006 (0.003) [0.087]	-0.035 (0.007) [0.000]	-0.103 (0.020) [0.000]
Landed value (industrial)	-0.254 (0.085) [0.003]	-0.002 (0.001) [0.057]	0.009 (0.003) [0.012]	0.024 (0.009) [0.006]
Mean (dep.var.)	0.296	0.120	0.427	0.455
Identifying observations	49,047	407,702	272,690	14,708
Singleton observations	0	0	0	0
Communities	5,954	24,301	17,371	9,027
Countries	14	32	26	26
Interview year range	2005–2016	1992–2018	2000–2018	2000–2018

Note. Estimates are based on equation (2.2). Dependent variables are reported in column headers and defined in Appendix A2.1. *Landed values* are defined in Section 2.3. The sample is restricted to communities in the coastal area (see Section 2.3). All specifications include location FEs using FAO fishing zone, interview month FEs, interview year FEs, fishing zone by interview month FEs, fishing zone by interview year FEs, and control variables (see Section 2.4; controls include ocean pH in the nearest waters; weather controls correspond to the year of interview). Standard errors (in parentheses) are clustered at the DHS community level, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Night-time and extractive fishing. Table B2.11 reports *p*-values of a test of heterogeneous effects by intensity of extractive and night-time fishing, produced estimating equation (2.2) on

the set of outcomes presented in Figure 2.5 by adding interaction terms between ocean pH while *in utero* and each of these variables and testing for joint equality to 0 of the coefficients on the interaction term(s). Table B2.12 shows descriptive statistics comparing areas with low versus high intensity of both types of fishing.

Table B2.11: Test of heterogeneous effects

Type of interaction:	Linear <i>p</i> -value (1)	Quadratic <i>p</i> -value (2)	Cubic <i>p</i> -value (3)
<i>NMR</i>			
Intensity of extractive fishing	0.296	0.044	0.096
Intensity of night-time fishing	0.637	0.850	0.625
<i>Physical development</i>			
Intensity of extractive fishing	0.166	0.535	0.847
Intensity of night-time fishing	0.642	0.619	0.490

Note. The table reports *p*-values for joint tests of equality to zero of the estimates on the interaction term(s). Estimates are based on equation (2.2) adding interaction terms between *pH in the nearest waters (in utero)* and the variables presented in the left column. *Linear* includes only the linear interaction, *quadratic* adds the quadratic interaction, and *cubic* adds the cubic interaction. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors are clustered at the ocean raster data point. All specifications include cluster fixed effects, birth year by birth month fixed effects, country by birth year fixed effects (local trend), country by birth month fixed effects (local seasonality), and time-varying controls (climatic/weather and demographic). The full list of controls is presented in Section 2.3. Observations are re-weighted to correct for oversampling of countries surveyed multiple times (see Appendix A2.1). *pH in the nearest waters (in utero)* is defined in Section 2.4. Appendix A2.1 provides further information on the variables and the list of surveys included in the study. We exclude DHS surveys for Peru as information for the intensity of night-time fishing is not available (see Appendix A2.1).

Table B2.12: Descriptive statistics by degree of extractive and night-time fishing

	Extractive fishing				Night-time fishing			
	High		Low		High		Low	
	Mean (1)	SD (2)	Mean (3)	SD (4)	Mean (5)	SD (6)	Mean (7)	SD (8)
A. Children								
Child is alive	0.92	0.27	0.93	0.26	0.92	0.26	0.92	0.27
Child is female	0.48	0.50	0.48	0.50	0.48	0.50	0.48	0.50
Birth order	2.51	1.79	2.56	1.81	2.53	1.77	2.54	1.83
Birth order in a multiple birth	0.04	0.25	0.03	0.22	0.04	0.24	0.03	0.23
Years since birth	12.22	7.91	12.34	7.84	12.16	7.82	12.39	7.91
Mother's age at birth of child	24.39	5.77	24.46	5.75	24.55	5.76	24.32	5.76
pH in the nearest waters (<i>in utero</i>)	8.05	0.03	8.05	0.03	8.05	0.03	8.05	0.03
B. Adult women								
Age at first delivery	20.92	4.31	20.85	4.18	21.00	4.32	20.79	4.18
Current age	30.68	9.73	30.63	9.88	30.69	9.70	30.63	9.89
Years of schooling	7.69	4.97	7.34	4.81	7.51	4.90	7.51	4.88
pH in the nearest waters (<i>in utero</i>)	8.07	0.03	8.07	0.03	8.07	0.03	8.07	0.03
At least primary schooling	0.69	0.46	0.68	0.47	0.68	0.47	0.69	0.46
Married	0.65	0.48	0.68	0.47	0.67	0.47	0.67	0.47
Working	0.49	0.50	0.47	0.50	0.47	0.50	0.49	0.50
Household head is female	0.23	0.42	0.21	0.40	0.22	0.42	0.21	0.41
Household head's age	46.01	13.03	46.18	13.19	45.64	13.07	46.48	13.14
Household members	5.65	3.25	5.59	2.79	5.60	3.00	5.63	3.05
Household wealth	3.62	1.28	3.20	1.37	3.54	1.33	3.29	1.34
Living in urban area	0.61	0.49	0.46	0.50	0.58	0.49	0.49	0.50
Distance from shore	22.54	27.10	39.55	30.68	23.04	27.09	38.16	30.96
Distance from another water body	36.82	50.36	57.17	132.93	35.46	49.23	57.14	129.88
Altitude	125.63	326.87	250.78	464.55	119.75	258.39	248.51	492.71
Temperature (° C)	26.15	3.70	26.47	2.83	26.53	2.87	26.13	3.59
Precipitations (mm)	1454.63	689.13	1660.70	644.30	1555.98	705.47	1563.89	647.25
Intensity of extractive fishing	0.23	0.49	0.00	0.00	0.20	0.49	0.04	0.16
Intensity of night-time fishing	0.16	0.28	0.10	0.35	0.29	0.43	0.01	0.01
C. Mortality rates								
Neonatal	27.87	164.61	26.95	161.94	28.41	166.15	26.53	160.70
Postneonatal	24.25	153.84	22.83	149.35	22.08	146.93	24.71	155.25
Child	23.65	151.94	19.56	138.50	21.00	143.37	21.92	146.43
Infant	51.54	221.10	49.33	216.56	49.94	217.82	50.76	219.51
Under-five	77.06	266.69	70.77	256.44	72.55	259.39	74.77	263.02

Note. Descriptive statistics of coastal areas by degree of extractive and night-time fishing. We define areas with above-median extractive or night-time fishing intensity as *high* intensity. We define areas with below-median extractive or night-time fishing intensity as *low* intensity. Coastal area includes all communities within 100 km from the ocean's shore (see Section 2.3). Means are reported in columns (1), (3), (5), and (7); standard deviations are reported in columns (2), (4), (6), and (8). *Years since birth* is measured at the time of the interview and is independent from the child being alive. *Mortality rates* are relative to 1,000 live births. *pH in the nearest waters (in utero)* is the average pH in the ocean grid cell closest to an individual's community during the 9 months before birth; it refers to the date of birth of the child in panel A and to the date of birth of the woman in panel B. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Seafood prices in the Philippines. We gather prices for the Philippines, a relevant setting in our context: its coastline is the 5th largest in the world, it is home to 9% of global coral reefs, and depends highly on fish. We gather monthly retail seafood prices at the province level for the period 1990–2018 from the [Philippine Statistics Authority \(2020\)](#). To compute real prices, data are supplemented with information about consumer prices indices and exchange rates from the [International Monetary Fund \(2025\)](#). We compute exposure to the average seafood price while *in utero* matching retail prices with individual information using the date and the province of birth. For identification, we rely on deviations in average retail seafood prices in logarithms from the spatially specific (and seasonally-adjusted) long-run trend by adding this variable in equation (2.2). Table B2.13 shows the results.


Table B2.13: Early-life exposure and market prices in the Philippines

Dependent variable:	NMR (deaths per 1,000 births)			
	(1)	(2)	(3)	(4)
Average seafood price (<i>in utero</i>)	5.768 (3.002) [0.056]	5.887 (3.000) [0.051]	5.791 (2.999) [0.055]	5.882 (2.983) [0.050]
pH in the nearest waters (<i>in utero</i>)		-5.147 (2.691) [0.057]	-5.797 (3.088) [0.062]	-5.687 (3.223) [0.079]
Mean (dep.var.)	15.410	15.410	15.410	15.412
Identifying observations	82,739	82,739	82,739	82,730
Singleton observations	9	9	9	9
Communities	2751	2751	2751	2751
Birth year range	1990–2017	1990–2017	1990–2017	1990–2017
Weather controls	-	-	Yes	Yes
Demographic controls	-	-	-	Yes

Note. Estimates are based on equation (2.2) using the benchmark specification. The sample is restricted to communities in the coastal area of the Philippines and to the period 1990–2018 (due to data availability; see Section 2.3). *pH in the nearest waters (in utero)* is defined in Section 2.4. *Average seafood price (in utero)* is the average fish price (including all available prices and reported in logarithms) in the province of birth of the child during the 9 months before birth. Standard errors (in parentheses) are clustered at the district by ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, district by birth year FEs, and district by birth month FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Seafood dependency. Figure B2.6 presents descriptive statistics for seafood dependency and estimates of the heterogeneous effect of ocean pH in the nearest waters (*in utero*) on NMR distinguishing by a country's fish dependency. As a separate measure of dependency on artisanal fishing, panel D focuses on heterogeneity by proximity to **coral reefs** (conditional on being at the shore).

Figure B2.6: Early-life exposure and neonatal mortality – by seafood dependency



Note. Panel A shows the average value of seafood proteins as share of total animal proteins by selected area (weighted by population) or by country, obtained from [FAO \(2019\)](#). In the right figure, vertical lines indicate the world's average (solid) and the average among the selected countries (dashed). Panel B shows heterogeneous effects of ocean pH while *in utero* on NMR by dependency on seafood proteins (*high* indicates the top tercile of the 1960–2013 sample distribution). Marginal effects are estimated using equation (2.2) restricting the sample to the corresponding group. Because the heterogeneity variable is at country-level, the specification includes community FEs, birth year by birth month FEs, $5^\circ \times 5^\circ$ grid cell by birth year FEs, $5^\circ \times 5^\circ$ grid cell by birth month FEs, and control variables (see Section 2.4). Panel C shows marginal effect of ocean pH in the nearest waters (*in utero*) on NMR as a function of shortest distance from a coral reef, estimated using equation (2.2) interacting ocean pH in the nearest waters (*in utero*) with distance from shore and distance from coral reefs and assuming 0km-distance from the shore. The specification includes community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). The dependent variable is an indicator variable equal to 1 if the child died within the first month of life and 0 if the child survived, multiplied by 1,000. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors are clustered at the ocean raster data point. In panel B, we report the 90% confidence interval. In panel C, the 90% confidence interval is indicated by dotted lines, beyond which the intervals are progressively shaded up to the 99% level. Within confidence bounds, color intensity reflects the relative density of observations across iterations. It is calculated by comparing the density in each iteration to a range between the lower bound (adjusted by 0.7) and the 99th percentile of densities across all iterations. These parameters were chosen to improve visibility. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.5 Selective migration

DHS surveys collect information on women's location at the time of the interview and, in only some of our selected surveys, the years a woman has lived in that location. If the respondent migrated, DHS does not provide information on the location before migration. We leverage data for the available countries and assess migration responses following the shock. Table B2.14 presents estimates of the effect of ocean acidity on the probability that the mother migrated to the community of the interview within the first five years following delivery. Mothers do not

adapt along this margin.

Table B2.14: Post-delivery selective migration

Dependent variable:	Mother migrated to community 0-4 years after delivery of child					
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	0.000 (0.002) [0.989]	-0.000 (0.002) [0.957]	0.000 (0.002) [0.963]	0.001 (0.003) [0.749]	0.002 (0.003) [0.518]	0.002 (0.003) [0.531]
Mean (dep.var.)	0.112	0.112	0.112	0.112	0.112	0.112
Identifying observations	1,016,246	1,016,246	1,015,068	1,016,242	1,016,242	1,015,064
Singleton observations	15	15	15	19	19	19
Communities	21,884	21,884	21,884	21,884	21,884	21,884
Countries	28	28	28	28	28	28
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes	-	Yes	Yes
Demographic controls	-	-	Yes	-	-	Yes
Seasonality	Country	Country	Country	Cell	Cell	Cell

Note. Estimates are based on equation (2.2). The dependent variable is an indicator variable equal to 1 if the mother of the child migrated to the community of the interview in the first 5 years of life of the child, and 0 otherwise. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, and control variables (see Section 2.4). Controls for local seasonality are either country by birth month FEs or 5°×5° cell by birth month FEs. Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Results presented in Figure B2.9 in Section B2.11 indicate that our effect is indeed specific to the location of the interview. To further address concerns about matching children with the correct exposure location, Table B2.15 estimates the effect of ocean acidity on neonatal mortality for different sub-samples selected based on whether the respondent was present in the location of the interview at the time of exposure. Estimates are based on our preferred specification (specification 3 in Table 2.3).

To address missing migration data in surveys, we identify respondents who were not present in the location of the interview at the time of exposure to ocean pH in the nearest waters using two approaches. First, we utilize information collected by DHS surveys. We refer to these surveys as *surveys with migration data*. Second, we expand information about migration to all surveys conducted for a country if at least one round of data collection includes the migration question. In this case, we impute missing information using (Poisson) LASSO regressions. In each regression, we consider as predictors numerical variables (various aspects of women’s characteristics, including current age, age at first birth, total number of births, number of children, number of living children, number of deceased children, age of the household head, and years of education), and factor variables (categorical variables that are included as indicators for each of their values, including a woman’s pregnancy status, year of birth, and whether she has ever been married, finished at least primary education, has access to electricity, owns a television, owns a radio, owns a refrigerator, and owns a car). For each country, we train the model on a random half of the sample with available information on migration, employing three different selection methods: cross-validation, Bayesian information criterion, and adaptive cross-validation. We predict the number of years a woman has spent in the location of the interview in the surveys with missing information about migration using the selection method

that performs best in out-of-sample prediction.

In columns (1)–(4), we focus on our full sample, but exclude children whose mothers moved to the location of the interview after the gestation period. Thus, we exclude children who, according to available information, were not in the location of the interview while *in utero*. In columns (1)–(2), we exclude those whose mothers migrated after the gestation period (assumed to be the year before the year of birth). In columns (3)–(4), we also exclude those who migrated after the year prior to the gestation period (assumed to be two years before the year of birth). Columns (1) and (3) use only survey data, while columns (2) and (4) integrate imputed data. Note that this approach combines surveys with and without information. Despite this limitation, estimates are, in absolute terms, larger but not statistically different compared to column (3) in Table 2.3. This suggests that excluding children who were not in the location of the interview while *in utero* for all surveys would lead to larger estimates.

In columns (5)–(6), we instead focus on the sample of migrants by restricting the sample to surveys with migration data and selecting children whose mothers reported having relocated to the location of the interview at some point during their life. In column (5), we include only children whose mothers migrated to the location of the interview in the year before the year of birth (during the gestation) or earlier. In column (6), to exclude temporary or short-term migrants, we further restrict the sample by including only long-term migrants, i.e., those who migrated to the location of the interview at least 17 years before (the median time a migrant has been living in the location).

Table B2.15: Table 2.3 – excluding children not in the interview location while *in utero*

Dependent variable: Sub-sample of children:	All	All	NMR (deaths per 1,000 births)			
			All	All	Mother is migrant	Mother is long-term migrant
	(1)	(2)	(3)	(4)	(5)	(6)
A. Benchmark specification						
pH in the nearest waters (<i>in utero</i>)	-1.903 (0.775) [0.014]	-1.938 (0.773) [0.013]	-1.893 (0.772) [0.015]	-1.916 (0.770) [0.013]	-1.967 (1.142) [0.086]	-2.631 (0.982) [0.008]
Mean (dep.var.)	30.069	29.674	30.157	29.644	27.624	30.406
Identifying observations	1,318,153	1,284,406	1,276,164	1,233,551	304,197	157,073
Singleton observations	115	116	149	153	871	832
Communities	31,236	31,233	31,175	31,169	19,631	15,115
Countries	36	36	36	36	28	28
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
B. Within-sibling specification						
pH in the nearest waters (<i>in utero</i>)	-2.371 (1.016) [0.020]	-2.476 (1.015) [0.015]	-2.448 (0.905) [0.007]	-2.539 (0.903) [0.005]	-2.036 (1.750) [0.246]	-2.800 (1.164) [0.017]
Mean (dep.var.)	31.383	31.042	31.568	31.125	30.109	31.431
Identifying observations	1,205,060	1,168,353	1,163,706	1,117,548	258,563	149,341
Singleton observations	114,768	117,725	114,115	117,659	46,819	8,710
Communities	30,987	30,982	30,767	30,759	18,863	14,867
Countries	36	36	36	36	28	28
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
Selected surveys:	All (Appendix Table A2.2)				With migration data	
Exclusion condition:	Mother migrated after gestation		Mother migrated after the year before gestation		Mother migrated after gestation	
Source of migration data:						
Survey	Yes	-	Yes	-	Yes	Yes
Survey + imputed	-	Yes	-	Yes	-	-

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to coastal areas (see Section 2.3), and to the sub-samples reported at the bottom of the table. *Surveys with migration data* includes all surveys in which information about migration is available. The *exclusion condition* indicates when the children is dropped from the sample and refers to the time when the mother of the child migrated to the location of the interview (if the condition is true, the observation is dropped). *Long-term migrants* are respondents that migrated to the location of interview at least 17 years before the interview (the median time in the location of interview among migrants). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, and country by birth month FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.6 Issues related to identification and measurement

Selection into identification. To verify the validity of our estimates of selection into identification, columns (1)–(3) in Table B2.16 present estimates of the benchmark specification restricting the sample to the identifying observations of the within-sibling specification, while columns (4)–(6) provide estimates of the effect using the identifying sample of the within-sibling specification and re-weighting as in Miller et al. (2021). To estimate the probability of being in the identifying sample of the within-sibling specification, we use a probit model and include mother and weather characteristics.

Table B2.16: The effect on neonatal mortality: identification checks

Dependent variable: Check:	NMR (deaths per 1,000 births)					
	Benchmark specification with within-sibling identifying sample			Re-weighting procedure		
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	-1.398 (0.737) [0.059]	-1.430 (0.724) [0.049]	-1.582 (0.706) [0.025]	-2.356 (0.870) [0.007]	-2.383 (0.857) [0.006]	-2.495 (0.849) [0.003]
Mean (dep.var.)	31.476	31.476	31.476	31.478	31.478	31.478
Identifying observations	1,474,945	1,474,945	1,474,945	1,474,353	1,474,353	1,474,353
Singleton observations	0	0	0	108,737	108,737	108,737
Communities	31,356	31,356	31,356	31,356	31,356	31,356
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes	-	Yes	Yes
Demographic controls	-	-	Yes	-	-	Yes

Note. In columns (1)–(3), estimates are based on equation (2.2) using the benchmark specification and restricting the sample to the identifying sample of the within-sibling specification. In columns (4)–(6), estimates are based on equation (2.2) using the within-sibling specification and the re-weighting procedure of Miller et al. (2021). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, and country by birth month FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Balance of mother characteristics. Table B2.17 presents estimates of equation (2.2) without control variables where the dependent variable is replaced by demographic controls.


Table B2.17: Placebo test: balance on observable characteristics

Dependent variable:	Age			Education		Other characteristics		
	First delivery (1)	Delivery (2)	Interview (3)	Years (4)	Primary or less (5)	Married (6)	Working (7)	Wealth (8)
pH in the nearest waters (<i>in utero</i>)	0.010 (0.016) [0.542]	0.002 (0.021) [0.938]	0.002 (0.021) [0.939]	0.014 (0.016) [0.379]	-0.001 (0.002) [0.426]	-0.000 (0.001) [0.807]	-0.001 (0.002) [0.641]	0.002 (0.003) [0.387]
Mean (dep.var.)	20.092	25.086	36.682	4.916	0.439	0.887	0.558	2.947
Identifying observations	1,583,706	1,583,706	1,583,706	1,583,065	1,583,630	1,583,705	1,454,950	1,339,312
Singleton observations	25	25	25	25	25	25	28	31
Communities	31,380	31,380	31,380	31,380	31,380	31,380	28,828	27,039
Countries	36	36	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018

Note. Estimates are based on equation (2.2) without control variables. The dependent variable is an indicator variable equal to 1 if the child died within the first month of life and 0 if the child survived, multiplied by 1,000. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. The full set of controls is reported in the bottom panel of the table, control variables are excluded. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Measurement error in the distance from the ocean. Figure B2.7 shows the distribution of the coefficients estimating iteratively equation (2.2) simulating a random error in the measurement of the distance from the shoreline of ± 10 , 30, and 50km.

Figure B2.7: The effect on neonatal mortality, by magnitude of measurement error




Note. Distribution of the marginal effect of ocean pH in the nearest waters (*in utero*) on NMR, estimated using equation (2.2) and introducing measurement error in the distance from the ocean. *pH in the nearest waters (in utero)* is defined in Section 2.4. The procedure performs 1,000 iterations. The vertical line represents our benchmark point estimate (column 3 in Table 2.3). The distribution fits are estimated non-parametrically using kernel density estimation and assuming an Epanechnikov kernel function. Bandwidths are estimated by Silverman’s rule of thumb. The sample is restricted to communities in the coastal area (see Section 2.3). Appendix A2.1 provides further information on the variables and the full list of surveys included in the study.

B2.7 Aggregate shocks

Comparison with droughts. We compare two shocks: one to agricultural productivity, as captured by the presence of a drought, and one to ocean pH. The first surely captures an income shock given the importance of rainfall on agriculture and the reliance of L&MICs on this economic activity. We generate a gridded dataset with a spatial resolution of $0.5^\circ \times 0.5^\circ$, and construct a yearly panel of night-time luminosity for the grid cells containing coastal DHS communities. We estimate equation (2.2) at the grid-cell level, matching each grid to its nearest point along the shoreline. We then compute drought and acidity exposure at this coastal location. To deal with zeros, we apply the inverse hyperbolic sine transformation on night-time luminosity. We define droughts using indicator variables taking value one when annual precipitations in the grid cell are at or below a defined percentile of the grid cell’s historical precipitations distribution (see, e.g., [Burke et al., 2015](#) for this approach). Because defining a shock depends on the long-run precipitation distribution, we use the University of Delaware Terrestrial Precipitation dataset ([UDEL, 2018](#)), provided by the National Oceanic and Atmospheric Administration (NOAA), which includes monthly data from 1900 to 2012. For comparability, we follow the same approach to define an *acidity shock*. These shocks do not occur simultaneously, with a raw contemporaneous correlation of -14.4%. Figure B2.8 presents the estimates of the effect on night-time luminosity of an acidity shock (panel A) and a drought (panel B) across various percentile bounds (reported in the horizontal axis) used to define these events.

Comparison with conflict. Using information about conflict events from the Uppsala Conflict Data Programme (UCDP) database at the $5^\circ \times 5^\circ$ resolution, we estimate equation (2.2) adding controls for the presence and the intensity of conflict while *in utero*. Table B2.18 presents estimates of the effect on NMR. Due to data availability, the birth year range is reduced to children born after 1984. For comparability, columns (3) and (6) are therefore restricted to the sample included in column (1) and (4), respectively.

Figure B2.8: Night-time luminosity: comparing acidification with droughts



Note. Estimates are based on equation (2.2). The dependent variable is the inverse hyperbolic sine-transformed satellite-based night-time luminosity at year t in grid cell i , ranging in levels between 0 (lowest) and 1 (highest). In panel A, an *acidity shock* is defined by a binary variable with a value of 1 when the yearly average pH in the nearest ocean grid cell is below the x^{th} percentile of the historical distribution of grid cell i , and 0 otherwise. In panel B, a *drought* is defined by a binary variable with a value of 1 when annual precipitations in grid cell i are below the x^{th} percentile of its historical precipitations distribution, and 0 otherwise. All specifications include grid cell FEs and year by macro-region FEs. Macro-regions are defined by $2.5^{\circ} \times 2.5^{\circ}$ resolution grid cells to guarantee sufficient variation over time of both acidity shocks and droughts. In panel A, standard errors are clustered at the ocean raster data point, and *controls* include the oxygen concentration in the nearest ocean grid cell, rainfall, temperature, and their interaction. In panel B, standard errors are clustered at the grid cell level, and *controls* include temperature levels. The sample is restricted to grid cells that contain one or more coastal DHS communities, where the median distance from the shoreline across all communities within a cell does not exceed 100 kilometers, and where water bodies cover no more than 50 percent of the cell area. (see Section 2.3). Further information on variables and the list of surveys is provided in Appendix A2.1.

Table B2.18: Comparing the effect size of ocean pH and conflict

Dependent variable:	NMR (deaths per 1,000 births)					
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	-1.044 (0.658) [0.113]	-1.048 (0.661) [0.113]	-1.048 (0.659) [0.112]	-1.609 (0.810) [0.048]	-1.613 (0.808) [0.046]	-1.619 (0.810) [0.046]
At least 1 violent event (<i>in utero</i>)	1.584 (1.085) [0.145]			1.586 (1.109) [0.153]		
Fatalities (<i>in utero</i>)		0.797 (0.414) [0.055]			0.812 (0.420) [0.054]	
Mean (dep.var.)	27.504	27.504	27.504	27.504	27.504	27.504
Identifying observations	1,227,711	1,227,711	1,227,711	1,227,703	1,227,703	1,227,703
Singleton observations	93	93	0	101	101	0
Communities	31,268	31,268	31,268	31,268	31,268	31,268
Countries	36	36	36	36	36	36
Birth year range	1989–2018	1989–2018	1989–2018	1989–2018	1989–2018	1989–2018
Seasonality	Country	Country	Country	Cell	Cell	Cell

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, p -values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, and control variables (see Section 2.4). Controls for local seasonality are either country by birth month FEs or $5^{\circ} \times 5^{\circ}$ cell by birth month FEs. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Labour supply. Table B2.19 estimates the effect of the contemporaneous ocean pH in the nearest waters on the labour supply of female respondents and their partners (when available). Columns (1) and (3) focus on any job, while columns (2) and (4) focus on agricultural and fishing jobs. Agriculture and fishing are distinct activities only in a few surveys, we grouped these activities together to maximize the number of observations.

Table B2.19: Contemporaneous exposure and labour supply

Dependent variable:	Respondent works in ...		Partner works in ...	
	Any job (1)	Agriculture/fishing (2)	Any job (3)	Agriculture/fishing (4)
pH in the nearest waters	-0.002 (0.003) [0.507]	-0.004 (0.002) [0.132]	0.002 (0.008) [0.842]	-0.005 (0.020) [0.782]
Mean (dep.var.)	0.473	0.089	0.897	0.234
Identifying observations	623,798	601,826	42,285	42,240
Singleton observations	3	3	0	0
Grid cells	821	821	143	143
Communities	28,865	28,846	2,989	2,989
Countries	36	36	7	7
Interview year	1990–2018	1990–2018	2015–2018	2015–2018

Note. Estimates are based on equation (2.2). Dependent variables are indicator variables equal to 1 if the respondent reports being working in the corresponding job at the time of the interview, and 0 otherwise. *pH in the nearest waters* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the female respondent's community in the month of the interview. The sample is restricted to coastal areas (see Section 2.3). Columns (1)–(2) use the full sample of women, whereas columns (3)–(4) restrict the sample to women who are either the household head or their partner. All specifications include location FEs using grid cells at the $1^\circ \times 1^\circ$ resolution, interview month FEs, interview year FEs, country by interview month FEs, country by interview year FEs, and control variables (see Section 2.4; weather controls correspond to the year of interview). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.8 Responses in food consumption

Table B2.20 shows estimates of the effect of ocean pH at the time of the interview on the probability of consuming different food items, considering macronutrients in columns (1)–(3), and micronutrients in column (4).

Table B2.20: Contemporaneous exposure and food consumption among women

Dependent variable:	Other proteins (1)	Carbohydrates (2)	Fats (3)	Other iron-rich food (4)
pH in the nearest waters	0.011 (0.013) [0.402]	-0.010 (0.006) [0.100]	-0.014 (0.016) [0.376]	0.004 (0.008) [0.639]
Mean (dep.var.)	0.774	0.870	0.689	0.870
Identifying observations	50,059	49,159	47,151	50,211
Singleton observations	2	2	2	2
Grid cells	255	239	239	255
Communities	6,120	5,956	5,588	6,122
Countries	15	14	14	15
Interview year range	2005–2016	2005–2016	2005–2015	2005–2016

Note. *Other proteins* and *other iron-rich food* includes only food items other than seafood. Estimates are based on equation (2.2). *pH in the nearest waters* is the average pH (multiplied by a factor of 100) in the ocean grid cell closest to the female respondent's community in the month of the interview. The sample is restricted to coastal areas (see Section 2.3). All specifications include location FEs using grid cells at the $1^\circ \times 1^\circ$ resolution, interview month FEs, interview year FEs, country by interview month FEs, country by interview year FEs, and control variables (see Section 2.4; weather controls correspond to the year of interview). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.9 Robustness to assuming a shorter gestation period

Table B2.21 replicates Table 2.3 assuming a gestation period of 8 months.

Table B2.21: Table 2.3 – assuming shorter gestation period

Dependent variable:	NMR (deaths per 1,000 births)					
	(1)	(2)	(3)	(4)	(5)	(6)
A. Benchmark specification						
pH in the nearest waters (<i>in utero</i>)	-0.887 (0.539) [0.100]	-0.846 (0.551) [0.125]	-0.893 (0.531) [0.093]	-1.793 (0.611) [0.004]	-1.733 (0.628) [0.006]	-1.742 (0.612) [0.005]
Mean (dep.var.)	30.473	30.473	30.474	30.474	30.474	30.475
Identifying observations	1,583,706	1,583,706	1,581,815	1,583,703	1,583,703	1,581,812
Singleton observations	25	25	25	28	28	28
Communities	31,380	31,380	31,380	31,380	31,380	31,380
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
B. Within-sibling specification						
pH in the nearest waters (<i>in utero</i>)	-1.404 (0.654) [0.032]	-1.379 (0.667) [0.039]	-1.472 (0.654) [0.025]	-2.128 (0.828) [0.010]	-2.119 (0.828) [0.011]	-2.268 (0.823) [0.006]
Mean (dep.var.)	31.476	31.476	31.476	31.476	31.476	31.476
Identifying observations	1,474,945	1,474,945	1,474,945	1,474,941	1,474,941	1,474,941
Singleton observations	108,786	108,786	108,786	108,790	108,790	108,790
Communities	31,356	31,356	31,356	31,356	31,356	31,356
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes	-	Yes	Yes
Demographic controls	-	-	Yes	-	-	Yes
Seasonality	Country	Country	Country	Cell	Cell	Cell

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4, but assuming a gestation period of 8 months instead of 9. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs. Seasonality is captured by either country by birth month FEs or 5°×5° cell by birth month FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.10 Early-life mortality rates

Table B2.22 presents estimates of the effect of ocean pH in the nearest waters (experienced *in utero*) on early-life mortality.

Table B2.22: The effect on early-life mortality rates

Dependent variables:	Post-neonatal (PMR)	Child (CMR)	Infant (IMR)	Under-5 (U5MR)
	(1)	(2)	(3)	(4)
pH in the nearest waters (<i>in utero</i>)	1.088 (0.505) [0.032]	0.164 (0.382) [0.667]	-0.243 (0.678) [0.720]	0.015 (0.976) [0.987]
Mean (dep.var.)	28.659	31.870	58.469	92.020
Identifying observations	1,468,359	1,140,070	1,514,850	1,215,597
Singleton observations	27	56	29	57
Communities	31,370	31,329	31,371	31,330
Countries	36	36	36	36
Birth year range	1972–2018	1972–2013	1972–2018	1972–2013

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and weather and demographic controls. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.11 Supplementary results on inference

Clustering. Table B2.23 shows estimates of equation (2.2) for NMR using different assumptions for the clustering of standard errors (reported in column).

Table B2.23: Robustness to clustering assumptions about standard errors


Dependent variable:	NMR (deaths per 1,000 births)					
Level of clustering:	None	1°x1° grid cell	Matched ocean cell	5°x5° grid cell	Country x survey year	Community
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	-1.493 (0.663) [0.025]	-1.493 (0.625) [0.017]	-1.493 (0.429) [0.001]	-1.493 (0.667) [0.025]	-1.493 (0.647) [0.023]	-1.493 (0.610) [0.015]
Mean (dep.var.)	30.474	30.474	30.474	30.474	30.474	30.474
Identifying observations	1,581,815	1,581,815	1,581,815	1,581,815	1,581,815	1,581,815
Singleton observations	25	25	25	25	25	25
Communities	31,380	31,380	31,380	31,380	31,380	31,380
Countries	36	36	36	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018	1972–2018

Note. Estimates are based on equation (2.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to the coastal area (Section 2.3). All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Standard errors are reported in parenthesis, *p*-values are reported in brackets. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Permutation-based inference. Focusing on Table 2.3, we implement three different tests. In the *birth dates within communities* test, birth dates are randomly reassigned within each community. In the *birth dates within countries* test, birth dates are randomly reassigned within each country, independently from the community and the survey. In the *across communities* test, mothers (and their children) are randomly allocated to different communities, indepen-

dently from the country and the survey. Figure B2.9 shows the distribution of estimates in each test and the empirical p -values.

Figure B2.9: The effect on neonatal mortality: permutation-based inference



Note. Distributions of marginal effects of ocean pH in the nearest waters (*in utero*) on NMR when birth dates are randomly reassigned. Tests are described in Appendix B2.11, and are based on 500 iterations. *pH in the nearest waters (in utero)* is defined in Section 2.4. Each graph depicts the empirical distribution of estimates using the specification in each of the columns in Table 2.3. In each iteration, marginal effects are estimated using equation (2.2). The sample is restricted to communities in the coastal area (see Section 2.3). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.12 Parental investments and adaptation


Table B2.24 shows estimates of the effect of *in-utero* exposure to the ocean pH on parental health investments and on health outcomes associated with poor contemporaneous nutrition. Figure B2.10 shows instead the effect of ocean pH on the probability of being underweight, distinguishing by the age of the child at the time of the measurement. The dependent variable is an indicator variable equal to 1 if the child has a weight-for-age z -score below negative 2 standard deviations, and 0 otherwise. Table B2.25 presents the same estimates presented in Table 2.4, but adding an interaction term between ocean pH and indicator variables capturing different dimensions. Panel A focuses on whether the household is part of the richest half of the sample, panel B on whether the mother is married, and panel C on whether the mother has at least primary school education.

Table B2.24: Parental investments and postnatal nutritional outcomes

Dependent variables:	Antenatal		Delivery		Nutrition	
	Number of visits (1)	w/ health professional (2)	In health centre (3)	w/ health professional (4)	Morbidity (5)	Anæmia (6)
pH in the nearest waters (<i>in utero</i>)	0.005 (0.010) [0.606]	-0.001 (0.004) [0.706]	0.000 (0.003) [0.980]	-0.003 (0.003) [0.215]	-0.002 (0.004) [0.659]	0.002 (0.006) [0.751]
Mean (dep.var.)	1.703	0.884	0.670	0.638	0.391	0.558
Identifying observations	269,330	156,134	176,787	267,900	339,407	114,370
Singleton observations	1,040	215	424	1,032	871	1,437
Communities	30,034	14,671	18,559	30,031	29,932	15,844
Countries	36	29	29	36	36	27
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018	1985–2018	1995–2018

Note. Estimates are based on equation (2.2). The dependent variables are reported in the column's header. *Morbidity* is an indicator variable equal to 1 if the child has experienced fever, cough or diarrhea in the weeks previous to the interview, and 0 otherwise. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. For cross-survey comparability, the samples are restricted to the last birth, independently from the child being alive. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Figure B2.10: Effect on the probability of being underweight



Note. Marginal effect of *pH in the nearest waters (in utero)* on the probability of the child to be underweight. *pH in the nearest waters (in utero)* is defined in Section 2.4. The dependent variable is an indicator variable equal to 1 if the child has a weight-for-age *z*-score below negative 2 standard deviations, and 0 otherwise. Confidence intervals are computed at 90% level. Estimates are based on equation (2.2) including community FEs, birth month by birth year FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Standard errors are clustered at the ocean raster data point. Appendix A2.1 provides further information on the variables and for the list of surveys included in the study.

Table B2.25: Early-life exposure and parental adaptation – heterogeneity


Dependent variables:	Antenatal	Delivery	Postnatal		
	(1)	(2)	Healthcare	Breastfed	Vaccinated
	(1)	(2)	(3)	(4)	(5)
A. By socio-economic status					
pH in the nearest waters (<i>in utero</i>)	-0.006 (0.008) [0.463]	-0.004 (0.005) [0.420]	0.004 (0.009) [0.644]	0.001 (0.003) [0.830]	-0.003 (0.007) [0.688]
× Richer household	0.007 (0.005) [0.180]	0.003 (0.004) [0.493]	0.001 (0.003) [0.805]	0.002 (0.001) [0.056]	0.002 (0.002) [0.426]
Richer household	0.066 (0.021) [0.002]	0.077 (0.012) [0.000]	0.019 (0.009) [0.032]	-0.001 (0.003) [0.789]	0.031 (0.008) [0.000]
Mean (dep.var.)	1.818	1.388	0.441	0.969	0.436
Identifying observations	140,538	153,196	101,075	172,692	177,763
Singleton observations	186	448	3,078	2,078	2,006
Communities	13,154	16,968	18,445	24,039	24,139
Countries	28	28	34	36	36
Birth year range	1998–2018	1998–2018	2002–2018	1999–2018	1999–2018
B. By marital status					
pH in the nearest waters (<i>in utero</i>)	-0.003 (0.007) [0.665]	-0.002 (0.005) [0.732]	0.008 (0.009) [0.397]	0.001 (0.003) [0.806]	-0.004 (0.007) [0.515]
× Married	0.002 (0.003) [0.373]	0.000 (0.004) [0.940]	-0.003 (0.002) [0.082]	0.000 (0.001) [0.665]	0.003 (0.002) [0.114]
Married	0.063 (0.010) [0.000]	0.036 (0.016) [0.023]	-0.004 (0.008) [0.621]	0.012 (0.003) [0.000]	0.036 (0.008) [0.000]
Mean (dep.var.)	1.808	1.359	0.441	0.972	0.435
Identifying observations	155,980	168,460	101,075	206,350	208,764
Singleton observations	217	481	3,078	2,336	2,269
Communities	14,669	18,481	18,445	28,029	27,887
Countries	29	29	34	36	36
Birth year range	1985–2018	1985–2018	2002–2018	1987–2018	1987–2018
C. By education					
pH in the nearest waters (<i>in utero</i>)	-0.003 (0.008) [0.724]	-0.004 (0.006) [0.518]	0.004 (0.009) [0.650]	0.000 (0.003) [0.943]	-0.003 (0.006) [0.650]
× At least primary schooling	0.002 (0.004) [0.625]	0.003 (0.005) [0.605]	0.000 (0.003) [0.865]	0.001 (0.001) [0.009]	0.002 (0.001) [0.143]
At least primary schooling	0.031 (0.016) [0.051]	0.086 (0.021) [0.000]	0.030 (0.011) [0.005]	0.002 (0.003) [0.593]	0.017 (0.007) [0.021]
Mean (dep.var.)	1.808	1.359	0.441	0.972	0.435
Identifying observations	155,980	168,460	101,075	206,350	208,765
Singleton observations	217	481	3,078	2,336	2,269
Communities	14,669	18,481	18,445	28,029	27,887
Countries	29	29	34	36	36
Birth year range	1985–2018	1985–2018	2002–2018	1987–2018	1987–2018

Note. Estimates are based on equation (2.2), adding an interaction term between ocean pH and different indicator variables. *Richer household* is an indicator variable equal for whether the household's wealth index is above the sample median. *Married* and *at least primary schooling* are indicator variables for whether the mother is married or has completed at least primary education, respectively. The dependent variables and sample restrictions are reported in Table 2.4. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures. Column (3) excludes the survey(s) for Indonesia and Morocco because information is not available in the corresponding surveys.

To test for the presence of adaptation, we introduce an additional check related to the effect on

neonatal mortality. To verify whether observable characteristics that could explain adaptation predict the effect on neonatal mortality, Figure B2.11 presents estimates of heterogeneous effects on neonatal mortality, distinguishing by children and mothers' demographics (panel A) and for location characteristics (panel B).

Figure B2.11: Ocean acidity and neonatal mortality – split sample heterogeneity



Note. Heterogeneous effects of pH in the nearest waters (in utero) on NMR by child and mother's demographics (panel A), and by location's characteristics (panel B). pH in the nearest waters (in utero) is defined in Section 2.4. Marginal effects are estimated using equation (2.2) restricting the sample to the corresponding group. For mother's age at birth, wealth index, agricultural land, population, we create an indicator variable indicating whether an observation is above or below the full sample's median of the variable of interest. Agricultural land and population are set at the 1970 level. Standard errors are clustered at the ocean raster data point. confidence intervals are computed at 90% level. All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

B2.13 Physical development and adulthood outcomes

Tables B2.26 and B2.27 replicate Table 2.5, but providing estimates separately for boys and girls and trimming observations in which z -scores are smaller than the 1st or larger than the 99th percentiles of the z -score distribution. Table B2.28 replicates Table 2.5 using anthropometrics among adults. For adults older than 18 years old, z -scores refer to standard reference curves at age 18, when physical development is assumed to be complete. Table B2.29 shows estimates on economic well-being. In column (1), we proxy economic well-being with a measure of wealth, computed as an asset-based index. Columns (2)–(6) focus on correlates of well-being, such as fertility (number of births), years of schooling, cognitive skills (determined by the ability to read a sentence), and labour supply. To account for family structure and capture intra-household dynamics, columns (1) and (6) select only women that are either a household head or their partner, while columns (2)–(5) refer to the full sample of women aged 15–49. Economic outcomes measured at household level or influenced by intra-household dynamics are more likely to be observed in women that are married and leading a family. Figure B2.12 shows heterogeneity of the effects on physical development and economic well-being by the intensity and type of fishing.

Table B2.26: Table 2.5 – estimates by sex at birth

Dependent variables:	z-scores		Indicators	
	W/h (1)	H/a (2)	Wasted (3)	Stunted (4)
A. Boys				
pH in the nearest waters (<i>in utero</i>)	-0.002 (0.018) [0.895]	-0.032 (0.020) [0.102]	0.009 (0.006) [0.141]	0.011 (0.007) [0.123]
Mean (dep.var.)	-0.338	-1.041	0.084	0.243
Identifying observations	115,472	115,528	115,472	115,528
Singleton observations	3,370	3,437	3,370	3,437
Communities	21,209	21,423	21,209	21,423
Countries	33	33	33	33
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018
B. Girls				
pH in the nearest waters (<i>in utero</i>)	-0.014 (0.019) [0.468]	0.024 (0.020) [0.226]	-0.004 (0.007) [0.594]	-0.013 (0.006) [0.037]
Mean (dep.var.)	-0.285	-0.942	0.076	0.227
Identifying observations	111,095	111,157	111,095	111,157
Singleton observations	3,508	3,577	3,508	3,577
Communities	20,843	21,052	20,843	21,052
Countries	33	33	33	33
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018

Note. Estimates are based on equation (2.2) estimated separately for boys (panel A) and girls (panel B). Dependent variables are reported in the column's header and defined in Appendix A2.1. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures. All panels exclude the survey(s) for Indonesia, Pakistan, and the Philippines because information is not available in the correspondent surveys.

Table B2.27: Table 2.5 – trimming z-scores

Dependent variables:	z-scores		Indicators	
	Weight-for-height (1)	Height-for-age (2)	Wasted (3)	Stunted (4)
A. Overall effect				
pH in the nearest waters (<i>in utero</i>)	-0.021 (0.016) [0.196]	-0.012 (0.015) [0.405]	0.006 (0.003) [0.090]	0.004 (0.004) [0.279]
Mean (dep.var.)	-0.309	-0.984	0.080	0.234
Identifying observations	232,339	232,575	232,339	232,575
Singleton observations	1,106	1,124	1,106	1,124
Communities	24,824	25,110	24,824	25,110
Countries	33	33	33	33
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018
B. Heterogeneity by sex				
pH in the nearest waters (<i>in utero</i>)	-0.001 (0.017) [0.969]	-0.033 (0.020) [0.100]	0.008 (0.006) [0.179]	0.011 (0.007) [0.147]
× female	-0.014 (0.020) [0.461]	0.057 (0.028) [0.047]	-0.011 (0.010) [0.303]	-0.023 (0.010) [0.022]
Mean (dep.var.)	-0.312	-0.993	0.080	0.236
Identifying observations	226,567	226,685	226,567	226,685
Singleton observations	6,878	7,014	6,878	7,014
Communities	23,979	24,248	23,979	24,248
Countries	33	33	33	33
Birth year range	1985–2018	1985–2018	1985–2018	1985–2018

Note. Estimates are based on equation (2.2). The sample includes observations trimmed to exclude z-scores below the 1st percentile or above the 99th percentile of the distribution. Dependent variables are reported in the column's header and defined in Appendix A2.1. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. Specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables. In panel B, FEs are sex specific. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures. All panels exclude the survey(s) for Indonesia, Pakistan, and the Philippines because information is not available in the correspondent surveys.

Table B2.28: Early-life exposure and physical development among adult women

Dependent variables:	z-scores		Indicators	
	W/h (1)	H/a (2)	Wasted (3)	Stunted (4)
pH in the nearest waters (<i>in utero</i>)	0.011 (0.007) [0.121]	0.010 (0.005) [0.072]	0.000 (0.001) [0.957]	-0.008 (0.003) [0.021]
Mean (dep.var.)	-0.310	-1.386	0.082	0.301
Identifying observations	324,160	327,124	324,160	327,124
Singleton observations	554	683	554	683
Communities	22,635	22,848	22,635	22,848
Countries	32	32	32	32
Birth year range	1972–2003	1972–2003	1972–2003	1972–2003


Note. Estimates are based on equation (2.2). Dependent variables are reported in the column's header. *W/h* (weight-for-height) and *h/w* (height-for-age) are z-scores from a reference scale. *Wasted* is an indicator variable equal to 1 for an for an abnormally low weight-for-height. *Stunted* is an indicator variable equal to 1 for an abnormally low height-for-age, and 0 otherwise. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. In Panels A and B, specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables. In panel C, specifications include community FEs, woman's birth year by woman's birth month FEs, country by woman's birth year FEs, country by mother's birth month FEs, and control variables (see Section 2.4). Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures. We exclude the survey(s) for Angola, Indonesia, Pakistan, and the Philippines because information is not available in the correspondent surveys.

Table B2.29: Early-life exposure and long-run economic well-being

Dependent variables:	Economic well-being	Fertility	Schooling	Cognitive skills	Labour supply	
	(1)	(2)	(3)	(4)	(5)	(6)
pH in the nearest waters (<i>in utero</i>)	0.016 (0.009) [0.062]	-0.008 (0.004) [0.046]	0.032 (0.035) [0.361]	0.000 (0.002) [0.924]	0.005 (0.004) [0.154]	0.014 (0.007) [0.040]
Mean (dep.var.)	3.096	1.552	7.183	0.771	0.425	0.513
Identifying observations	212,741	497,982	433,480	414,000	429,270	190,645
Singleton observations	1,161	536	538	794	549	2,255
Communities	25,432	30,429	27,878	26,824	27,860	24,720
Countries	36	36	36	36	36	36
Birth year range	1972–2003	1972–2003	1972–2003	1972–2003	1972–2003	1972–2003
Sample:	Head/partner	All	All	All	All	Head/partner

Note. Estimates are based on equation (2.2). The dependent variables are reported in the column's header. *Economic well-being* is a household-level asset-based index which ranges from 1 (poorest) to 5 (richest). *Fertility* is the number of births per woman. *Schooling* is the number of completed years of education. *Cognitive skills* is an indicator variable equal to 1 if the respondent is able to read a whole sentence in her native language or has completed at least secondary schooling, and 0 otherwise. *Labour supply* is an indicator variable equal to 1 if the respondent is working at the time of the interview, and 0 otherwise. *pH in the nearest waters (in utero)* is defined in Section 2.4. The sample is restricted to coastal areas (see Section 2.3), and in columns (1)–(6) to women in the household that are household head or their partner. Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, woman's birth year by woman's birth month FEs, country by woman's birth year FEs, country by woman's birth month FEs, and control variables (see Section 2.4). Column (2)–(4) have a reduced number of observations because, for comparability of estimates, we include only the random sub-sample of women that completed both the education and the work modules. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Figure B2.12: Acidity shocks and resource exploitation – effects on adults



Note. Estimated impacts of a one-standard-deviation increase in acidity on long-run indicators as a function of intensity of fishing. Intensities range between 0 (no presence) and 1 (high). Estimates are based on equation (2.2) introducing interaction terms between ocean pH and a quadratic polynomial in the corresponding intensity. *Economic well-being* is a household-level asset-based index which ranges from 1 (poorest) to 5 (richest). *Physical development* is the average z-score of available anthropometric measures. The sample is restricted to communities in the coastal area (see Section 2.3). We exclude surveys for Peru as information for the intensity of night-time fishing is not available (see Appendix A2.1). All specifications include community FEs, birth year by birth month FEs, country by birth year FEs, country by birth month FEs, and control variables (see Section 2.4). Standard errors are clustered at the ocean raster data point. Confidence intervals are computed at 90% level. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Chapter 3

The Future Impact of Ocean Acidification on Child Health

This chapter is co-authored with Alex Armand (NOVA SBE)

Abstract

While short-term variations in ocean acidity are known to adversely affect neonatal health in coastal communities, the long-term consequences and the potential for human adaptation remain poorly understood. This paper first provides empirical tests for long-run adaptation to ocean acidification by examining whether historical exposure has led to increased resilience in 36 low- and middle-income countries. Our analysis finds no evidence of historical adaptation; on the contrary, exposure to either historically more acidic waters or faster rates of acidification appears to amplify vulnerability. Building on this finding, we quantify the historical mortality burden, identifying coastal hotspots where the impact has been greatest. We then use these results to project the impact of ocean acidification on neonatal mortality through 2100 under various climate and societal scenarios. Projections indicate that, absent adaptation, ocean acidification could cause 77.2 million neonatal deaths by 2100 under a high-emissions scenario (RCP8.5), a figure that is halved under a low-emissions pathway (RCP2.6). These findings quantify the substantial public health co-benefits of climate change mitigation and underscore the need for proactive adaptation policies, as endogenous adaptation appears insufficient to counter the threat.

Keywords: Climate Change; Ocean; Acidification; Health; Mortality; Development; Projections

JEL Codes: I15, Q20, Q54, O10

3.1 Introduction

Since the Industrial Revolution, the world's oceans have absorbed a significant portion of anthropogenic carbon dioxide emissions, functioning as the planet's largest carbon sink. This absorption has altered the ocean's chemical balance, leading to a 26% increase in surface water acidity—a phenomenon known as ocean acidification (Doney et al., 2020). This ongoing chemical shift poses a threat to marine life. While the effects vary by species, scientific studies show that the biological impacts are generally large and negative (Kroeker et al., 2010; Alter et al., 2024), a finding supported by laboratory experiments (Wittmann and Pörtner, 2013). This knowledge gap is emphasized by IPCC (2019), which states in its latest report about the ocean that there is no direct evidence on how climate change may have already contributed to malnutrition by decreasing seafood supply in coastal communities.

Chapter 2 of this thesis established a causal link between short-term variations in ocean acidity and adverse child health outcomes in the coastal regions of 36 low- and middle-income countries. The analysis uses data from all available Demographic and Health Surveys (DHS) (Croft et al., 2018) from countries with a coastline for which coordinates are available, resulting in a repeated cross-section of more than 1.5 million births and 500,000 women from 1972 to 2018. The findings show that *in utero* exposure to increased ocean acidity significantly raises the risk of neonatal mortality. This effect is primarily driven by maternal malnutrition, as higher acidity reduces the catch of small-scale fisheries, which in turn increases seafood prices and lowers the consumption of essential nutrients by pregnant women. The analysis further suggests that these nutritional shocks were not severe enough to trigger an observable adaptive response in maternal or child health investments.

The findings from Chapter 2 raise important questions about the long-term implications of ocean acidification. First, do communities adapt as ocean acidity gradually worsens over decades? Second, what has been the historical burden of ocean acidification on neonatal mortality in coastal communities? Third, what is the potential scale of this health problem in the future? This chapter addresses these questions by first testing for evidence of long-run adaptation to ocean acidification. Based on those results, we then quantify the historical burden of ocean acidification on neonatal mortality and subsequently project its future impacts through the end of the 21st century. Our analysis continues to focus on the coastal populations of the 36 low- and middle-income countries studied previously, a region that is home to 41% of the world's coastal population and is particularly vulnerable to disruptions in marine resources.

Our empirical analysis finds no discernible evidence of long-run adaptation to ocean acidification. Coastal communities that were initially exposed to higher levels of acidity do not experience a smaller mortality impact from new shocks; on the contrary, the effect is larger, which suggests that long-term exposure may exacerbate vulnerability rather than foster resilience. Furthermore, we test whether adaptation is more pronounced in locations where the long-term

decline in pH has been steepest and thus more noticeable. Consistent with our previous result, we find that the mortality impact is more severe in these areas. This finding suggests that a greater salience of the environmental threat does not trigger effective adaptation and instead points to increased vulnerability.

Building on these empirical insights, our chapter also quantifies the historical mortality burden attributable to ocean acidification between 1975 and 2018. We reveal that ocean acidification, on average, has contributed positively to neonatal mortality across all 36 countries, but this impact is highly heterogeneous both across and within regions. We find the adverse impact is generally more significant in communities situated closer to the shore, consistent with the observed reliance on small-scale fisheries in these areas. While some locations show a negative attributable impact, in the vast majority of areas the mortality burden is substantial, with some communities enduring an attributable mortality burden exceeding 10 deaths per 1,000 live births. This spatial variation underscores the need to identify specific coastal hotspots, such as the extensive coastlines of the Philippines, India, East Africa, and parts of Central and South America, where the burden has been particularly high.

These empirical results directly inform our projection methodology. Based on our findings, we use a *no adaptation* case as our benchmark scenario. However, recognising that this represents a stark, though empirically grounded, baseline, we also model several more optimistic scenarios that allow for future adaptation and mitigation. This approach enables us to project a range of potential outcomes, from a worst-case trajectory to futures where policy interventions and societal changes may lessen the mortality burden.

Our projections indicate a substantial mortality burden from future ocean acidification. Under a high-emissions scenario (RCP8.5), and assuming no adaptation, ocean acidification is projected to cause a cumulative 77.2 million neonatal deaths in the studied countries by 2100. The analysis highlights the significant benefits of climate change mitigation; a global shift to a low-emissions pathway (RCP2.6) would reduce this figure to 38.0 million, averting an estimated 39.2 million deaths. The difference between these outcomes underscores the sensitivity of future child health to current energy policy choices. Furthermore, the limited evidence of long-run adaptation suggests that realising additional reductions in mortality will require proactive and targeted policy interventions, as endogenous adaptation appears insufficient.

This chapter makes two primary contributions. First, we contribute to the literature that links short-run environmental shocks to long-run adaptation, a central challenge in projecting future climate impacts. While previous work has tested for long-run adaptation to changes in temperature (Deschênes and Greenstone, 2011; Dell et al., 2012; Burke et al., 2015), we provide the first empirical tests for long-run adaptation to ocean acidification. By examining whether communities that were initially exposed to more acidic waters show any signs of increased resilience, our results provide an empirical grounding for the assumptions used in our projection

model.

Our second and main contribution is to provide the first quantitative projections of the cost of ocean acidification in terms of human lives. This work adds to the emerging literature on the impacts of ocean acidification on human communities (e.g., [Colt and Knapp 2016](#)) and extends the broader climate change literature, which has generally focused on the future impacts of rising temperatures (e.g., [Moore and Diaz 2015](#); [Burke et al. 2015](#)). By translating future emissions scenarios into a tangible human cost—the number of neonatal lives at risk—we provide a new input into the global calculus of climate change damages (e.g., see [Auffhammer 2018](#)).

3.2 Data

The empirical analysis in this chapter builds upon the comprehensive dataset constructed for Chapter 2. For a detailed description of the primary data sources, variable construction, and sample selection criteria, readers are referred to Section 2.3. This section provides a summary of the core data used for the analysis on long-run adaptation and introduces the new data sources incorporated for the projection exercise.

Demographic and Health Surveys. We use data from 92 surveys of the Demographic and Health Surveys (DHS) conducted between 1990 and 2018 across 36 low- and middle-income countries with a coastline. The DHS provides nationally representative data on population and health, with a specific focus on maternal and child health. A key feature of the DHS is the detailed birth history recorded for each female respondent, which allows for the precise calculation of early-life mortality rates. Our primary outcome variable is the Neonatal Mortality Rate (NMR), defined as the number of deaths per 1,000 live births that occur within the first month of life. The analysis is restricted to communities located in coastal areas, defined as those within 100 kilometres of the shoreline, calculated as the shortest straight-line distance ([United Nations, 2003](#)). The dataset also includes a range of control variables for the mother, child, and household used in our regression models.

Ocean Acidity. We use two distinct sources for ocean pH data to link our empirical analysis with future projections. For the analysis of long-run adaptation, we use historical ocean pH data from the Hadley Global Environment Model 2 - Earth System (HadGEM2-ES). This model provides monthly global raster data at a $1^\circ \times 1^\circ$ resolution for the period 1972–2018, which we match to DHS communities using a nearest-proximity criterion. For the projections, we instead use global forecasts of ocean acidification from the IPCC’s Sixth Assessment Report ([Tagliabue et al., 2022](#)). These data provide the annual average surface ocean pH from 2006 to 2100 under two different emissions scenarios: a low-emissions scenario (RCP2.6) and a high-emissions scenario (RCP8.5).

Population and Birth Projections. A critical component of this chapter is the projection of the number of live births in the specific coastal areas under study. To the best of our knowledge, such projections do not exist. We, therefore, construct a yearly time series from 1975 to 2100 through a multi-step process. First, we use projections of country-level population and crude birth rates from the UN’s medium-fertility scenario ([United Nations, 2024](#)). This scenario, which assumes a gradual decline in global fertility rates toward replacement levels, serves as the UN’s baseline. Second, because our focus is on coastal areas, we assume a distribution of future populations between coastal and inland areas within each country. Factors such as trade, inequality, tourism, migration, and coastal management can influence population dynamics between coastal and inland areas. To address this, we draw on data from [Merkens et al. \(2016\)](#), which provides a gridded population dataset every five years from 2005 to 2100 for different scenarios. We consider the SSP1 scenario, which assumes a sustainable development path, and the SSP5 scenario, which anticipates rapid economic growth driven by fossil fuels. For each country and the world, we compute the coastal population share for each available year, applying these shares to project population and birth dynamics in coastal areas based on the UN’s estimates. Between 2005 and 2020, the global coastal population is estimated to have grown from 2.54 billion to 3 billion. Under SSP1, it is projected to decrease slightly to 2.98 billion by 2100, while under SSP5, it is expected to rise to 3.2 billion. Although RCPs and SSPs are not directly interchangeable, we assume RCP2.6 aligns with SSP1 and RCP8.5 with SSP5.

Weather Controls. To account for potentially correlated climatic factors in our historical analysis, we incorporate yearly weather controls from the PRIO-GRID v2.0 database ([Tollefsen et al., 2012](#)). This provides data at a $0.5^\circ \times 0.5^\circ$ resolution, which we spatially merge to DHS communities. Specifically, we control for total annual precipitation (in millimetres), drawn from the GPCP v.2.2 dataset, and mean annual temperature (in degrees Celsius), drawn from the GHCN/CAMS dataset, and their interaction.

3.3 Methodology

This section details the two-part approach used in this chapter. We first outline the empirical strategy used to test for long-run adaptation to ocean acidification. The strategy uses a multi-way fixed effects model to examine whether initial exposure or the salience of long-term trends has led to any adaptation responses. We then describe the methodology used to project the future neonatal mortality burden, which builds upon the findings from the adaptation tests to define a set of credible scenarios.

3.3.1 Testing for long-run adaptation

The core identification strategy exploits short-term, plausibly exogenous temporal and geographical variations in ocean pH near coastal communities. We match communities to the nearest ocean data point and link this environmental data to individual birth histories from DHS. Identification relies on isolating the impact of these shocks from confounding factors such as long-term local trends and seasonality. Conditional on this demanding structure, these short-run deviations in pH are plausibly exogenous, analogous to shocks in weather patterns. The benchmark estimating equation is:

$$y_{ic,mt} = \beta R_{c,mt} + \mathbf{X}_{ic,mt}\gamma + \Omega_{c,mt} + \epsilon_{ic,mt} \quad (3.1)$$

The key independent variable is the ocean pH, $R_{c,mt}$, matched to each child's community c , and month m and year t of birth. As pH is a logarithmic scale, a 0.01-unit decrease corresponds to an approximate 2.3% increase in acidity; a shock of this magnitude is equivalent to the median within-year variation for a given location and roughly one-third of a standard deviation of in-utero exposure in our sample. Our identification comes from a multi-way fixed effects model that isolates the impact of short-run pH deviations. The fixed effects structure includes DHS community fixed effects (location effects), birth year-by-month fixed effects (non-spatial time effects), country-by-birth month fixed effects (region-specific seasonality) and country-by-birth year fixed effects (region-specific trends). The vector of controls, $\mathbf{X}'_{ic,mt}$, further accounts for child characteristics (sex, birth order, twins), maternal and household variables (mother's age and its square, age at birth, education, household size), and local weather conditions (annual temperature, precipitation, and their interaction).

To extend the analysis from short-run shocks to long-run adaptation, this chapter performs two empirical tests.

First, we check whether communities initially exposed to higher levels of ocean acidity show any signs of increased resilience to new shocks (Dell et al., 2014). The intuition behind this test is that communities adapted to a more challenging historical baseline may have already developed economic structures, technologies, or social norms that make them less vulnerable. For example, a city with a historically hot climate is likely better adapted to a heatwave than a cooler city. Similarly, communities in historically more acidic waters may be less reliant on pH-sensitive species and thus more resilient to further increases in acidity.

Second, we test whether communities in areas with a steeper long-term decline in pH experience a smaller mortality impact. This test is motivated by the idea of adaptation as a dynamic learning process. A more rapid and noticeable environmental change provides a stronger, more salient signal of the threat. If communities learn and adapt to this signal, we would expect to see the smallest negative impacts in the very places where the problem has been most apparent

over the last several decades.

For the first test, we interact the contemporaneous pH shock with a measure of initial exposure. Let $\bar{R}_{c,0}$ denote the average pH in the nearest waters for community c during the baseline period of 1972–1975, which we standardise to ease interpretation. The estimating equation is then:

$$y_{ic,mt} = \beta_1 R_{c,mt} + \beta_2 (R_{c,mt} \times \bar{R}_{c,0}) + \mathbf{X}'_{ic,mt} \Gamma + \Omega_{c,mt} + \epsilon_{ic,mt} \quad (3.2)$$

The coefficient of interest for this test is β_2 , which captures the differential effect of a pH shock based on initial conditions. Evidence of adaptation would be a negative and statistically significant β_2 . This would suggest that the adverse impact of a decrease in $R_{c,mt}$ (higher acidity) is attenuated in locations that were initially more acidic (i.e., had a lower $\bar{R}_{c,0}$). Conversely, a positive β_2 would imply a compounding effect, where communities with worse initial environmental conditions are more vulnerable to subsequent shocks.

For the second exercise, we compute the long-term trend in ocean pH for each ocean raster cell, defined as the slope of the linear regression of pH on time over the period 1972–2018, corrected for seasonality. We then create an indicator, T_c , that equals one if a location's rate of acidification is above the sample median (i.e., exhibits a faster rate of acidification), and zero otherwise. The variation in these long-term trends across our sample is non-negligible. For example, the trend at the 5th percentile of the distribution (-0.0017 pH units/year) is markedly different from the trend at the 95th percentile (-0.0013 pH units/year). Over a 30-year period, this difference implies that a location at the 5th percentile of the trend distribution becomes approximately 2.3% more acidic than a location at the 95th percentile, solely due to the difference in their long-term acidification rates. This absolute change of 0.012 pH units is also large in a spatial context, corresponding to approximately 32% of one cross-sectional standard deviation in pH levels (0.037).

To test for adaptation based on the salience of the trend, we employ a fully flexible specification where the effects of covariates and fixed effects are also allowed to vary with the trend indicator, T_c . This demanding specification controls for any observable and unobservable differences between high- and low-trend areas that might also have differential time trends. The estimating equation for this test is:

$$y_{ic,mt} = \beta_1 R_{c,mt} + \beta_2 (R_{c,mt} \times T_c) + \mathbf{X}'_{ic,mt} \Gamma + (\mathbf{X}'_{ic,mt} \times T_c) \Delta + \Omega_{c,mt,T_c} + \epsilon_{ic,mt} \quad (3.3)$$

In this fully interacted model, β_1 represents the effect of ocean acidity in locations with a below-median rate of acidification ($T_c = 0$). The additional effect for locations with an above-median rate of acidification ($T_c = 1$) is given by β_2 . Evidence of adaptation would be a positive and statistically significant β_2 , as this would indicate that the total effect for high-trend

areas ($\beta_1 + \beta_2$) is less negative than the effect for low-trend areas (β_1). This specification is particularly robust as it allows the effects of all control variables and fixed effects to differ between the two groups.

3.3.2 Estimating historical impacts and projecting future burden

Our projection methodology begins with an estimation of the historical neonatal mortality rate NMR attributable to ocean acidification. This estimate is constructed from the community level upwards, using birth-level data from the DHS.

First, we predict the birth-level NMR under observed pH conditions for all births within 100 kilometres of the shore. To increase predictive power, this estimation uses a flexible version of equation (3.1) which allows the effect of *in utero* exposure to ocean pH to vary with distance from the shore via interaction terms: distance, distance squared, and cubed. Next, we calculate the counterfactual NMR for each child born at time mt in community c , $\text{NMR}_{ic,mt}^{CF}$, by applying the same model but holding the *in utero* pH exposure at its 1975 level for that location and month. This procedure preserves natural seasonality while removing long-term trends and shock-related variations post-1975. The difference between the predicted and counterfactual NMR yields the mortality rate attributable to acidification for each child. This value is then averaged at the community level to obtain NMR_c^{OA} .

These community-level estimates, however, are confined to the discrete locations surveyed by the DHS. To create a continuous surface of the mortality impact across each country's entire coastal zone, interpolation is required. For this purpose, we employ ordinary kriging to impute NMR_c^{OA} at all unmeasured locations within 100km of a country's coastline. Ordinary kriging is a geostatistical method that predicts values at unmeasured locations by fitting a spatial auto-correlation function—a semi-variogram—to the data (Cuadros et al., 2017). The process is implemented using the `autoKrig` function in the `automap` R package (Hiemstra et al., 2009), which automates model selection. It fits several semi-variogram models, such as spherical, exponential, and Gaussian, and selects the one that minimises the residual sum of squares. From this best-fitting model, it determines three key parameters: the nugget, representing the variance at zero distance; the range, defined as the distance at which spatial correlation becomes negligible; and the sill, which is the total variance of the process. This fitted model is then used to predict values across the continuous coastal surface. Interpolation is performed separately for each country, though due to limited coastal variation, the Democratic Republic of Congo and Namibia are estimated jointly with Angola, and India is estimated jointly with Pakistan.

Once values at unmeasured locations are imputed, we use the LandScan Global 2023 database (Lebakula et al., 2024), which provides a distribution of population density across the globe at approximately 1km $\times$ 1km spatial resolution, to create a population mask. We replace interpo-

lated values with missing values in areas with no population.

To construct the historical time series for our analysis, we aggregate the original birth-level estimates to the country-year level. This yields two key series: the predicted total NMR for country k in year t , denoted $\text{NMR}_{k,t}$, and the counterfactual NMR, $\text{NMR}_{k,t}^{CF}$. The portion of mortality attributable to ocean acidification, $\text{NMR}_{k,t}^{OA}$, is defined by their relationship:

$$\text{NMR}_{k,t} = \text{NMR}_{k,t}^{CF} + \text{NMR}_{k,t}^{OA} \quad (3.4)$$

While this process provides a historical time series from 1975 to 2018, projecting to 2100 requires further steps. First, we project $\text{NMR}_{k,t}^{CF}$ forward by fitting an exponential decay curve to the 1975–2018 series and extrapolating to 2100; this model achieves an adjusted R^2 of 99.5%. For projections of $\text{NMR}_{k,t}^{OA}$, we use ocean acidification forecasts from the IPCC Sixth Assessment Report (IPCC, 2022), considering two scenarios for 2006–2100: a low-emissions scenario, RCP2.6, and a high-emissions scenario, RCP8.5. We use the historical simulation for the 1975–2005 period for consistency.¹

We estimate the future attributable mortality, NMR_t^{OA} , using the following expression:

$$\text{NMR}_t^{OA} = \left[\frac{d \text{NMR}}{d R} \right]_t \times \frac{\Delta \text{pH}_{1975,t}}{0.01} \quad (3.5)$$

where $\left[\frac{d \text{NMR}}{d R} \right]_t$ is the time- t effect on NMR of *in-utero* exposure to ocean pH, as obtained from equation (3.1), and $\Delta \text{pH}_{1975,t}$ is the change in pH from 1975 to time t . The denominator scales the effect to a 0.01 unit change in pH, consistent with our main analysis; see Section 2.3.

Finally, by combining the projected NMR^{CF} and NMR^{OA} series with coastal population projections, we compute the cumulative number of neonatal deaths attributable to ocean acidification under several adaptation scenarios. First, we consider *no adaptation*, where the effect of ocean acidification on NMR is constant over time at the level estimated in our main specification from column 3 of Table 2.3. This, combined with the high-emissions scenario, forms our worst-case baseline. Second, we model adaptation through internal migration, assuming the coastal population decreases linearly by either 20 percent under a *migration away from coast—low* scenario or 50 percent under a *migration away from coast—high* scenario between 2024 and 2100.² We also model other adaptation forms, for example dietary changes or health investments, by assuming the marginal effect of pH on NMR decreases linearly over time. This effect either halves by 2100 in a *decreasing effect—slow* scenario or reaches zero by 2100 in a *decreasing effect—fast* scenario. Our most *optimistic adaptation* scenario combines the high migration assumption with the fast-decreasing effect.

¹These series reflect the global average pH, with coastal areas representing only a small portion of the oceans. We assume that coastal pH follows, on average, a similar trend.

²For some countries, like the Philippines, this assumption is not relevant because the whole or most of the country is considered a coastal area.

3.4 Results

This section presents the main findings of the chapter. We begin by detailing the results of our empirical tests for long-run adaptation to ocean acidification. We then present the results of our projection exercise, quantifying the future neonatal mortality burden under various emissions and adaptation scenarios.

3.4.1 Evidence on long-run adaptation

A central prerequisite for projecting future impacts is to understand whether communities adapt to long-term environmental change. If initial exposure to a worse climatic conditions, i.e. a more acidic ocean, fosters resilience, then projections that ignore this adaptation may be overly pessimistic. This subsection presents the results of two empirical tests designed to detect such adaptation. The findings consistently show a lack of observable long-run adaptation to ocean acidification in our sample.

Our first test examines whether communities in areas with more acidic initial conditions have developed a greater resilience to subsequent pH shocks. Table 3.1 presents the results of this test. The coefficient on the interaction term between the contemporaneous pH shock and the historical baseline pH is positive and statistically significant across all specifications. This result runs directly counter to the adaptation hypothesis. A positive coefficient indicates that the negative health impact of a decrease in ocean pH is *more* severe in locations that were initially more acidic. For instance, in our preferred specification (Column 3), a one-standard-deviation lower baseline pH worsens the impact of a contemporaneous pH shock by 1.28 deaths per 1,000 live births. Rather than fostering resilience, adverse initial conditions appear to increase vulnerability, a finding that underscores that not only the *change* in pH matters for NMR, but also its baseline *level*.

A second possibility is that adaptation is triggered not by the level of exposure, but by the salience of the environmental change. A large literature on learning suggests that adaptation is more likely when the negative signal is strong and clear. In our context, this would imply that adaptation should be more pronounced in locations where the long-term decline in pH has been steepest and thus most noticeable. To test this, we interact the pH shock with an indicator for whether a location's long-term acidification trend is above the median. The results, presented in Table 3.2, once again fail to support the adaptation hypothesis. The coefficient on the interaction term is statistically significant and negative. This indicates that the adverse mortality impact of a pH shock is significantly larger in areas with a faster rate of acidification. Specifically, in these locations, the impact of a pH shock on neonatal mortality is more than twice as large as in locations with a slower rate (-1.133 vs -3.782 deaths per 1,000 live births). The fact that a more noticeable environmental threat is associated with worse, not better, outcomes

provides strong evidence against the idea that the salience of the problem has spurred effective adaptation.

Taken together, these two empirical tests paint a clear and consistent picture. Over the past five decades, we find no evidence that coastal communities in these 36 countries have spontaneously adapted to the growing threat of ocean acidification in a way that mitigates its impact on neonatal mortality. This lack of historical adaptation provides a strong empirical justification for using a “no adaptation” scenario as a plausible, if concerning, baseline for our future projections.

3.4.2 The historical impact of ocean acidification

This subsection aims to quantify the historical mortality burden attributable to ocean acidification between 1975 and 2018. The methodology, thoroughly detailed in Section 3.3.2, involves a critical decomposition of the observed NMR at the community level. We separate this observed rate into two distinct components: a counterfactual NMR, representing what would have occurred in the complete absence of ocean acidification, and a specific component directly attributable to the impacts of ocean acidification.

The comprehensive results of this historical decomposition are visually presented in Figure 3.1. In this figure, the total height of each blue bar precisely represents our estimate of the average NMR for the entire 1975–2018 period. Conversely, the accompanying red bar illustrates the counterfactual NMR, which signifies the rate that would have prevailed had ocean acidification not impacted the environment since 1975. Consequently, the gap between these two bars represents the historical mortality burden directly attributable to ocean acidification. The figure reveals two insights. First, the contribution of ocean acidification to neonatal mortality is positive across all 36 countries included in our study, consistently indicated by the blue bar being taller than the red bar. Second, this contribution exhibits significant geographical heterogeneity. The impact is most pronounced in countries such as the Philippines, where the substantial gap between the actual and counterfactual bars indicates an attributable burden of approximately 9 deaths per 1,000 live births. In stark contrast, the gap is considerably smaller in countries like the Democratic Republic of Congo, where the impact is estimated at only about 2 deaths per 1,000 live births. It is also important to note that our estimates of total NMR, which range from 14.7 in the Philippines to 49.4 in the DRC, show strong consistency with established official statistics, such as those provided by (UNICEF, 2024), thereby lending credibility to our decomposition approach.

However, moving beyond country-level averages, a large degree of heterogeneity becomes apparent across communities when viewed spatially. Figure 3.3 illustrates this spatial variation, presenting the impact as a function of distance from the coast and disaggregated by broader geographical regions. Two key patterns distinctly emerge. First, the adverse impact of ocean

acidification is generally more significant in communities situated closer to the shore. This observed pattern is consistent with the findings of Chapter 2, which highlight the greater reliance of coastal communities on small-scale fisheries, rendering them inherently more susceptible to the detrimental effects of ocean acidification. Second, substantial variation exists not only within but also across different regions concerning the impact of ocean acidification. While in some specific locations the estimated impact may even be negative, suggesting that the counterfactual NMR would have been higher than the actual observed NMR, in the vast majority of areas the mortality burden is substantial, with some communities enduring an attributable mortality burden exceeding 10 deaths per 1,000 live births.

This spatial heterogeneity highlights the need to precisely identify areas that have been subjected to a higher burden from ocean acidification. By identifying and focusing on regions that have experienced the most significant impacts, we can considerably improve the strategic targeting of potential mitigation and adaptation policies. Given that we only observe NMR and its counterfactual at specific surveyed locations, we address this data limitation by imputing the NMR attributable to ocean acidification at unmeasured locations through the application of ordinary kriging (Cuadros et al., 2017). This method allows us to generate maps of areas with higher vulnerability. The results of this spatial interpolation are presented in Figure 3.2, which presents a continuous map of the estimated historical NMR attributable to ocean acidification across the entire coastal zones of the countries included in our sample. For analytical purposes, we define *hotspot* as those areas in the top 25% of the entire sample's distribution, as derived from our kriging results.

The resulting map from this interpolation reveals significant sub-national heterogeneity in the mortality burden, providing a more nuanced understanding than country-level averages alone. Hotspots of particularly high attributable mortality are clearly evident in specific coastal regions. For example, the extensive coastlines of the Philippines, India, and East Africa consistently show broad areas experiencing a high mortality burden (Figures 3.4 and 3.6). In Central and South America, there is also evidence of significant hotspots, particularly along the Caribbean coast (Figure 3.5). In contrast, the impact along West African coastlines appears more concentrated within distinct, smaller locations.

3.4.3 Projections of future mortality burden

Building on these historical estimates, we project the NMR components to 2100. Figure 3.7 illustrates the evolution of the counterfactual NMR (NMR^{CF}) and the ocean acidification-induced NMR (NMR^{OA}) over time. The counterfactual series shows a steady decline, reflecting expected improvements in public health and economic development. The ocean acidification component, however, shows a steady increase, diverging significantly under the low- and high-emissions scenarios. Panel A, which reflects our benchmark *no adaptation* assumption, shows

a stark future where the gains from development are partially offset by the growing burden of ocean acidification. Panels B and C illustrate more optimistic scenarios where the marginal effect of ocean pH on NMR decreases over time, either halving by 2100 (Panel B) or falling to zero (Panel C). These panels demonstrate that even under optimistic assumptions about adaptation, ocean acidification will continue to pose a significant threat to neonatal health for decades to come.

Finally, we translate these projected mortality rates into the total number of deaths attributable to ocean acidification. Figure 3.8 shows the cumulative number of deaths over time, and Table 3.3 summarises the total projected deaths by 2100. Our central estimate, under the high-emissions (RCP8.5) and no-adaptation scenario, projects that ocean acidification will be responsible for a cumulative 77.21 million neonatal deaths in the coastal regions of the 36 countries studied between 1975 and 2100. The projections also highlight the life-saving benefits of aggressive climate change mitigation. Shifting from the high-emissions pathway to the low-emissions pathway (RCP2.6) dramatically reduces the projected mortality burden to 37.98 million deaths, averting nearly 40 million deaths. The increasingly optimistic adaptation scenarios further reduce the burden, with the most optimistic scenario combining high migration and a fast-decreasing effect projecting only 12.86 million deaths under RCP2.6 and 18.14 million under RCP8.5. These results underscore the critical importance of both climate change mitigation and proactive adaptation strategies in safeguarding child health in the face of ocean acidification.

It is important to acknowledge that these projections are subject to several sources of uncertainty. The primary sources are the statistical uncertainty in our estimated effect of ocean pH on neonatal mortality, the uncertainty in future climate projections, and the uncertainty in future population growth. The confidence intervals presented in Table 3.3 account only for the first of these, based on the 90% confidence interval of our main regression coefficient from Table 2.3. The true range of potential outcomes is therefore likely wider. Despite this uncertainty, the central estimates provide a clear and compelling case that ocean acidification represents a major, and heretofore unquantified, threat to global child health.

3.5 Conclusion

This chapter provides the first projections of the future neonatal mortality burden attributable to ocean acidification, while also quantifying its substantial historical impact.

Our empirical tests consistently reveal little evidence to support the assumption that communities will adapt spontaneously to ocean acidification in a way that mitigates its health impacts. On the contrary, areas with worse initial conditions or more rapid rates of acidification appear more, not less, vulnerable to further shocks. This concerning finding suggests that effective


adaptation will necessitate deliberate and targeted policy interventions, such as strengthening public health systems, promoting dietary diversity, and managing fisheries sustainably.

Building on these insights, our analysis further demonstrates that ocean acidification has already imposed a substantial and highly heterogeneous mortality burden on coastal communities between 1975 and 2018. This impact has been consistently positive across all 36 countries in our study, yet it exhibits significant spatial variation. Higher attributable neonatal mortality rates are observed closer to the shore, consistent with the heightened reliance of these communities on small-scale fisheries. Distinct hotspots of high attributable mortality are clearly evident in specific coastal regions, such as the Philippines, India, and East Africa, with impacts along West African coastlines appearing more concentrated. This established past and ongoing burden underscores the concrete human cost of ocean acidification, providing a critical empirical foundation for understanding the scale of the future challenge.

The future projections, informed by this lack of historical adaptation, are stark: the business-as-usual emissions scenario could potentially lead to millions of additional neonatal deaths by the end of the century in the world's most vulnerable coastal regions due to this single, often overlooked, consequence of climate change.

However, our projections also make it clear that adaptation efforts alone are insufficient to address the full scale of the projected burden. The most effective policy for saving lives is decarbonization. The significant gap in projected deaths between the high- and low-emissions scenarios represents a clear and quantifiable health dividend from global climate action. Policies aimed at emissions reduction are, in this light, among the most impactful public health interventions available to the global community in the 21st century.

Figure 3.1: Estimates of $NMR_{k,t}$ and $NMR_{k,t}^{CF}$, 1975–2018 average




Note. The figure presents the 1975–2018 country-level average of NMR and of (counterfactual) NMR in absence of any ocean acidification from 1975. The sample is restricted to the countries used in the study and to their coastal area (see Appendix A2.1). Details about the methodology followed to compute these statistics are presented in Section 3.3.2.

Figure 3.2: Estimates of NMR_c^{OA}




Note. Acidification-related NMR hotspots, defined as the top 25% of the entire sample's distribution resulting from kriging (Cuadros et al., 2017). The map shows the estimated NMR attributable to ocean acidification at unmeasured locations within 100km of the coast for the 36 countries in our sample. The estimates are obtained by ordinary kriging, which interpolates values at unmeasured locations based on the spatial auto-correlation of observed data. The methodology followed to compute these estimates is detailed in Section 3.3.2.

Figure 3.3: Estimates of NMR_c^{OA} : By distance and regions



Note. Panel A presents a kernel-weighted local polynomial regression of NMR_c^{OA} on distance from shore. Panel B shows a kernel density of NMR_c^{OA} by regions. The sample is restricted to the countries used in the study and to their coastal area (see Appendix A2.1). Details about the methodology followed to compute these statistics are presented in Section 3.3.2.


Figure 3.4: Estimates of NMR_c^{OA} and hotspots: Africa



Note. Panel A displays the geographical distribution of NMR_c^{OA} . Panel B illustrates *Hotspot* locations, defined as the top 25% of the kriging distribution for the entire sample. *Coastal areas* are those within 100 km of the coastline. *No data* indicates countries not included in the study, inland areas, or unpopulated regions as reported by the LandScan Global 2023 database (Lebakula et al., 2024). Further details on the methodology used to compute these statistics are provided in Section 3.3.2.

Figure 3.5: Estimates of NMR_c^{OA} and hotspots: Central and South America

A. NMR_c^{OA}




B. Hotspots




Note. Panel A displays the geographical distribution of NMR_c^{OA} . Panel B illustrates *Hotspot* locations, defined as the top 25% of the kriging distribution for the entire sample. *Coastal areas* are those within 100 km of the coastline. *No data* indicates countries not included in the study, inland areas, or unpopulated regions as reported by the LandScan Global 2023 database (Lebakula et al., 2024). Further details on the methodology used to compute these statistics are provided in Section 3.3.2.

Figure 3.6: Estimates of NMR_c^{OA} and hotspots: South and Southeast Asia

A. NMR_c^{OA}




B. Hotspots




Note. Panel A displays the geographical distribution of NMR_c^{OA} . Panel B illustrates *Hotspot* locations, defined as the top 25% of the kriging distribution for the entire sample. *Coastal areas* are those within 100 km of the coastline. *No data* indicates countries not included in the study, inland areas, or unpopulated regions as reported by the LandScan Global 2023 database (Lebakula et al., 2024). Further details on the methodology used to compute these statistics are provided in Section 3.3.2.

Figure 3.7: Counterfactual and ocean-acidification-induced NMR, 1975–2100



Note. The figure shows the evolution of NMR_t^{CF} and $NMR_{k,t}^{OA}$ over time. For NMR_t^{CF} , we show both the values estimated using DHS data and the fitted series. Section 3.3.2 details the assumptions. Panel A shows the values assuming the effect of ocean pH experienced while *in utero* on NMR is constant over time. In panel B, we assume the effect gradually diminishes linearly over time halving by 2100. In panel C, we assume the effect gradually diminishes linearly over time reaching 0 by 2100. The *low-emissions scenario* is the IPCC RCP2.6 scenario, targeting global warming limits of around 1.5°C–2°C by 2100 through strong mitigation efforts. The *high-emissions scenario* is the IPCC RCP8.5 scenario, a worst-case high-emissions scenario with rising emissions potentially increasing temperatures by 4°C–5°C or more by the end of the century. Assumptions and the methodology followed to compute these estimates are detailed in Section 3.3.2.

Figure 3.8: Neonatal deaths attributable to ocean acidification, 1975–2100



Note. Cumulative number (in millions) of neonatal deaths attributable to ocean acidification from 1975 to 2100, by year. The *low-emissions scenario* is the RCP2.6 scenario, targeting global warming limits of around 1.5°C–2°C by 2100 through strong mitigation efforts. The *high-emissions scenario* is the RCP8.5 scenario, a worst-case high-emissions scenario with rising emissions potentially increasing temperatures by 4°C–5°C or more by the end of the century. Scenarios are obtained from [IPCC \(2022\)](#). Each panel makes alternative assumptions concerning adaptation, detailed in Section 3.3.2, ranging from lowest (panel A) to highest (panel F). Section 3.3.2 details the methodology followed to compute these estimates. Table 3.3 reports the values in the year 2100, including the confidence intervals accounting for uncertainty in the estimate of the effect of ocean acidity experienced while *in utero* on neonatal mortality.

Table 3.1: Ocean acidity and neonatal mortality – heterogeneity by initial conditions

	Dependent variable: NMR (deaths per 1,000 births)		
	(1)	(2)	(3)
pH in the nearest waters (<i>in utero</i>)	-2.044 (0.719) [0.005]	-2.091 (0.699) [0.003]	-2.279 (0.687) [0.001]
× initial conditions	1.092 (0.316) [0.001]	1.089 (0.319) [0.001]	1.283 (0.313) [0.000]
Mean (dep.var.)	30.473	30.473	30.474
Identifying observations	1,583,706	1,583,706	1,581,815
Singleton observations	25	25	25
Communities	31,380	31,380	31,380
Countries	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes
Demographic controls	-	-	Yes

Note. Estimates are based on equation (3.2). *pH in the nearest waters (in utero)* is defined in Section 2.4. *Initial conditions* refer to a location's (standardised) average between 1972–1975. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth month FEs, country by birth year FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 3.2: Ocean acidity and neonatal mortality – heterogeneity by rate of acidification

Dependent variable: NMR (deaths per 1,000 births)			
	(1)	(2)	(3)
pH in the nearest waters (<i>in utero</i>)	-1.179 (0.686) [0.086]	-1.152 (0.663) [0.083]	-1.133 (0.644) [0.079]
× faster acidification	-2.596 (1.206) [0.032]	-2.602 (1.212) [0.032]	-2.649 (1.187) [0.026]
Mean (dep. var.)	30.473	30.473	30.474
Identifying observations	1,583,689	1,583,689	1,581,798
Singleton observations	42	42	42
Communities	31,380	31,380	31,380
Countries	36	36	36
Birth year range	1972–2018	1972–2018	1972–2018
Weather controls	-	Yes	Yes
Demographic controls	-	-	Yes

Note. Estimates are based on equation (3.3). *pH in the nearest waters (in utero)* is defined in Section 2.4. *Initial conditions* refer to a location's (standardised) average between 1972–1975. The sample is restricted to communities in the coastal area (see Section 2.3). Standard errors (in parentheses) are clustered at the ocean raster data point, *p*-values are reported in brackets. All specifications include community FEs, birth year by birth month FEs, country by birth month FEs, country by birth year FEs. The full list of controls is presented in Section 2.4. Appendix A2.1 provides detailed information on variables, selected surveys, and weighting procedures.

Table 3.3: Total neonatal deaths attributable to ocean acidification, 1975–2100

	Low-emissions scenario			High-emissions scenario			$\% \Delta_{\text{low-high}}$
	Estimate	90% CI bounds		Estimate	90% CI bounds		
		Lower	Upper		Lower	Upper	
Assumptions about adaptation	(1)	(2)	(3)	(4)	(5)	(6)	(7)
No adaptation	37.98	10.10	65.83	77.21	20.54	133.80	50.81
Migration away from coast (low)	34.48	9.17	59.75	67.83	18.04	117.55	49.17
Migration away from coast (high)	30.02	7.98	52.02	55.67	14.81	96.47	46.08
Decreasing effect (slow)	26.30	7.00	45.58	49.55	13.18	85.87	46.92
Decreasing effect (fast)	14.61	3.89	25.33	21.90	5.82	37.95	33.29
Optimistic adaptation	12.86	3.42	22.29	18.14	4.83	31.44	29.11

Note. Total number (in millions) of neonatal deaths attributable to ocean acidification from 1975 to 2100. Assumptions about adaptation are detailed in Section 3.3.2. The *low-emissions scenario* is the IPCC RCP2.6 scenario, targeting global warming limits of around 1.5°C–2°C by 2100 through strong mitigation efforts. The *high-emissions scenario* is the IPCC RCP8.5 scenario, a worst-case high-emissions scenario with rising emissions potentially increasing temperatures by 4°C–5°C or more by the end of the century. Assumptions and the methodology followed to compute these estimates are detailed in Section 3.3.2. Column (7) reports the percent reduction in the estimate from the high-emissions to the low-emissions scenario. The evolution over time of the cumulative number of neonatal deaths attributable to ocean acidification is reported in Figure 3.8.

Bibliography

- Alter, K., J. Jacquemont, J. Claudet, M. E. Lattuca, M. E. Barrantes, S. Marras, P. H. Manríquez, C. P. González, D. A. Fernández, M. A. Peck, et al. (2024). Hidden Impacts of Ocean Warming and Acidification on Biological Responses of Marine Animals Revealed through Meta-Analysis. *Nature Communications* 15(1), 2885.
- Auffhammer, M. (2018). Quantifying Economic Damages from Climate Change. *Journal of Economic Perspectives* 32(4), 33–52.
- Burke, M., S. M. Hsiang, and E. Miguel (2015, nov). Global Non-linear Effect of Temperature on Economic Production. *Nature* 527(7577), 235–239.
- Colt, S. G. and G. P. Knapp (2016). Economic Effects of an Ocean Acidification Catastrophe. *American Economic Review* 106(5), 615–19.
- Croft, T. N., A. M. J. Marshall, and C. K. Allen (2018). Guide to Dhs Statistics. Demographic and Health Surveys Program.
- Cuadros, D. F., J. Li, A. J. Branscum, A. Akullian, P. Jia, E. N. Mziray, and F. Tanser (2017). Mapping the Spatial Variability of Hiv Infection in Sub-Saharan Africa: Effective Information for Localized Hiv Prevention and Control. *Scientific Reports* 7(1), 9093.
- Dell, M., B. F. Jones, and B. A. Olken (2012). Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics* 4(3), 66–95.
- Dell, M., B. F. Jones, and B. A. Olken (2014). What Do We Learn from the Weather? the New Climate-Economy Literature. *Journal of Economic Literature* 52(3), 740–98.
- Deschênes, O. and M. Greenstone (2011, October). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics* 3(4), 152–85.
- Doney, S. C., D. S. Busch, S. R. Cooley, and K. J. Kroeker (2020). The Impacts of Ocean Acidification on Marine Ecosystems and Reliant Human Communities. *Annual Review of Environment and Resources* 45(1), 83–112.
- Hiemstra, P. H., E. J. Pebesma, C. J. Twenhöfel, and G. B. Heuvelink (2009). Real-Time Automatic Interpolation of Ambient Gamma Dose Rates from the Dutch Radioactivity Monitoring Network. *Computers and Geosciences* 35(8), 1711 – 1721.
- IPCC (2019). Changing Ocean, Marine Ecosystems, and Dependent Communities. In H.-O. Pörtner, D. Roberts, V. Masson-Delmotte, P. Zhai, M. Tignor, E. Poloczanska, K. Mintenbeck, A. Alegría, M. Nicolai, A. Okem, J. Petzold, B. Rama, and N. Weyer (Eds.), *IPCC*

Special Report on the Ocean and Cryosphere in a Changing Climate, pp. 447–587. Cambridge, UK and New York, NY, USA: Cambridge University Press.

IPCC (2022). AR6 Srocc Data for Figure Spm.1: Past and Future Changes in the Ocean and Cryosphere. MetadataWorks.

Kroeker, K. J., R. L. Kordas, R. N. Crim, and G. G. Singh (2010). Meta-Analysis Reveals Negative yet Variable Effects of Ocean Acidification on Marine Organisms. *Ecology letters* 13(11), 1419–1434.

Lebakula, V., J. Epting, J. Moehl, C. Stipek, D. Adams, A. Reith, J. Kaufman, J. Gonzales, B. Reynolds, S. Basford, A. Martin, W. Buck, A. Faxon, A. Cunningham, A. Roy, Z. Barbose, J. Massaro, S. Walters, C. Woody, A. Iman, A. Wilkins, E. Powell, and M. Urban (2024). Landscan silver edition. LandScan Global, October 09, 2024, CY 2023, raster digital data, Global.

Merkens, J.-L., L. Reimann, J. Hinkel, and A. T. Vafeidis (2016). Gridded Population Projections for the Coastal Zone under the Shared Socioeconomic Pathways. *Global and Planetary Change* 145, 57–66.

Moore, F. C. and D. B. Diaz (2015, feb). Temperature Impacts on Economic Growth Warrant Stringent Mitigation Policy. *Nature Climate Change* 5(2), 127–131.

Tagliabue, A., A. Ekaykin, B. Mazon, C. Derkse, N. Abram, R. Hock, R. van Waal, T. Frölicher, M. Aschwanden, and E. Lambert (2022). AR6 Srocc Data for Figure Spm.1: Past and Future Changes in the Ocean and Cryosphere.

Tollefsen, A. F., H. Strand, and H. Buhaug (2012). Prio-Grid: A Unified Spatial Data Structure. *Journal of Peace Research* 49(2), 363–374.

UNICEF (2024). Levels and Trends Child Mortality-Report 2023: Estimates Developed by the United Nations Inter-Agency Group for Child Mortality Estimation. Technical report, UNICEF, World Health Organization, World Bank, United Nations.

United Nations (2003). Ecosystems and Human Well-Being: A Framework for Assessment. United Nations, Island Press, Washington DC.

United Nations (2024). World Population Prospects 2024. Department of Economic and Social Affairs, Population Division.

Wittmann, A. C. and H.-O. Pörtner (2013). Sensitivities of Extant Animal Taxa to Ocean Acidification. *Nature climate change* 3(11), 995–1001.