

ABSTRACT

Title of Dissertation: **ESSAYS ON GENDER INEQUALITIES AND THE
ECONOMIC IMPACTS OF CLIMATE CHANGE**

Fatima Najeeb
Doctor of Philosophy, 2025

Dissertation Directed by: **Professor Erich Battistin**
Department of Agricultural and Resource Economics

Climate change is accelerating the frequency and intensity of extreme weather events such as floods, heatwaves, and droughts. In 2019 alone, 4.5 billion people were exposed to extreme weather events, marking a 12% increase from 2010, with South Asia as the hardest hit region (Doan et al., 2023). These climate shocks affect various dimensions of a household's well-being, such as finances, livelihoods, and health, with the potential to reverse hard-won social and economic gains. They have particularly severe consequences for women, who are often at the frontline of climate crises due to existing social and economic inequalities.

This dissertation examines the ways in which climate change intersects with gender inequalities across multiple dimensions: within the household, during critical life stages, and through policy efforts to bridge gender gaps in educational attainment. Specifically, the three chapters explore how climate shocks affect the distribution of household resources among members (Chapter 1), how extreme heat experienced by mothers during pregnancy affects children's early-life health (Chapter 2), and how government programs can close gender gaps in secondary education

attainment (Chapter 3). Chapters 1 and 2 directly investigate the economic and health impacts of climate change, while Chapter 3 focuses on gender inequality in education—a barrier that affects women’s economic opportunities and may indirectly influence their capacity to adapt to climate shocks in the future. Together, these chapters highlight the interconnectedness of climate change, economic vulnerability, and gender inequality in climate-vulnerable settings.

Chapter 1 examines how climate shocks can have heterogeneous effects on household members’ well-being by examining the impact of floods on intra-household resource allocation in Bangladesh, one of the countries most vulnerable to climate change. While much of the existing climate-economy literature examines the effect of climate shocks on household-level consumption (Karim and Noy, 2016), this chapter demonstrates that women’s and children’s material well-being deteriorates relative to that of men in the same household as the share of the household budget available to them declines, while the share allocated to men increases. As the division of the household budget among members is not directly observable, I specify a structural model to estimate the shares of the household budget allocated to men, women, and children and leverage quasi-random variation in flood exposure to examine how these shares respond to a flood shock over time. Using panel data on households combined with geospatial data on flooding, I show that exposure to a flood leads to a reallocation of resources away from women and children—primarily women—toward men, with effects persisting even four years after the shock. As a result, women and children face a greater risk of falling below the poverty line compared to men living under the same roof.

This reallocation of resources is driven by a decline in employment opportunities for women, leading to a decrease in their earnings relative to men’s after the flood, as women are less able to transition to alternative occupations. While both men and women are less likely to work in

farming after the flood, men can compensate by working as day laborers, an option less available to women in this context. This limited adaptability for women highlights gendered economic vulnerabilities that climate shocks exacerbate, underscoring the importance of integrating gender-sensitive labor market support into climate resilience efforts to mitigate these hidden intrahousehold inequalities.

Chapter 2 examines how exposure to extremely hot temperatures during pregnancy affects children's early-life health, which can contribute to long-term economic and social inequalities. This chapter also focuses on the Bangladeshi context, and uses exogenous variation in prenatal temperature experienced by cohorts of pregnant mothers who give birth in the same district and month but in different calendar years. Combining household and child-level survey data with reanalysis temperature data, I find that short-lived spikes in temperature during pregnancy do not significantly impact health in early childhood. However, more prolonged exposure to hot temperatures, especially during the third trimester, decreases their height-for-age z-score, a measure of long-term undernourishment, by 0.186 and increases their probability of stunted growth, a signal of chronic undernourishment, by 6.9 percentage points. This chapter demonstrates that one channel through which the economic and health consequences of climate change materialize is women's exposure to extreme heat during pregnancy.

Chapter 3 shifts from the direct impacts of climate shocks to examining a policy intervention aimed at reducing gender inequalities in education. In climate-vulnerable settings, improving girls' education can be an important step toward enhancing future resilience, as education expands women's economic opportunities and access to information. This co-authored chapter evaluates a government-administered cash transfer program in Pakistan that offered households a monthly stipend conditional on girls' regular attendance in secondary school (grades 6 to 10). Using

administrative school-level data and a regression discontinuity design that exploits the program's eligibility cutoff based on district-level adult literacy rates, we find that the program increased girls' secondary school enrollment consistently over more than a decade (2004-2015). These positive effects persisted despite the real value of the cash benefit decreases by more than 60%, suggesting that the program may have influenced household attitudes or contributed to shifting social norms around girls' education. While this potential climate resilience benefit is not directly tested in this chapter, the policy's role in reducing gender gaps in education is directly relevant to the broader inquiry into gender inequality in South Asia.

These three chapters collectively illustrate the multifaceted relationship between climate change, economic impacts, and gender inequalities. Chapters 1 and 2 directly show how climate shocks—floods and heat—worsen gendered economic and health vulnerabilities within and across generations. Chapter 3 complements this by showing how policy interventions can reduce the gender gap in education, which may have long-term consequences for women's economic and adaptive capacity in climate-vulnerable settings. Taken together, my dissertation highlights that addressing gender inequality is a necessary element of effective climate adaptation strategies in contexts where climate change and gender disparities are intertwined.

ESSAYS ON GENDER INEQUALITIES AND THE
ECONOMIC IMPACTS OF CLIMATE CHANGE

by

Fatima Najeeb

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2025

Advisory Committee:

Professor Erich Battistin, Chair
Professor Kenneth Leonard
Professor Cory Smith
Professor Louis Preonas
Professor Sebastian Galiani

© Copyright by
Fatima Najeeb
2025

Acknowledgments

I am incredibly grateful to my advisors, Professor Erich Battistin, Professor Ken Leonard, and Professor Cory Smith, for the countless ways in which they helped improve this dissertation, for their generosity with their time, and for consistently engaging with my work. This has been an incredible learning experience, made possible by their thoughtful mentorship, and it has been an absolute pleasure getting to know them and learning from such extraordinary researchers.

I am also grateful to other faculty at the University of Maryland whose teaching and feedback have shaped my thinking and helped me develop my research. In particular, I would like to thank Professor Louis Preonas, Professor Sebastian Galiani, Professor Jorge Holzer, and Professor Robert Chambers for their feedback and support at various points in my PhD journey.

I am thankful for the support of my family and friends, without whom it would not have been possible to complete this dissertation. I owe a special thanks to my parents, Najeeb Tariq and Tanveer Fatima; my brothers, Omar and Ahmad; my sisters-in-law, Aleena and Sumrana; and my absolutely remarkable niece and nephews — Fareeha, Arham, and Ibrahim — for their constant love, encouragement, and silliness.

I would also like to thank my friends, both within and beyond the University of Maryland, for their emotional support throughout. I am lucky to have had some spectacular people around me who gave me hope and provided much-needed comic relief and tea. In particular, I would like to thank Saman, Sara, Maja, Tommy, Varun, Aqila, and Fatima — and a shoutout to Momo for

helpful comments and suggestions. A special thanks also to Ben and Mike, for letting me stay in their coffee shop for as long as I wanted — most of this dissertation was written there.

I also want to acknowledge the help and support of several staff members. Thank you to Dany Burns and Billie Ray for making the administrative side of things so much easier. I would also like to thank the University of Maryland Graduate School for awarding me the Ann G. Wylie Dissertation Fellowship, which provided the financial support I needed to complete this dissertation during my final year. I am also grateful to the International Food and Policy Research Institute (IFPRI) for sharing some data on households that was used in this dissertation. IFPRI bears no responsibility for the analyses or interpretations of the data presented here.

Table of Contents

Acknowledgements	ii
Table of Contents	iv
List of Tables	vi
List of Figures	vii
Chapter 1: Deepening Intrahousehold Inequalities? Evidence from Floods in Bangladesh	1
1.1 Introduction	1
1.2 Related literature	7
1.3 Intra-household allocation of resources	10
1.3.1 The efficient collective household model	11
1.3.2 Engel curves of private assignable goods	15
1.3.3 Identification of resource shares	17
1.3.4 Estimation of resource shares	19
1.4 Data and sample selection	22
1.4.1 Flooding data	22
1.4.2 Household level data	23
1.4.3 Exposure to floods	24
1.4.4 Ex-ante probability of flooding	25
1.4.5 Villages with similar ex-ante flood risk	27
1.4.6 Descriptive statistics for the 2014 sample	30
1.5 Results	33
1.5.1 Effect of floods on intra-household resource allocation	33
1.5.2 Shadow budget and individual-level poverty	41
1.5.3 Effect of the flood on the household's total budget	44
1.6 Mechanisms	48
1.7 Robustness checks	53
1.7.1 Adjusting the model	54
1.7.2 Effect of a flood on resource shares in out-of-sample villages	55
1.7.3 Stricter criteria for unflooded villages	56
1.7.4 Stricter criteria for flooded villages	56
1.8 Conclusion	61
Chapter 2: Impact of Exposure to Hot Temperatures In-Utero on Child Health Outcomes	65

2.1	Introduction	65
2.2	Data	70
2.3	Descriptive analysis	76
2.4	Empirical strategy	79
2.5	Empirical findings	86
2.6	Conclusion	93
Chapter 3: Effects over the Life of a Program: Evidence from an Education Conditional Cash Transfer Program for Girls		96
3.1	Introduction	96
3.2	Related literature	101
3.3	Program context and design	103
3.4	Data and empirical strategy	107
3.4.1	Data	107
3.4.2	Empirical strategy	109
3.5	Results	112
3.5.1	Validity of the local randomization-based RD window	112
3.5.2	Program effects on secondary school enrollment	113
3.5.3	Falsification tests	117
3.5.4	Program cost-effectiveness	118
3.6	Conclusion	121
Appendix A: Supplementary details for Chapter 1		133
A.1	Illustration of identification of resource shares from Engel curves	133
A.2	Intuition behind identification of the change in resource shares in treated and untreated villages	135
A.3	Identification in a fully specified example	138
A.4	Additional details on data and sample selection	141
A.4.1	Cleaning household panel and categorizing individuals	141
A.4.2	Additional geographic data	142
A.4.3	Narrowing sample based on ex-ante flood risk	144
A.4.4	Trends in rainfall in flooded and unflooded villages	145
Appendix B: Additional tables and figures for Chapter 1		148
B.1	Additional tables and figures for Chapter 2	168
B.2	Additional tables for Chapter 3	175
Bibliography		177
Bibliography		177

List of Tables

1.1	Summary of household characteristics	31
1.2	Effect of floods on resource shares of men and women	34
1.3	Resource shares in reference households and the effect of floods	36
1.4	Shadow budget of men, women and children in 2017 PPP\$	42
1.5	Effect of the flood on household's financial indicators	46
1.6	Effect of the flood on employment and earnings	49
2.1	Descriptive statistics for children under age 5	75
2.2	Effects of prenatal heat exposure	87
2.3	Heterogeneous effects of prenatal heat exposure by trimester	88
2.4	Results for subsample of local births	91
2.5	Exploring the agricultural productivity channel (controlling for mother's education)	92
3.1	School participation rate by gender, age group, and district program status in Punjab, 2003 and 2014	129
3.2	District program assignment status	130
3.3	Validity check of window length	131
3.4	Program effects on girls' secondary school enrollment	132
B.1	Additional descriptive statistics for 2014 sample	162
B.2	The determinants of resource shares of men and women	163
B.3	Effect of number of floods on resource shares	164
B.4	Resource shares in reference households and the effect of number of floods	165
B.5	Resource shares of all households	165
B.6	Other mechanisms	166
B.7	Effects on the 2016 and 2018 cohorts	167
B.8	Effect of interventions on children's height	172
B.9	Main results without standardizing temperature variables	172
B.10	Exploring the agricultural productivity channel (controlling for NDVI)	173
B.11	Main results with Conley standard errors	174
B.12	Balance test for school characteristics at baseline (2003)	175
B.13	Validity check of window length for 30% false cutoff	176

List of Figures

1.1	Cumulative number of floods across villages from 2000 - 2018	24
1.2	Cohorts of villages flooded from 2012-2018	26
1.3	Number of villages across cohorts	27
1.4	Standardized annual rainfall in flooded and unflooded villages	29
1.5	Resource shares by number of floods	35
1.6	Resource shares of men and women over time	39
1.7	Shadow budgets per individual by flood exposure	43
1.8	Effect of the flood on main occupations of men and women	52
1.9	Effect on resource shares by share of village flooded	58
1.10	Effect on women's and children's earnings relative to men by share of village flooded	59
2.1	Relationship between the two measures of temperature	74
2.2	Correlation between HAZ and height	76
2.3	Variation in number of heat-wave days by district	78
2.4	Variation in height by district	78
2.5	Residual variation in temperature variables	82
2.6	Share of births by season	83
2.7	Factors possibly affecting fertility decisions	85
2.8	Randomization inference results	93
3.1	Net enrollment rate by province, Pakistan, 2014/15	124
3.2	School participation rate for girls' ages 11–15 by district, Punjab, 2014	124
3.3	School participation rate by age, gender, and district program status, Punjab, 2014	125
3.4	Regression discontinuity plot, girls' secondary school enrollment effect, 2003–05	125
3.5	Program effects on girls' secondary school enrollment	126
3.6	Program effects on boys' secondary school enrollment	126
3.7	Program effects on enrollment girls using false cutoffs	127
3.8	Program effects on the number of teacher positions filled	127
3.9	Real value of annual program cash benefit, 2004–15	128
A.1	Example of Engel curves for men and women	134
A.2	Illustrative example of Engel curves across treated and untreated villages	137
A.3	Factors predicting probability of flooding	143

A.4	Differences in standardized annual rainfall in full sample	146
A.5	Standardized annual temperature	147
B.1	Administrative divisions	148
B.2	Flood map from Global Flood Database	149
B.3	Share of village flooded during 2012-2018	149
B.4	Frequency of flooding from 2000 - 2018	150
B.5	Actual floods in 2012-2018 vs. ex-ante probability of flooding	150
B.6	Selected treated and untreated villages	151
B.7	Balance check at baseline	152
B.8	Main occupations of men and women at baseline	153
B.9	Food budget shares across survey years	154
B.10	Resource shares and total household expenditure	155
B.11	Shadow budgets by treatment status	156
B.12	Differences in the price of food across flooded and unflooded villages	157
B.13	Robustness of main results across models	158
B.14	Reducing allowable difference in ex-ante risk of floods	159
B.15	Effect of the flood on men's standardized monthly earnings	160
B.16	Effect of the flood on presence of non-earning children	161
B.17	Correlation between HAZ and height, full	168
B.18	Spatial variation in average temperature during pregnancy, by season of birth	168
B.19	Trends in residualized height-for-age z-score (HAZ) by birth year	169
B.20	Fertility trends by season	169
B.21	Exploring adaptive behavior in fertility decisions.	170
B.22	Residual variation in temperature variables, local births only	171

Chapter 1: Deepening Intrahousehold Inequalities? Evidence from Floods in Bangladesh

1.1 Introduction

Climate change has increased the frequency and intensity of weather extremes, including heatwaves, floods, and droughts (Seneviratne et al. 2021). In 2019, 4.5 billion people were exposed to extreme weather events—12% more than in 2010—highlighting the growing global vulnerability to climate-induced events (Doan et al. 2023). While the climate-economy literature addresses the broader economic impacts of these climate shocks on the household’s overall budget and consumption (Karim and Noy 2016), less is known about their heterogeneous impacts across individuals living under the same roof. Traditional measures of per-capita consumption often overlook the realities of intra-household resource distribution, assuming equality where it may not exist. This generalization is particularly problematic in regions where entrenched cultural norms disadvantage certain types of household members. In such contexts, it becomes important to study how resources are divided across members, and how these resource shares respond to various shocks.

Within the cultural context of South Asia, women are particularly vulnerable to the effects of climate shocks. Qualitative studies in the region show that these shocks place acute pressure

on women’s daily workloads, reduce labor force participation, and increase dependency (Memon 2020; Rezwana and Pain 2020). Furthermore, entrenched gender norms can restrict their decision-making autonomy and freedom of movement after these events, further hindering their recovery (Nellemann, Verma and Hislop 2011). These constraints can weaken women’s bargaining power, potentially reducing their resource share—defined as the proportion of the total household budget spent on goods and services they consume. Such a reduction can, in turn, diminish their material well-being relative to other household members.

I investigate how climate shocks change the intra-household allocation of resources and, by extension, the material well-being of individuals living under the same roof. I study this question in the context of flood shocks in Bangladesh.¹ I seek to answer whether flood events divert resources away from women in favor of others, and if they are likely to become poorer than others as a result. Building on recent work showing that Bangladeshi women on average have 6 percentage points or 16% lower resource shares than men (Brown, Calvi, and Penglase 2021), this paper further examines whether exposure to a negative climate shock exacerbates existing within-household gender inequalities up to four years after the event.

To account for potential differences in flood anticipation, I focus on villages that are equally at risk of flooding ex-ante, where the timing of the first flood during the study period (2012 to 2018) is as-good-as-random. Bangladesh experienced 20 floods during this period, affecting 99 of the 275 villages in my sample. The quasi-random timing of floods, among villages with similar ex-ante risk, yields cohorts of villages that experienced their first flood in different years. I focus on the largest cohort: villages that experienced their first flood in 2014. Using panel data from

¹Globally, Hallegatte et al. (2016) document a stronger association between flood exposure and extreme poverty compared to other climate shocks. Bangladesh, being one of the most climate-vulnerable countries in the world, is projected to experience “continuously increasing” exposure to floods (Tellman et al. 2021), making it an appropriate context for this study.

the Bangladesh Integrated Household Survey, conducted in 2011 (prior to the flood), in 2015 (six months post-flood), and in 2018 (four years later), my research design allows me to examine both the immediate effects six months post-flood and the longer-term impacts four years afterward.

I specify a structural model of collective households to estimate the share of the total household budget allocated to men, women, and children in unflooded villages, given that the division of household resources among members is not directly observable. Grounded in collective household theory (Chiappori 1988, 1992), the model uses limited information on individual-level consumption to estimate the share of the *total* household budget that is claimed by men, women and children in the household. Specifically, it uses variation in expenditures on private and assignable goods – goods consumed exclusively by a known member – and total household expenditures, and recovers intra-household resource allocation by comparing the slopes of Engel curves for these goods across different types of household members. This approach has gained traction since the seminal paper by Dunbar, Lewbel and Pendukar (2013) who identify resource shares in a nuclear family in a single period. I extend this approach to a longitudinal setting to estimate the evolution of resource shares in unflooded households. I then overlay exogenous variation in flooding to estimate any *changes* in intra-household resource allocation with exposure to a flood in the short and longer run.

I find that exposure to a flood triggers a reallocation of household resources away from women —and, to a lesser extent, children—towards men of the household. Six months after the flood, women’s resource shares decrease by 8.5 percentage points (30%) relative to women in unflooded households. Children’s resource shares also decrease, though by a smaller 2.7 percentage points. These resources are redistributed to men, whose share increases by 11.2 percentage points, significant at the 5% level. Four years later, women and children continue

to experience decreases in resource shares: women's resource shares are 6.9 percentage points lower (22%) compared to those in unaffected households, while children's shares decline by 4.6 percentage points. Men's resource shares remain significantly higher, increasing by 11.5 percentage points (38%). This reallocation of resources away from women and towards men becomes more pronounced with the number of floods experienced.

Using predictions of resource shares from the structural model, I compute the material well-being of men, women, and children by multiplying their resource shares by the observed total household expenditure, yielding the shadow budget — the unobserved budget allocated to each household member to maximize their individual utility. I find that men in flooded households see a boost in their daily shadow budget, while women and children experience a decrease. Six months after the flood, the daily shadow budget for a woman in flooded households averages \$3.33 (in terms of 2017 PPP\$), compared to \$6.59 for a man, resulting in a daily difference of \$3.26. Given that the international extreme poverty line is \$2.15 per day, an average woman in a flooded household is closer to extreme poverty compared to a man. On the other hand, in unflooded households during the same period, the average daily shadow budget for a man is \$5.89 and the gender gap is only \$0.61 per day. Four years later, the gender gap in flooded households is \$3.18 per day, compared to just \$0.16 per day in unflooded households. Importantly, men in flooded households have a shadow budget that is \$1.41 *higher* than men in unflooded households.

To better contextualize these individual-level effects, I examine the flood's impact on the household's total budget. I find that exposure to a flood reduces the total household expenditure by approximately 14% in the months following the event, with this reduction persisting even four years later. Therefore, the increase in men's shadow budget in flooded households in both the short and longer run occurs even though the household's total budget decreases.

I show that the reallocation of household resources is driven by a decline in women's earnings relative to men after the flood, as women are less able to transition into alternative occupations following the shock. In the months after the flood, women's earnings relative to men's earnings decrease by about 4 percentage points (53%). In the longer run, this decline reduces but remains at about 2 percentage points (24%). This decrease in women's relative earnings is linked to a reduction in female employment post-flood. I find that women are 8.8 percentage points less likely to be employed six months after the flood, leading to a significant drop in their earnings. In contrast, men's monthly earnings remain unaffected. This disparity arises in my sample because women, who are less likely to work in farming or salaried positions after a flood, do not find employment in other sectors to compensate for their loss. Meanwhile, men, though less likely to be employed in farming and self-employed activities after the flood, significantly increase their work as day laborers. As pointed out by Afridi, Mahajan and Sangwan (2022), this ability to substitute occupations helps men protect their income sources, while women's earnings decline. Consequently, the relative earnings gap between women and men widens, which drives a shift in household resources from women to men.²

My structural estimates are robust to the inclusion of different household covariates, changes to the sample of villages, and extensions of the analysis to villages in the 2016 and 2018 cohorts. Additionally, short-run results remain stable across variations in flood intensity: as the intensity of the flood increases—measured by a larger share of the village being inundated—the effects on women's resource shares remain consistent. However, in the long run, these effects diminish with higher flood intensity. The same pattern is observed in the relative earnings of women compared

²I do not find any significant or economically meaningful changes in male migration or fertility in my sample as a result of the flood.

to men within the household. As flood intensity rises, men's earnings also decline in the longer run, reducing the disparity between women's and men's earnings. This reduced disparity drives the weakening of the long-term redistribution of household resources away from women and toward men. Therefore, when the impact of the flood on women's relative earnings subsides, the impact on women's resource shares also diminishes in the longer run.

This paper expands and connects two strands of literature. First, it advances the literature on the impact of climate shocks. To my knowledge, this is the first study, along with parallel work by Aminjonov (2025), to disaggregate the impact of a climate shock on the intra-household division of resources among household members, as well as their individual-level expenditures, which are not directly observable. Existing research in the climate-economy literature primarily investigates the effects of climate shocks on consumption for the household as a unit (Anttila-Hughes and Hsiang 2013; Poaponsakorn, Meethom, and Pantakua 2015; Noy, Nguyen, and Patel 2021), and does not delve into these effects within the household. Other research by Christian et al. (2018) examines the effects of a climate shock on a limited set of observable private good expenditures. However, it does not account for potential shifts in unobserved consumption across a broader set of goods, and thus does not fully capture changes in an individual's total consumption. In contrast, I use a structural model to recover the share of total household budget allocated to men, women and children, and leverage exogenous variation in flood exposure to provide a holistic understanding of its effects on resource shares and individual-level expenditures.

This paper also contributes to the literature on resource share estimation by examining not only how resource shares change in response to a climate shock but also how these effects can persist over several years. Some existing studies use resource share estimation to measure individual-level poverty in a cross-section (e.g., Lechene, Pendukar, and Wolf 2022), while

others—such as Calvi (2020), Tommasi (2019), Sokullu and Valente (2021), and Hernandez-de-Benito (2024)—apply it to analyze the effects of programs or policies, with the latter focusing on local crime. This research extends these previous approaches by focusing on the impact of climate-induced shocks on intra-household resource allocation, examining not only the immediate effects but also how they evolve over time. While Aminjonov (2025) examines a related question, this paper differs in key ways—most notably in tracing how resource shares evolve over time following a climate shock, and in its focus on flood shocks in rural Bangladesh rather than rainfall shocks in Malawi.

The remainder of this paper is organized as follows. Section 2 discusses the background and contributions of this study to the existing literature. Section 3 outlines the theory behind the structural model used to identify resource shares over time and explains the estimation approach. In Section 4, I describe the household and flood data, explain how I predict the ex-ante probability of flooding at the village level, and use this measure to narrow my sample to villages with similar flood risk. I also demonstrate that, within this sample, exposure to a flood can be considered as-good-as-random. Section 1.5 presents the main results from the structural estimation of resource shares and shadow budgets, including results by the number of floods. Section 1.6 explores potential mechanisms driving the results. Section 1.7 provides robustness checks and examines how flood intensity affects the stability of the results. Section 1.8 concludes.

1.2 Related literature

The collective household model has become the main framework through which household allocation decisions are studied (see, for example, Browning, Chiappori, and Weiss 2014). This

model characterizes the allocation problem for a collection of individuals assuming they all have their own objective functions and bargain with each other to yield Pareto efficient allocations (Chiappori 1988, 1992). A few studies have developed methods based on this model to estimate the proportion of the total household budget allocated to each individual: their resource share. Lewbel and Pendakur (2008), Browning, Chiappori, and Lewbel (2013), and Dunbar, Lewbel, and Pendakur (2013) are some prominent studies in this area.³ The first two studies identify resource shares in a child-free setting, whereas the latter allows for nuclear families.

Specifically, Dunbar, Lewbel, and Pendakur (2013) identify individuals' resource shares by using Engel curves of private assignable goods and imposing semi-parametric restrictions on preferences of household members. This model has been extended in recent years to allow for complex household types and economies of scale (Calvi et al. 2023), linear re-framing of previous non-linear models (Lechene, Pendakur and Wolf 2022), and relaxed restrictions on preferences (Brown, Calvi, and Penglase 2021; Sokullu and Valente 2021).⁴

To my knowledge, Sokullu and Valente (2021) is the only paper that studies the evolution of resource shares in a longitudinal setting, over one year.⁵ Their structural model relaxes the semi-parametric restriction on preferences in Dunbar, Lewbel, and Pendakur (2013), but introduces an identifying assumption that resource shares change from baseline values. Therefore, if resource

³Intra-household resource distribution not only gives us insight into the consumption dynamics within a household, but also informs us of the bargaining power of individuals. Browning, Chiappori, and Lewbel (2013) show that there is a monotonic correspondence between the fraction of total household expenditure allocated to an individual (i.e. their resource share) and their bargaining power.

⁴Most of these papers use resource shares as a means to estimate individual-level poverty at a point in time. For example, Dunbar, Lewbel, and Pendakur (2013) show that children have higher rates of poverty than their parents in Malawi. A few studies utilize resource share estimation as a way to evaluate the effects of policies that improve the status of women in the household. Calvi (2020) found that strengthening women's inheritance rights in India increased the share of resources devoted to them. Similarly, Tommasi (2019) found that access to PROGRESA redistributed resources away from fathers towards mothers in beneficiary households in Mexico. Hernandez-de-Benito (2024) shows that an increase in local crime in Mexico decreases women's resource shares in the household.

⁵Sokullu and Valente (2021) evaluate the effect of PROGRESA over 13 months. Contrary to Tommasi (2019), they find that after one year, mother's share of household resources declined in treated households.

shares change immediately after a climate event and then revert to baseline values in the next period then the identification of resource shares would not be possible. This makes it less suitable for analyzing how resource shares evolve several years after a temporary climate shock, as in this study. Therefore, I adopt a different set of identifying assumptions that I argue are better suited for studying the effects of a temporary climate event over the longer-run. Specifically, I apply the Similarity Across People semi-parametric restriction, as in Dunbar, Lewbel, and Pendukar (2013), but extend it over different time periods to capture the evolution of resource shares over four years.

Qualitative field research in South Asia suggests that floods could disproportionately affect the status of women in the household. Cultural norms in the region dictate that women are responsible for household work such as providing clean drinking water and food in any climate stressed situation. Memon (2020) interviews women in flood prone areas of Pakistan and finds that when women are unable to fulfil their household duties in a satisfactory manner after a flood, they can be met with emotional or physical violence. In Bangladesh, Rezwana and Pain (2020) discuss how women's household work and responsibilities increase after a flood, making it difficult for them to search and take-up new jobs. This increases their dependence on other household members. Similarly, Mitchel, Tanner and Lussier (2007) find that women's livelihoods are particularly affected by erratic monsoon patterns, with women expressing losses in crops, animal fodder, poultry and livestock.

The new climate-economy literature is limited in its ability to characterize distributional impacts of climate-induced shocks, such as floods, *within* the household. Existing papers study the effects on consumption and income at the household level (Anttila-Hughes and Hsiang 2013; Poaponsakorn, Meethom, and Pantakua 2015) rather than for individuals. A few studies look into

how effects could vary by sector of economic activity (Noy, Nguyen, and Patel 2021) and time since the flood shock (De Alwis and Noy 2019; Mueller and Quisumbing 2011). Some papers delve into risk-coping strategies employed by affected households. For example, Gianelli and Canessa (2022) find an increase in male migration in households exposed to a flood.⁶

A number of studies have explored how these shocks impact proxies for women’s status, but this paper goes a step further by investigating how these changes in status trigger a reallocation of household resources. For example, Afridi, Mahajan, and Sangwan (2022) show that exposure to climate shocks reduced women’s labor market opportunities more than men’s, largely due to gender norms that restricted women’s ability to transition to other types of employment. In contrast, Gianelli and Canessa (2021) find that a climate shock led to an increase in women’s labor supply as a coping strategy. Hoddinott (2006) provides evidence of lower body mass index among women relative to male household members following a climate shock, while Miguel (2005) finds an increase in violence against elderly women. Dessy et al. (2021) highlight that in some contexts, climate shocks have resulted in polygamous marriages, which can affect women’s bargaining power within households.⁷ Building on this literature, I focus specifically on how shifts in women’s status within the household following a climate shock translate into tangible changes in their material well-being.

1.3 Intra-household allocation of resources

This section outlines the theory and estimation strategy for identifying the fraction of total household consumption expenditure allocated to men, women, and children, referred to as their

⁶See Karim and Noy (2016) and Hallegatte et al. (2020) for a review of the literature on the effects of climate shocks on household consumption and poverty.

⁷See Fruttero et al. (2023) for a review of the literature on gendered effects of climate change.

resource share. Since household members consume both private goods and shared goods among them, with the degree of sharing unknown, individual level consumption in households with multiple members is not directly observed. I will describe the structural model developed in Dunbar, Lewbel, and Pendukar (2013) to identify resource shares in the cross-section and extend it to identify resource shares in a panel data setting. For simplicity, I focus on households that have one man, woman, and child. The model can be generalized to nuclear households with multiple children and extended households. Households can differ in terms of observable characteristics such as their composition, socio-economic attributes, and exposure to the “treatment”, i.e., a flood. These differences could affect preferences and how resources are distributed within the household. However, these covariates are not *required* for identification of resource shares, so I leave them implicit to simplify notation. These covariates will enter explicitly into the model in the empirical estimation in Section 1.3.4 to allow for some flexibility in preferences and resource shares across households.

1.3.1 The efficient collective household model

In any given period, all household members J come together to purchase K goods at market prices $p = (p^1, \dots, p^K)$. The vector of total observable quantities of goods purchased by the household is $z = (z^1, \dots, z^K)$, and the unobservable quantities of goods consumed by household member j is given by $x_j = (x_j^1, \dots, x_j^K)$. For ease of exposition, I omit time subscripts for now.

Since goods can be sharable or nonsharable, the quantities purchased by the household are not necessarily equal to the sum of individual consumption. A linear consumption technology transforms household-level purchased goods z to individual-level consumption x_j , taking into

account economies of scale and complementarities in consumption for all goods in the household, except for those privately consumed by individuals. This technology is represented through a $K \times K$ matrix A such that $z = A \sum_{j=1}^J x_j$.⁸ Therefore, if there is one man, one woman and one child, $z = A(x_m + x_w + x_c)$. Matrix A also allows us to express the unobserved shadow prices of goods as $A'p$. This means that for shared goods with economies of scale like gasoline or heating, the shadow price of that good will be less than the market price. If goods are not shared by members, then the shadow price should equal to the market price. Intuitively, if we could solve for the bundle of goods x_j and price it at shadow prices $A'p$, we will get the individual-level expenditure for person j . Dividing by the total household expenditure would give us the share of total household consumption allocated to that individual.

Following the literature, I assume a collective household model in which every household member j has their own utility function $U_j(x_j)$, the allocation of goods within the household allocation is Pareto efficient, and the household does not suffer from money illusion.⁹ $U_j(x_j)$ is twice continuously differentiable, monotonically increasing and strictly quasi-concave over the vector of consumption goods x_j . The household maximizes a joint social welfare function U_h in which the preferences of each individual are weighted. These weights are called “Pareto weights” and denoted as μ_j .¹⁰ To allow for caring preferences, each member’s utility function is weakly

⁸ A is a block diagonal matrix that allows for sharing or jointness of consumption for all good except a private assignable good – a good that is consumed exclusively by a known member of the household. For any private assignable good, there will be a 1 in the elements corresponding to this good and all off-diagonal rows and columns for this good will be 0. This is because these goods do not have any economies of scale in consumption (since they are not shared). The off-diagonal elements being 0 implies that there are no complementarities in consumption for these goods. For all other goods, including any shared goods like shelter or gasoline, there are no restrictions on complementarities and economies of scale.

⁹ Recent work by Brown, Calvi and Pengele (2021) does not reject the assumption of Pareto efficiency in the Bangladeshi context using a formal test based on distribution factors, which are variables affecting claims to resources but not preferences for goods. Lewbel and Pendukar (2024) allow for some types of consumption inefficiencies in collective households and show that estimates of resource shares change at most by only 2.7 percentage points.

¹⁰ $\mu_j = \mu_j(p)$ is homogeneous of degree zero, so $\mu_j(p_t) = \mu_j(p_t/y_t)$ (Dunbar, Lewbel and Pendukar 2013).

separable over their own utility as well as the utility of other household members. Therefore, utility for person j takes the following form: $\tilde{U}_j = \tilde{U}_j(U_j(x_j), \dots, U_J(x_J))$. Household members maximize the Bergson-Samuelson social welfare function:

$$U_h(\tilde{U}_j(x_j), \dots, \tilde{U}_J(x_J), p, y) = \sum_{j=1}^J \mu_j \tilde{U}_j, \quad (1.1)$$

where y is total household expenditure and J is the total number of members. Browning, Chiappori, and Lewbel (2013) show that there is a one-to-one correspondence between resource shares and Pareto weights. Resource shares are often understood as measures of intra-household bargaining power, but they are also influenced by altruism, particularly in the case of shares claimed by children (Dunbar, Lewbel, and Pendukar 2019). If exposure to a climate shock increases a member's Pareto weight, it will increase the weight of that person's utility in the household's joint social welfare function. This will change the way the household optimally allocates resources amongst its members. In what follows I will be using Pareto weights interchangeably with bargaining power.

The household solves the following optimization problem:

$$\begin{aligned} \max_{z, x_1, \dots, x_N} \quad & \sum_{j=1}^J \mu_j \tilde{U}_j(x_j), \\ \text{s.t.} \quad & z = A \sum_{j=1}^J x_j \text{ and } y = z'p, \end{aligned} \quad (1.2)$$

The solution of this optimization problem gives the private quantities for the vector x_j for each member j . Pricing x_j at shadow prices $A'p$ and scaling by y yields the resource share of individual j , which is denoted as η_j below. This is the parameter of interest.

By the second welfare theorem, the solution to the household's allocation problem above can be equivalently represented as a two-step process. Household members first allocate resources so that the shadow budget constraint $\eta_j y$ for each household member can be determined; and then each member j chooses bundle x_j that maximizes their utility subject to shadow prices $A'p$ (which are the same for all household members), and a shadow budget constraint $\eta_j y$ (which are different across members).

The household's demand for a generic good k can be expressed as follows:

$$z^k = A^k \sum_{j=1}^J h_j^k(\eta_j y, A'p), \quad (1.3)$$

where $h_j^k(\eta_j y, A'p)$ is individual j 's Marshallian demand for good k subject to the shadow price and their shadow budget constraint. For example, if a household has one man, one woman, and one child, the household demand for good k would be:

$$z^k = A^k (h_m^k(\eta_m y, A'p) + h_w^k(\eta_w y, A'p) + h_c^k(\eta_c y, A'p)). \quad (1.4)$$

If a good j is privately consumed by only person j , then the household's demand for that good simplifies because there are no economies of scale and good j factors into only the utility function for person j .

$$z^j = h_j^j(\eta_j y, A'p). \quad (1.5)$$

In what follows, to simplify notation, the Marshallian demand for a good consumed exclusively by person j is represented as $h_j(\eta_j y, A'p)$. Equations (1)-(5) apply in each time period t .

1.3.2 Engel curves of private assignable goods

Dunbar, Lewbel, and Pendukar (2013) show that resource shares for a person j where $j = m, w, c$ can be identified in the cross-section by comparing the slopes of the household's Engel curves for a private good that is assignable to each person j . An Engel curve captures the relationship between the share of total household budget spent on a good and total household expenditure, holding prices fixed. A private assignable good (PAG) is a good that is not shared and is consumed exclusively by a known member of the household. Examples include food, clothing or footwear, among others. The PAG in this paper is the food expenditure of person j .¹¹ Food is private in the sense that a unit of food eaten by one person cannot also be consumed by someone else. It is also assignable because the Bangladesh Integrated Household Survey (BIHS) – the household data I am using for this analysis – has information on food expenditure of each household member.¹²

The Marshallian demand for a private good j that is assignable to person type j can be derived from that individual's indirect utility function. Following Dunbar, Lewbel and Pendukar (2013) I assume that resource shares are independent of total expenditures in poor households¹³ and that each person has PIGLOG preferences; however, I extend their cross-sectional approach to multiple periods, with each person j having their own PIGLOG indirect utility function in each

¹¹There are several reasons for preferring food over clothing expenditure shares. One key reason is that food is more private and assignable compared to clothing, which can be shared among different household members. Additionally, food expenditures are likely to be larger than clothing expenses because food is less durable, which can contribute to relatively steeper Engel curves. For a detailed discussion of these advantages, see Lechene, Pendukar, and Wolf (2022).

¹²If the BIHS did not have information on individual level food consumption, food would still have been a private good but not assignable and could not have been used for estimation.

¹³This assumption implies that, for low levels of total expenditure, if a household's total expenditure increases, resource shares across household members would not change. This is a strong assumption, but there is some empirical support for it (Menon, Perali, and Pendakur 2012). Most importantly, this restriction does not apply to variables closely related to household expenditure, such as household income, or member's wages, or measures of wealth (Dunbar, Lewbel and Pendukar 2013).

period t :

$$V_{jt}(p_t, y_t) = \psi_{jt} \left[\ln \left[\ln \left(\frac{y_t}{G_{jt}(p_t)} \right) \right] + F_{jt}(p_t), \tilde{p}_t \right], \quad (1.6)$$

where p_t is the price of all goods including the PAGs, $\tilde{p}_t$ is the price of all goods except the PAGs. G_{jt} is an individual-specific expenditure deflator function, that is non-zero, differentiable and homogenous of degree 1. F_{jt} is individual-specific, differentiable, and homogenous of degree 0. ψ_{jt} is differentiable and strictly monotonically increasing. The price of the PAG for person j is p_{jt} .

To derive person j 's Marshallian demand for the PAG in a given time period, apply Roy's identity, such that $h_{jt}(p_t, y_t) = -\frac{\partial \psi_{jt}(\cdot)}{\partial p_{jt}} / \frac{\partial \psi_{jt}(\cdot)}{\partial y_t}$. This gives the following expression (surpressing arguments of G_{jt} and F_{jt} for conciseness):

$$h_{jt}(p_t, y_t) = y_t \left(\frac{G'_{jt}}{G_{jt}} + F'_{jt} \ln G_{jt} \right) - F'_{jt} y_t \ln y_t, \quad (1.7)$$

where G'_{jt} and F'_{jt} are first derivatives with respect to p_{jt} . As an individual makes a consumption decision based on their shadow budget constraint $\eta_{jt} y_t$, and shadow prices $A' p_t$, the $h_{jt}(p_t, y_t)$ in equation (1.7) can be expressed as $h_{jt}(A' p_t, \eta_{jt} y_t)$.

To express as budget shares of the PAG, multiply both sides by $\frac{p_{jt}}{y_t}$, and rearrange terms:

$$W_{jt} = \eta_{jt} \alpha_{jt} + \beta_{jt} \eta_{jt} \ln \eta_{jt} + \beta_{jt} \eta_{jt} \ln y_t, \quad (1.8)$$

where $\alpha_{jt} = p_{jt} \left(\frac{G'_{jt}}{G_{jt}} + F'_{jt} \ln G_{jt} \right)$, and $\beta_{jt} = -p_{jt} F'_{jt}$. W_{jt} is the share of the total household budget that is spent on good j . α_{jt} and β_{jt} can be interpreted as preference parameters because they determine the intercept and slope of this curve. β_{jt} can be loosely interpreted as the marginal

propensity to consume the PAG as household expenditure increases.

Importantly, the budget share spent on a PAG for person j , W_{jt} , and their resource share, η_{jt} , capture different information and one cannot substitute for the other. For instance, differences in preferences for goods can lead to cases where an individual holds the highest W_{jt} but has the lowest resource share in the household. I provide an example in Appendix A.1 to illustrate this argument.

1.3.3 Identification of resource shares

To estimate the resource shares of person j at any given time t , I would *ideally* like to estimate a system of Engel curves, one for each person j :

$$\begin{aligned} W_{wt} &= \left[\eta_{wt}\alpha_{wt} + \beta_{wt}\eta_{wt} \ln\eta_{wt} \right] + \eta_{wt}\beta_{wt} \ln y_t, \\ W_{mt} &= \left[\eta_{mt}\alpha_{mt} + \beta_{mt}\eta_{mt} \ln\eta_{mt} \right] + \eta_{mt}\beta_{mt} \ln y_t, \\ W_{ct} &= \left[\eta_{ct}\alpha_{ct} + \beta_{ct}\eta_{ct} \ln\eta_{ct} \right] + \eta_{ct}\beta_{ct} \ln y_t. \end{aligned} \tag{1.9}$$

With data on budget shares and total household expenditure, we could regress W_{jt} on y_t , and the slope would estimate $\eta_{jt}\beta_{jt}$, while the constant would estimate the term in the brackets in equation (1.9). However it is not possible to identify η_{jt} from this system, as there are only 6 moment conditions (3 slopes and 3 intercepts) and 9 unknowns for any period t . If we account for the restriction that $\eta_{mt} + \eta_{wt} + \eta_{ct} = 1$, it brings the number of unknown parameters down to 8, but the model is still under identified.

Therefore further restrictions are needed on the unknown parameters to identify η_{jt} . One possible assumption is that $\beta_{jt} = \beta_t$, which implies the marginal propensity to consume the

private assignable good varies across time, but is the same across individuals within a given time period. This is equivalent to applying the Similarity Across People (SAP) restriction in Dunbar, Lewbel, and Pendukar (2013) for each time period t .

Since $\beta_{jt} = -p_{jt} F'_{jt}$, this restriction has two components: first, the price of the PAG (p_{jt}) should not vary across household members, and second, that the derivative of $F_{jt}(p_t)$ with respect to p_{jt} must not vary with person j . In this paper, the private assignable good is food expenditure, and food prices do not vary based on who consumes the good, satisfying the first condition. The latter restriction implies that the shape of the Engel curves for the PAG should be the same across j . It does not imply that preferences for the PAG are identical across j , as they can still vary across individuals through an intercept shift (i.e. through the α_{jt}). Importantly, this restriction on preferences is only for the PAG, not any other good. In Appendix A.3, I provide a fully specified example of an indirect utility function that satisfies this requirement.

With $\beta_{jt} = \beta_t$, the number of unknown parameters reduces to 6 in a given time period t and matches the number of moment conditions, allowing the identification of η_{jt} by taking the ratio of the slopes and solving for η_{jt} . If the slope of the Engel curve is steeper for men relative to women, then it means they have a higher share of household resources devoted to them.

There are alternative assumptions that can be considered on β_{jt} . For example, Sokullu and Valente (2021), who also use longitudinal data to see how resource shares evolve over one year, propose a Similarity Over Time (SOT) restriction that implies β_{jt} can vary across individual, but not over time – i.e. $\beta_{jt} = \beta_j$. However, a key identifying assumption in this approach requires some change in the resource shares of household members since baseline.¹⁴ If resource shares do

¹⁴Under SOT, assumption 4 in Sokullu and Valente (2021) simplifies to the requirement that the following matrix needs to be non-singular, where η_{j2}/η_{j1} is the resource share of person j in $t = 2$ relative to baseline (i.e. $t = 1$), and η_{j3}/η_{j1} is the corresponding ratio for $t = 3$.

not change over time, or if they change and then revert back to original levels, then SOT cannot be imposed. Given that I am studying the effects of a flood up to four years after the event, and considering the temporary nature of this shock, it is possible that resource shares change immediately after the flood and then revert back. This possibility makes the SOT restriction less viable in the context of this study. The SAP restriction, on the other hand, is agnostic about resource shares changing over time.

1.3.4 Estimation of resource shares

I follow Calvi (2020) and adjust the system of Engel curves in equation (1.9) to allow for extended households, pool data from all three rounds of the Bangladesh Integrated Household Survey, and estimate the following system of equations for a household h :

$$\begin{aligned} W_{wht} &= \eta_{wht} \alpha_{wht} + \beta_{ht} \eta_{wht} \ln\left(\frac{\eta_{wht}}{F_h}\right) + \eta_{wht} \beta_{ht} \ln y_{ht} + \epsilon_{wht}, \\ W_{mht} &= \eta_{mht} \alpha_{mht} + \beta_{ht} \eta_{mht} \ln\left(\frac{\eta_{mht}}{M_h}\right) + \eta_{mht} \beta_{ht} \ln y_{ht} + \epsilon_{mht}, \\ W_{cht} &= \eta_{cht} \alpha_{cht} + \beta_{ht} \eta_{cht} \ln\left(\frac{\eta_{cht}}{C_h}\right) + \eta_{cht} \beta_{ht} \ln y_{ht} + \epsilon_{cht}. \end{aligned} \tag{1.10}$$

$$\begin{pmatrix} 1 & 1 & 1 \\ \eta_{m2}/\eta_{m1} & \eta_{w2}/\eta_{w1} & \eta_{c2}/\eta_{c1} \\ \eta_{m3}/\eta_{m1} & \eta_{w3}/\eta_{w1} & \eta_{c3}/\eta_{c1} \end{pmatrix}$$

If resource shares do not change since baseline, or if they change back to original levels in $t = 3$, then this matrix would be singular. Also, if they do not change substantially, the matrix can be near-singular, which can cause numerical instability and lead to difficulties in achieving convergence and reliable parameter estimates in iterative procedures (Yuan and Chan, 2008).

Results from my estimates in Table 1.3 show that, in the reference household, resource shares for women do not change substantively since baseline, and the determinant of the matrix above is near-zero (-0.128). Therefore, when I tried to estimate resource shares using the SOT restriction, the model produced unreliable results.

F_h , M_h , and C_h refer to the numbers of women, men, and children in the household, respectively. They are added to allow for the possibility of an extended household, with multiple male and female adults, as is common in many developing countries like Bangladesh. If there is just one man, one woman, and one child, this model reverts to equation (1.9). Since exposure to a flood could potentially affect the number of men, women, and children in the household, F_h , M_h , and C_h are numbers at baseline. I categorize adults as those aged 16 or above at baseline in 2011, adjusting the threshold in subsequent surveys to ensure the same individuals are consistently defined as men, women, or children across all survey rounds. Further details are provided in Appendix A.4.

As errors could be correlated across equations, I estimate the model using Nonlinear Seemingly Unrelated Regression (NLSUR). To allow preferences and resource shares to vary flexibly across households, I model each unobservable parameter— η_{jht} , β_{ht} , and α_{jht} —as a combination of observable household socioeconomic characteristics, time variables, and exposure to flooding. Specifically, the resource share η_{jht} for person type $j = m, w, c$ at time t in household h are modeled as:

$$\begin{aligned} \eta_{jht} = & \eta_{0j} + \eta_{1j}Midline + \eta_{2j}Endline + \eta_{3j}Flood_h \times Midline \\ & + \eta_{4j}Flood_h \times Endline + \eta_{5j}X_h + \eta_{6j}W_{ht}, \end{aligned} \quad (1.11)$$

where $Flood_h$ is an indicator variable for whether the household was exposed to a flood, $Midline$ and $Endline$ are indicator variables referring to the second and third rounds of the BIHS survey (the first survey conducted in 2011/12 is the baseline), X_h refers to household covariates that are time invariant, and W_{ht} are covariates that can vary with time. Standard errors allow for the

arbitrary correlation for error terms within the village, which is the level of treatment assignment. It should be noted that η_{jht} represents the resource share that is allocated collectively to people of that category. Within each person category, resource shares are assumed to be evenly distributed, as in Calvi (2020). For instance, if the household allocated 30% of its resources to children and there are two children in the household, then each child would get 15%.

In the preferred specification, X_h includes the the numbers of men, women and children in the household (centered around the mean at baseline), whether the household had a woman with at least primary education, and a child under the age of 5 – all measured at baseline. W_{ht} includes six binary variables for whether the numbers of men, women, and children in the household increased or decreased since baseline. Given the findings in Calvi (2020) about how women in India approaching post-reproductive ages experience a decrease in the household resources allocated to them, W_{ht} also includes a binary variable for whether the household has a woman over the age of 40.¹⁵ I check the robustness of the main estimates by adding additional covariates in Section 1.7.

The parameters of interest, η_{3j} and η_{4j} , along with the constant term η_{0j} , provide key insights into resource allocation. The constant term reflects the resource share for person type j in a reference household at baseline—specifically, an unflooded household with the average number of men, women, and children, where all binary covariates in equation (1.11) are zero. Meanwhile, η_{3j} and η_{4j} indicate whether resources shift towards or away from person type j after exposure to a flood in 2014, with the former capturing the short-run effect and the latter the longer-run effect compared to unaffected households. Appendix A.2 provides a graphical illustration of how these

¹⁵Calvi (2020) shows that in households with children, women’s resource shares begin to decrease approximately after the age of 40.

parameters are identified within the Engel curve framework.

Since exposure to the flood could affect preferences, I model β_{ht} and α_{jht} similarly and I allow them to change with time, exposure to a flood, and the household covariates in equation (1.11). This approach is intended to account for the findings in De Rock, Potoms, and Tommasi (2022), which provide evidence of changing preferences following the receipt of a cash transfer program.¹⁶

1.4 Data and sample selection

1.4.1 Flooding data

My source for flood maps is the Global Flood Database (GFD) which uses daily satellite imagery at 250-metre resolution to estimate flood extent for 913 large flood events documented by the Dartmouth Flood Observatory (DFO) from 2000 to 2018 (Tellman et al. 2021). For each flood event, the GFD provides a map that shows the maximum observed surface-water extent, the start and end date of the flood, as well as an estimate for the total population that was exposed. The DFO catalogue is compiled largely from news reports, so flood event maps provide a lower bound of maximum flood extent.

From the GFD, I map 97 floods that affected Bangladesh between 2000 to 2018.¹⁷ As an example, Appendix Figure B.2 shows the flood map for one event that occurred from August 20 to September 8 in 2014 that impacted around 11 million people. The GFD also allows me to generate a satellite-based flood plain that outlines all areas that have experienced a flood from

¹⁶Put differently, I do not use flood exposure as a distribution factor – defined as variables that affect the distribution of resource but not preferences for goods —because a flood, as a significant event, could influence preferences, invalidating it as a distribution factor.

¹⁷91 of these floods were due to heavy rain, 2 due to dams, and 4 due to a tropical storm.

2000 to 2018. This is one of the variables I use to predict the probability of flooding in villages that will be described in more detail later in Section 1.4.4.

1.4.2 Household level data

I use all three waves of the Bangladesh Integrated Household Survey (BIHS), which is a household panel survey nationally-representative of rural Bangladesh conducted by International Food Policy Research Institute (IFPRI) in 2011/12, 2015, and 2018/19 (Ahmed, 2013; IFPRI, 2016; IFPRI, 2020).¹⁸ Households are located in all of Bangladesh’s 64 districts and in 275 villages.¹⁹ The dataset provides detailed socio-economic information, including demographics, household expenditures, employment, agricultural activities, and migration.

The BIHS stands out because it contains information on food intake of each individual household member, unlike most household panel surveys. As explained in Section 3, this disaggregated information, even if it is just for the food module, provides crucial information on a private assignable good (PAG) – a good consumed exclusively by a known member of the household – that is necessary for estimating resource shares. I focus on a balanced panel of 1,596 households that had at least one man, one woman, and one child in each round of the survey, and had usable information on their food expenditure shares. Details about data cleaning are in Appendix A.4.

¹⁸The BIHS defines a household as a group of people who live together and take food from the “same pot”. If two brothers stay in the same house with their families but they do not share food costs and they cook separately, then they are considered two separate households.

¹⁹There are 323 total geo-located villages across all rounds of the survey, but households were repeatedly surveyed in 275 of them.




Figure 1.1: Number of times each village has been hit by a flood from 2000 to the final round of the BIHS survey in 2018/19. Villages in green-shaded unions did not get hit by any floods from 2000 to 2018. Gray lines show inland perennial water. The unions in which the villages are located are presented in this graph instead of village GPS coordinates, to protect the identity of individual villages.

1.4.3 Exposure to floods

To determine whether households in my sample were exposed to floods, I obtain data from IFPRI on average village GPS coordinates, with households located within 5km. I create a buffer of 5km around these village GPS coordinates, overlay it with each of the 97 flood maps for Bangladesh from 2000 to 2018, and compute the share of area within 5km of the village that was inundated during each flood event. I binarize this treatment as equal to 1 if more than 15% of the area around the village is flooded and 0 otherwise. Treatment is assigned at the village level, and all households in the village are considered flooded (or unflooded).²⁰ The threshold is set at 15% because during the study period – i.e. the time between the baseline and endline surveys – 15.3% is the average area around the village that was inundated by water across all flood events

²⁰I do not have access to household GPS coordinates, which are kept confidential.

(see Figure B.3). I perform sensitivity checks to this threshold in Section 1.7.4.

From 2000 until the last BIHS survey conducted in 2018, villages experienced up to 42 floods. Figure 1.1 illustrates the cumulative number of floods experienced by the 275 villages in my sample. To protect the identity of individual villages, the graph depicts the union (a level 4 administrative boundary) in which each village is located. About 65% of the villages were exposed to at least one flood, with an average of 11.02 flood occurrences per village from 2000-2018 (see Figure B.4).

Between the first and last BIHS survey, there were 20 flood events in Bangladesh. These floods affected 99 villages in the sample at various points in time. Figure 1.2 shows villages based on the year they experienced their *first* flood since the baseline BIHS survey was conducted. Among them, 84 villages faced their initial flood event in 2014, 2016, or 2018. Specifically, in 2014, 51 villages were exposed to their first flood event since the baseline BIHS survey, forming the largest cohort of flooded villages.

1.4.4 Ex-ante probability of flooding

Villages do not have a uniform risk of flooding. This can be seen in Figure 1.1, which shows that villages located in certain parts of the country experience a higher frequency of flooding than others. For example, villages closer to inland waterbodies experienced a greater number of floods relative to areas further away. Villages also differ in terms of other geographic characteristics that can influence their probability of being exposed to a flood, such as rainfall, topography and vegetation cover, among others (USGS). Figure A.3 illustrates differences in rainfall, Normalized Difference Vegetation Index (NDVI), flood zones, and elevation across villages in my sample (see




Figure 1.2: Cohorts of villages based on the year they were hit by their first flood since the baseline BIHS survey. 99 villages were hit by floods from 2012-2018 at different times. The unions in which the villages are located are presented in this graph, to protect the identity of individual villages.

Appendix A.4 for details).

I estimate a probit model to predict a village's probability of flooding at baseline. The outcome variable is one if a village experienced any flood by the time of the first household survey in 2011 and zero otherwise. Village-level regressors include time variant and invariant characteristics. The former category includes village-level trends for annual rainfall, NDVI, and land surface temperature from 2001 until 2011. Time invariant characteristics include the village's elevation, climate zone, distance to inland waterbodies, administrative division, whether the village is located in a flood plain, is close to a major river or stream, or in an area prone to tidal, flash, or river flooding.

To see how well this measure of ex-ante flood risk is associated with future floods, I examine its relationship with the actual number of floods experienced by villages during the study period between 2012 and 2018. Figure B.5 illustrates that villages with higher ex-ante




Figure 1.3: Number of villages (flooded and unflooded) in each cohort after narrowing sample based on ex-ante flood risk and geographic area.

probabilities experienced a greater number of floods during this period.

1.4.5 Villages with similar ex-ante flood risk

I limit my sample to households located in villages that have similar geographic characteristics and ex-ante flood risk. To accomplish this, within each cohort, I narrow down the sample to flooded and unflooded villages that are located within the same geographic area and have a maximum allowable difference in ex-ante flood risk of 0.3. For instance, if a flooded village has an ex-ante flood probability of 0.9 then selected unflooded villages in the same geographic area must have a baseline probability of 0.6 or higher. This ensures a comparison of flooding effects between villages with similar baseline risk of flooding, differing only in treatment status, while maintaining geographic similarity among villages. Further details can be found in Appendix A.4. In Section 1.7.3, I demonstrate that the main results are robust to different thresholds of baseline

flood risk.

Based on these selection criteria, I narrow the sample from 275 to 118 villages, of which 49 experienced a flood during the study period (2012 to 2018), and 69 shared a similar ex-ante flood risk and were geographically close but did not experience a flood during the study period. The 2014, 2016 and 2018 cohorts constitute 75% of the villages in this sample, as shown in Figure 1.3.

I focus exclusively on the 2014 sample for the main results, as it is the largest cohort of villages (see Figure 1.3), and draw on the 2016 and 2018 samples only for robustness checks. Households in these cohorts are balanced in terms of socio-economic characteristics at baseline – including household composition, education attainment, total household expenditure, and food expenditure shares, as illustrated in Appendix Figure B.7. They are also balanced in terms of ex-ante flood probability and geographic region, and are equally likely to be in a flash flood or river flood zone. Furthermore, these covariates are not jointly significant in determining treatment status.²¹

To illustrate further that the timing of the first flood during the study period is arguably random, I plot the difference in annual rainfall between flooded and unflooded villages in the selected sample in the years preceding and following the flood (Figure 1.4). High rainfall was cited as the cause for 19 of the 20 floods that occurred in Bangladesh during the study period, but after narrowing down the sample based on ex-ante flood risk and geographic location, there are no pre- or post- trends in annual rainfall in flooded compared to unflooded villages (Panel A in Figure 1.4). The graph is flat, and the estimates are statistically indistinguishable from zero. For example, the point estimate at the time of the event is 0.23 standard deviations, which amounts to

²¹F statistic: 0.64; pvalue 0.85




Figure 1.4: Differences in standardized annual rainfall in flooded and unflooded villages in the years before and after their first flood during the study period. Panel A shows these differences for the 2014, 2016, and 2018 cohorts in a selected sample of villages that have a similar ex-ante probability of flooding and are located in the same 100km grid. Households in this sample of villages are used to produce the main results. Panel B is based on a further restricted sample of villages, excluding three treated villages (and associated control villages) in the Sylhet division that experienced some of the highest rainfall during the study period relative to surrounding villages. This reduced sample is used for robustness checks only. Shaded area represents 95% confidence intervals.

only 14.7 cm of rainfall in a year. Given that the 20-year average annual rainfall for a village is 235.34 cm, this difference is economically insignificant. Further details are provided in Appendix Section A.4, including the specification used to produce Figure 1.4.²²

To address any lingering concerns about rainfall differences, I restrict the sample further by removing seven villages, including three flooded villages that experienced exceptionally high

²²Prior to narrowing the sample based on ex-ante flood risk and geographic location, there existed systematic differences in rainfall between flooded and unflooded villages in the full sample, see Figure A.4 in Appendix.

rainfall compared to their surrounding villages. This subsample is used only for robustness checks. Panel B presents rainfall differences in this robustness sample, where the point estimates along the event timeline are even closer to zero. Figure A.5 in the Appendix plots the differences in land surface temperature between flooded and unflooded villages, showing similar results.

1.4.6 Descriptive statistics for the 2014 sample

Only the 2014 sample in Figure 1.3 is used to produce the main results in Section 1.5. These households are observed in the baseline survey conducted in 2011/2012, then six months after their first flood in 2014 (corresponding to the midline survey in 2015), and finally, four years post-flood, in the endline BIHS survey conducted in 2018/19. From the initial flood in 2014 until the endline survey in 2018, villages experienced up to 11 floods.

The average household is nuclear with two children, as shown in Table 1.1. Children have the highest food budget share at 21%, followed by men at 18% and women at 16%.²³ In terms of employment and earnings, 86% of households had a working woman, 98% had a working man, and children's labor participation increased from 12% at baseline to 43% by 2018.²⁴ Also, note that men's average monthly income is substantially higher than that of women and children. On average, women's earnings are 8.26%, and children's earnings are 6.12% of men's monthly earnings.

These villages are not representative of rural Bangladesh and are likely poorer and more flood-prone. The average baseline risk of flooding is 95% and flood-affected households experienced an average of 2.1 floods since their initial flood in 2014 until the time of the midline survey and

²³Table B.1 in the Appendix has other household characteristics and shows that the mean age of both women and men increases from baseline to endline, which is encouraging as it reinforces the idea that the same individuals are being tracked in each person category.

²⁴The legal age of employment in Bangladesh is 14 years.

Table 1.1: Summary of household characteristics

	All	Baseline	Midline	Endline
Number of men (16+ years at baseline)	1.16 (0.42)	1.19 (0.46)	1.16 (0.43)	1.12 (0.38)
Number of women (16+ years at baseline)	1.20 (0.44)	1.21 (0.46)	1.20 (0.43)	1.17 (0.43)
Number of children (≤ 15 years at baseline)	2.16 (0.97)	2.03 (0.98)	2.20 (0.96)	2.26 (0.97)
Men's food budget share	0.18 (0.08)	0.19 (0.09)	0.19 (0.08)	0.16 (0.07)
Women's food budget share	0.16 (0.07)	0.16 (0.07)	0.17 (0.08)	0.14 (0.07)
Children's food budget share	0.21 (0.11)	0.18 (0.11)	0.24 (0.10)	0.23 (0.10)
Total annual household exp. (1000 Taka)	178.56 (123.95)	113.03 (76.21)	181.69 (112.74)	240.97 (139.39)
1(Household below extreme poverty line)	0.16 (0.37)	0.27 (0.44)	0.12 (0.33)	0.09 (0.29)
1(Woman worked in last 7 days)	0.86 (0.35)	0.82 (0.39)	0.87 (0.34)	0.90 (0.30)
1(Man worked in last 7 days)	0.98 (0.15)	0.98 (0.12)	0.97 (0.16)	0.97 (0.17)
1(Child worked in last 7 days)	0.26 (0.44)	0.12 (0.33)	0.20 (0.40)	0.43 (0.50)
Women's monthly earnings (Taka)	423.03 (956.90)	333.78 (768.14)	434.09 (968.31)	500.69 (1,099.15)
Men's monthly earnings (Taka)	7,858.52 (5,746.82)	5,594.24 (3,987.57)	7,843.52 (5,468.43)	10173.33 (6,569.12)
Children's monthly earnings (Taka)	309.21 (1,276.41)	63.97 (486.72)	276.36 (1,177.02)	537.49 (1,698.67)
Ratio of women's to men's monthly earnings (%)	8.26 (27.00)	10.56 (38.84)	6.47 (14.78)	7.75 (21.25)
Ratio of children's to men's monthly earnings (%)	6.12 (32.73)	2.00 (18.55)	8.53 (46.29)	7.24 (25.03)
1(Access to relief program)	0.06 (0.24)	0.09 (0.29)	0.06 (0.23)	0.03 (0.17)
1(Flooded)	0.32 (0.47)	0.32 (0.47)	0.32 (0.47)	0.32 (0.47)
Number of floods if treated	2.59 (2.95)	0.00 (0.00)	2.12 (0.76)	5.65 (3.03)
Annual rainfall in cm (20 year average)	235.34 (63.70)	235.34 (63.75)	235.34 (63.75)	235.34 (63.75)
Probability of flood risk at baseline	0.95 (0.08)	0.95 (0.08)	0.95 (0.08)	0.95 (0.08)
Observations	1173	391	391	391

Notes: Descriptive statistics for households in selected flooded and unflooded villages in the 2014 sample. Mean with standard deviations in parentheses.

5.65 floods by the endline survey. Additionally, 27% of households in the sample live in extreme poverty, based on the World Bank's international extreme poverty line (\$2.15 per day per capita in PPP terms). The poverty headcount ratio is approximately 28.1%, which is higher than the World Bank's estimate of 18% one year prior to the baseline survey.²⁵ Furthermore, across all survey rounds, an average of only 6% of households reported having access to a flood relief program.

There is also a clear division of labor on gender lines at baseline, with women's main occupations concentrated in farming, while men's are spread across a wider range of sectors. Figure B.8 illustrates the percentage distribution of main occupations for men and women, excluding non-earning roles such as students or housewives. I grouped earning occupations into 10 broad categories, five of which are related to farming (including raising livestock, raising poultry, farming one's own land, sharecropping, and other types of farming). About 90% of women with earning occupations were engaged in farming, predominantly in raising poultry, livestock, or dairy farming, while the remaining 10% worked in salaried jobs, non-agricultural day labor, or in trade and production. On the other hand, around 40% of men were involved in farming, but primarily as sharecroppers or cultivating their own land.²⁶ The division of labor is particularly evident in farming, where men do not engage in livestock, poultry, or dairy farming, and women have minimal participation in sharecropping or cultivation of own land. Furthermore, nearly 15% of men worked as agricultural day laborers, a category where no women were represented, and around 20% of men were self-employed, with another 20% employed in trade, production, or salaried positions.

²⁵<https://data.worldbank.org/indicator/SI.POV.DDAY?locations=BD>

²⁶Other farming includes less common types such as fish farming or homestead farming, where small-scale farming is done for personal use or local trade.

1.5 Results

1.5.1 Effect of floods on intra-household resource allocation

I estimate the model outlined in Section 1.3.4 and find that exposure to a flood led to a redistribution of household resources. In Columns 1 and 2 of Table 2, I show the effects of exposure to a flood on the resource shares of men and women, as shown in equation (1.11). The interaction between the binary treatment and post-treatment time dummies reveals how resource shares shift when households are hit by a flood. Six months following the flood, women's resource share decreases by 8.5 percentage points (hereafter pp), on average, compared to women in unaffected households (Column 1). This effect is statistically significant at the 10% level. Four years later, the decrease in women's resource share reduces marginally to 7 pp. In contrast, men's resource shares increase by 11.2 pp in the short run and by 11.5 pp four years later, with both estimates remaining statistically significant (Column 2). The model estimates parameters directly for men and women only, and estimates for children can be backed out given that resource shares for all household members add up to one in a time period.

Notice that in the absence of a flood, men's resource shares decrease over time, as indicated by the negative, large, and significant time fixed effects for men. The implications of this will be explored further in Table 1.3, but it suggests that men gradually cede around 13 pp of their resource shares, primarily reallocating them to children as they grow older. Table B.2 in the Appendix reports estimates for all resource share parameters in equation (1.11), including the constant and changes in resource shares with the other observable household characteristics.²⁷

²⁷Figure B.10 validates the structural model by illustrating that the estimated resources shares are independent of total household expenditures, among this sample of low-income households.

Table 1.2: Effect of floods on resource shares of men and women

	(1) Women	(2) Men	(3) Women	(4) Men
1(Midline)	0.001 (0.043)	-0.134*** (0.040)	0.046 (0.048)	-0.188*** (0.049)
1(Endline)	0.036 (0.045)	-0.134*** (0.048)	0.011 (0.057)	-0.164** (0.064)
1(Flood)×1(Midline)	-0.085* (0.046)	0.112** (0.049)		
1(Flood)×1(Endline)	-0.069 (0.053)	0.115* (0.066)		
#Flood ₂₀₁₅ ×1(Midline)			-0.069*** (0.012)	0.088*** (0.019)
#Flood ₂₀₁₈ ×1(Endline)			-0.015*** (0.003)	0.025*** (0.005)
Observations	1173	1173	1173	1173
Total parameters	102	102	102	102

Notes: Estimates from the system of Engel curves from Section 1.3.4, controlling nonlinearly for 17 covariates: a constant, time fixed effects, interactions of the flood variable with post-treatment time dummies, the demeaned number of men, women, and children at baseline, six dummies for changes in household composition since baseline, and binary variables indicating whether the household had a child under five at baseline, whether at least one woman had completed primary education at baseline, and whether at least one woman was over the age of 40. The sample size is 391 households across 59 villages. A total of 102 parameters are estimated. The constant, taste parameters and changes in resource shares with other household observables in eq (1.11) are omitted from this table for brevity. Standard errors are clustered at the village level and are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Conclusions based on the number of floods further support the main findings from Columns 1 and 2, indicating a similar diversion of household resources from women to men. This specification is the same as in equation (1.11), replacing the binary flood variable with the number of floods experienced by these survey periods. As the timing of the first flood is as-good-as-random for this sample of villages, so would be the number of floods accumulated up until a set point in time. Columns 3 and 4 in Table 1.2 show that for each additional flood experienced before the midline survey in 2015, women lose 6.5 pp of their household resources while men gain 8.8 pp, both of which are statistically significant at 1% level. Four years after the first flood, women's resource




Figure 1.5: Average resource shares of men, women and children, as predicted by the model, at the time of the endline survey in households that experienced the number of floods indicated on the x-axis. Estimates at zero correspond to resource shares in untreated households at the time of the endline survey. Shaded areas represent 95% confidence intervals. Sample size: 391 households.

share decreases by 1.5 pp and men's increases by 2.5 pp, which remain statistically significant at 1%. Therefore, if a village in the 2014 cohort experienced just one flood during the study period, it would have a relatively larger effect in the short-run than in the longer run, on average. If there are additional floods, then the effect on resource shares would be additive.²⁸ These shifts in resource shares are not driven by the most flood-exposed villages, as results remain robust even after excluding villages with the highest flood exposure (see Table B.3).

As flood exposure accumulates, the redistribution of household resources becomes increasingly skewed in favor of men. Figure 1.5 plots the average resource shares of men, women, and children

²⁸Table B.4 in the appendix shows that if the treatment effects in Columns 3 and 4 Table 1.2 are multiplied by the average number of floods experienced by midline and endline surveys (2.1 and 5.6, respectively), the resulting resource shares for flooded villages are similar to those reported in Columns 4 and 5 of Table 1.3.

Table 1.3: Resource shares in reference households and the effect of floods

	Reference household			Flooded household			(7) Δ Children's share
	(1) Women	(2) Men	(3) Children	(4) Women	(5) Men	(6) Children	
Baseline	0.272*** (0.037)	0.438*** (0.052)	0.289*** (0.042)	0.272*** (0.037)	0.438*** (0.052)	0.289*** (0.042)	
Short-run	0.273*** (0.041)	0.305*** (0.034)	0.422*** (0.046)	0.189*** (0.038)	0.416*** (0.058)	0.394*** (0.050)	-0.027 (0.040)
Long-run	0.308*** (0.050)	0.305*** (0.035)	0.387*** (0.041)	0.239*** (0.035)	0.420*** (0.068)	0.341*** (0.050)	-0.046 (0.037)

Notes: Estimates for resource shares of men, women, and children in reference households and flooded households. A reference household is located in a control village, with one man, one woman and two children, and all binary household covariates in Table B.2 set to zero. Columns 4-6 report the resource shares in a reference household in a treated village. Column 7 reports the effect of exposure to a flood on children's resource shares. Standard errors are clustered at the village level and are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

at the time of the endline survey (as predicted by the model) against the cumulative number of floods experienced by the time of the last household survey. Each point represents the average resource share allocated to a person type in villages that experienced the cumulative number of floods on the x-axis. Estimates at zero correspond to resource shares in unflooded households at the time of the endline survey. It shows that women's resource shares start to decline consistently after the first flood, but children's resource shares start to drop after the sixth flood. This drop in resource shares of women and children is absorbed by men of the household. In households that experienced 11 floods by the endline survey, men have a resource share of 63.4%, which is 22.7 pp higher than that of men in households with only one flood.

In discussing the remaining results, I will focus on Columns 1 and 2 in Table 1.2 (and Table B.2), as these offer a more concise view of the overall effects of flood exposure. It follows that the effects would be more pronounced in villages that experienced a greater number of floods.

I show the predicted resource shares of men, women, children in reference households over

time in Columns 1 to 3 in Table 1.3. A reference household is a household in an unflooded village where the numbers of men, women and children are set at the sample mean and binary household covariates, mentioned in equation (1.11), are set to zero. Therefore, in terms of household composition, the reference household is nuclear with two children, and its composition remains unchanged over time. In each time period, the resource shares for women (men) in reference households can be derived using estimates of the constant term and the corresponding parameter values for time effects for women (men). At baseline, the resource shares for women, men, and children in a reference household are 27.2%, 43.8% and 29%, respectively. The gap between men and women's resource shares at baseline is 16.6 pp. Over time, the disparity between men and women in these reference households narrows. This is partly due to a marginal increase in women's resource shares in unflooded villages from baseline values, but the primary reason is that men transfer around 13 pp of their resource shares to children as they age.

I then bring in exogenous variation from floods and show the predicted resource shares if a reference household experiences this climate shock in Columns 4 to 5. Resource shares in flooded villages can be derived using predicted resources shares from Columns 1 to 3 and incorporating the average effect of a flood on men and women's resource shares from Columns 1 and 2 in Table 1.2. Six months after the flood, women in flooded households have a resource share of 18.9% which is 30.8% less than that of women in a reference household. Four years after the flood, women experience a drop of 22.4% in their resource shares. Men's resource shares improve by about 37% in both post-flood periods.

The resource shares for children can be backed-out as all resource shares must sum to one in each time period. The effect of exposure to the flood on children's shares is reported in Column 7. In the short-run, children in flooded households lose 2.7 pp of their resource shares. This

effect increases to 4.6 pp four years after the first flood, although these effects are statistically insignificant.

Taken together, the results show that women and children experience a decrease in the share of household resources allocated to them after a flood, but these effects are stronger for women. It is important to note that resource shares are assumed to be evenly distributed among individuals within each person category. On average, households have two children, so the reduction in children's total resource shares is distributed across two individuals. In contrast, the decline in women's resource shares is borne by one individual, as the average household has one woman. This drop in resource shares of women and children is absorbed by men of the household, who experience a sustained and statistically significant increase in their resource shares after the first flood.

Importantly, Table 1.3 suggests that even if there are no changes in household composition since baseline, there is still evidence for a reallocation of household resources away from women and children towards men of the household after a flood. This helps alleviate concerns about the potential impact of changes in household composition on the results. A more detailed discussion of the mechanisms behind these shifts will be provided in Section 1.6.

The stark reallocation of household resources away from women toward men is evident in Figure 1.6, which plots the distributions of the predicted resource shares of men, women, and children in all flooded and unflooded households at each time period. The distributions at baseline are quite similar for flooded and unflooded households – any minor differences observed at baseline are attributed to differences in observable household characteristics used to model heterogeneity in equation (1.11). After exposure to the flood, the distribution of women's resource shares in flooded households is shifted to the left compared to those in unflooded




Figure 1.6: Distribution of resource shares of men, women, and children in all households in the 2014 sample, obtained as predictions from equation (1.11). Blue and gray histograms correspond to flooded and unflooded households, respectively. Sample size: 391 households.

households, indicating a decrease in resource shares for women in treated households. The distribution for children's resource shares follows a similar pattern, but the shift is less prominent. In contrast, the distribution of men's resource shares in flooded households is shifted to the right compared to those in unflooded households, suggesting an increase in resource shares for men in flood-affected households. Appendix Table B.5 reports descriptive statistics for the predicted resource shares in flooded and unflooded households.

A closer look at Figure 1.6 shows that women in flooded households appear to be providing insurance for men's consumption by compensating them for their transfers to children. Notice that in flooded households, men's resource shares remain stable between 44-46% in all years, whereas in unflooded households men yield about 13 pp of their resources to their children as they grow older, which can be attributed to their growing needs. This is also reflected in the time fixed effects in Table 1.2. On the other hand, women's resource shares in unflooded villages remain consistently between 32-35% across years, but in flooded households, women experience a drop in their resources of around 10 pp compared to their average baseline values. It is these transfers from women that allow men to maintain their resource shares in flooded households while women end up with a lower share.

Note that food expenditure shares remain unchanged across flooded and unflooded households, and the underlying effects of the flood on intra-household resource distribution become apparent only through changes in the slopes of Engel curves for person-level food expenditure. Appendix Figure B.9 shows that the distributions of observed food expenditure shares for men, women, and children remain consistent in flooded and unflooded households in all time periods. This suggests that naively using only observational data on food expenditures would likely show no discernible effect from the floods, potentially leading to the mistaken conclusion that the flood had no impact

on intra-household inequalities.

Additionally, a higher share of food expenditure does not necessarily correspond to a higher resource share, as these are distinct entities and one cannot proxy for the other. Appendix Figure B.9 and Table 1.1 show that children consistently hold the largest food expenditure share, around 20% in all periods, but Figure 1.6 and Table 1.3 show that men have the largest resource share, exceeding 40%. This emphasizes the need to use the structural model to uncover the effects of the flood shock on the distribution of the household budget, rather than relying solely on limited observational data for a private good.

1.5.2 Shadow budget and individual-level poverty

Using the predicted resource shares from the structural model, I calculate the shadow budget by multiplying the predicted shares with the observed total household expenditure at time t —i.e., $\hat{\eta}_{jht} \times y_{ht}$. The shadow budget represents the unobserved individual-level budget allocated to men, women, or children, which they use to maximize their utility. Comparing these values to benchmarks such as the international extreme poverty line of \$2.15 per day (in 2017 PPP\$) allows for assessing how close individuals are to reaching the minimum threshold required for a decent standard of living.

The redistribution of resources away from women and children in the household results in an increase in the material well-being of men in flooded households. Table 1.4 presents the descriptive statistics for the shadow budget. Panel A shows the total shadow budget allocated to each person category. Six months after the flood, men in flooded households had a mean daily shadow budget of \$7.46, which was about \$3.55 higher than the average for women and \$1.92

Table 1.4: Shadow budget of men, women and children in 2017 PPP\$

	Unflooded			Flooded		
	Women	Men	Children	Women	Men	Children
<i>Panel A: By person type</i>						
Baseline	4.71 (3.66)	6.59 (4.79)	3.16 (2.52)	4.95 (3.26)	6.94 (4.92)	3.36 (2.80)
Short-run	6.29 (4.41)	6.77 (5.54)	6.81 (4.43)	3.91 (2.45)	7.46 (4.35)	5.54 (3.26)
Long-run	7.39 (4.80)	7.35 (5.67)	6.31 (3.88)	5.20 (3.02)	8.49 (5.68)	4.66 (3.04)
<i>Panel B: By individual</i>						
Baseline	3.91 (2.91)	5.75 (4.27)	1.58 (1.20)	4.25 (2.84)	6.07 (4.20)	1.69 (1.10)
Short-run	5.28 (3.04)	5.89 (4.21)	3.33 (2.09)	3.33 (1.53)	6.59 (3.38)	2.72 (1.61)
Long-run	6.31 (3.45)	6.47 (4.13)	2.99 (1.97)	4.70 (2.78)	7.88 (5.06)	2.27 (1.43)
Observations	798	798	798	375	375	375

Notes: Descriptive statistics of the shadow budget of men, women, and children in all households in the selected sample. Panel A presents the shadow budgets by person type, whereas panel B presents the shadow budget by individual in each person type category. Means are reported in 2017 PPP\$ terms. Standard deviations are in parentheses. Sample size: 391 households.

higher than the average allocated to children in these households. It was also approximately \$0.69 higher than the average budget for men in unflooded households. By the time of the final household survey, four years after the first flood, the shadow budget allocated to men in flooded households was \$1.14 higher than the average for men in unflooded households, while the budgets for women and children were \$2.19 and \$1.65 lower, respectively.

To compare these shadow budgets to the \$2.15-a-day international extreme poverty line using PPP exchange rates, I compute the shadow budget per individual within each category. The descriptive statistics are shown in Panel B of Table 1.4. The median household typically consists of one man, one woman, and two children. Within each group, the shadow budget is




Figure 1.7: Kdensity of the daily shadow budget per man, woman and child in each time period. All shadow budgets are in terms of 2017 PPP\$. Vertical dashed line is the current extreme poverty line of \$2.15 a day in 2017 PPP terms. Sample size: 391 households.

assumed to be evenly split among members if there are multiple individuals of the same type. When comparing flooded and unflooded villages, the average daily shadow budget per individual is lower for women and children in flooded areas, while the shadow budget per man is relatively higher. This indicates that women and children are particularly vulnerable to falling below the extreme poverty line.

This is illustrated further in Figure 1.7, which plots the distributions of daily shadow budgets per man, woman, and child compared to the international poverty line. At baseline, the distributions are similar across both treated and untreated households, with children being the most likely to fall below the extreme poverty line. This observation is consistent with previous estimates from Bangladesh and other lower-income countries, which indicate that children are the most vulnerable to extreme poverty (Lechene, Pendukar, and Wolf, 2022). Following a flood, within a given time period, women and children in flooded households are more likely to fall below the extreme poverty line compared to their counterparts in untreated households, while the distribution for men shifts slightly to the right.²⁹

1.5.3 Effect of the flood on the household's total budget

To contextualize the changes in individual-level budgets and to provide a clearer understanding of what is happening to the household as a whole, I explore the effects of exposure to a flood on

²⁹Figure B.11 in the appendix plots the distributions of the total shadow budget allocated to each person type. By the endline survey, the distributions for the shadow budget allocated to all women, men, and children in unflooded households align more closely. This alignment occurs primarily because, in the absence of a flood, resources are transferred from men to children, without a compensatory adjustment from women.

the household's total budget. I use the following specification, which mirrors equation (1.11):

$$Y_{ht} = \lambda_0 + \lambda_1 Midline + \lambda_2 Endline + \lambda_3 Flood_h \times Midline + \lambda_4 Flood_h \times Endline + \lambda_5 X_h + \lambda_6 W_{ht} + \epsilon_{ht} \quad (1.12)$$

Since exposure to a flood is as-good-as-random, λ_3 and λ_4 represent the effects of the flood on the outcome six months and four years after the shock. X_h and W_{ht} are household observable characteristics that are not required for identification of causal effects, but are included to improve precision. As villages experienced up to 11 floods between their first flood in 2014 until the endline survey in 2018, I additionally study effects based on the number of floods.³⁰ As the timing of the first flood is arguably random, so is the number of floods experienced by any given time period. For all outcome variables, I report results both with and without household-level controls (X_h and W_{ht}), with the former being the preferred specification. As a robustness check, I also reproduce results for the more restrictive set of villages in Panel B in Figure 1.4. Standard errors are clustered at the village level since that is the level of treatment assignment.

I find that exposure to a flood had a negative and lasting effect on the total household budget. Column 2 in Table 1.5 presents results from equation (1.12). Six months after the flood, the total household budget in flood-affected households decreased by 0.23 standard deviations, on average, relative to households not exposed to a flood, as shown in Column 2 in Panel A. This represents a decrease of about 14% relative to the average household budget in the midline survey.³¹ This negative effect endures over time, with affected households having a total budget

³⁰To do this, I modify equation (1.12), replacing the binary flood indicator with the number of floods experienced by the midline and endline surveys.

³¹Table 1.1 shows the mean total household expenditure is 181,690 Taka at midline with standard deviation of 112,740 Taka.

Table 1.5: Effect of the flood on household's financial indicators

	Total expenditure			Poverty			Cons. assets			Ag. assets			Other prod. assets		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Panel A															
Flood $\times$ <i>Midline</i>	-0.229* (0.111)	-0.232*** (0.086)	-0.243*** (0.090)	0.090* (0.047)	0.119** (0.047)	0.115** (0.050)	-0.102 (0.124)	-0.084 (0.073)	-0.119 (0.071)	-0.063 (0.104)	-0.087 (0.091)	-0.115 (0.092)	-0.171 (0.108)	-0.166 (0.105)	-0.175 (0.110)
Flood $\times$ <i>Endline</i>	-0.259* (0.147)	-0.264** (0.116)	-0.259** (0.122)	0.021 (0.037)	0.050 (0.032)	0.049 (0.034)	-0.153 (0.166)	-0.114 (0.123)	-0.107 (0.128)	-0.047 (0.115)	-0.126 (0.103)	-0.150 (0.104)	-0.180 (0.177)	-0.171 (0.168)	-0.233 (0.173)
Panel B															
# Flood ₂₀₁₅ $\times$ <i>Midline</i>	-0.089* (0.047)	-0.090** (0.037)	-0.093** (0.039)	0.039* (0.022)	0.048** (0.022)	0.046** (0.023)	-0.035 (0.051)	-0.044* (0.025)	-0.057** (0.024)	-0.021 (0.050)	-0.019 (0.037)	-0.030 (0.039)	-0.085 (0.051)	-0.073 (0.053)	-0.077 (0.055)
# Flood ₂₀₁₈ $\times$ <i>Endline</i>	-0.035** (0.017)	-0.034** (0.014)	-0.033** (0.015)	0.006 (0.005)	0.008** (0.004)	0.008* (0.004)	-0.006 (0.022)	-0.017 (0.013)	-0.016 (0.013)	-0.006 (0.014)	-0.021* (0.012)	-0.024** (0.012)	-0.021 (0.026)	-0.016 (0.026)	-0.024 (0.028)
Observations	1173	1173	1125	1173	1173	1125	1162	1160	1112	1162	1156	1108	1162	1162	1114
Controls		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Robustness sample			Yes			Yes			Yes			Yes			Yes

Notes: Estimates from eq (1.12) for households in the 2014 sample only. Except for “Poverty,” all outcome variables listed in the headers are standardized. “Poverty” is an indicator variable denoting whether the household budget *per capita* is less than \$2.15 per day in PPP terms. Panel A presents estimates using a binary treatment variable, while Panel B shows results based on the number of floods experienced by the midline or endline surveys. Controls include the number of men, women, and children at baseline, the baseline outcome variable, a binary variable indicating whether a woman had completed at least primary education at baseline, and whether she was over the age of 40. Robustness sample refers to the villages used to produce Panel B in Figure 1.4. Standard errors are clustered at the village level and are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

that is 0.26 standard deviations below that of unflooded households even four years post-flood. These findings, when combined with those in Section 1.5.2, show that men in flooded households experienced an increase in their individual-level budget, even though the total household budget declined.

The observed decrease in the total household budget is reflected in an increased incidence of extreme poverty and a decrease in the value of consumption assets owned by the household. Using the World Bank's current extreme poverty line of \$2.15 per day per person (in 2017 PPP terms) as a benchmark, Column 5 shows that experiencing a flood increases the likelihood of a household falling below the poverty line by 11.9 pp in the short run, a result that is statistically significant at the 5% level. In the long run, the likelihood increases by 5 pp, but is no longer statistically significant.³² The value of consumption assets decreases by 0.08 standard deviations in the months following a flood and by 0.114 standard deviations in the longer-run (Column 8 in Panel A). Panel B shows that experiencing one additional flood by the midline survey reduces the value of consumption assets by 0.04 standard deviations in the short run, which is statistically significant at the 10% level, and by about 0.02 standard deviations in the long run.

The persistent reduction in the overall household budget may be linked to a decline in productive assets, though the evidence is suggestive. Column 11 shows that the value of agricultural assets decreases by approximately 0.09 standard deviations shortly after the flood and by 0.126 standard deviations four years later. Non-agricultural productive assets decrease by around 0.17 standard deviations in both periods. The direction of these effects is corroborated in Panel B, which examines the impact of experiencing an additional flood by the midline and endline surveys

³²If a household has 3 members, for instance, then the total household budget needs to be at least \$6.45 ($\2.15×3) per day to not be categorized as extremely poor.

on household budget and assets.

1.6 Mechanisms

I find that women's employment is disproportionately affected after the flood, while male employment remains relatively stable. Table 1.6 focuses on changes in female and male employment and earnings within the household. Using equation (1.12), Column 2 shows that flood exposure decreases the likelihood of women being employed by 8.8 pp in the short run, significant at the 10% level. This effect reduces to 1.2 pp in the longer-run. In contrast, Column 5 shows that men's employment is not significantly affected: they are 1.5 pp less likely to be employed immediately after the flood and 1.6 pp more likely in the long run, but both effects are small and statistically insignificant.

This disparity arises because women, who are less likely to work in farming or salaried positions after a flood, do not find employment in other sectors to compensate for their loss. Figure 1.8 shows the effect of a flood on women and men holding the main occupations listed on the y-axis. As described in Section 1.4.6, there is a clear division of labor at baseline: women primarily work in poultry and livestock farming, while men are more concentrated in sharecropping, farming their own land, and a broader range of non-farming occupations. Figure 1.8 illustrates that men are 5.2 pp less likely to work in sharecropping and own farming and 4.6 pp less likely to work in self-employed activities after a flood. However, they are able to compensate by increasing their wage labor, particularly as agricultural day laborers (6.04 pp more likely, significant at 1% level). In contrast, women face a 11.5 pp reduction in raising poultry and a 2 pp drop in salaried work, both significant at the 10% level, yet they do not move

Table 1.6: Effect of the flood on employment and earnings

	Woman employed			Man employed			Female earnings			Male earnings			Female-to-male earnings		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Panel A															
Flood $\times$ <i>Midline</i>	-0.079 (0.055)	-0.088* (0.052)	-0.064 (0.051)	-0.021 (0.023)	-0.015 (0.020)	-0.017 (0.021)	-0.168 (0.113)	-0.198* (0.103)	-0.172 (0.106)	-0.020 (0.129)	-0.012 (0.127)	-0.006 (0.135)	-4.204*** (1.514)	-4.455** (1.722)	-4.302** (1.809)
Flood $\times$ <i>Endline</i>	-0.002 (0.034)	-0.012 (0.037)	0.010 (0.033)	0.010 (0.017)	0.016 (0.016)	0.012 (0.017)	-0.099 (0.141)	-0.138 (0.133)	-0.109 (0.140)	-0.035 (0.143)	-0.020 (0.140)	0.003 (0.145)	-2.039 (2.416)	-2.092 (3.167)	-2.102 (3.319)
Panel B															
# Floods ₂₀₁₅ $\times$ <i>Midline</i>	-0.041 (0.028)	-0.042 (0.027)	-0.033 (0.026)	-0.005 (0.009)	-0.002 (0.008)	-0.002 (0.008)	-0.103*** (0.039)	-0.112*** (0.039)	-0.103** (0.040)	-0.042 (0.040)	-0.037 (0.041)	-0.036 (0.043)	-1.882*** (0.623)	-2.048** (0.818)	-1.985** (0.848)
# Floods ₂₀₁₈ $\times$ <i>Endline</i>	-0.001 (0.004)	-0.002 (0.005)	0.001 (0.004)	0.003** (0.001)	0.004** (0.002)	0.003* (0.002)	-0.026 (0.019)	-0.031 (0.019)	-0.028 (0.020)	-0.005 (0.016)	-0.004 (0.017)	-0.000 (0.017)	-0.344 (0.358)	-0.366 (0.440)	-0.369 (0.456)
Observations	1172	1170	1122	1173	1173	1125	1161	1149	1101	1162	1160	1113	1106	1076	1035
Controls		Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Robustness sample			Yes			Yes			Yes			Yes			Yes

Notes: Estimates for households in flooded and unflooded villages in the 2014 sample only, following eq (1.12). Outcome variables mentioned on the headers are: a binary variable for whether any woman in the household was engaged in any economic activity, a binary variable for whether any man in the household was engaged in any economic activity, standardized monthly income of women in the household, standardized monthly income of men in the household, and female earnings as a percentage of male earnings. Panel A presents estimates using a binary treatment variable, while Panel B shows results based on the number of floods experienced by the midline or endline surveys. Controls include the number of men, women, and children at baseline, the baseline outcome variable, a binary variable indicating whether a woman had completed at least primary education at baseline, and whether she was over the age of 40. Reduced sample refers to a subset of the main sample used to produce Panel B in Figure 1.4. Standard errors are clustered at the village level and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

to roles like agricultural day labor, which is could be more labor-intensive and may be considered inappropriate for women in this context depending on the nature of the work. Recall no women had agricultural day labor as their main occupation at baseline. This lack of substitution provides some context behind the observed shock to women's earnings in the short run, as they do not replace the loss of farming and salaried work with another occupation. This aligns with findings in Afridi, Mahajan, and Sangwan (2022), which emphasize how cultural barriers limit women's employment flexibility after a climate shock.

The reduction in women's employment after the flood also has a direct impact on their earnings, which has been shown to influence women's intra-household bargaining power (see Doss 2013). Column 8 indicates that women's monthly earnings in flooded villages decreased by 0.2 standard deviations in the months following the first flood, compared to women in unflooded villages. This effect is statistically significant at the 10% level. Four years after the first flood, the reduction in women's monthly earnings decreases to 0.138 standard deviations. Although this effect is statistically insignificant in the long run, it represents a decrease of approximately 30% of the mean monthly earnings of women at the time of the endline survey (as reported in Table 1.1).³³ In contrast, male monthly earnings decrease by a smaller magnitude, with reductions of 0.012 standard deviations in the short run and 0.02 standard deviations in the long run. These reductions are both statistically and economically insignificant.

As women lose out on earnings more than men, there is a change in the monthly earnings of women relative to men in the household. This is the most likely explanation for what is driving the change in resource shares, and this has been used as a proxy for bargaining power

³³Given the relatively high standard deviation of women's monthly earnings, a change of 0.13 standard deviations corresponds to a reduction of 142 Taka, compared to an average monthly income of 500.69 Taka at the endline survey.

in existing research (e.g., Hoddinott and Haddad 1995). Column 14 shows that exposure to a flood leads to a decrease of 4.2 pp in female earnings as a percentage of male earnings, on average. In this context, the average ratio of female to male earnings is 10.56% at baseline. A 4.2 pp decrease represents a 40% reduction in the average ratio prior to the flood, highlighting a significant negative impact on gender earnings parity. In the longer-run, the magnitude of this effect reduces to 2 pp, which is close to a 20% reduction in the average ratio at baseline.


Alternative mechanisms, such as transportation and housing expenses, are unlikely to explain the reallocation of household resources. Column 8 in Appendix Table B.6 shows that men are approximately 7 pp less likely to work outside the village six months after the flood and 4.1 pp less likely in the long run. This reduction coincides with a decrease in transportation expenses (Column 11). Additionally, housing-related expenses (such as rent and maintenance costs) decrease by 0.157 standard deviations in the short run and 0.246 standard deviations in the long run (Column 14), further indicating that these factors do not drive the changes in resource allocation.

I also consider whether changes in household composition, such as male migration or fertility, drive the observed effects on intra-household resource allocation, but these factors appear unlikely.³⁴ Column 2 in Appendix Table B.6 shows no significant increase in male migration in flooded households.³⁵ Similarly, Column 5 indicates no noticeable changes in fertility, as the

³⁴Recall that the structural model already controls for changes in household composition since the baseline. Results considering a reference household as shown in Table 1.3 indicate that even in the absence of changes in household composition, a decrease in female resource shares is still observed. Consequently, channels such as fertility and male migration are examined here largely because they are commonly considered in the literature, although no significant findings are expected from these channels.

³⁵This is contrary to findings in Gianelli and Canessa (2022), who show that floods in Bangladesh led to significant increases in male migration. One possible explanation for this difference is that the selected sample in this paper is restricted to households in flood-prone areas that have at least one man, one woman, and one child, whereas Gianelli and Canessa (2022) studied a broader rural population in Bangladesh, which may have included more diverse household structures and migration patterns.

Women



Men




Figure 1.8: Effect of exposure to a flood on a binary variable for whether a household has a man or woman whose main occupation is listed on the y-axis, following equation (1.12). The specification controls for household composition, the outcome variable, and woman's educational attainment (completion of at least primary education)—all measured at baseline—as well as whether she was over the age of 40. Bars represent 90% confidence intervals, with standard errors clustered at the village level. Sample size: 391 unique households.

age of the youngest child in flooded households is marginally higher, suggesting no increase in fertility rates post-flood. Thus, neither migration nor fertility seem to account for the observed redistribution of resources.³⁶

Lastly, there may be concerns that the flood could affect agricultural markets by altering food prices, but I show that they remain unchanged. Figure B.12 compares food prices for over 300 commodities between flooded and unflooded villages across all three survey rounds. It can be seen that price differences are economically and statistically insignificant across food groups and time periods. For prepared foods—i.e., food purchased outside the home—there is more variation in prices between flood affected and unaffected villages, likely reflecting differences in the types of prepared foods available. However, these items constitute only 0.58%, 0.69%, and 1.07% of food costs for men, women, and children, respectively, making it unlikely that they significantly affect the estimates. In Section 1.7, I re-estimate the model excluding prepared foods from the dependent variable (food expenditure shares), and the results remain consistent.

1.7 Robustness checks

This section presents results from a series of robustness checks that examine how the estimates change when additional covariates are included in the model, when the model is estimated using data from a completely different set of villages, or when stricter criteria are used for categorizing flood-affected and unaffected villages.

³⁶Regarding concerns that the flood may have led to fatalities, there were 34 households in the sample where a death occurred during the study period. Of these, 5 households experienced the death of a child, while most of the adult deaths involved men and women over the age of 60.

1.7.1 Adjusting the model

To keep the model parsimonious, I did not include household covariates that were not essential. However, I check the robustness of the effects of the flood by adding optional covariates. In Figure B.13, Model 1 reflects the preferred estimates from Columns 1 and 2 from Table 1.2. In Models 2-4, I introduced one of the following binary variables: whether the household had only sons, whether the household head was female, and whether the village experienced more than the mean number of floods before 2011.³⁷ Model 5 incorporates region fixed effects. Model 6 re-estimates the preferred model by adjusting food expenditure shares to exclude food prepared outside the home. Lastly, Model 7 uses the preferred specification but using the robustness sample in Panel B of Figure 1.4.

Short-run results remain consistent across all specifications, as shown in Figure B.13. In some cases, standard errors increase, such as when a binary variable indicating whether all children are male is introduced. This variable does not significantly influence the determination of resource shares in this context, adding noise to the model. This finding aligns with Brown, Calvi, and Penglase (2021), who note that son preference is less prominent in the rural Bangladeshi context. Additionally, including region fixed effects does not alter the estimates, suggesting that the observed effect is not influenced by regional factors. Even when prepared foods are excluded or when the sample of villages itself is reduced, the short-run estimates remain unchanged.

In the long run, effects are robust for men and remain between 10 to 13 pp. Effects for women also remain negative and stable across Models 1 through 6, ranging from 5 to 7 pp. In Model 7, however, while the effects for women remain negative, the magnitude decreases to

³⁷From 2000, the first year for which the Global Flood Database provided data, until the first household survey, villages in the sample experienced an average of 8 floods.

3.3 percentage points. Section 1.7.4 will discuss in more detail the robustness of the effects on women's resource shares in the long run and the factors that contribute to its reduction in some cases.

1.7.2 Effect of a flood on resource shares in out-of-sample villages

To determine whether the changes in resource shares observed in the 2014 cohort are unique to that group or consistent across other flood-prone villages, I re-estimate the model for an out-of-sample group of villages that experienced their first flood in 2016 or 2018. This sample does not overlap with the 2014 sample and is smaller, comprising 225 households after combining the 2016 and 2018 cohorts.³⁸ To reduce the number of parameters to be estimated, the model utilizes only the midline and endline surveys, with the midline survey serving as the base, as it is the most recent survey prior to households in the 2016 and 2018 cohorts being exposed to a flood.

Estimates for this new sample of villages are presented in Table B.7 and show a continued redistribution of household resources, primarily away from women and towards men. Exposure to a flood reduces women's resource shares by 6.6 pp, while men's resource shares increase by 8.1 pp. Note that the magnitude of this effect is very similar to the short-run effects observed for one additional flood in Columns 3 and 4 in Table 1.2.³⁹ These results remain consistent even when the sample is restricted to villages used in Panel B of Figure 1.4.

³⁸There are 18 flooded villages in this sample, out of which 8 are from the 2018 cohort and 10 from 2016.

³⁹Flooded households in the 2016/2018 sample experienced an average of 1.2 floods by the time of the endline survey.

1.7.3 Stricter criteria for unflooded villages

I examine how the effects change when the threshold for the allowable difference in ex-ante probability of flooding between flooded and unflooded villages is reduced. The main sample used for the estimates of the 2014 cohort permits a maximum difference of 0.3 in ex-ante flood risk. I investigate whether the effects of a flood are robust to tightening this threshold, gradually reducing it from 0.3 to 0.1, and excluding unflooded villages that no longer meet the revised criteria. If a flood-affected village does not have a nearby unaffected village that satisfies the stricter probability cutoff, then that flooded village is also excluded from the sample.

The results, presented in Figure B.14, show that both the short-run and long-run estimates remain stable. At the most stringent cutoff of 0.1, 9 villages are removed from the main sample (1 flooded and 8 unflooded).

1.7.4 Stricter criteria for flooded villages

Lastly, I examine how the estimates change when the criterion for defining flooded villages is made more stringent. In the main sample used to produce the main results for the 2014 cohort, a village is categorized as “flooded” if at least 15% of the area around the village is inundated during a flood event. I gradually increase this threshold and remove flooded villages that no longer meet the revised criteria. For example, if the new threshold is set at 20%, villages that initially had between 15% and 20% of their area flooded will be excluded from the sample. I re-estimate the model if the threshold for the first flood experienced by a village in the 2014 cohort increases from 15% to 33%, with results presented in Figure 1.9.⁴⁰ Panel A reports the

⁴⁰After 33%, the model ceases to converge at the starting values used to produce the main set of estimates in Columns 1 and 2 in Table 1.2.

robustness of short-run effects based on the new threshold used to define flooded villages, while Panel B presents the long-run estimates.

The short-run effects of flood exposure on women's resource shares remain stable, but as the flooding threshold rises, the positive effect on men's resource shares increases. For instance, when the threshold is set at 25%, men's resource shares increase by 20 pp compared to 11 pp when the threshold was 15% in the main results. This suggests that the short-run effect on women's resource shares is stable regardless of the intensity of the initial flood, but men gain more resources as the flood intensity increases. The difference is compensated by children, who give up more resources as the severity of the flood increases, relative to their counterparts in unflooded villages. In the long run, the effect on men's resource shares is stable at around 11 pp as flood intensity increases, but the impact on women is less consistent. At lower levels of flood intensity, women experience a decrease in their resource shares with exposure to a flood, in line with the main results in Columns 1 and 2 in Table 1.2. But as the intensity of the initial flood increases, women no longer experience a negative effect on their resource shares.

The stability of the flood's impact on women's resources in the short run, contrasted with the variation in the long run, is closely linked to changes in women's earnings relative to men as the flood intensity increases. Panel A in Figure 1.10 highlights this relationship, showing that, six months after the flood, as more of the village area is flooded, women's monthly earnings as a percent of men's earnings consistently drop by around 4 pp relative to men (as in Table 1.6). This consistent drop suggests that regardless of flood intensity, women's relative earnings are negatively impacted. However, in the longer run, as the intensity of the flood increases, the relative earnings gap between women and men decreases. Once this gap diminishes, the negative effect of the flood on women's resource shares also fades.




Figure 1.9: Effect of a flood on men's and women's resource shares. Results vary with the threshold used to categorize a flooded village. As the threshold for the minimum village area flooded increases, the number of treated villages decrease. There are 391 households in 59 villages when the threshold is set at 15%. This drops to 338 households in 51 villages at 33%. Beyond 33% the model ceases to converge. Standard errors are clustered at the village level. Shaded area represents 90% confidence intervals.




Figure 1.10: Effect of the flood on monthly earnings of women and children relative to men, following specification in equation (1.12). The outcome variable in Panel A is women's monthly earnings as a percent of men's monthly earnings in the household. In Panel B, its children's monthly earnings as a percent of men's monthly earnings in the household. Controls to increase precision include the number of men, women, and children at baseline, the baseline outcome variable, a binary variable indicating whether a woman had completed at least primary education at baseline, and whether she was over the age of 40. Bars are 90% confidence intervals.

The flood's decreasing effect on the women-to-men earnings ratio is partially attributable to the fact that as the intensity of the initial flood increases, men's monthly earnings also begin to decline more. At lower levels of flood intensity, exposure to a flood does not have a substantial impact on men's earnings, as shown in Table 1.6. However, with greater flood intensity, men's earnings start to decrease, as shown in Figure B.15 in the Appendix.

The increasing impact of floods on the resource shares of children is similarly reflected in their relative earnings compared to men in the household, as shown in Panel B of Figure 1.10. Recall that children can be up to age 15 at the time of the baseline survey, age 18 by the midline, and up to the age of 21 by the last household survey. Therefore, it is possible that some children in the household are of working age, as the legal age of employment in Bangladesh is 14 years.⁴¹ In the short run, exposure to a flood decreases the relative earnings of children by 4.8 to 8.8 pp as flood intensity increases. In the long run, this effect remains at around a 10 pp decrease, regardless of the flood's intensity.⁴²

Overall, the patterns observed in the reallocation of household resources as the intensity of the first flood increases, as illustrated in Figure 1.9, closely mirror the relative earnings of women and children compared to men in the household. This reinforces the idea that changes in relative earnings due to flood exposure are changing the status of women and children in the household, ultimately triggering a redistribution of household resources.

⁴¹<https://www.npr.org/sections/goatsandsoda/2016/12/07/504681046/study-child-laborers-in-bangladesh-are-working-64-hours-a-week>

⁴²This reduction in children's relative earnings, particularly in the long run, is likely due to a decrease in their employment, as illustrated in Figure B.16 in the Appendix. As flood intensity increases, Panel A in Figure B.16 shows that the probability of a child being employed reduces by about 20 pp in the long run. Although changes are observed across multiple occupations, self-employment experiences the most significant decline (Panel B).

1.8 Conclusion

In 2019, 87% of South Asia’s 2 billion people were exposed to climate-induced shocks, making it the hardest-hit region globally (Doan et al. 2023). Within this vulnerable context, I estimate how a climate shock affects the distribution of household resources among men, women, and children living together, as well as their unobserved individual-level expenditures. In rural Bangladesh, where existing gender disparities and cultural norms are already skewed against women, I find that women bear the brunt of a climate shock.

I specify a structural model of collective households to estimate the shares of the total household budget allocated to men, women, and children living under the same roof and how they change in the years after a flood shock. To account for differences in flood anticipation, I restrict my sample to villages with similar ex-ante risk of flooding, where flood exposure can be considered as-good-as-random. I show there is a reallocation of household resources away from women—and, to a lesser extent, children—towards men following a flood shock. Six months post-flood, women’s share of the household budget drops by 8.5 pp (children by 2.7 pp) relative to counterparts in unflooded households, and four years later, the decline is 6.9 pp for women (4.6 pp for children), while men consistently experience an increase of about 11 pp. Additionally, I show that as the number of floods experienced increase, so does the gender divide within the household.

To connect these changes in resource shares to material well-being, I use the predicted resource shares and calculate each person’s unobserved individual-level budget within the household. Six months post-flood, the daily shadow budget for a woman in a flooded household averages \$3.33 (in 2017 PPP terms), compared to \$6.59 for a man—a daily gap of \$3.26. With the extreme

poverty line at \$2.15/day, these numbers reveal a greater vulnerability for women in flooded households. For context, in unflooded households, the difference between men's and women's shadow budgets in this period is just \$0.61. Four years later, this gap narrows only slightly to \$3.18 in flooded households, while it remains negligible in unaffected ones. In both post-treatment periods, men in flooded households have higher shadow budgets than men in unflooded households. This disparity underscores how climate shocks increase women's vulnerability, pushing them closer to poverty, while men in flooded households move further from it.

The widening gap in intra-household resource allocation is driven by the decline in women's relative earnings compared to men after a flood. Although both men and women suffer losses in farming-related occupations, men are able to transition into agricultural day labor to protect their earnings, but women are unable to substitute into other employment. As a result, women's employment falls by 8.8 pp in the short run, while men's employment remains largely stable. Women's earnings as a percentage of men's decrease by 4 pp, deepening the gap in earnings between them, which ultimately drives the shift to a more unequal division of the total household budget.

While this is the first study to use resource share estimation to examine the intra-household effects of a climate shock, a few other studies have used this methodology to evaluate the effects of government programs or local crime. To put these results in context, Calvi (2020) found that an inheritance policy for women in India improved women's resource share by 0.8-2.3 pp. In Mexico, Tommasi (2019) found that access to PROGRESA increased the resource share of mothers by 2.6 pp, and Hernandez-de-Benito (2024) shows that an average increase in local crime reduced women's resource shares by 5-8 pp. These studies do not examine the evolution of resource shares over time, but the redistribution effects of a flood are towards the higher end

of these estimates.

As Bangladesh is one of the 32 countries in the world where flood exposure is projected to “continuously increase” in the near future (Tellman et al. 2021), the costs and impact of these shocks on the material well-being of women and children are likely to multiply. To explore this further, I conduct a counterfactual exercise, calculating the poverty headcount ratio – the percentage of people whose daily shadow budget is less than the \$2.15 a day international extreme poverty line – if *all* villages in the 2014 sample experienced a flood in 2014. I predict resource shares assuming all households had been flooded and calculate their individual-level budgets. Based on these shadow budgets, I compute the total number of people living on less than \$2.15 per day, divided by the total sample population. Under the actual treatment status, the poverty rate was 22% at the midline survey in 2015 and 25% at the endline survey. In the counterfactual scenario, where all villages experienced a flood, the poverty rate increases to 26.5% in the short run and 30% in the long run. This increase is driven by more women and children falling into extreme poverty.

This paper makes a strong case to incorporate intra-household inequalities into climate adaptation policies and financing goals. Current policy approaches generally overlook this important dimension of inequality. In the case of Bangladesh, while their National Adaptation Plan and Climate Change and Gender Action Plan acknowledge that climate change exacerbates gender disparities, the primary focus is community-level inequalities (MOEF 2022; BCCT and UN Women Bangladesh, 2024). I show that, on average, women and children lose \$1.95 and \$0.61 per day, respectively, six months after a flood, and \$1.61 and \$0.72 four years later. For individuals living near or below the international extreme poverty line, these losses are substantial, severely undermining women’s and children’s ability to meet even their most basic needs and further

entrenching their vulnerability.

Given the significant impact of climate shocks on intra-household resource allocation and its growing importance, it is important to develop policies that address these inequalities. Since the decline in women's relative earnings drives these disparities, diversifying women's employment opportunities could help mitigate the income losses they face after a flood. Before the flood, women in my sample were predominantly engaged in poultry, livestock, and dairy farming, with limited involvement in other sectors, making them particularly vulnerable to income loss. I show that this loss in women's employment is largely driven by a reduction in poultry farming after the flood. Expanding employment in other sectors could provide greater resilience against future shocks. Additionally, there is a need to strengthen existing livelihoods to make them more secure in the face of climate shocks. For example, building flood-protected shelters to safeguard poultry, livestock, and animal fodder could help them maintain their income during extreme weather events (Mitchell, Tanner, and Luisser, 2007).

The policy response after these shocks needs to account for the changes in bargaining power within households, as demonstrated in this paper. Implementing cash transfer programs targeted specifically toward women could be promising, as they may strengthen women's income during climate shocks. However, there is a risk, as suggested by Solluku and Valente (2021) in their analysis of PROGRESA, that these transfers could reduce men's altruistic behavior toward women. Therefore, their impact on resource allocation remains uncertain, as reduced husbands' altruism could counteract the transfer's intended benefits. Further research is needed to evaluate the effectiveness and implications of these interventions in mitigating the adverse effects of climate shocks on intra-household inequalities.

Chapter 2: Impact of Exposure to Hot Temperatures In-Utero on Child Health Outcomes

2.1 Introduction

Increasing temperatures pose serious risks to human health as climate change progresses (Mora et al., 2017). Given the potentially lasting consequences for health and economic productivity, there has been a growing interest in economics to study the effects of *direct* exposure to extreme heat on a number of outcomes, including mortality (Deschenes and Greenstone, 2011), worker productivity (Adhvaryu et al., 2020), labor supply (Somanathan et al., 2021) and cognitive performance (Garg et al., 2020; Graff Zivin et al., 2018).

However, less is understood about the perils of *indirect* exposure to heat stress on health and socioeconomic outcomes. For example, only a few papers have explored the causal impact of in-utero heat exposure on child health outcomes *at birth*, such as birth weight (Deschenes et al., 2009; Andalón et al., 2016), dehydration (Kim et al., 2019), length and Apgar score¹ (Andalón et al. 2016). Some studies estimate the effects of heat stress on long-run adult outcomes including earnings (Fishman et al. 2019; Isen et al. 2017), education attainment (Wilde et al. 2017) and height (Hi and Li, 2019). To my knowledge, there is only one existing study on the causal effects

¹The Apgar score is a measure of the child's overall health recorded 5 minutes after birth.

of prenatal exposure to heat on the health of children under the age of five (Bratti et al., 2021).²

The effects of prenatal heat exposure on early childhood health warrant further attention. A large number of epidemiological studies document a positive association between prenatal exposure to heat stress and likelihood of premature and low birth weight deliveries (Strand et al. 2011). These infants are at a greater risk of poor health and developing disabilities over time (Goldenberg and Culhane, 2007). Preterm birth and low birth weight are also strong predictors for reduced height (Farooqi et al., 2006; Stein et al, 2013), which has been linked to poor cognitive performance, low adult wages, and serves as a channel for inter-generational transmission of poverty (Grantham-McGregor et al., 2007). To make matters worse, catch-up growth in height is unlikely after a child's second birthday, and whether such growth meaningfully reverses the functional consequences of early-life nutritional deficiencies is unclear (Leroy et al., 2015). Taken together, early childhood health is a pathway through which disparities in economic well-being persist across generations, contributing to the persistence of inequality.

This study contributes to the literature by examining the causal effect of in-utero heat stress on children's height-for-age and the probability of stunted growth, defined as having sufficiently low height-for-age, in early years of life. I study this question in Bangladesh, which is among the countries most vulnerable to climate change, and focus on height-for-age of children under five as it is a signal of long-term malnutrition and likely to capture effects of elevated temperatures during the pregnancy. Alternative outcomes such as weight-for-age are signals for short-run malnutrition. To my knowledge, there are only two economic studies on the causal effects of prenatal heat exposure on measures of height in the developing world: Hu and Li (2019) on China,

²Other related studies include Skoufias and Vinha (2012) and Baker and Anttila-Hughes (2020), but these do not focus on prenatal heat exposure.

and Bratti et al. (2021) on Sub-Saharan Africa. The former investigates the long-run effects into adulthood while the latter focuses on a region that is quite different from Bangladesh in many dimensions, including female labor force participation, quality of prenatal health services, and climate.³

More broadly, this paper builds on a long-standing literature on the fetal origins hypothesis on how conditions in the womb can be attributed to chronic conditions later in life (Barker, 1990). Economic literature on this topic exploits random and unanticipated changes in a pregnant woman's environment to disentangle the causal effects of those environmental factors on the health and well-being of the child. For example, Almond et al. (2009) study the effects of prenatal exposure to radiation from the Chernobyl nuclear disaster on cognitive ability of children, Almond and Mazumder (2011) on diurnal fasting in the month Ramadan on birth weight, and Currie and Rossin-Slater (2013) on prenatal exposure to hurricanes on abnormal conditions of the newborn.⁴

The identification strategy in this study relies on exogenous differences in a pregnant woman's environment. The ideal experiment to examine the effects of in-utero exposure to heat stress would compare women giving birth to a child within the same narrowly-defined geographic area around the same time in the calendar year, but who experience different temperatures during their pregnancy. I implement this idea empirically by exploiting variation in temperature experienced by cohorts of pregnant mothers who give birth in the same district and month but in different calendar years. This variation in temperature within the same district and month is plausibly unrelated to maternal characteristics or household decisions, making it "as good as random". While I cannot test this assumption explicitly, I provide support for it using randomization

³Randell et al. (2020) study the association between heat stress during pregnancy on the height of children in Ethiopia, but their specification does not account for time-invariant seasonal factors that could be different across geographic areas, such as seasonality in birth weight that could vary by geographic region (Strand et al., 2011).

⁴Please see Almond et al., (2018) for recent review of this literature.

inference.

Unlike other studies on the effects of in-utero heat exposure, I distinguish between the effects of relatively short-lived “spikes” in temperature and more prolonged exposure to elevated ambient temperatures during the course of the pregnancy. A number of studies explore the effects of an additional hot day during the entire pregnancy and find that it can have significant effects on socioeconomic outcomes later in life (Isen et al., 2017; Hu and Li, 2019; Kim et al., 2019; Deschenes et al., 2009). Following this literature, one measure of heat exposure in this study is the number of heat-wave days during a pregnancy, as it would help elucidate how vulnerable a child’s health is to short spells of high temperature. On the other hand, Wilde et al. (2017) and Fishman et al. (2019) use average temperatures during the pregnancy to gauge impacts on education attainment and adult earnings, respectively. This measure reflects a more persistent increase in the temperature, and it is the second indicator I use to measure heat exposure.

I use data from the Bangladesh Integrated Household Survey, a household survey that contains detailed information on health, parental characteristics, and other household attributes. From this, I construct a child-level dataset, which I merge with temperature data from the ERA5 reanalysis dataset. Contrary to studies largely conducted in the US and China, I find that one additional hot day does not have a statistically significant effect on child health in Bangladesh.⁵ However, more prolonged exposure to heat stress, manifested by a higher average temperature during the pregnancy, does have a statistically significant impact on a child’s height-for-age

⁵This could be due to differences in the context. One prominent factor could be differences in female labor force participation. Over the past two decades, the US and China have maintained an average female labor force participation rate of approximately 58% and 68%, respectively, whereas the average in Bangladesh has been only 29% (World Bank, 2021). In my data, only 30% of mothers are working. Therefore, it is possible that mothers in my sample have more flexibility to adapt to short-term spikes in temperature by, for example, rescheduling outdoor tasks around the hottest time of the day so that they are less exposed to heat. However, this would be harder to do for more prolonged increases in temperature, and this could explain why there is a significant effect when there is a one standard deviation increase in average temperature.

and probability of stunted growth. On average, a one standard deviation increase in average temperature during the gestation period resulted in children having a height-for-age that was 0.186 standard deviations lower than what it would have been without the temperature increase. Additionally, their probability of stunting increased by 6.9 percentage points. Examining whether these effects are heterogeneous based on the timing of exposure during pregnancy, I find that they are primarily driven by heat stress in the third trimester.

I explore the potential role of various confounding factors to further establish that the observed effects on children are indeed caused by heat stress during pregnancy. One potential confounder is changes in fertility patterns, but I find no evidence supporting this explanation. There are no differences in the observed characteristics of mothers or households for children born in different seasons. Additionally, I find no evidence that mothers adapt the timing of their pregnancies to avoid heat stress during pregnancy. I also examine whether heat stress affects children's health outcomes through reduced agricultural yields, but I do not find support for this channel either.

My results are robust to sensitivity checks. I repeat the analysis on a subsample of children for whom birth location data is available, and the suggestive evidence from this smaller sample aligns with my main findings. To ensure that the model is correctly specified and not detecting a result by chance, I conduct a placebo test by randomly reassigning the “treatment” (i.e., average in-utero temperature) 1,000 times across children within the same district-birth month cell. The results further strengthen my interpretation that the effects are causal.

The rest of this paper proceeds as follows. Section 2.2 describes the data on households and daily weather in Bangladesh as well as the different definitions of heatwaves applied, followed by Section 2.3 which graphs the variation of temperature and its correlation with childhood stunting

in the Bangladeshi context. Section 2.4 outlines the identification strategy and investigates the relevance of possible confounding factors. Section 2.5 presents the main findings of the paper on heatwaves and the prevalence of stunting among children under 5 years, followed by the conclusion in Section 2.6.

2.2 Data

Household and child level data

I use all three waves of the Bangladesh Integrated Household Survey (BIHS), which is a panel survey nationally-representative of rural Bangladesh conducted by International Food Policy Research Institute (IFPRI) in 2011/12, 2015, 2018/19 covering around 6,500 households (Akhtar, 2013; IFPRI, 2016; IFPRI, 2020). Households are located in all of Bangladesh's 64 districts and in 259 of 544 sub-districts. It contains detailed information on anthropometry of all children under the age of 5, their age, child nutrition practices before and after birth, and other household characteristics.

I first assembled the panel so I could track each household over time, then using the mother's ID for each child, I was able to match a mother's characteristics (such as her education and age) to each child. As joint families are the norm in Bangladesh where different sets of parents could reside in the same household, I required the parents' ID to match each child to the right mother and father. Father's ID was available in the data only from 2015 onward, so I was not always able to match a child to their father's characteristics because of the prevalence of missing values. Therefore, I match the household head's characteristics to each child to complement the mother's characteristics.

Once I matched each child to the characteristics of their mother and household head, I constructed a child-level panel of 5,916 children born between 2007 and 2018. Using information on a child's height, age, and sex, I computed their height-for-age z-score (HAZ) for children under the age of 5 using the WHO Child Growth Standards (WHO, 2006), which is the standard in the literature on child nutrition. This is my primary outcome of interest as measures of height are used to reflect long-standing malnutrition, whereas measures of weight signal shorter-term nutritional deficiencies. Therefore, if there are detrimental effects of high temperatures during the pregnancy, they are likely to be picked up in measures of height rather than weight. The height-for-age z-score compares a child's height to the median height of a reference population of children of the same gender and age, weighted by the standard deviation of the same reference population. A child is considered stunted if their height-for-age z-score is two standard deviations below the median of the reference population.

I dropped children who had a height-for-age z-score that was 6 standard deviations below or above the median of the reference population,⁶ as well as children who had a birth weight greater than 6,500 g, as these values are medically implausible. I also dropped children who showed inconsistencies in the reported sex in the panel data, or whose date of birth was missing. This left me with a resulting sample of 5,837 children, 2,061 of whom were followed in all three waves of the panel survey. For this analysis, however, I use only the cross-section of the 5,837 children. See Table 2.1 for descriptive statistics.

I obtained licensed confidential data from IFPRI on each child's exact date of birth, so that I could best approximate the pregnancy period, as well as the village GPS coordinates that would

⁶As well as other standard nutrition indicators such as weight-for-age and weight-for-height z-scores that were 6 standard deviations above or below the median.

allow me to merge this household survey with data on temperature and precipitation which I will describe next.⁷

Temperature and rainfall data

Given the presence of only 35 ground weather stations that provide sparse ground coverage for Bangladesh's 64 districts, I collect daily temperature and precipitation data from ERA5 (Herbach et al., 2020). This is a type of reanalysis data constructed by the European Centre for Medium-Term Weather Forecasting that uses observational data combined with a physics-based model to produce estimates for weather/climate on a 31 km grid. The use of such reanalysis products has been supported in the literature where observational data is sparse or of low quality (Auffhammer et al. 2013) and are commonly used in economics (see for example, Garg et al., 2020 and Adhvaryu et al., 2020). From this data I am able to compute variables for temperature and precipitation at the district level from 1980 to 2020, by weighting a pixel based on how much it covers (or “masks”) the geographic area of a particular district.⁸ Weather variables are matched to households based on the district of residence.

Unlike Wilde et al. (2017) and Fishman et al. (2019) that use monthly temperatures, I extract *daily* rather than monthly temperature (and precipitation) data to maximize the variability of the heat exposure variables across pregnant mothers within a district-month cell for a given birth year.

Definitions of heat exposure

I construct two measures for heat exposure nine months before a child's date of birth, as

⁷This data was provided by the IFPRI. IFPRI bears no responsibility for the analyses or interpretations of the data presented here.

⁸If a district is covered by 60% of pixel 1, 80% of pixel 2, 20% of pixel 3, and 3% of pixel 4, the mean for variable *X* will be calculated for the district using the relative weights of each pixel. The pixel that masks the district the most will be given the highest weight.

well as during each trimester of the pregnancy. Figure 2.1 illustrates the non-linear relationship between these two measures.

The first measure is the number of “heat-wave days” during the pregnancy. Many studies on temperature shocks report the effect of an additional hot day on long-term outcomes. Different criteria are used to define what constitutes a hot day depending on the context. For example, Hu and Li (2019) use a threshold of 29.4°C (85°F), and Isen et al. (2017) use multiple cutoffs from 24 to 32+°C. In a study conducted specifically on Bangladesh, Nissan et al. (2017) are the first to define a heat wave in the context of Bangladesh and find that “elevated minimum and maximum daily temperatures over the 95th percentile for 3 consecutive days” have significant impacts on human health. Keeping in line with this definition, I categorize a day as a “heat-wave day” if it is the third consecutive day when the mean temperature that day exceeds the 95th percentile of daily values from 1980 to 2020, which is 29.8°C.⁹ Heat-wave days could occur consecutively or separately and would reflect the effect of a relatively short-lived heat shock during the course of the pregnancy. Women in my sample experience a varying number of heat-wave days, ranging from 0 to as high as 109.

The second measure is the average temperature during the gestation period in degrees Celsius, following Wilde et al. (2017) and Fishman et al. (2019). This measure would signify the effects of a more persistent change in the temperature experienced during pregnancy.

Table 2.1 shows that, on average, a pregnant woman in the sample experiences roughly 30 heat-wave days during the gestation period, and an average temperature of 25.5°C. In terms of variation in the temperature variables, the standard deviation in average temperature in my setting

⁹This threshold of 29.8°C applies to the whole country. To determine this threshold for the whole country, I used temperature data at a national level, not the district level. For everything else in this study, I use district-level daily temperature.




Figure 2.1: Scatter depicts mean number of heatwave days by average temperature during the pregnancy, where each bin is 0.02°C. Each point reflects the average number of heat-wave days that fall in each bin. Dotted lines indicate the 95% confidence interval. The dash line represents a Lowess smoothing line.

(1.43°C) is generally more than that reported by Wilde et al. (2017) and Fishman et al. (2019) in Sub-Saharan Africa and Ecuador, respectively. But the variation in the number of heat-wave days is lower than in other studies such as Hu and Li (2019). That could be because I employ a relatively stronger measure for what constitutes a hot day: unlike other papers, temperatures need to be consistently high for three *consecutive* days in my setting to be categorized as a heat-wave day.

The mean height-for-age z-score (HAZ) is less than 0, implying that the average child in the sample is shorter than the median height of the reference population of the same age and sex, as determined by the WHO Child Growth Standards. Around 40% of the children less than 5 are stunted, which is close to the 41.3% stunting rate in 2011 (World Bank, 2021).

In terms of mother's activities, only 30% of the mothers in the sample worked.¹⁰ This number is not far from Bangladesh's overall female labor force participation rate which ranged

¹⁰81% of the mothers that worked were involved in raising poultry or livestock.

Table 2.1: Descriptive statistics for children under age 5

	Mean	Min	Max	SD	N
<i>Weather variables</i>					
No. of heat-wave days during pregnancy	29.47	0.00	109.00	18.33	5,837
No. of heat-wave days in first trimester	9.35	0.00	69.00	11.35	5,837
No. of heat-wave days in second trimester	10.38	0.00	65.00	11.85	5,837
No. of heat-wave days in third trimester	9.75	0.00	66.00	11.71	5,837
Total precipitation during pregnancy (cm)	18.41	4.48	58.49	8.11	5,837
Avg daily temperature during pregnancy (°C)	25.48	21.45	29.04	1.43	5,837
Avg daily temperature in first trimester (°C)	25.27	16.45	31.18	3.53	5,837
Avg daily temperature in second trimester (°C)	25.73	16.55	31.50	3.45	5,837
Avg daily temperature in third trimester (°C)	25.44	16.32	31.30	3.54	5,837
Avg daily precipitation during pregnancy (cm)	0.07	0.02	0.22	0.03	5,837
<i>Outcome variables</i>					
Height-for-age z-score	-1.63	-5.97	5.88	1.42	5,605
Stunted (0/1)	0.40	0.00	1.00	0.49	5,605
<i>General child and household-level characteristics</i>					
Male (0/1)	0.51	0.00	1.00	0.50	5,837
Age in months	25.50	0.03	59.97	15.24	5,832
Muslim (0/1)	0.91	0.00	1.00	0.29	5,837
Household head's years of education	3.70	0.00	16.00	3.91	5,825
Household size	5.26	2.00	20.00	2.01	5,837
No. of girls under 5 years	0.66	0.00	4.00	0.66	5,837
No. of boys under 5 years	0.68	0.00	4.00	0.63	5,837
Mother's years of education	5.49	0.00	16.00	3.62	5,714
Mother's age at first birth	19.70	13.00	37.00	2.96	4,051
Mother works (0/1)	0.30	0.00	1.00	0.46	5,713
Monthly per-capita food expenditure (Taka)	1,747.65	215.11	11,893.68	953.08	5,837
Monthly per-capita nonfood expenditure (Taka)	1,271.25	144.00	11,452.21	994.14	5,837
Access to water supply (0/1)	0.85	0.00	1.00	0.36	5,837
Access to electricity (0/1)	0.62	0.00	1.00	0.49	5,837
Avg NDVI during pregnancy	0.54	0.06	0.80	0.08	4,212
Geographic place of delivery	1.41	1.00	3.00	0.58	2,635

Note: These summary statistics are reported for the full sample of 5,837 children under the age of five years. Any difference between the number of observations for a variable and the sample size is due to missing values.

from 30-36% between 2011 to 2019 (World Bank, 2021).

Information on geographic place of delivery was available for only 2,635 children in my sample, and 1,683 of these children were *certainly* born in the same district where the household was located. The remaining children mainly reported they were born in their mother's hometown,




Figure 2.2: Correlation between height-for-age z-score (HAZ) and heat exposure during pregnancy, for children under 2 years. Scatter depicts mean HAZ by number of heat wave days in-utero, where each bin is 2 days. Each point reflects the average HAZ of children under age 2 that fall in each bin. Dotted lines indicate 95% confidence interval. The dash line represents a Lowess smoothing line. See Figure B.17 in the appendix for the average HAZ over the full range of heat-wave days.

but it is unclear if that is within the same district or not. The main analysis in this paper uses information available for all 5,837 children under the age of 5, but I supplement my findings by restricting to this sub-sample of 1,683 children later in the paper.

2.3 Descriptive analysis

Figure 2.2 below plots the relationship that exists in the raw data between the height-for-age z-score for children under the age of two and the number of heat-wave days during the pregnancy. The relationship is negative, as expected, which provides credence to the hypothesis that hot temperatures during pregnancy may have a role to play in a child's nutritional status in early years of life. The magnitude of the slope is modest, which is also in line with expectations. For comparison, nutrition interventions and conditional cash transfer programs lasting 6 to 24

months have been shown to change height-for-age z-scores by 0.161 to 0.41 (see Table B.8 in the appendix for examples). Since the slope reflects the change in z-score for each additional heat-wave day, it is reasonable that the change would be relatively small.

In-utero heat exposure would depend on season (or month) of birth. The winter season in Bangladesh stretches from October to February, during which temperatures are low and there are few hot days. Children born in March, for example, would be exposed to less heat during the pregnancy relative to children born in September, as the latter would be exposed to the summer heat. Figure 2.3 illustrates the number of heat-wave days a woman was exposed to during her pregnancy, across districts, separated by the child's season of birth. The premonsoon/summer season stretches from March to May, monsoon from June to September, followed by the winter season.¹¹ Figure 2.3 illustrates variation in prenatal heat exposure across districts and shows that children born in the winter were exposed to a greater number of heat-wave days than those born in the summer.¹²

Children born in the winter season experience stronger heat exposure in-utero, which may be associated with their height. To explore this, Figure 2.4 plots the relationship between the height (in cm) of children under the age of two and their season of birth.¹³ The difference is quite subtle, but it appears that overall the height for children born in the winter, who generally experience a greater number of heat-wave days as shown in Figure 2.2, is slightly less than that of children born in the pre-monsoon season. The magnitude of the effect is small, but this is in line with expectations, as Table B.8 in the appendix shows that other interventions typically change height by less than 2 cm. The fact that children born in the winter appear to be relatively shorter in

¹¹Categorizing months into season is based on Nissan et al (2017).

¹²A similar pattern can be seen using average temperatures instead. See Figure B.18 in the appendix.

¹³I used height (measured in cm) instead of height-for-age z-score for illustrative purposes so that this subtle relationship is easier to see graphically.




Figure 2.3: Variation in number of heat-wave days across Bangladesh's 64 districts, by season of birth.




Figure 2.4: Variation in height (cm) across Bangladesh's 64 districts, by season of birth. Gray and blue error bars reflect the 95% confidence intervals corresponding to heights of children born in the pre-monsoon and winter seasons, respectively. Overall, across all districts, the average height for children (less than two years) born in the pre-monsoon was 71.9 cm, while it was 68.2 cm for children born in the winter season.

the raw data adds more credibility to the hypothesis that hot temperatures could negatively impact a child's nutritional outcomes, even a few years after birth. To test this hypothesis rigorously, I develop the identification strategy described next.

2.4 Empirical strategy

I estimate the following reduced-form linear relationship between heat exposure during the pregnancy and the outcome Y_{idmy} , for child i in district d , born in month m and year y :

$$Y_{idmy} = \alpha + \beta Temp_{idmy} + \delta Rainfall_{idmy} + X_i' \gamma + \phi_{d,m} + \mu_y + \epsilon_{idmy} \quad (2.1)$$

The main outcome variables are the height-for-age z-score for children under five and a dummy variable indicating whether a child is stunted – that is, whether their height-for-age is at least 2 standard deviations below the median of the reference population for the same gender and age. $Temp_{idmy}$ refers to one of two measures of heat exposure – either the number of heat-wave days or the average temperature (in degrees Celsius) during the pregnancy. β is the coefficient of interest.

I control for precipitation in a household's district during the gestation period, $Rainfall_{idmy}$. As temperature and rainfall have been shown to be highly correlated (Auffhammer et al. 2013) and rainfall around the time of birth can impact later-life health (Maccini and Yang, 2009), omitting it could bias the coefficient on $Temp_{idmy}$. To be consistent with the literature, when the temperature measure used is the number of heat-wave days, $Rainfall_{idmy}$ is the total precipitation (in cm) in a household's district during the gestation period. In specifications where I use average temperature, $Rainfall_{idmy}$ is the average precipitation during the gestation period. X_i is a vector

of characteristics of child i such as child's age (in months), sex, number of boys and girls in the household under the age of 5 years, household head's educational attainment, and access to water and electricity at the time of the survey.

The ideal experiment to address this question would compare women giving birth to a child within the same narrowly-defined geographic area around the same time in the calendar year, but who have somehow been exposed to different temperatures during their pregnancy. I take this idea to the data and exploit variation in temperature experienced by cohorts of pregnant mothers who give birth in the same district and month but in different calendar years.

As parents can sort into different parts of the country because of their preferences for hot or cold temperatures that could be associated with the health of their children, comparing women across districts would allow these unobservable differences to bias the coefficient of interest. Similarly, the season or month of birth is associated with multiple socioeconomic outcomes (Buckles and Hungerman, 2013) and parents can plan to conceive and give birth in certain seasons. For example, women may be aware of the consequences of high temperatures during pregnancy and prefer to give birth in the pre-monsoon season, as they would be exposed to less heat during the gestation period. These preferences could make them systematically different from women who prefer to give birth in the winter season in ways that could be correlated with the outcomes of interest. The extent of this type of sorting could vary across different parts of the country – women who experience harsher summers could have more reason to plan pregnancies this way, than women who reside in cooler parts of the country (Wilde et al., 2017; Bratti et al., 2021). To mitigate these sources of bias, I include district-by-month fixed effects, $\phi_{d,m}$, to control for time-invariant factors within a district-month cell, and exploit variation in heat exposure *within* this cell over different calendar years. I replace this with district-by-season fixed

effects in some of the less stringent specifications.

I would also want to control for trends in the dependent variable over time, to better identify the changes in the outcome variables that can be attributed only to prenatal hot temperatures. As trends in the primary outcome of interest (height-for-age) are highly non-linear within a district-month cell (see Figure B.19 in the appendix) and cannot be captured accurately by linear or quadratic trends, I include year fixed effects in equation (2.1), μ_y . Figure 2.5 illustrates residual variation in the temperature variables, after netting out the fixed effects.

The identifying assumption is that within a district-season or district-month cell, and controlling for time-varying differences in the outcome variables that are common in all parts of the country, exposure to higher temperatures is exogenous. While I cannot test this assumption explicitly, I provide support for it later using randomization procedures. Standard errors are clustered at the level of treatment assignment, following Abadie et al. (2021). In this case, the treatment is assigned at the district level, since everyone residing in the same district at a given point in time is exposed to the same temperature. As a robustness check, I also estimate the same specification using Conley standard errors, which account for potential spatial correlation in the error terms.

To study whether the timing of heat exposure matters, as suggested by the epidemiological literature (Strand et al., 2011), I estimate the following regression where I breakdown $Temp_{idmy}$ by trimester:

$$Y_{idmy} = \alpha + \sum_{T=1}^3 \beta^T Temp_{idmy}^T + \delta Rainfall_{idmy} + X_i' \gamma + \phi_{d,m} + \mu_y + \epsilon_{idmy}, \quad (2.2)$$

where T refers to the trimester, $Temp^T$ refers to heat exposure during each trimester.

Possible confounding factors




Figure 2.5: Variation in the two measures of in-utero heat exposure (the number of heat-wave days and average temperature °C during the pregnancy) after taking fixed effects into account.

I examine the role of potential confounding factors—fertility patterns, the possibility that heat affects child health through reduced agricultural yields, and uncertainty about pregnancy start dates—to assess whether they raise concerns for interpreting the results.

1. **Fertility:** If parents become more cognizant about the effects of prenatal heat stress on child health, there could be a change in fertility trends over time in a given season.¹⁴ To check whether this is a concern in my data, I plot the share of all births in a given calendar year, by season (see Figure 2.6). For example, in 2007 around 20% of all births were in the pre-monsoon season, compared to 42% in the winter season. Keeping the season fixed, the share of births remains fairly constant across birth years in the sample, showing there was

¹⁴Barreca et al. (2018) find evidence in the US from 1931-2010 suggesting that temperature shocks can alter birth rates in certain seasons.




Figure 2.6: Share of births by season. Shares calculated by dividing the number of births in a particular season in a given birth year, by the total number of births in that birth year.

no consistent change in fertility trends in a season.

There is a difference in levels, but it could simply be because the pre-monsoon season is three months long, whereas winter is from October to February. In Figure B.20 in the appendix, I adjust the duration of all seasons to four months. The difference in levels narrows significantly but does not disappear completely. I explore differences in observable characteristics across households and mothers that could account for this minor difference in levels. Figure 2.7 graphs household observable characteristics that could influence fertility decisions: monthly per-capita food and non-food expenditure¹⁵ (as a measure of household income), mother's age at first birth, and mother's education, by season and year of birth. Changes in these characteristics over time is similar across birth season – children born in a given year seem to have a similar family background, regardless of season of birth.

¹⁵I use information on monthly per-capita expenditure if it was collected in the same year as the birth year of the child, or one year before or after, to be able to approximate of financial well-being close to the time of birth.

To add to this evidence, I check for adaptive behavior in the timing of births – specifically, whether heat stress during pregnancy influences the timing of subsequent pregnancies. If women experience high temperatures during a pregnancy, they might plan their next child to avoid the hot summer months, for example by timing the next birth for the pre-monsoon season instead of winter. To explore this, I transform the child-level BIHS data into a mother-level dataset, tracking up to five births per woman over different years. Panel A in Figure B.21 shows that heat exposure remains fairly constant across births by the same woman, and Panel B shows no evidence of adaptive behavior. This brings me to conclude that fertility behavior is not a cause of concern in this setting.

2. **Agricultural incomes:** As rural Bangladesh is heavily reliant on agriculture, a potential concern is that higher temperatures could reduce agricultural yields, thereby affecting household incomes, which could in turn influence children’s health outcomes. Since agricultural productivity itself may be affected by hotter temperatures, it is not explicitly controlled for in the specifications. However, resource-poor households are likely to be more vulnerable to this channel. To assess the importance of this channel in this setting—without introducing a bad control—a proxy for household income or resources is used: mother’s education.¹⁶

$$Y_{idmy} = \alpha + \beta Temp_{idmy} + \lambda Temp_{idmy} \times LowEduc_i + \delta Rainfall_{idmy} \quad (2.3)$$

$$+ X_i' \gamma + \phi_{d,m} + \mu_y + \epsilon_{idmy}$$

LowEduc_i is a dummy variable indicating whether child *i*’s mother’s education was below

¹⁶In Table B.10 in the appendix, I report results for equation (2.3) but explicitly control for agricultural productivity during the pregnancy with the normalized difference vegetation index (NDVI) in the household’s vicinity, instead of mother’s education.




Figure 2.7: Factors possibly affecting fertility decisions. Dotted lines represent the upper and lower bounds of the 95% confidence intervals. Panels report the mean per-capita household expenditure, mother's age at first birth, and mother's education for all children born in a particular year. Expenditures were deflated to the baseline value.

the median in the sample: 5 years. λ would indicate whether resource-poor households face any additional impact of heat exposure, possibly through the agricultural income channel.

3. **Uncertainty about start date of pregnancy:** Unfortunately, the survey does not include information on the start date of the pregnancy. Therefore, similar to many other studies on the impact of temperature shocks during pregnancy, in-utero heat exposure is measured by assuming a full-term pregnancy, with the approximate date of conception reverse-engineered from the reported date of birth. This raises two complications. First, there could be some measurement error in heat exposure if a child was born prematurely. However, this is unlikely to pose a major concern in this study. In rural Bangladesh, approximately 20% of births are preterm, but around 70% of these preterm births occur between 8.5 to 9 months (Baqui et al., 2013). For these children, the measurement error would be less than two weeks. Applying these rates to the sample suggests that only about 6% of the children are likely to have been born before 8.5 months.

Secondly, medical literature indicates that the likelihood of a preterm delivery can increase with exposure to adverse environmental conditions, including high temperatures (Strand et al., 2011). Unfortunately, there is no way in the data to ascertain the actual start of the pregnancy to determine the extent to which this could be an issue in this study. This must be taken into account when interpreting the results.

2.5 Empirical findings

Estimates of equation (2.1) showing effects on height-for-age z-score and stunting are reported in Table 2.2. Panel A reports effects of short-lived “spikes” in temperature, whereas

Table 2.2: Effects of prenatal heat exposure

	HAZ			Stunted		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Heat-wave days during pregnancy	-0.024 (0.037)	-0.037 (0.043)	-0.031 (0.045)	0.001 (0.013)	0.002 (0.016)	0.000 (0.017)
Panel B						
Avg temperature during pregnancy	-0.010 (0.044)	-0.197* (0.111)	-0.186* (0.109)	-0.005 (0.014)	0.072* (0.037)	0.069* (0.038)
Observations	5602	5602	5590	5602	5602	5590
District-birthmonth FE	No	Yes	Yes	No	Yes	Yes
District-season FE	Yes	No	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	No	Yes	No	No	Yes
Sample	Full	Full	Full	Full	Full	Full

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Table reports results of equation 2.1. Temperature variables are standardized to mean 0 and unit standard deviation. Standard errors clustered at the district level. All regressions control for rainfall during the pregnancy, the child's age in months and sex. Additional covariates include number of boys and girls under the age of 5, the household head's years of education, and the household's access to water and electricity at the time of the survey.

Panel B reflects the impact of a more prolonged exposure to elevated temperatures during the full nine months. For ease of interpretation, both measures of heat exposure have been standardized to have mean 0 and unit standard deviation. All regressions control for rainfall during the pregnancy, the child's age and sex, as well as year fixed effects and district-by-season or district-by-month fixed effects. Columns (3) and (6) include additional controls to improve precision, including the number of boys and girls under the age of 5, the household head's years of education, and the household's access to water and electricity at the time of the survey.

Contrary to studies done in other countries measuring the effects of an additional hot day (see, for example, Hu and Li (2019) for the effects observed in China), results from panel A show that a one standard deviation increase in the number of heat-wave days (~ 18 days) during

Table 2.3: Heterogeneous effects of prenatal heat exposure by trimester

	HAZ			Stunted		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Heat-wave days in first trimester	-0.015 (0.034)	0.025 (0.044)	0.027 (0.046)	0.008 (0.013)	-0.000 (0.015)	0.000 (0.015)
Heat-wave days in second trimester	-0.000 (0.026)	-0.024 (0.044)	-0.019 (0.044)	-0.012 (0.010)	-0.017 (0.017)	-0.018 (0.017)
Heat-wave days in third trimester	-0.043 (0.036)	-0.074 (0.046)	-0.071 (0.047)	0.019 (0.011)	0.027 (0.017)	0.026 (0.017)
Panel B						
Avg temperature in first trimester	-0.084 (0.059)	-0.138 (0.128)	-0.138 (0.124)	0.015 (0.017)	0.054 (0.043)	0.054 (0.043)
Avg temperature in second trimester	0.006 (0.039)	-0.091 (0.104)	-0.080 (0.105)	-0.004 (0.014)	0.036 (0.041)	0.033 (0.041)
Avg temperature in third trimester	-0.036 (0.044)	-0.302** (0.132)	-0.285** (0.135)	0.001 (0.013)	0.101** (0.039)	0.097** (0.040)
Observations	5602	5602	5590	5602	5602	5590
District-birthmonth FE	No	Yes	Yes	No	Yes	Yes
District-season FE	Yes	No	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	No	Yes	No	No	Yes
Sample	Full	Full	Full	Full	Full	Full

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Table reports results of equation 2.2. Temperature variables are standardized to mean 0 and unit standard deviation. Standard errors clustered at the district level. All regressions control for rainfall during the pregnancy, the child's age and sex. Additional covariates include number of boys and girls under the age of 5, the household head's years of education, and the household's access to water and electricity at the time of the survey.

the pregnancy does not have a statistically significant effect on the height-for-age z-score, or probability of stunted growth. On the other hand, panel B indicates that, on average, a one standard deviation increase in average temperature ($\sim 1.4^\circ\text{C}$) during the course of the pregnancy decreases the height-for-age z-score by 0.186, and increases the probability of a child being stunted by 0.069 (or 6.9 percentage points). These estimates become more precise after adjusting

for spatial correlation using Conley standard errors (see Table B.11 in the appendix).¹⁷ As the average age of child in the sample is 25-26 months, this effect translates into a decrease in height of approximately 0.6 cm relative the median height of children that age. To give this result some context, the *Familias en Accion* conditional cash transfer program in Colombia improved the height-for-age z-score by 0.161 and decreased the probability of stunting by 0.069 for children less than 2 years, after being exposed to the intervention for one year (Attanasio et al., 2005). Taken together, the results in Table 2.2 suggest that child health outcomes in Bangladesh are sensitive to a persistent increase in temperature during the course of the pregnancy rather than temporary spikes in temperature.

As epidemiological literature suggests the effect of environmental factors could depend when they occur during a pregnancy (Strand et al., 2011), I compute heat exposure in each trimester and report results of equation (2.2) in Table 2.3. As before, all temperature variables have been standardized. Comparing the size of the coefficient across trimesters, both panels A and B indicate that child health outcomes are most sensitive to heat in the third trimester. However, only elevated average temperatures in the third trimester have a statistically significant impact: a one standard deviation increase in the average temperature during this time decreases the height-for-age z-score of children by 0.29, and increases the probability of being stunted by 9.7 percentage points.

While there is no consensus in the literature about *when* elevated temperatures matter the most (Strand et al., 2011), there are some studies documenting the detrimental effects of heat stress later in the pregnancy. Deschenes et al. (2009) find the effects of higher temperatures on low birth weight to be the strongest in the second and third trimesters when the fetus undergoes

¹⁷To see effects without standardizing the temperature variables, see Table B.9.

rapid growth. In a more recent study, Ha et al. (2017) find that heat stress in the third trimester, but not in the first or second trimesters, significantly increases the risk of a low birth weight delivery. As low birth weight is a predictor of reduced height in later years (Farooqi et al., 2006; Stein et al., 2013), it is consistent for the effects in the last trimester to be the strongest.

As a sensitivity check, I restrict the sample to the 1,683 children who reported they were born in their district of residence.¹⁸ Children who were reportedly born close to where their house is located are more likely to be present there during the third trimester when the effects of heat exposure appear to be the greatest. As expected, the absolute value of the coefficients in both panels A and B increase, but the estimates are noisier because of the reduced sample size.¹⁹ Even though the results from Table 2.4 are not conclusive, they provide suggestive evidence in line with the main results in Table 2.2.

A potential concern is that hotter temperatures could affect pregnant mothers through changes in agricultural productivity, rather than just by the temperature itself. As rural Bangladesh is largely agrarian, changes in agricultural productivity could affect household incomes and food security, and ultimately fetal development. One way to proxy for local agricultural productivity is the normalized difference vegetation index (NDVI) in a household's vicinity, to provide a measure of plant greenness.²⁰ However, NDVI, or any other measure of agricultural productivity, would itself be a potential outcome so it would be problematic to control for this variable directly. But if this channel is important, it will most likely affect resource-poor households the most. I use a proxy of permanent household income or resources – mother's education – which is unlikely to

¹⁸As mentioned in section 2.2, there was information on geographic place of delivery in the data for 2,635 of the children in my sample, and 1,683 of these children were born close to their home. For the remaining children, nothing can be said with certainty about their district of birth.

¹⁹Figure B.22 in the appendix plots the residual variation in the temperature variables for this reduced sample.

²⁰NDVI is measured on a scale from -1 to 1. Areas with more dense vegetation would have a higher value for NDVI.

Table 2.4: Results for subsample of local births

	HAZ			Stunted		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Heat-wave days during pregnancy	-0.086 (0.066)	-0.149 (0.120)	-0.135 (0.117)	0.024 (0.019)	0.043 (0.039)	0.039 (0.039)
Panel B						
Avg temperature during pregnancy	-0.230* (0.121)	-0.301 (0.283)	-0.282 (0.272)	0.069** (0.028)	0.083 (0.074)	0.079 (0.073)
Observations	1666	1666	1661	1666	1666	1661
District-birthmonth FE	No	Yes	Yes	No	Yes	Yes
District-season FE	Yes	No	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	No	Yes	No	No	Yes
Sample	Local	Local	Local	Local	Local	Local

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Table reports results of equation 2.1, but only for a sub-sample of children whose location of birth was the same as their district of residence. Temperature variables are standardized to mean 0 and unit standard deviation. Standard errors clustered at the district level. All regressions control for rainfall during the pregnancy, the child's age and sex. Additional covariates include number of boys and girls under the age of 5, the household head's years of education, and the household's access to water and electricity at the time of the survey.

be affected by the “treatment”. The interaction of heat exposure with mother's education would help shed light on whether this channel is important in the context of this study.²¹

Table 2.5 reports the results of equation (2.3). Temperature variables are again standardized and have a mean of 0. Panels A and B show that mothers who are less educated in my sample but are exposed to mean levels of heat exposure are more likely to have children who are malnourished, on average. Effects on children of women who are relatively more educated but experienced a one standard deviation increase in average temperatures are very similar to those reported in Table 2.2 on the full sample. The coefficient on the interaction is small and statistically insignificant across

²¹In Table B.10 in the appendix, I perform the same analysis but interact with the average NDVI in a household's vicinity during the pregnancy period instead of mother's education, and the results are consistent with those in 2.5.

Table 2.5: Exploring the agricultural productivity channel (controlling for mother's education)

	HAZ		Stunted	
	(1)	(2)	(3)	(4)
Panel A				
Heat-wave days	-0.038 (0.042)	-0.051 (0.049)	0.013 (0.014)	0.016 (0.017)
Low Educ	-0.287*** (0.048)	-0.306*** (0.049)	0.085*** (0.014)	0.090*** (0.015)
Heat-wave days $\times$ Low Educ	0.029 (0.042)	0.024 (0.045)	-0.025* (0.014)	-0.026 (0.016)
Panel B				
Avg. temp	0.010 (0.046)	-0.193* (0.115)	-0.002 (0.016)	0.076* (0.039)
Low Educ	-0.286*** (0.049)	-0.307*** (0.049)	0.084*** (0.014)	0.089*** (0.016)
Avg. temp $\times$ Low Educ	-0.041 (0.030)	-0.034 (0.036)	-0.008 (0.011)	-0.010 (0.012)
Observations	5484	5484	5484	5484
District-birthmonth FE	No	Yes	No	Yes
District-season FE	Yes	No	Yes	No
Year FE	Yes	Yes	Yes	Yes

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Table reports results of equation 2.3. Temperature variables are standardized to mean 0 and unit standard deviation. Standard errors clustered at the district level. *LowEduc* is a binary variable equal to 1 if mother's education was below the median (5 years). All regressions control for rainfall during the pregnancy, the child's age and sex.

specifications, indicating that the effects of heat are *not* compounded by whether a household is resource-poor. These results provide suggestive evidence that the agricultural yield is unlikely to be a primary driver of the observed effects in this context.

To ensure the model has been specified correctly and the effects of prenatal average temperatures do not occur by chance, I randomly reassigned prenatal average temperature and rainfall across children in a district-month cell. Therefore, a child is randomly matched to the in-utero "treatment" received by another child, who could be born in a different year. This random reassignment




Figure 2.8: Distribution of the coefficient on average temperature in equation 2.1, after randomly reassigning prenatal average temperature and rainfall across children within a district-month cell 1000 times. Vertical lines represent the coefficients in columns (3) and (6) in Table 2.2 for height-for-age (left) and stunting (right), respectively.

was repeated 1000 times, and each time I computed the coefficient of interest, β in equation (2.1). Figure 2.8 plots the distribution of these coefficients using the false temperature and rainfall measures assigned to each child. The vertical lines correspond to the coefficient estimates reported in columns (3) and (6) of Table 2.2. The distribution of the β s using the false “treatment” is centered around 0. Moreover, the results reported in Table 2.2 appear highly unlikely to have occurred by chance: the fraction of the 1000 coefficient estimates that were more extreme than the results in Table 2.2 were 0.007 for height-for-age and 0.004 for stunting. This strengthens confidence that the observed relationships are not driven by spurious correlations.

2.6 Conclusion

This study contributes to understanding how climate change can reinforce existing inequalities through its effects on early-life health. Building on the long-standing literature on the fetal origins

hypothesis, which links conditions in the womb to chronic health outcomes later in life (Barker, 1990), I investigate whether hot temperatures during pregnancy affect the health of children up to five years after birth. I use variation in temperature experienced by cohorts of pregnant mothers who give birth in the same district and month but in different calendar years, after netting out year specific effects. Because this residual variation in temperature within a narrowly defined district-month cell over time is plausibly exogenous, I am able to estimate the causal effects of prenatal heat exposure on children's height-for-age z-score and their probability of stunted growth—two well-established indicators of long-term undernourishment with consequences for later-life well-being.

The results show that short-lived spikes in temperature during the pregnancy – the number of heat-wave days – do not have a statistically significant impact on the health of children in early years of life. However, more prolonged exposure to hot temperatures – the average temperature during the gestation period – significantly reduced the height-for-age z-score. On average, mothers who experienced a one standard deviation increase in average temperature during pregnancy ($\sim 1.4^{\circ}\text{C}$) gave birth to children that had a height-for-age that was 0.186 standard deviations below what it would have been had they not been exposed to this increase in temperature. The probability that these children are stunted also increased by 0.069. To give this result some context, the *Familias en Accion* conditional cash transfer program in Colombia decreased the probability of stunting by the same magnitude in children younger than 24 months after families had been exposed to the program for a year (Attanasio et al., 2005).²²

This finding is in contrast to other studies primarily in the US and China that find statistically

²²But it came with a cost – families with children less than 6 years received a minimum stipend of \$188.16 during this time (2002\$).

significant effects of one additional hot day during the pregnancy on child or adult outcomes (including birth weight, and long-term adult earnings and height). However, this inconsistency could be due to differences in the context. One prominent factor could be differences in female labor force participation. Over the past two decades, the US and China have maintained an average female labor force participation rate of approximately 58% and 68%, respectively, whereas the average in Bangladesh has been only 29% (World Bank, 2021). In my data, only 30% of mothers are working. Therefore, it is possible that mothers in my sample have more flexibility to adapt to short-term spikes in temperature by, for example, rescheduling outdoor tasks around the hottest time of the day so that they are less exposed to heat. However, this would be harder to do for more prolonged increases in temperature, and this could explain why there is a significant effect when there is a one standard deviation increase in average temperature.

This study has important implications for policymakers in the wake of changing global temperatures under climate change, particularly in resource-constrained countries. The fact that children born in unusually hot years have a greater probability of being stunted and that the biological time-frame for catch-up growth is only two years (Leroy et al. 2015), cohorts of children born in hot years could be permanently disadvantaged, unless parents and policymakers find sustainable ways to adapt. Since stunting has been linked to lower educational attainment and lower wages later in life (Grantham-McGregor et al., 2007), these climate-related health shocks have the potential to reinforce existing inequalities. While increasing access to air conditioners has been found to reduce the adverse effects of heat on human health in high-income countries like the US (Barreca et al., 2016), this may not be a viable solution in some low-income countries with limited access to reliable electricity. This highlights the need for research alternative adaptation strategies, such as early heat warning systems.

Chapter 3: Effects over the Life of a Program: Evidence from an Education Conditional Cash Transfer Program for Girls

Fatima Najeeb, Dhushyanth Raju, Esha Chhabra

3.1 Introduction

How do the effects of a program evolve over many years of its life, over successive cohorts of beneficiaries? We examine this question in relation to an education conditional cash transfer (CCT) program targeted to girls of secondary school age in Punjab province, Pakistan. Specifically, we examine how the effects on girls' secondary school enrollment evolved over a decade of the program's existence, from its launch in 2004 until 2015.

While there is extensive evaluative literature on education CCT programs in developing countries, most studies examine effects on the likelihood of school participation—and, to a lesser extent, the likelihood of class attendance—one or two years after the program has been introduced (Fiszbein et al. 2009; Baird et al. 2014; Garcia and Saavedra 2017). We are unaware of any evaluations of education-only interventions in developing countries that examine how effects on enrollment outcomes evolve over an extended period of a program's implementation and across multiple cohorts of beneficiaries.¹

¹Most evaluations of long-term education CCT programs (e.g., in Cambodia, Malawi, and Colombia) assess

The question is of interest for two reasons. First, effects on beneficiary cohorts could change over time because of changes in relevant environmental, institutional, or intervention-specific factors. In the CCT program that we examine, a key intervention-specific factor that changed was the real value of the program benefit, which was not indexed to inflation. Indeed, the value of the benefit declined substantially over the study's observation period not only in real terms but relative to average household consumption spending, weakening the program's economic incentive as observed by households. As a result, standard economic theory would predict that program effects should fall over the study's observation period.

Second, the effects of a program could change after the first few years of implementation as different agents have time to reoptimize their behavior in response to the program's introduction—in the case of an education program, such agents as households, school staff, and education system administrators (Glewwe and Muralidaran 2016). Most evaluations that examine effects in the first few years after a program's introduction are unlikely to capture such reoptimizing of behavior, which can dampen or amplify program effects (Das et al. 2013). Our investigation, by looking at the evolution of enrollment effects over the extended life of the program, allows us to converge on the program's "full" effect.

The CCT program that we study, the Punjab Female School Stipend Program, is aimed at stimulating household demand for education where that demand is particularly weak, by targeting cash transfers on the basis of gender, age, and location. The school participation rate (the percentage of children attending school) is lower for girls than for boys, and this gender gap widens marginally when children reach preadolescence. In addition, the gender gap

individuals who were exposed for a limited period and follow them years later, rather than tracking successive cohorts throughout a program's implementation (see Molina Millán et al. 2019 for a review).

in school participation is greater in areas of the province with poorer overall socioeconomic outcomes. The program, initiated by the Punjab government in 2004, offers households cash benefits tied to girls' regular attendance in secondary grades in government schools. The program was targeted to districts in the province with low adult literacy, which is correlated with poor overall socioeconomic status. And because the program offers the CCT benefit only to students in government schools (which provide free education), it implicitly also targets poorer households. Richer households in the province are more likely than poorer ones to send their children to fee-based private schools (Nguyen and Raju 2015).

Assignment of the program is based on a clear rule. Of the 34 districts in Punjab at the time of its launch in 2004, the program was offered to the 15 districts with the lowest adult literacy rates, which results in a cutoff literacy rate of 40 percent. Data from the 1998 census, six years before the program's conception, were used for the assignment variable. Thus, there is no risk of manipulation of the assignment variable by program administrators or beneficiaries. And to our knowledge, this variable (or correlated variables) and, importantly, the cutoff have not been used in assigning any other programs.

Exploiting the program assignment process that generates a cutoff, we evaluate program effects on the basis of a sharp regression discontinuity (RD) design. Because the program assignment variable is not continuous but discrete, with 34 mass points (given 34 districts), local randomization-based RD (Cattaneo, Idrobo, and Titiunik, 2020, 2024) is a better fit than the parametric and continuity-based nonparametric methods typically used in RD applications. Local randomization-based RD uses the smallest window that yields two consecutive mass points (two districts), one on either side of the cutoff. We apply this approach to government administrative data from annual censuses of government schools. These data are independent of those used

by the government to administer the CCT program. Constructing a school-level panel with the census data, we examine program effects in the RD program district on school-level changes in girls' enrollment in secondary grades across different subperiods from 2003 (before the program) until 2015, as well as potential spillover effects on boys.

The Punjab Female School Stipend Program has been evaluated in multiple studies (Population Council 2007; Hasan 2010a, 2010b; Chaudhury and Parajuli 2010; Alam, Baez, and Del Carpio 2011). These evaluations have used various sources of data, covering periods extending as far as 2009. The evaluations tend to find significant, positive effects on girls' secondary school enrollment, among other measures. While our research questions align with those in previous studies, a key distinction is that our observation period extends several years further, evaluating effects up to 2015 following the program's launch in 2004.

Another key difference is that our study employs an identification strategy more appropriate to the context and data. Chaudhury and Parajuli (2010) use a triple-differences RD model, comparing outcomes before and after program onset, between boys' and girls' schools, and between program and nonprogram districts. However, Hasan (2010a) provides evidence that boys may not be a valid counterfactual for evaluating the effects of the program on girls' school enrollment. In a different study, Hasan (2010b) estimates effects using difference-in-differences, comparing outcomes before and after program onset and between program and nonprogram districts, but relies on nonrepresentative data. Moreover, all previous evaluations of the program use parametric or continuity-based nonparametric methods, which are susceptible to bias—parametric estimations incorporate observations far from the cutoff, while nonparametric estimations treat the assignment variable as continuous, though it is not in this case.

Our local randomization-based RD estimations suggest that the program had significant,

positive effects on girls' secondary school enrollment throughout the observation period. The average gain in the number of girls per school in the RD program district ranges from 33 to 53, equivalent to an average gain in girls' secondary enrollment at the school level of 25–41 percent. These effects are on the higher end of the range of effects documented in previous evaluations of the program. There is also a positive spillover effect on boys, with schools in program districts experiencing an additional 33–48 boys per school in secondary enrollment across different sub-periods, compared to schools in non-program districts.

Program effects remain stable over the study's observation period, though measured with substantial imprecision, yet remain statistically significant. The results suggest that effects were robust to a large decline in the real value of the program benefit over the observation period and to potential reoptimization by agents in later years of the program. Our calculations of cost-effectiveness, using the estimated program effect over the full observation period (2003–15), indicate that for every \$100 spent annually, the program induced the enrollment of 0.2–2.3 beneficiary girls.²

As suggested, the stable effects on girls' secondary school enrollment over time despite the large decline in the real value of the program benefit are inconsistent with what standard economic theory would have predicted. The observed pattern does suggest potential behavioral explanations, however. For example, as Benhassine et al. (2015) argue in explaining the effects they find in evaluating an unconditional cash transfer program in Morocco, “labeling” a program as support for education can affect perceptions of the importance of schooling. Similarly, the

²As additional questions, we examine whether the effects on girls' secondary school enrollment numbers are reflected in girls' secondary school participation rate and whether they translate into an effect on female education attainment. To answer these questions, we use data on school participation status for adolescent girls and years of education for young women from district-representative household sample surveys for 2003 and 2014. But we fail to find balance in means for selected covariates between program and nonprogram RD districts, a key validity check for applying local randomization-based RD estimation.

program we evaluate may have signaled the government’s commitment to girls’ secondary education—or to secondary education more generally—making it more salient to households and communities and inducing girls’ enrollment independent of the benefit size. This mechanism may also explain the positive spillover effects on boys.

Alternatively, greater exposure of households and communities to girls’ secondary education in the early years of the program, induced by the economic incentive of the program benefit, may have generated a durable, positive change in the “taste” for girls’ secondary education or helped correct for preexisting biases in beliefs or preferences (Kremer, Rao, and Schilbach 2019). And these adjustments in beliefs and preferences may have catalyzed a positive shift in social norms around girls’ secondary education.

3.2 Related literature

Most evaluations of education CCT programs in developing countries assess program effects on the likelihood of school participation, typically in the first few years after the program has been initiated (Fiszbein et al. 2009; Baird et al. 2014; Garcia and Saavedra 2017). A meta-analysis of the effects of education CCT programs finds a significant, positive average effect on the likelihood of school participation, at both the primary and secondary school levels (Garcia and Saavedra 2017). The study also finds a larger average effect on the likelihood of participation at the secondary level than at the primary level. Effects found by studies included in the meta-evaluation sample vary substantially, however.

Among education CCT programs that, like the one we evaluate, are targeted to girls and condition cash benefits on secondary school participation, almost all those that have been evaluated

are in Asia. Here too, evaluations generally find that the programs have significant, positive effects on the likelihood of school participation. For example, the Japan Fund for Poverty Reduction scholarship program in Cambodia, which awarded scholarships to girls from poor households who were in their last grade of primary school, increased the likelihood of school participation by 30 percentage points (Filmer and Schady 2008). Similarly, a CCT program in Turkey increased the likelihood of secondary school participation for girls from very poor households by 11 percentage points (Ahmed et al. 2007). And the Female Secondary School Assistance Project in Bangladesh, targeted to rural girls of secondary school age, increased the likelihood of girls' school participation by 8 percent (Khandker, Pitt, and Fuwa 2003). Similar evidence is available for Pakistan. For example, Cheema et al. (2016) examine the effects of a CCT program introduced in 2012 for beneficiaries of the Benazir Income Support Program (BISP). The CCT program benefit for households was conditioned on regular school attendance by children of primary school age. Using household survey data collected in 2016, and propensity score matching of BISP households in non-CCT program areas to BISP+CCT program households in CCT program areas, they find that the CCT program increased the likelihood of school participation.

The program that we study also has been the subject of earlier evaluations (Population Council 2007; Chaudhury and Parajuli 2010; Hasan 2010a, 2010b; Alam, Baez, and Del Carpio 2011), including some based on the same data that we use—annual rounds of government school censuses. In contrast to our study, however, these evaluations examine effects on girls' education outcomes only in the early years of the program.

For example, using government school census data and a parametric triple-differences RD approach,³ Chaudhury and Parajuli (2010) find that the program increased female enrollment in

³The triple-differences approach compares outcomes before and after program onset, between boys' and girls'

grades 6–8 in government schools by an average of 6 girls per school (or by 9 percent relative to initial enrollment) over the period 2003–05. Using government school census data and difference-in-differences estimation,⁴ Hasan (2010a) finds that the program increased female enrollment in grades 6–8 in government schools by an average of 11 girls per school over the period 2003–08. Hasan also finds that the program increased corresponding male enrollment by an average of 10 boys per school. In another study, using nonrepresentative panel household sample survey data for three districts (one program district, two nonprogram districts) and within-family difference-in-differences estimation,⁵ Hasan (2010b) finds that the program had no effect on the likelihood of government school participation by girls ages 10–14 in 2005 or 2006 but had a significant positive effect of 3 percentage points on this outcome in 2007. Finally, using government school census data from 2003–09 and parametric RD estimation, Alam, Baez, and Del Carpio (2011) find that the program had significant positive effects on girls’ enrollment in grades 6–8 from 2005 to 2009. They find average gains in girls’ enrollment by grade ranging from 10 to 32 percent.

3.3 Program context and design

Pakistan, a lower-middle-income economy with a per capita income in 2018 of \$4,928 (in 2011 PPP\$), is the sixth most populous country in the world (World Bank 2019). The country is divided into four large provinces (Balochistan, Khyber Pakhtunkhwa, Punjab, and Sindh) as well as small territories. Punjab is a major province. It accounted for 53 percent of Pakistan’s population of almost 208 million in 2017 (Pakistan Bureau of Statistics 2017a) and for 55 percent middle schools (grades 6–8), and between program and nonprogram districts.

⁴The difference-in-differences approach compares outcomes before and after program onset and between program and nonprogram districts.

⁵Here, within-family differences are those between siblings of different genders.

of the country's annual output in 2015/16 (PERI 2017). Sixty-three percent of the province's population lives in rural areas (Pakistan Bureau of Statistics 2017b).

Government schools in the province provide education at five levels: preprimary, primary (grades 1–5), middle (grades 6–8), high (grades 9–10), and intermediate (grades 11–12). In this study, we define secondary school as grades 6–10 (middle and high school), and secondary school age as ages 11–15. Punjab tends to perform better than the other provinces in girls' and boys' school participation rates (see Figure 3.1 for net enrollment rates in 2014/15). Nevertheless, within Punjab, there is a large gender gap in school participation rates. In 2014, near the end of our study's observation period, the school participation rate for girls of secondary school age (66 percent) was 7 percentage points lower than that for boys (73 percent) (Table 3.1).

Government schools provide free education, and most students in Punjab attend government schools. In 2014, among students of secondary school age, 65 percent of boys and 67 percent of girls attended government schools, while most of the rest attended private schools. School participation rates vary substantially across districts within the province. In 2014, for girls of secondary school age, they ranged from a low of 28 percent in Rajanpur to a high of 89 percent in Gujrat (Figure 3.2). While the levels differ across time, preprogram patterns in school participation are qualitatively similar to the patterns discussed here.

The Punjab government initiated the Female School Stipend Program in 2004 with the aim of increasing the rate of girls' participation in secondary education in parts of the province with the poorest education outcomes.⁶ The selection of districts for the program was based on the literacy rate (among those age 10 and above), as measured by the 1998 Pakistan Population and

⁶The first distribution of cash benefits was initiated in January 2004, based on enrollment and attendance information for the October–December 2003 quarter.

Housing Census. Of the 34 districts that Punjab had at the time, the 15 with literacy rates below 40 percent were selected for the program (Table 3.2). Over time, as the number of districts in Punjab increased from 34 to 36,⁷ the number of program districts increased to 16, and the number of nonprogram districts to 20. The program assignment status in Table 3.2 remains constant during the study period.

Districts selected for the program had poor school participation outcomes (see Table 3.1). Measured in 2003, preprogram school participation rates for children of primary and secondary school ages were lower in program districts than in nonprogram districts, for both boys and girls. Moreover, the gender difference in preprogram school participation rates was greater in program than in nonprogram districts, particularly among children of secondary school age. In program districts, 59 percent of boys of secondary school age were enrolled in school, while it was 18 percentage points lower for girls. In nonprogram districts, the corresponding gender difference was only 4 percentage points. By 2014, near the end of our observation period, school participation rates had increased for both boys and girls and in both program and nonprogram districts, but large gender differences in school participation rates remained for children of secondary school age in program districts (Figure 3.3).

When initiated in 2004, the program covered girls enrolled in government schools in grades 6–8. In 2005, it was expanded to also cover girls enrolled in grades 9–10. Under the program, households were offered a cash benefit of 200 Pakistan rupees (PRs) a month (3.40 US\$) for each girl enrolled in grades 6–10 in a government school, conditional on an attendance rate of at least 80 percent in the preceding academic quarter.⁸ At baseline, the value of the per-girl benefit was

⁷Nankana Sahib was formed from the division of Sheikhpura in 2005, and Chiniot from the division of Jhang in 2009.

⁸Currency conversion is at the average exchange rate of PRs 58 per U.S. dollar in 2004. The benefit amount was estimated to cover safe travel to a government school with secondary grades, which in rural areas was typically

14 percent of average monthly household consumption spending per adult equivalent in Punjab (or 15 percent of the corresponding average in rural Punjab).⁹ Program benefits were transferred to households on a quarterly basis, with benefits for a given quarter delivered in the following quarter. The maximum individual transfer per quarter was PRs 600 (= PRs 200 × 3 months), offered for four quarters in a year.

Key steps in program administration began at the government school, where the head teacher was responsible for recording female students' attendance for the quarter in a standard register as well as for determining their eligibility and the amount of their benefit. The teacher also filled in a money order form for each female student beneficiary. The school sent a copy of the completed register and the individual money order forms to the district education department, which transmitted the money orders to the government post service. The post service then delivered cash to beneficiary girls and households, typically at either the girls' school or their home. In addition, district education departments would transmit data and documentation on beneficiary girls and cash delivery to the provincial education department.

A sample-based process evaluation conducted during the first year of the program by a government-contracted consulting firm finds that while program implementation suffered from delays and was burdensome for different parties (including beneficiaries), all sample teachers and virtually all sample students and parents were aware of the attendance condition for program eligibility, and all sample beneficiary households reported receiving the full benefit amount (Innovative

located some distance from the girl's village. While the program benefit is small in absolute terms compared with those in samples of CCT programs in developing countries that have been evaluated (Garcia and Saavedra 2017), the level of the benefit as a percentage of household consumption spending is within the range of benefit sizes in these samples (Baird et al. 2014; Benhassine et al. 2015).

⁹Based on statistics from the 2004/05 Pakistan Household Income and Expenditure Survey. We are unable to estimate the program benefit level as a percentage of average household consumption spending per adult equivalent in program districts because available household sample surveys that gather data on income and consumption spending are not representative at the district level.

Development Consultants 2005).

In 2017, the Punjab government adjusted the value and delivery of program benefits to households. Renaming the program Khadim-e-Punjab Zewar-e-Taleem, it raised the benefit value, which had remained constant in nominal terms since the program's introduction, to PRs 1,000 a month (\$9.50).¹⁰ The new benefit level is equivalent to 18 percent of average household consumption spending per adult equivalent in Punjab (or 21 percent of the corresponding average for rural Punjab).¹¹ The government also began modifying how program benefits are delivered to households by requiring parents of beneficiary girls to register for an ATM card called the Khidmat Card. This card allows them to withdraw program benefits from branches of the Bank of Punjab, which is a commercial bank. Our study's observation period ends in 2015, before the introduction of these adjustments to the program.

3.4 Data and empirical strategy

3.4.1 Data

We use data from the Punjab government's Annual School Census from 2003 to 2015. These data are independent of those used by the government for administering the girls' stipend program. Initiated in 2003, the census is an annual field survey of government schools that gathers basic information on schools and teachers, including student enrollment by grade and gender, with a reference date of October 31.¹² Its main purpose is to monitor the status of government schools.

¹⁰At the average exchange rate of PRs 105 per U.S. dollar in 2016.

¹¹Based on the 2015/16 Pakistan Household Income and Expenditure Survey.

¹²The government school year runs from April to March.

The provincial education department has primary responsibility for managing all stages of the census and sole responsibility for designing the census questionnaire. At the school level, teachers complete the questionnaires delivered to them. These questionnaires are collected by district government officials and returned to the provincial education department, where the information is digitized.

Using the annual censuses, we construct a school-level panel.¹³ On the basis of this panel, we construct the outcome measure—the absolute change in school-level girls’ enrollment in secondary grades for six subperiods of different lengths originating from baseline (2003): 2003–05, 2003–07, 2003–09, 2003–11, 2003–13, and 2003–15. Table B.12 presents baseline comparisons (in 2003, one year before the program’s launch) between schools in program and non-program districts. On average, total enrollment in secondary grades is approximately 83 students higher per school in non-program districts compared to schools in program districts. This pattern holds for both boys’ and girls’ secondary school enrollment. Schools in non-program districts also tend to have more teachers, a lower student-to-teacher ratio, and better infrastructure. For example, they are more likely to have access to electricity and to be enclosed by a boundary wall, which could be important for girls’ safety. For the empirical strategy, it is therefore important to compare districts that are similar. This motivates the use of a regression discontinuity (RD) design, which is described in the next section.

¹³While government schools are assigned unique identification (ID) numbers, school ID protocols were adjusted by the provincial and district education departments in certain years and assigned school IDs were not strictly adhered to during data collection, entry, and processing. Consequently, to construct an accurate school-level panel, we worked closely with provincial education department officials, specifically the Programme Monitoring and Implementation Unit (PMIU), with intimate working knowledge of the Annual School Census data and ground-level knowledge of schools.

3.4.2 Empirical strategy

We estimate program effects on the basis of a sharp RD design. One of two outcomes is possible for school S , conditional on the district literacy rate X_D . Schools in districts with a literacy rate below 40 percent are covered by the program. These schools have a potential girls' secondary school enrollment outcome of $Y_S(1)$. Schools in districts with a literacy rate equal to or above 40 percent are not covered by the program and have a potential outcome of $Y_S(0)$. In theory, the expected outcomes, given X_D , are two functions: $\mathbb{E}[Y_S(1) \mid X_D]$ if $X_D < 40$ and $\mathbb{E}[Y_S(0) \mid X_D]$ if $X_D \geq 40$. The sharp RD program effect at the cutoff, τ^{RD} , reflects the difference between the potential outcomes:

$$\tau^{RD} \equiv \mathbb{E}[Y_S(1) - Y_S(0) \mid X_D = 40] \quad (3.1)$$

Parametric methods approximate the global shape of these functions, but observations far from the cutoff can distort the approximation close to the cutoff and misrepresent the RD effect (Gelman and Imbens, 2014; Cattaneo, Idrobo, and Titiunik, 2020, 2024).

A common nonparametric solution is a continuity-based RD design, which approximates the conditional expectation functions as polynomials within an optimal bandwidth close to the cutoff. This design is less sensitive to extreme features of the data generation process or outliers far from the cutoff (Cattaneo and Vazquez-Bare, 2016; Imbens and Kalyanaraman, 2012).

However, the continuity-based RD design does not fit key defining features of our data. The design requires the program assignment variable—the district literacy rate—to be a continuous random variable. Because program assignment is at the district level while our outcome indicator

is at the school level, our assignment variable is discrete, as schools within the same district have the same assignment score. With 34 districts at baseline in 2003, the score has that many “mass points.”

Considering these data characteristics, we adopt the local randomization-based RD design as the only practical and valid approach for estimation and inference, following recommendations in Cattaneo, Idrobo, and Titiunik (2024), as there are only 34 mass points. This design formalizes the idea that treatment is as-if randomly distributed within a small window close to the cutoff, by imposing randomization-type assumptions that are stronger than the assumptions under the continuity-based RD design. It requires the existence of a window W_0 close to the cutoff where the distribution of the program assignment variable is known and uncorrelated with potential outcomes. Within W_0 , potential outcomes do not depend directly on the program assignment variable; they are associated only through program assignment. In other words, the design requires $\mathbb{E}[Y_S(1) \mid X_D]$ and $\mathbb{E}[Y_S(0) \mid X_D]$ to be constant functions of X_S within W_0 , and the vertical distance between these functions would reflect the local randomization-based RD effect, τ_{LRD} (Cattaneo, Titiunik, and Vazquez-Bare 2017; Cattaneo, Idrobo, and Titiunik 2020, 2024).

Following Cattaneo, Idrobo, and Titiunik (2024), because we have only 34 mass points, the window that we select for local randomization is the smallest window, W_{LRD} , with two consecutive mass points, one on either side of the cutoff. The two districts within this window are Khushab and Khanewal (see Table 3.2).

Within W_{LRD} , we assume complete randomization—or fixed-margins randomization—under which the numbers of program and nonprogram schools are fixed. Using this mechanism, we apply two inference methods.

The first is *Fisherian inference*, which tests the sharp null hypothesis of no program effect

for any given school and leads to correct inference even when the number of observations within W_{LRD} is very small. Fisher's p -value is calculated using randomization inference through the permutation of the program assignment status of observations under the sharp null hypothesis of no program effect.

The second is *Neyman inference*, which is based on the large-sample approximate behavior of the statistic and tests the null hypothesis of no program effect. While the Fisherian framework focuses on testing the sharp null hypothesis, one of the main purposes of the Neyman framework is to estimate the sample average program effect. The difference in average outcomes between program and nonprogram observations is an unbiased estimator of this parameter, assuming complete randomization within W_{LRD} . Because our sample of schools within W_{LRD} is sufficiently large, we use the Neyman method to obtain unbiased estimates of the sample average program effect.

Manipulation of program assignment status can undermine the identification of program effects. In our case, program assignment is based on district literacy rates obtained from 1998 data, while the program was initiated in 2004. Thus, we can consider districts to be plausibly exogenously sorted around the cutoff of 40 percent.

Migration between program and nonprogram districts could pose a threat to identification; however, available evidence suggests this is unlikely. Since the Annual School Census does not include information on children's addresses, we examine migration patterns using the 2012–13 Pakistan Labor Force Survey (LFS), conducted near the end of the evaluation period. Among male household heads, 86.2% reported residing in their district of birth. The primary reasons for migration were employment or returning home. Only 1.3% cited education, though it is unclear whether this referred to their own or their family members' education. Given the low incidence

of education-related migration, systematic bias in our estimates due to migration across program and nonprogram districts is unlikely.¹⁴

Identification of program effects can also be undermined if the program assignment variable and cutoff were also used for other programs. To our knowledge, no other government program, in education or in any other sector, has used district literacy rates and this program's cutoff for assignment.

Finally, while the enrollment data in government school censuses may not be fully accurate,¹⁵ we have no reason to suspect that the quality of data varies systematically between program and nonprogram observations within W_{LRD} .

3.5 Results

3.5.1 Validity of the local randomization-based RD window

The validity of local randomization-based RD estimation requires that there be at least one covariate whose preprogram value is correlated with the program assignment variable (the district literacy rate) everywhere except within W_{LRD} . This test is comparable to the baseline balance test in randomized controlled trials.

We select two important school-level covariates whose preprogram values should be associated with the district literacy rate: number of teachers and number of classes. We test the null hypothesis of no effect on the program assignment variable for each covariate. The window

¹⁴Other major household surveys, including the Multiple Indicator Cluster Survey (MICS) (2011, 2014) and the Pakistan Social and Living Standards Measurement Survey (PSLM) (2014), do not contain migration-related variables, limiting further analysis.

¹⁵Reported enrollment data may not be accurate if school enrollment records are not up to date as of the reference date for the government school census or if reported enrollment data become high stakes (and thus subject to manipulation) because they are used by government education authorities for decisions affecting the interests of stakeholders in the public education sector.

would be invalid if the null hypothesis is rejected for at least one covariate within it.

We fail to reject the null hypothesis within the smallest window W_{LRD} of 0.6 percentage points (as well as within the next smallest window of 2.6 percentage points), which contains our RD program and nonprogram districts, Khanewal and Khushab (Table 3.3). The results indicate that the preprogram values for the covariates are balanced within the selected small window around the cutoff, but not in larger windows.

3.5.2 Program effects on secondary school enrollment

Effect on girls' enrollment

Before discussing the local randomization-based RD results, we illustrate the RD effect on girls' secondary school enrollment for the subperiod 2003–05 in Figure 3.4, according to continuity-based nonparametric RD estimation. While this estimation method does not yield a valid estimate of the program effect because of the structure of our data (as discussed in Section 3.4), it does allow us to graph the program effect. The figure fits fourth-order local polynomials to bin-specific averages of our school-level measure of the change in girls' secondary school enrollment in the subperiod 2003–05 on either side of the cutoff (depicted by the vertical line). Program districts are located to the right of the cutoff. The figure clearly shows a positive program effect on girls' secondary school enrollment at the cutoff.¹⁶

The local randomization-based RD results are presented in Table 3.4. We find that the program produced statistically significant gains in girls' secondary school enrollment across all subperiods between 2003 and 2015, whether we apply Fisherian or Neyman inference. Program

¹⁶Continuity-based nonparametric RD plots for other subperiods show similar results and are available upon request.

effects over the different subperiods range from an average gain of 33 to 53 girls per school. Given an average preprogram secondary school enrollment of 130 girls per school in the RD program district, the effects are equivalent to an average gain of 25–41 percent in girls’ secondary school enrollment. These effects are on the higher end of the range of effects previously documented for the program we examine, which may reflect, among other things, differences across the studies in the data, sample, and estimation strategy.

We also aimed to examine differences in program effects between rural and urban areas.¹⁷ While the sample of schools in rural areas within W_{LRD} is large enough to apply the Neyman approach to estimate the program effect, the corresponding sample of schools in urban areas is too small. The Fisherian approach could provide valid inference for such a small sample, but not the program effect. Therefore, we do not examine program effects in urban areas. In rural areas, results show that the average gain in girls’ secondary school enrollment is statistically significant across all subperiods and ranges from 10 to 31 girls per school (see column 2 of Table 3.4).

While overall program effects differ somewhat in size, they are statistically similar across all subperiods, according to an examination of Fisherian confidence intervals (Figure 3.5). The size and statistical significance of the program effects are sustained over the study’s observation period despite a decline in the real value of the annual program benefit of more than 60 percent over the period, from PRs 2,400 in 2004 to PRs 860 in 2015 (Figure 3.9). As a result of this decline in the real value of the benefit, along with an increase in average real household income over the period, the relative benefit level decreased from an amount equivalent to 14 percent of

¹⁷On the one hand, program effects may be larger in rural communities than in urban areas because the economic incentive (in terms of benefit size relative to average household consumption spending) is marginally stronger in rural communities. On the other hand, program effects may be smaller in rural communities because government secondary schools that girls can attend are fewer and on average farther away, or because social stigma may be associated with travel by socially disadvantaged girls to schools outside their own rural community (Jacoby and Mansuri 2015).

average household consumption spending per adult equivalent in Punjab to 4 percent.¹⁸

Effect on boys' enrollment

With the program inducing parents to send their daughters to school, spillover effects on boys may have occurred as well. Figure 3.6 presents the estimated local randomization-based RD effects of the program on boys' secondary school enrollment across different subperiods. As in the analysis for girls, this estimation relies on local randomization within a small window around the literacy rate cutoff, where program assignment is treated as as-if random. The outcome variable measures the absolute change in boys' secondary enrollment between the 2003 baseline and each subsequent period. However, as shown in Table 3.3, the window validity check for boys indicates marginal covariate imbalance in the smallest window. This should be kept in mind when interpreting the results.

The results show a consistent and positive impact throughout the decade following the program's launch. Between 2003 and 2005, schools in program districts experienced an average increase of 40 boys per school in secondary enrollment compared to schools in non-program districts. These effects remained stable over time, ranging from 33 to 48 additional boys per school across different subperiods. The findings indicate that, while the program was specifically designed to improve girls' secondary school enrollment, it also had a positive spillover effect on boys' enrollment, possibly due to increased household awareness of education and shifts in social norms.

Program anticipation effects for girls in government primary schools

While the program provides cash benefits to girls enrolled in government secondary schools, it can indirectly induce girls to complete their primary education and transition to secondary

¹⁸Based on statistics from the 2004/05 and 2015/16 Pakistan Household Income and Expenditure Surveys.

education. We aimed to examine the existence and evolution of “program anticipation effects,” by estimating the local randomization-based RD effects on the school-level change in girls’ enrollment in grade 5 (the terminal primary grade) in government primary schools in each of the subperiods.¹⁹ But when performing our window validity check, we find that two of the covariates for primary schools with enrolled girls—the number of teachers and the number of classrooms—had preprogram values that were significantly correlated with the program assignment variable within our selected window, W_{LRD} , precluding the use of our estimator.

Program effects on girls’ education attainment

Gains in girls’ secondary school enrollment may not necessarily translate into gains in education attainment if girls encounter institutional or environmental hurdles that are difficult to overcome. Rigorous evidence on the long-term education and economic effects of development interventions, including education CCT programs, is scarce but growing (Bouguen et al. 2018). A recent meta-evaluation of the effects of education CCT programs finds that the average effect on the likelihood of school completion is significant and positive, although most of the studies included in the meta-evaluation sample do not find significant effects on this outcome (Garcia and Saavedra 2017).

Examining the program in our study, Alam, Baez, and Del Carpio (2011) find that it increased the likelihood of girls ages 15–16 completing middle school (grade 8) and transitioning to high school (grade 9). Their analysis applied difference-in-differences and parametric and nonparametric RD estimations to data from the 2003 and 2007/08 Punjab Multiple Indicator Cluster Surveys, which are large-scale household sample surveys representative at the district

¹⁹Chaudhury and Parajuli (2010) and Hasan (2010a) also examine program anticipation effects on girls in government primary schools in the first few years of the program.

level. As discussed earlier, however, these empirical methods are not well suited for identifying the effects of this program because of the structure of the program assignment variable.

We aimed to examine program effects on girls' secondary school participation rate and education attainment on the basis of the 2003 and 2014 Punjab Multiple Indicator Cluster Surveys. But as with the attempt to examine program anticipation effects on girls' grade-5 enrollment, we find that several household-level covariates for girls of secondary school age or for young women had preprogram values that were correlated with the program assignment variable within our selected window, W_{LRD} , for the randomization-based RD estimation. The covariates we tested include self-reported household income; number of rooms in the dwelling; connection to such utilities as gas, electricity, telephone, cable, or internet; and possession of such consumptive or productive assets as a bicycle, a television, or farmland. As a result, we were unable to examine this research direction.

3.5.3 Falsification tests

To assess the validity of the local randomization-based RD design, we conducted two falsification tests using alternative program assignment cutoffs at 30% and 50% instead of the actual 40% literacy rate cutoff. Under these alternative thresholds, districts would have been included in the program if their 1998 adult literacy rate was below the respective cutoff. These changes alter the set of districts near the false assignment threshold, as there is substantial variation in adult literacy rates across districts (Table 3.2). Given these alternative specifications, we reselect the window for local randomization, with Table B.13 providing validity checks for

window length.²⁰

Figure 3.7 presents the results of these falsification tests, plotting the estimated program effects on girls' secondary school enrollment using the false cutoffs. Across all sub-periods, despite constructing the outcome variable identically to the main analysis, the estimates at these false cutoffs are not statistically significant. The absence of significant effects at these false cutoffs provides additional evidence that the estimated program effects are not driven by arbitrary threshold selection. We also explored conducting this falsification test for boys' secondary school enrollment. However, as shown in Table B.13, the window validity check is not satisfied in this case, as covariates are statistically significant within the smallest window. As a result, we only report the falsification test results for girls, where the window validity check holds under the false cutoffs.

As an additional falsification test, we examine the effect of the program on a placebo outcome—the number of teacher positions filled (see Figure 3.8). Since the program did not target teacher hiring, any detected effect would suggest spurious RD estimates. However, the estimated effect is small and statistically indistinguishable from zero across all sub-periods, confirming that the program's estimated effects are not driven by unrelated changes in school characteristics.

3.5.4 Program cost-effectiveness

We assess the program's cost-effectiveness in relation to girls' secondary school enrollment in the RD program district, following Dhaliwal et al. (2013) and costing guidelines from the

²⁰Table B.13 presents the window validity checks for the 30% false cutoff. Similar tables for the 50% false cutoff are available upon request.

Abdul Latif Jameel Poverty Action Lab (J-PAL n.d.). Per beneficiary, apart from the cost of the benefit transfer (PRs 2,400 a year), the government paid a fee to the government post service of PRs 50 per money order (which amounts to PRs 200 a year given quarterly benefit transfers) and what the government refers to as an “incidental expense” of 1 percent of the benefit amount (PRs 24 a year).²¹ These values remained the same over the study’s observation period. Thus, the total annual per-beneficiary cost was PRs 2,624. This per-beneficiary cost almost perfectly matches the value from official government accounting records of actual annual outlays and beneficiaries for the program.

Note that the PRs 2,624 value reflects recurrent costs and does not capture fixed costs or costs incurred in designing and establishing the program. Moreover, for recurrent costs, the value does not capture the cost of time and other resources spent by provincial and district education department officials or school teachers in carrying out their program responsibilities. Nor does it capture any costs borne by households in ensuring that they were present when and where needed to receive the benefit transfer or in pursuing any grievances with the authorities relating to problems with receiving the benefit. While such omissions from cost calculations are not unusual in impact evaluation studies for developing countries (Dhaliwal et al. 2013), the PRs 2,624 value should be considered an underestimate of the actual annual per-beneficiary cost of the program.

In 2015, there were 28,559 girls in secondary grades in government schools in the RD program district. Comparing enrollment numbers from the school census rounds and program beneficiary numbers from government administrative data, we find that about 90 percent of

²¹This maximum administrative cost of PRs 224 per beneficiary per year is about 9 percent of the maximum benefit transfer per beneficiary per year. This percentage is on par with the modal ratio of per-beneficiary annual administrative cost to benefit transfer in a sample of CCT programs in developing countries that have been evaluated (Garcia and Saavedra 2017).

enrolled girls receive program benefits.²² Using these numbers, we estimate that the government spent PRs 67,496,341 on the program in the RD program district in 2015.

For the calculation of cost-effectiveness, we arbitrarily choose the program effect for the subperiod 2003–15, an average gain of 41 girls per school in the RD program district. This effect, which lies roughly in the middle of the range of program effects across subperiods, is equivalent to 32 percent of average preprogram girls' enrollment in secondary grades in government schools in the RD program district. Given a total preprogram enrollment of 25,676 girls in secondary grades in the RD program district, this effect translates into 8,109 girls being induced by the program to enroll in school.

From the total estimated outlay and the total number of girls induced to enroll in 2015, we estimate an annual cost per beneficiary girl induced to enroll of PRs 8,324 (or \$81).²³ Accounting for imprecision in the estimated program effect, we find that this annual cost ranges from PRs 57,106 (\$554) at the lower bound of the 90 percent confidence interval for the estimate of the program effect to PRs 4,450 (\$43) at the upper bound for the same interval.²⁴ As an alternative measure of cost-effectiveness, we estimate that for every \$100 spent annually, the program induced the enrollment of 1.2 beneficiary girls (or, after accounting for the imprecision in the estimated program effect, 0.2–2.3 beneficiary girls).²⁵

²²That the number of beneficiary girls is lower than girls' enrollment is presumably due to the application of the regular attendance condition for program benefit eligibility.

²³At the average exchange rate of PRs 103 per U.S. dollar in 2015.

²⁴The lower- and upper-bound estimates of program effects are 6 and 77 girls per school.

²⁵Although we use the standardized approach proposed by Dhaliwal et al. (2013), we are unable to compare the cost-effectiveness values for this program with those for other CCT programs in developing countries because of differences in outcome indicators. The available cost-effectiveness values are based on program effects on school participation rates, attendance rates, dropout rates, and education attainment (see, for example, Garcia and Saavedra 2017; and J-PAL 2019a and 2019b), while our outcome measure is school enrollment numbers.

3.6 Conclusion

Most rigorous evaluations of development interventions, including of CCT programs, examine effects in the first few years of a program’s implementation. Yet a program’s effects on later cohorts could differ from those on early cohorts—because agents have had the time to reoptimize their behavior in response to the program or because there may have been changes in key features of the program’s design or implementation. This study contributes to the evaluative literature for developing countries by examining the evolution of effects of a public education CCT program targeted to girls of secondary school age over the first decade of the program’s life, from its inception in 2004 until 2015. This examination is made possible by a sharp RD design based on an observed, numerical program assignment variable (the district adult literacy rate) and government administrative data from annual censuses of government schools that capture school-level data on girls’ enrollment in secondary grades. We apply a local randomization-based RD approach, which is better suited to the discrete nature of the assignment variable—comprised of 34 mass points—than the parametric and nonparametric methods used in previous evaluations (Cattaneo, Idrobo, and Titiunik 2024).

We find that the program had statistically significant, large positive effects on girls’ enrollment in secondary grades in government schools across cohorts throughout the period of observation—and that the effects were of similar size. Moreover, there was also a positive spillover effect on boys’ secondary school enrollment. These program effects were observed even though the program’s economic incentive effect on households weakened substantially over the observation period, as a result of a marked decline in the real value of the program benefit (which was not indexed to inflation) accompanied by a rise in average real household income. They were also observed

despite potential reoptimizing behavior by various agents in later years of the program.^{26 27}

To ensure that our estimates capture the true program effects, we conducted falsification tests using false cutoffs for the treatment assignment threshold. The results show no significant effects on girls' secondary school enrollment at these false cutoffs. Additionally, there is no detectable impact on the placebo outcome of the number of teacher positions filled. These findings reinforce the validity of our estimates, suggesting that the observed program effects are not driven by arbitrary threshold selection or unrelated changes in school characteristics.

The finding of sustained program effects on girls' secondary school enrollment over the study's observation period despite a loss in economic incentive suggests potential behavioral explanations. The program may have led to durable changes in household and community behavior regarding girls' secondary education—or secondary education more generally—as shown by improvements in boys' enrollment as well. For example, independent of the benefit amount, the program may have increased the salience of secondary education, shifting household and community perceptions about its acceptability and value. More generally, greater exposure to girls' secondary education induced by the program may have helped to correct for preexisting biases in beliefs and preferences of households and communities—or in their decision-making processes—in relation to girls' secondary education. As another potential explanation, consistent with a multiple-equilibria characterization of the setting, program benefits may have functioned as an initial catalytic agent, triggering communities to ratchet up from a low girls' secondary education equilibrium to a higher one (a shift in social norms).

²⁶Note that the documented RD effects, which relate to the RD program district of Khanewal, may not apply to program districts more generally.

²⁷As additional questions, we aimed to examine whether the program produces anticipation effects on girls in government primary schools and whether the effects on girls' secondary school enrollment numbers are reflected in girls' secondary school participation rates and translate into gains in female education attainment. But our window validity condition for using local randomization-based RD failed to hold for these additional lines of investigation.

We see two relatively straightforward directions for further evaluative research on the program. The first is to examine the effects on girls' enrollment of the higher benefit amount (along with the new method for delivering benefits) introduced in 2017. Relative to average household consumption spending per adult equivalent in the province, the benefit amount increased more than fourfold, from 4 to 18 percent. Does the restored—and enhanced—economic incentive induce greater gains in girls' secondary school enrollment? In due course, enough rounds of government annual school censuses should be available to examine the early effects of the program modification on girls' secondary school enrollment.

The second direction is to examine whether the program improves girls' academic achievement as measured by test scores, evidence that has been lacking for this program. Under the management of the Punjab Examination Commission, the provincial and district education departments have been administering standardized exams to students in grades 5 and 8 annually. Validity and reliability appear to be a concern for the existing rounds of exams. But with enough improvement in the design and administration of the exams, future rounds may serve as a credible source of data for analyzing the program's effects on girls' academic achievement.




Figure 3.1: Net enrollment rate by province, Pakistan, 2014/15. The net enrollment rate for a given set of grades is the percentage of children in the corresponding age group who attend those grades, by province. KPK = Khyber Pakhtunkhwa. Data from Pakistan Bureau of Statistics (2016).




Figure 3.2: School participation rate for girls' ages 11–15 by district, Punjab, 2014. The school participation rate for a given age group is the percentage of children in that age group who attend school. Estimates are adjusted for survey sampling weights. Own estimates based on data from the 2014 Punjab Multiple Indicator Cluster Survey.




Figure 3.3: School participation rate by age, gender, and district program status, Punjab, 2014. The school participation rate for a given age is the percentage of children of that age who attend school. Estimates are adjusted for survey sampling weights. Own estimates based on data from the 2014 Punjab Multiple Indicator Cluster Survey.




Figure 3.4: Regression discontinuity plot, girls' secondary school enrollment effect, 2003–05. The gray dots depict sample averages of the outcome variable within each evenly spaced bin. The red vertical line depicts the cutoff, with program districts to the right of the cutoff. The solid black curves depict fourth-order polynomial fits separately for program and nonprogram districts. Own estimates based on data from the Punjab government's Annual School Census for 2003 and 2005.




Figure 3.5: Program effects on girls' secondary school enrollment. Fisherian 95 percent confidence intervals are obtained using the Stata software package `rdrandinf` developed by Cattaneo, Titiunik, and Vazquez-Bare (2016). Own estimates based on data from the Punjab government's Annual School Census for various years.




Figure 3.6: Program effects on boys' secondary school enrollment. Fisherian 95 percent confidence intervals are obtained using the Stata software package `rdrandinf` developed by Cattaneo, Titiunik, and Vazquez-Bare (2016). Own estimates based on data from the Punjab government's Annual School Census for various years.




Figure 3.7: Program effects on enrollment for girls using false cutoffs. Fisherian 95 percent confidence intervals are obtained using the Stata software package `rdrandinf` developed by Cattaneo, Titiunik, and Vazquez-Bare (2016). Own estimates based on data from the Punjab government’s Annual School Census for various years.




Figure 3.8: Program effects on the number of teacher positions filled. Fisherian 95 percent confidence intervals are obtained using the Stata software package `rdrandinf` developed by Cattaneo, Titiunik, and Vazquez-Bare (2016). Own estimates based on data from the Punjab government’s Annual School Census for various years.




Figure 3.9: The base year for deflation is 2004, the year the program was launched. The program's cash benefit remained fixed in nominal terms from 2004 to 2015. Yearly consumer price index data from Global Economic Monitor, World Bank, <https://datacatalog.worldbank.org/dataset/global-economic-monitor>.

Table 3.1: School participation rate by gender, age group, and district program status in Punjab, 2003 and 2014

Districts	Gender	School participation rate	
		2003 (1)	2014 (2)
Secondary school age (ages 11–15)			
All	Male	67	73
	Female	57	66
Program	Male	59	65
	Female	41	52
Nonprogram	Male	72	80
	Female	68	77
Primary school age (ages 6–10)			
All	Male	72	85
	Female	66	79
Program	Male	63	76
	Female	51	65
Nonprogram	Male	81	92
	Female	76	90

Note: The school participation rate for a given age group is the percentage of children in that age group who attend school. Estimates are adjusted for survey sampling weights. Own estimates based on data from the 2003 and 2014 Punjab Multiple Indicator Cluster Surveys.

Table 3.2: District program assignment status

District	Adult literacy rate (1998)	Recentered literacy rate (cutoff 40%)	Covered by program?
	(1)	(2)	(3)
Rajanpur	20.7	19.3	Yes
Muzaffargarh	28.5	11.5	Yes
Lodhran	29.9	10.1	Yes
Dera Ghazi Khan	30.6	9.4	Yes
Rahim Yar Khan	33.1	6.9	Yes
Bhakkar	34.2	5.8	Yes
Pakpattan	34.7	5.3	Yes
Bahawalpur	35.0	5.0	Yes
Bahawalnagar	35.1	4.9	Yes
Kasur	36.2	3.8	Yes
Vehari	36.8	3.2	Yes
Jhang	37.1	2.9	Yes
Okara	37.8	2.2	Yes
Layyah	38.7	1.3	Yes
Khanewal	39.9	0.1	Yes
Khushab	40.5	-0.5	No
Hafizabad	40.7	-0.7	No
Mianwali	42.8	-2.8	No
Multan	43.4	-3.4	No
Sahiwal	43.9	-3.9	No
Sargodha	46.3	-6.3	No
Mandi Bahauddin	47.4	-7.4	No
Sheikhupura	47.8	-7.8	No
Attock	49.3	-9.3	No
Toba Tek Singh	50.5	-10.5	No
Faisalabad	51.9	-11.9	No
Narowal	52.6	-12.6	No
Gujranwala	56.5	-16.5	No
Chakwal	56.7	-16.7	No
Sialkot	58.9	-18.9	No
Gujrat	62.2	-22.2	No
Jhelum	63.9	-23.9	No
Lahore	64.7	-24.7	No
Rawalpindi	70.4	-30.4	No

Note: The recentered literacy rate is computed by subtracting the district's adult literacy rate from 40 percent, the value we set as the program assignment cutoff. District adult literacy rates from the 1998 Pakistan Population and Housing Census.

Table 3.3: Validity check of window length

Window length	Minimum balance test p-value	Nonprogram schools n	Program schools n
(1)	(2)	(3)	(4)
Girls' Schools			
0.6	0.602	74	197
2.6	0.325	123	492
4.6	0	550	992
6.6	0	850	1,554
8.6	0	1,162	1,756
10.6	0	1,490	1,909
12.6	0	2,025	2,016
14.6	0	2,025	2,016
16.6	0	2,242	2,016
Boys' Schools			
0.6	0.073	217	122
2.6	0.135	510	175
4.6	0.025	1,052	690
6.6	0.032	1,623	933
8.6	0.023	1,864	1,280
10.6	0.003	2,094	1,569
12.6	0	2,247	1,997
14.6	0	2,247	1,997
16.6	0	2,247	2,197
18.6	0	2,247	2,377

Note: Balance test results using preprogram values for two covariates: number of teachers and number of classes. Column 1 presents the window length in which the balance test is performed. Column 2 presents the minimum p-value across the covariates. Columns 3 and 4 show the number of nonprogram and program schools in each window. Balance tests use the Stata package `rdrandinf` (Cattaneo et al., 2016). Own estimates based on data from the Punjab government's Annual School Census for 2003.

Table 3.4: Program effects on girls' secondary school enrollment

	Overall (1)	Rural (2)
2003–05		
Program effect	32.884	10.278
Fisher's p-value	0.000	0.013
Large-sample p-value	0.009	0.017
2003–07		
Program effect	53.485	26.251
Fisher's p-value	0.000	0.000
Large-sample p-value	0.000	0.000
2003–09		
Program effect	52.727	30.926
Fisher's p-value	0.000	0.000
Large-sample p-value	0.001	0.000
2003–11		
Program effect	49.345	23.992
Fisher's p-value	0.001	0.000
Large-sample p-value	0.007	0.000
2003–13		
Program effect	35.340	12.909
Fisher's p-value	0.049	0.061
Large-sample p-value	0.108	0.066
2003–15		
Program effect	41.162	18.447
Fisher's p-value	0.027	0.017
Large-sample p-value	0.072	0.019
Nonprogram schools n	74	64
Program schools n	197	178

Note: Program effects are obtained using local randomization methods with two mass points, assuming a complete randomization mechanism. Randomization p-values are obtained using 1,000 permutations. Fisherian p-values correspond to a test of the sharp null hypothesis of no program effect. Large-sample p-values are based on Neyman inference methods and correspond to a test of the null hypothesis of no program effect. Local randomization results are generated using the Stata software package `rdlocrand` developed by Cattaneo, Titiunik, and Vazquez-Bare (2016). Own estimates based on data from the Punjab government's Annual School Census for various years.

Appendix A: Supplementary details for Chapter 1

A.1 Illustration of identification of resource shares from Engel curves

Figure A.1 demonstrates the intuition behind identification of resource shares at a time period t . Recall that information on floods or any other household socio-economic attributes (except for their total household expenditure and budget shares on a private assignable good) is not required for identification of η_{jt} for a set of households in time t . Identification of η_{jt} comes from the slopes of Engel curves of a private assignable good and the identifying assumptions in Section 1.3.2.

For simplicity, I only consider households with one man and one woman in this example, therefore $\eta_{mt} + \eta_{wt} = 1$. The horizontal axis plots the logarithm of total household expenditure and the vertical axis plots the budget share spent on a private assignable good, such as food consumed by men and women denoted as W_{mt} and W_{wt} . If the budget shares for a private assignable good and total household expenditure are observable, then we can draw these Engel curves for a hypothetical homogeneous set of households.

The slopes for men and women are $\eta_{mt}\beta_{mt}$ and $\eta_{wt}\beta_{wt}$ (equation 1.9), and these will be estimated as the coefficients on $\ln y$ when it is regressed on W_{mt} and W_{wt} , respectively. Under the shape-similarity assumption that $\beta_{mt} = \beta_{wt} = \beta_t$, resource shares can be backed out by taking




Figure A.1: Hypothetical Engel curves for a private assignable good (food in this case) for men and women at a time period t across a group of households that is homogeneous except for their total household expenditure and food expenditure shares. Children are excluded from this example. The Engel curve for women is three times as steep as that for men. Both Engel curves are downward sloping in accordance with Engel's Law.

the ratio of the slopes of men and women, setting it equal to the ratio of the coefficients on $\ln y$ and solving for η_{wt} and η_{mt} . After the cancellation of the β_t terms, two unknown parameters remain, along with two equations: the ratio of the slopes and the condition that resource shares sum to one (i.e., $\eta_{mt} + \eta_{wt} = 1$). Hence, the model is just identified.

Note that Figure A.1 also shows that budget shares for a private assignable good are not necessarily indicative of the share of total household resources allocated to that individual. In this hypothetical example, men have a higher food budget share which could reflect their different tastes. However, the *slope* of their Engel curve is one third that of women's, therefore women have a higher resource share than men.

A.2 Intuition behind identification of the change in resource shares in treated and untreated villages

In this section, I will go through a hypothetical example of how the causal effect of floods on resource shares is identified through changes in the slopes of Engel curves. I only consider households with men and women for now, but the intuition can be extended to households with children. For simplicity, assume for now that floods hit villages randomly, and within a time period t , households are homogeneous except for 3 sources of variation:

1. Total household expenditure;
2. Food expenditure shares for $j = m, w$;
3. Exposure to flooding (random).

For this group of homogeneous households, the resource share for $j = m, w$ at time t , is modeled as the following *within* the system of Engel curves estimated through NLSUR:

$$\begin{aligned}\eta_{jt} = & \eta_{0j} + \eta_{1j}Survey2 + \eta_{2j}Survey3 \\ & + \eta_{3j}Flood_h \times Survey2 + \eta_{4j}Flood_h \times Survey3,\end{aligned}\tag{A.1}$$

For the 2014 cohort, η_{3j} and the η_{4j} parameters are important in determining the relationship between floods and resource shares of men and women in the short and longer-run, respectively.¹ Figure A.2 below provides a graphical illustration to explain the intuition and sources of variation needed to identify these effects.

¹The baseline period corresponds to the first household survey conducted in 2011, the second survey was done 6 months after the 2014 floods, and the third survey was conducted 4 years after the 2014 flood.

In each of the six graphs in Figure A.2 resource shares are identified under the assumption that $\beta_{mt} = \beta_{wt} = \beta_t$. The slopes for the Engel curves for men and women in time t are $\eta_{mt}\beta_t$ and $\eta_{wt}\beta_t$, respectively, with $\eta_{mt} + \eta_{wt} = 1$. By taking the ratio of the slopes, the β_t cancel out and we can solve for η_{mt} and η_{wt} . The parameters η_{3j} and η_{4j} are identified through a sort of decomposition of η_{jt} across flooded and unflooded villages within time t .

In Figure A.2, the top panel represents the hypothetical Engel curves across three time periods in flooded villages only. The bottom panel illustrates the hypothetical Engel curves in untreated villages. The steeper the Engel curve, greater is the resource share of that person type.

At baseline, graph (a) shows that the Engel curve for men is steeper than that for women, implying that men have a higher resource share than women. Since floods are randomly assigned in this example, graph (d) shows the corresponding Engel curves at baseline for households in untreated villages, and mimic the same pattern as treated villages.

By the second time period, the Engel curve for women becomes flatter in flooded villages, while it gets steeper for men, implying that women's resource share decreases. In unflooded villages, the opposite happens: women's Engel curve becomes steeper relative to baseline and men's Engel curve gets relatively flatter and they are both parallel to one another implying equal resource shares. Comparing j 's resource share in (b) and (e) identifies η_{3j} .

By the time of the last survey, the slope for women's Engel curve decreases further in flooded villages, while it increases for men, implying that women's resource share declines. In unflooded villages, the opposite happens: women's Engel curve becomes steeper relative to the short-run and men's Engel curve gets relatively flatter. In fact, in this hypothetical example, women have higher resource share than men by the third time period. Comparing j 's resource share in (c) and (f) identifies η_{4j} .




Figure A.2: Hypothetical Engel curves for men and women (indexed as j) across households in flooded and unflooded villages in the 2014 cohort, for three time periods. In each graph, the resource shares of men and women are identified under the assumption that $\beta_{jt} = \beta_t$. Comparing the resource share for j across (b) and (e) identifies η_{3j} . Similarly, comparing the resource share for j across (c) and (f) identifies η_{4j} .

A.3 Identification in a fully specified example

I build on the fully specified model in Dunbar, Lewbel and Pendukar (2013) and Sokullu and Valente (2021), but specify $F_{jt}(p_t)$ differently. Assume the following indirect utility function:

$$V_{jt}(p_t, y_t) = \exp \left[\ln \left[\ln \left(\frac{y_t}{G_{jt}(p_{jt}, \tilde{p}_1, \tilde{p}_{2t})} \right) \right] + \underbrace{\frac{p_{jt}}{e^{a'_1 \ln \tilde{P}_1}} + \frac{\tilde{p}_{21t}}{e^{a'_{2jt} \ln \tilde{P}_{2t}}}}_{F_{jt}(p_t)} \right], \quad (\text{A.2})$$

where p_{jt} is the price of the private assignable good (PAG) for person j at time t , all other goods are split into two categories: $\tilde{p}_1$ denotes the prices of K_1 goods that have time-invariant prices, and $\tilde{p}_2$ denotes prices of goods K_2 whose prices vary across time. $\tilde{P}_1$ is a vector of time-invariant prices, while $\tilde{P}_{2t}$ is a vector of prices that vary with time. As $F_{jt}(p_t)$ is homogenous of degree 0, a'_1 and a'_{2jt} are constant vectors with all elements adding up to 1. $G_{jt}(p_t)$ is a person specific expenditure deflator function that is homogenous of degree 1. Given that $F_{jt}(p_t)$ and $G_{jt}(p_t)$ are homogenous of degree 0 and homogenous of degree 1, respectively, it means that if all prices double (without any change in y_t), for example, then that change will affect $V_{jt}(p_t, y_t)$ only through the expenditure deflator function.

As mentioned in equation (1.8), the Engel curves for the PAG for person j takes the following form:

$$W_{jt} = \eta_{jt} \alpha_{jt} + \beta_{jt} \eta_{jt} \ln \eta_{jt} + \beta_{jt} \eta_{jt} \ln y_t,$$

where $\alpha_{jt} = p_{jt} \left(\frac{G'_{jt}}{G_{jt}} + F'_{jt} \ln G_{jt} \right)$ and $\beta_{jt} = -p_{jt} F'_{jt}$. Recall that G'_{jt} and F'_{jt} are first derivatives with respect to the price of the private assignable good for person j . From the expressions of α_{jt} and β_{jt} we can see that a change in price of the non-PAG goods *over time* will affect the Engel

curve for the PAG only through person j 's expenditure deflator G_{jt} and captured in α_{jt} (sort of an income effect).

As mentioned in Section 1.3.2, I impose the a restriction on the functional form of β_{jt} , such that $\beta_{jt} = \beta_t$, to identify η_{jt} in the system of Engel curves in equation (1.9). This is equivalent to applying the Similarity Across People (SAP) restriction in Dunbar, Lewbel, and Pendukar (2013). Given the functional form of $F_{jt}(p_t)$ in equation (A.2), the derivative with respect to the price of the PAG does not vary with person j :

$$\frac{\partial F_{jt}(p_t)}{\partial p_{jt}} = \frac{1}{e^{a'_1 \ln \tilde{P}_1}} \quad (\text{A.3})$$

Secondly, person-level food expenditure serves as the designated PAG for men, women, and children. Given that food items have the same price regardless of the consumer, p_{jt} varies over time (t) but remains the same across persons (j). Therefore, in the context of this study, it is plausible that $\beta_{jt} = \beta_t$.

The household maximizes the following problem:

$$\begin{aligned} \max_{z, x_1, \dots, x_N} \quad & \sum_{j=1}^J \mu_{jt} [U_{jt}(x_{jt}) + \rho_{jt}], \\ \text{s.t.} \quad & z_t = A \sum_{j=1}^J x_{jt} \quad \text{and} \quad y_t = z'_t p_t. \end{aligned} \quad (\text{A.4})$$

where $\mu_{jt} = \mu_{jt}(p)$ and $\rho_{jt} = \rho_{jt}(p)$ are homogeneous of degree zero, so, for example, $\mu_{jt}(p_t) = \mu_{jt}(p_t/y_t)$. μ_{jt} is the Pareto weight and ρ_{jt} is a utility transfer function, allowing individuals to receive utility from another household member's utility.

Given Pareto efficiency, the household allocation problem is decentralized. In the first

step, household members allocate the total budget among themselves, creating “shadow” budget constraints for each member. In the second step, members use these shadow budgets, along with their individual preferences, to demand a vector of consumption quantities that maximizes their utility (at shadow prices accounting for the extent to which goods are shared). Because the allocation problem is decentralized in this way, we can reformulate the utility maximization problem as below. This reframing makes it more evident how changes in Pareto weights due to climate shocks alter the household’s objective function, ultimately affecting the optimal allocation of resources.

$$\max_{\eta_f, \eta_m, \eta_c} \quad \rho_t + \sum_{j=m,w,c} \mu_{jt} V_{jt}(A'p_t, \eta_{jt} y_t) \quad \text{s.t.} \quad \eta_m + \eta_w + \eta_c = 1 \quad (\text{A.5})$$

where $\rho_t = \sum_{j=1}^J \mu_{jt} \rho_{jt}$ and $V_{jt}(A'p_t, \eta_{jt} y_t)$ is a person specific indirect utility, at shadow price $A'p_t$, and shadow budget $\eta_{jt} y_t$. The expression in equation (A.2) can be re-written as:

$$V_{jt}(A'p_t, \eta_{jt} y_t) = \ln\left(\frac{\eta_{jt} y_t}{G_{jt}(A'p_t)}\right) \exp(p_{jt} e^{-a'_1 \ln \tilde{A}_1 \tilde{P}_1}) \exp(A_1 \tilde{p}_{21t} e^{-a'_{2jt} \ln \tilde{A}_{2t} \tilde{P}_{2t}}) \quad (\text{A.6})$$

Plug this into the maximization problem in equation (A.5):

$$\max_{\eta_f, \eta_m, \eta_c} \quad \rho_t + \sum_{j=m,w,c} \tilde{\mu}_{jt} \ln\left(\frac{\eta_{jt} y_t}{G_{jt}}\right) \quad \text{s.t.} \quad \eta_m + \eta_w + \eta_c = 1 \quad (\text{A.7})$$

where $\tilde{\mu}_{jt} = \mu_{jt} \exp(p_{jt} e^{-a'_1 \ln \tilde{A}_1 \tilde{P}_1}) \exp(A_1 \tilde{p}_{21t} e^{-a'_{2jt} \ln \tilde{A}_{2t} \tilde{P}_{2t}})$ and $\rho_t = \sum_{j=1}^J \mu_{jt} \rho_{jt}$.

After deriving the first order conditions, the optimal allocation of η_{jt} is:

$$\eta_{jt}^* = \frac{\tilde{\mu}_{jt}}{\tilde{\mu}_{mt} + \tilde{\mu}_{wt} + \tilde{\mu}_{ct}} \quad (\text{A.8})$$

Equation (A.8) shows that a shock to the Pareto weights – which can be considered a measure of a person’s bargaining power – will manifest in a change in the optimal share of household resources assigned to that person.

A.4 Additional details on data and sample selection

A.4.1 Cleaning household panel and categorizing individuals

Utilizing all three rounds of the BIHS, I initially had a baseline sample of 5,543 households, of which 4,603 were surveyed across all three rounds. Around 8% of the baseline sample was lost after either round 1 or round 2, 1.6% of the baseline sample was interviewed in rounds 1 and 3 but missing in round 2, and 7.7% of households in round 1 split into different households by time of the last survey. After removing outliers, such as households with more than five children and those with total household expenditures above the 99th percentile, as well as dropping households with missing values for food budget shares – which will be the dependent variable when estimating the system of Engel curves in equation 13 – the sample reduced to 3,010 households located in 275 villages. From this subset, I focus on the 1,596 households containing at least one man, one woman, and one child.

To consistently track resource shares of the same individuals over time, I define adults as those aged 16 or above at baseline (2011) and children as those aged 15 or below at baseline. By the time of the midline survey in 2015, this threshold changes, categorizing children as individuals below the age of 18. By the endline survey in 2018, children are defined as those up to 21 years old. The reason for adjusting the age threshold across surveys is to ensure that I am categorizing the same individuals as men, women, and children in each survey round.

If I were to use a flat 15-year cutoff for children across all survey rounds, individuals categorized as children at baseline could age out of this threshold, leading their resource shares to be mixed with those of their parents. This would contaminate the results for adult men and women in the household by introducing a different demographic.

Although it is theoretically possible to track individuals over time using member ID numbers assigned in the household survey, this approach is impractical here. Household composition can change over time, not only due to new births but also due to other relatives coming to stay permanently in the household. In this context of extended households, multiple adult men and women might reside together. Therefore, as household composition changes across different survey rounds, I needed a uniform way to categorize men, women, and children, ensuring that new permanent members of the household are correctly included in a person type category.²

A.4.2 Additional geographic data

Figure A.3 covers four main climatic and geographic features relevant to flooding, to show the variation in these factors across the country. Panel A depicts the average annual rainfall at the village level using data over 20 years (2000 to 2020). Villages on the eastern side of the country experience relatively higher rainfall per year. Rainfall data were collected from the Climate Hazards Group InfraRed Precipitation with Station (CHIRPS) database that combines satellite and station-level information to provide measures of rainfall over a 5 km grid.

As a measure of vegetation cover, average monthly Normalized Difference Vegetation Index (NDVI) at the village level is presented in panel B of Figure A.3. The NDVI is a measure of

²This is one of the reasons I use food expenditures, rather than clothing expenditures, to estimate resource shares. Food expenditures are available in the household survey at the individual level, allowing for adjustments based on changing age thresholds over time.




Figure A.3: Panel A shows average annual rainfall for each village from 01/2000 to 12/2020. Values in the legend denote percentiles. Panel B shows the average monthly Normalized Difference Vegetation Index (NDVI) at the village level as a proxy for vegetation cover from 1/2000 to 12/2018. Values in the legend denote percentiles. Panel C depicts the elevation of the country in meters. Panel D highlights areas that are prone to tidal, flash and river flooding as well as the degree of severity. High, moderate and low risk areas are denoted by the letters “H”, “M” and “L” respectively. Areas that are not prone to flooding are highlighted in pastel yellow. Gray lines depict inland perennial water. Panels A and B show unions, not villages.

greenness, so a higher NDVI should correspond to areas with greater vegetation. Monthly NDVI was collected using satellite imagery from January 2000 to December 2018 at 5 km resolution from the NOAA AVHRR Surface Reflectance product.³

Panel C shows the elevation in meters.⁴ The relatively lower elevation of the northeastern region compared to surrounding areas is of particular importance as it forms a saucer shaped shallow depression, making it prone to flash flooding. This geographic feature could explain why this region appears flooded in Figure B.2.

Lastly, panel D shows a flood risk map developed by the Bangladesh Agricultural Research Council (BARC) which outlines the areas prone to river flooding, flash flooding or tidal flooding. The map also details whether this risk is high, medium or low.⁵

A.4.3 Narrowing sample based on ex-ante flood risk

The criteria to refine the sample to a subset of villages with comparable flood risk is as follows: the maximum allowable difference in ex-ante flood risk between flooded and unflooded villages is 0.3, they should be located within the same geographic area (i.e. within a 100km grid). Each flooded villages can be associated with up to 3 similar unflooded villages. The resulting sample narrows down to 118 villages, 49 of which experienced a flood during 2012-2018 and 69 villages that had a similar risk of flooding and shared the same geographic area.

The 2014, 2016, and 2018 cohorts form the largest cohorts in the selected sample. This makes sense as figure 1.2 shows that these cohorts together make up 75% of all flooded villages in this sample, which is expected since 84 of the 99 flooded villages in the full sample belonged

³https://developers.google.com/earth-engine/datasets/catalog/NOAA_CDR_AVHRR_NDVI_V5#description

⁴Obtained from <https://data.humdata.org/dataset/bangladesh-contour-lines>

⁵Obtained from <https://data.humdata.org/dataset/bangladesh-hazards>

to these cohorts.

A.4.4 Trends in rainfall in flooded and unflooded villages

The equation used to produce the event study graphs in Figure 1.4 is below:

$$Rainfall_{vct} = \phi_0 + \sum_{t=2011}^{2018} \phi_t (Flood_{vc} \times Year_t) + \gamma_t + \epsilon_{vct} \quad (A.9)$$

$Rainfall_{vct}$ represents the standardized annual rainfall in village v in year t for cohort c , where c can be 2014, 2016, or 2018. γ_t are year fixed effects, and $Flood_{vc}$ is a binary variable equal to one if village v belonging to cohort c experienced a flood since 2011 until the endline survey in 2018. The interaction of $Flood_{vc}$ with a binary $Year_t$ variable allows us to estimate the difference in rainfall in flooded and unflooded villages in the years preceding and following the first flood.

All parameters are estimated separately for each cohort. Estimates at a given point on the event timeline are computed by averaging the ϕ_t parameter across cohorts, weighted by the number of flooded villages in that cohort. For example, the effect observed at event time 0 is the weighted average of ϕ_{2014} for the 2014 cohort, ϕ_{2016} for the 2016 cohort, and ϕ_{2018} for the 2018 cohort, with greater weight given to estimates from the cohort with a larger number of flooded villages.

In the full sample, there are significant differences in rainfall between flooded and unflooded villages in the years preceding and following the first flood. Villages that experienced a flood had been experiencing increasingly higher rainfall than unflooded villages in the years before the flood, shown in Figure A.4, and this continues after the event.

Once the sample is narrowed to a subset of villages that have a similar ex-ante risk of

flooding and that are located within the same 100km grid, this pre- and post-trend is eliminated (Panel A in Figure 1.4). The point estimates become even closer to zero if I remove three flooded villages (and four associated unflooded villages) in Sylhet division that experienced exceptionally high rainfall relative to surrounding villages (Panel B in Figure 1.4). Narrowing the sample helps to remove pre-existing differences in other weather variables too, such as land surface temperature, shown in Figure A.5.




Figure A.4: Differences in standardized annual rainfall in flooded and unflooded villages in the years before and after their first flood during the study period, for villages in the 2014, 2016, and 2018 cohorts prior to narrowing down sample based on ex-ante flood risk. Shaded area represents 95% confidence intervals.




Figure A.5: Differences in standardized annual temperature in flooded and unflooded villages in the years before and after their first flood during the study period. Panel A shows these differences in the full sample for the 2014, 2016, and 2018 cohorts. Panel B shows these differences in a refined subsample of villages that have a similar ex-ante probability of flooding and are located in the same 100km grid. Households in this sample of villages are used to produce the main results. Panel C contains a further reduced sample of villages, excluding three treated villages (and associated control villages) in the Sylhet division that experienced some of the highest rainfall during the study period relative to surrounding villages. This reduced sample is used in a robustness check. Shaded area represents 95% confidence intervals.

Appendix B: Additional tables and figures for Chapter 1



Figure B.1: Regions in Bangladesh. Bangladesh is divided into 8 level 1 administrative divisions. I grouped these administrative divisions into four broad regions.



Figure B.2: Blue shading depicts the maximum flood extent for one event from 8/20/2014 to 9/08/14 from Global Flood Database.




Figure B.3: Distribution of the share of village area flooded for all 20 flood events that occurred since the start of the baseline survey to the endline. Vertical dashed line is the mean at 15.35%.




Figure B.4: Frequency of floods in the 179 villages that experienced at least one flood from 2000 to the third round of the BIHS survey in 2018/19. The mean number of floods for these “treated” villages is 11.02.




Figure B.5: Average number of times villages in each bin were treated during 2012-2018. Error bars indicate 95% confidence intervals.




Figure B.6: Villages that actually experienced a flood in 2012-2018 and were successfully paired with unflooded villages shown in green based on the selection criteria. To protect village confidentiality, unions in which villages are located are depicted instead of village GPS coordinates to protect village confidentiality.




Figure B.7: Balance check for flooded and unflooded villages in 2014, 2016, and 2018 cohorts at baseline. Figure shows the point estimates of the binary flood/treatment variable from a series of regressions of the village-level outcome variables listed on the vertical axis on the treatment variable. Household-level variables are continuous and have been standardized. Flood probability is continuous and estimated using the probit model described in Section 1.4.4. Region and flood-prone indicators (flash or river flooding) are binary variables. Error bars indicate 95% confidence intervals.




Figure B.8: Percentage distribution of men and women across their main earning occupations, excluding non-earning categories (e.g., student, housewife, jobless, and retired). The survey provided data on 72 earning occupations, which I grouped into 10 categories: salaried jobs, agricultural and non-agricultural labor, five farming-related categories (e.g., livestock, poultry, own farm, sharecropping, other farming), self-employment, and trade/production. Percentages are calculated relative to the total number of men and women in earning occupations, with bars in navy for men and red for women. Sample size: 391 households.




Figure B.9: Distribution of food expenditure shares of men, women and children in each time period. Blue and gray histograms refer to flooded and unflooded households, respectively. Sample size: 391 households.




Figure B.10: Estimates of resource shares of men, women and children against deciles of total household expenditure across all rounds. Shaded area represents 95% confidence intervals. Sample size: 391 households.




Figure B.11: Kdensity of the total shadow budget allocated to all men, women and children in each time period (not by individual). All shadow budgets are in terms of 2017 PPP\$. Sample size: 391 households.




Figure B.12: Difference in the price (in logs) of the commodity on the y-axis across treated and untreated villages in 2014 sample only. Sample size: 391 households. Standard errors are clustered at the village level. Bars represent 90% confidence intervals.




Figure B.13: Effect of a flood on resource shares of men and women across models with varying household covariates. Sample size for models 1-6 is 391 households, and 375 households for model 7. Main estimates in Table B.2 are those corresponding to model 1. Models 1,6,7 estimate 102 parameters. Models 2-4 estimate 108 parameters. Model 5 estimates 126 parameters. In models 2 - 4, I add one of the following binary variables: whether the household only had sons, whether the household head was female, whether the village experienced more than the mean number of floods prior to 2011. Model 5 includes four region dummies. Model 6 adjusts the food expenditure shares to exclude prepared foods. Model 7 utilizes data from a subset of the main sample used to produce panel B in Figure 1.4. Standard errors are clustered at the village level. Bars represent 90% confidence intervals.




Figure B.14: Effect of a flood on resource shares varies with the maximum allowable difference in ex-ante probability of experiencing a flood. As the difference reduces, the criteria for untreated villages becomes more strict and the number of villages included in the estimation decreases. There are 391 households in 59 villages when the allowable difference is 0.3. This drops to 337 households in 50 villages when the allowable difference is 0.1. The model did not converge when the difference was set to 0.18-0.20 (at the starting values used in the preferred model in Table B.2). Standard errors are clustered at the village level. Shaded region represents 90% confidence intervals.




Figure B.15: The outcome variable is standardized men's monthly earnings. Controls to increase precision include the number of men, women, and children at baseline, the baseline outcome variable, a binary variable indicating whether a woman had completed at least primary education at baseline, and whether she was over the age of 40. Bars are 90% confidence intervals.




Figure B.16: The outcome variable in panel A is a binary variable for whether any child in the household is employed. In Panel B, it is a binary variable for whether any child in the household does self-employed work. Controls to increase precision include the number of men, women, and children at baseline, the baseline outcome variable, a binary variable indicating whether a woman had completed at least primary education at baseline, and whether she was over the age of 40. Bars are 90% confidence intervals.

Table B.1: Additional descriptive statistics for 2014 sample

	All	Baseline	Midline	Endline
Household size	4.51 (1.16)	4.43 (1.14)	4.56 (1.14)	4.56 (1.19)
Fraction of female children	0.47 (0.37)	0.47 (0.39)	0.46 (0.36)	0.47 (0.35)
Mean age of women	36.54 (9.02)	33.27 (8.67)	36.65 (8.77)	39.71 (8.47)
Mean age of men	41.63 (10.10)	38.25 (9.58)	41.58 (9.67)	45.06 (9.90)
Mean age of children	9.13 (4.19)	7.10 (3.52)	9.18 (3.81)	11.11 (4.21)
Max years of education of men	3.39 (4.05)	3.40 (4.05)	3.47 (4.11)	3.30 (3.99)
Max years of education of women	3.49 (3.66)	3.51 (3.66)	3.47 (3.59)	3.50 (3.73)
Max years of education of children	4.23 (3.45)	2.43 (2.61)	4.45 (3.28)	5.81 (3.51)
Value of cons. assets (1000 Taka)	38.92 (49.19)	26.09 (40.48)	36.80 (47.46)	54.04 (54.67)
Value of prod. assets (1000 Taka)	17.27 (30.60)	2.12 (9.56)	18.16 (26.79)	31.84 (39.70)
Value of ag. assets (1000 Taka)	3.28 (6.96)	2.39 (7.31)	3.24 (5.82)	4.20 (7.52)
Observations	1173	391	391	391

Notes: Descriptive statistics for households in selected flooded and unflooded villages in the 2014 sample. Mean with standard deviations in parentheses. Computed from the 2011/12, 2015, and 2018/19 BIHS surveys.

Table B.2: The determinants of resource shares of men and women

	(1) Women	(2) Men
Constant	0.272*** (0.037)	0.438*** (0.052)
1(Midline)	0.001 (0.043)	-0.134*** (0.040)
1(Endline)	0.036 (0.045)	-0.134*** (0.048)
1(Flood)×1(Midline)	-0.085* (0.046)	0.112** (0.049)
1(Flood)×1(Endline)	-0.069 (0.053)	0.115* (0.066)
1(Child under 5 at baseline)	0.055** (0.022)	0.002 (0.020)
Number of children at baseline	-0.034*** (0.011)	-0.058*** (0.012)
Number of adult men at baseline	-0.037* (0.021)	0.060*** (0.022)
Number of adult women at baseline	0.101*** (0.028)	-0.095*** (0.029)
1(Children increase)	-0.027 (0.017)	-0.029 (0.020)
1(Children decrease)	0.045 (0.030)	0.068** (0.031)
1(Men increase)	-0.110** (0.046)	0.205*** (0.042)
1(Men decrease)	0.076 (0.051)	-0.157*** (0.055)
1(Women increase)	0.103** (0.049)	-0.068 (0.063)
1(Women decrease)	-0.152*** (0.045)	0.227*** (0.051)
1(Woman educated at baseline)	0.016 (0.022)	0.039* (0.023)
1(Woman of age 40+)	-0.005 (0.023)	0.039 (0.028)
Observations	1173	1173

Notes: Estimates for the determinants of resource shares of men and women in the household, omitting estimates of the taste parameters for brevity (a total of 102 parameters are estimated in this model). The results are derived from the estimation of the system of Engel curves from section 1.3.4, controlling nonlinearly for 17 covariates: a constant, time fixed effects, interactions of the binary flood variable with post-treatment time dummies, the demeaned number of men, women, and children at baseline, six dummies for changes in household composition since baseline, and binary variables indicating whether the household had a child under five at baseline, whether at least one woman had completed primary education at baseline, and whether at least one woman was over the age of 40. The sample size is 391 households across 59 villages. Standard errors are clustered at the village level and are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.3: Effect of number of floods on resource shares

	All treated villages		Without top quantile	
	(1) Women	(2) Men	(3) Women	(4) Men
Constant	0.289*** (0.046)	0.456*** (0.062)	0.296*** (0.041)	0.424*** (0.056)
1(Midline)	0.046 (0.048)	-0.188*** (0.049)	0.029 (0.045)	-0.177*** (0.041)
1(Endline)	0.011 (0.057)	-0.164** (0.064)	-0.014 (0.052)	-0.135** (0.053)
#Flood ₂₀₁₅ × 1(Midline)	-0.069*** (0.012)	0.088*** (0.019)	-0.063*** (0.014)	0.087*** (0.024)
#Flood ₂₀₁₈ × 1(Endline)	-0.015*** (0.003)	0.025*** (0.005)	-0.011 (0.007)	0.018** (0.008)
Observations	1173	1173	1095	1095

Notes: Main parameter estimates for men and women in the household. A total of 17 covariates were included: a constant, time fixed effects, interactions of the number of floods with each post-treatment time dummy, the demeaned number of men, women, and children at baseline, six dummies for changes in household composition since baseline, and binary variables indicating whether the household had a child under five at baseline, whether at least one woman had completed primary education at baseline, and whether at least one woman was over the age of 40. Columns 1 and 2 report estimates for all treated villages in the selected sample (391 households in 59 villages). Columns 3 and 4 report estimates for a truncated sample excluding villages that experienced 9 to 11 floods, which form the top quantile (365 households in 56 villages). A total of 102 parameters are estimated in both set of estimates. Standard errors are clustered at the village level and are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.4: Resource shares in reference households and the effect of number of floods

	Reference household			Flooded household			(7) Δ Children's share
	(1) Women	(2) Men	(3) Children	(4) Women	(5) Men	(6) Children	
Baseline	0.289*** (0.046)	0.456*** (0.062)	0.255*** (0.043)	0.289*** (0.046)	0.456*** (0.062)	0.255*** (0.043)	
Short-run	0.335*** (0.036)	0.268*** (0.033)	0.397*** (0.048)	0.190*** (0.031)	0.453*** (0.051)	0.357*** (0.048)	-0.039 (0.036)
Long-run	0.300*** (0.037)	0.291*** (0.035)	0.408*** (0.048)	0.215*** (0.039)	0.434*** (0.045)	0.339*** (0.055)	-0.057** (0.027)

Notes: Resource shares of men, women, and children in reference households and flooded households. A reference household is located in a control village, with one man, one woman and two children, and all binary household covariates are set to zero. Columns 4-6 report the resource shares in a reference household in a treated village if they experienced the mean number of floods by the survey year (2.1 floods by 2015 and 5.6 floods by 2018). Column 7 reports the effect of the average number of floods on children's resource shares. Standard errors are clustered at the village level and are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Table B.5: Resource shares of all households

	Unflooded			Flooded		
	Women	Men	Children	Women	Men	Children
Baseline	0.32 (0.06)	0.46 (0.07)	0.22 (0.10)	0.33 (0.06)	0.45 (0.08)	0.22 (0.10)
Short-run	0.32 (0.06)	0.33 (0.08)	0.35 (0.09)	0.23 (0.06)	0.44 (0.08)	0.33 (0.10)
Long-run	0.35 (0.07)	0.34 (0.09)	0.31 (0.09)	0.28 (0.05)	0.46 (0.09)	0.26 (0.09)
Observations	798	798	798	375	375	375

Notes: Descriptive statistics of resource shares of men, women, and children in all households in the selected sample. Means are interpreted as the average resource share of a person type in a given time period. Standard deviations are in parentheses. Sample size: 391 households.

Table B.6: Other mechanisms

	Change in # of men			Age youngest child			Works outside village			Transport exp.			Housing exp.		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Panel A															
Flood $\times$ <i>Midline</i>	0.006 (0.025)	0.002 (0.023)	0.002 (0.023)	-0.120 (0.613)	0.387 (0.383)	0.306 (0.402)	-0.091 (0.064)	-0.073 (0.063)	-0.055 (0.066)	-0.113 (0.099)	-0.142 (0.127)	-0.160 (0.135)	-0.150 (0.100)	-0.157 (0.098)	-0.140 (0.101)
Flood $\times$ <i>Endline</i>	-0.016 (0.031)	-0.021 (0.029)	-0.025 (0.027)	-0.432 (0.846)	0.064 (0.691)	-0.065 (0.724)	-0.059 (0.073)	-0.041 (0.060)	-0.046 (0.060)	-0.123 (0.127)	-0.216* (0.120)	-0.270** (0.114)	-0.224** (0.110)	-0.246** (0.113)	-0.249** (0.117)
Panel B															
#Floods ₂₀₁₅ $\times$ <i>Midline</i>	0.001 (0.012)	-0.002 (0.011)	-0.002 (0.011)	0.086 (0.263)	0.236 (0.141)	0.207 (0.147)	-0.053* (0.029)	-0.046 (0.030)	-0.039 (0.031)	-0.069* (0.035)	-0.070 (0.048)	-0.076 (0.050)	-0.055 (0.042)	-0.053 (0.043)	-0.046 (0.044)
#Floods ₂₀₁₈ $\times$ <i>Endline</i>	-0.002 (0.004)	-0.003 (0.004)	-0.004 (0.003)	0.012 (0.128)	0.053 (0.092)	0.039 (0.099)	-0.004 (0.009)	-0.003 (0.008)	-0.004 (0.008)	-0.021 (0.019)	-0.032* (0.019)	-0.039* (0.020)	-0.018 (0.023)	-0.019 (0.024)	-0.018 (0.025)
Observations	1173	1173	1125	1173	1173	1125	1173	1173	1125	1094	1020	975	1162	1160	1112
Controls		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Robustness sample			Yes			Yes			Yes			Yes			Yes

Notes: Estimates for households in treated and untreated villages in the 2014 sample only. Outcome variables mentioned on the headers are: change in the number of men in the household since baseline, age of the youngest child in the household (in years), a binary variable for whether any man in the household worked outside the village, standardized monthly transportation expenses of the household, and standardized annual housing-related expenses of the household. Panel A presents estimates using a binary treatment variable, while Panel B shows results based on the number of floods experienced by the midline or endline surveys. Controls include the number of men, women, and children at baseline, the baseline outcome variable, a binary variable indicating whether a woman had completed at least primary education at baseline, and whether she was over the age of 40. Reduced sample refers to a subset of the main sample used to produce Panel B in Figure 1.4. Standard errors are clustered at the village level and are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.7: Effects on the 2016 and 2018 cohorts

	All treated villages		Robustness sample	
	(1) Women	(2) Men	(3) Women	(4) Men
Constant	0.222*** (0.046)	0.485*** (0.067)	0.233*** (0.055)	0.465*** (0.063)
1(Endline)	0.19*** (0.059)	-0.11* (0.060)	0.17*** (0.058)	-0.11* (0.060)
Flood $\times$ 1(Endline)	-0.067** (0.030)	0.081** (0.032)	-0.047 (0.029)	0.076** (0.030)
Observations	450	450	432	432

Notes: Main parameter estimates for men and women in households located in villages in the 2016 and 2018 cohorts with their corresponding untreated villages that do not overlap with the main sample. This estimation uses only data from the midline and endline surveys, using 2015 as the base year. In both set of estimates, 90 parameters are estimated. A total of 15 covariates were included: a constant, one post-treatment dummy, interaction of the binary flood variable with the post-treatment time dummy, the demeaned number of men, women, and children at baseline, six dummies for changes in household composition since baseline, and binary variables indicating whether the household had a child under five at baseline, whether at least one woman had completed primary education at baseline, and whether at least one woman was over the age of 40. Columns 1 and 2 report estimates for all treated villages in the selected sample (225 households in 37 villages). Columns 3 and 4 report estimates for a reduced sample of villages used in Panel B of Figure 1.4 (216 households in 35 villages). Standard errors are clustered at the village level and are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.1 Additional tables and figures for Chapter 2




Figure B.17: Correlation between height-for-age z-score (HAZ) and heat exposure during pregnancy, for children under 2 years. Scatter depicts mean HAZ by number of heat wave days in-utero, where each bin is 2 days. Each point reflects the average HAZ of children under age 2 that fall in each bin. Dotted lines indicate 95% confidence interval associated with each point. The dash line represents a Lowess smoothing line.




Figure B.18: Variation in average temperature (in °C) during pregnancy across Bangladesh's 64 districts, by season of birth.




Figure B.19: Trends in residualized height-for-age z-score (HAZ) by birth year. Residualized HAZ obtained by taking residuals of the regression of HAZ on number of heat-wave days and precipitation, child's age in months and sex, as well as district-month fixed effects.




Figure B.20: Share of births by season, after adjusting length of seasons to four months each. Shares calculated by dividing the number of births in a particular season in a given birth year, by the total number of births in that birth year.




Figure B.21: Exploring adaptive behavior in fertility decisions. Panel A: Variation in mean number of heatwave days and mean average temperature during the gestation period across the first five births. Gray and blue error bars reflect the 95% confidence intervals. Panel B: Share of births in each season, based on whether the previous birth was in the winter season, which has the highest exposure to heat in-utero. A mother-level dataset was created for both graphs, and only mothers who gave birth to more than one child and whose first birth was covered by the BIHS survey were included. Therefore, these graphs are generated using a restricted sample of 1,763 children born to 776 mothers.




Figure B.22: Residual variation in temperature variables after netting out fixed effects, but only for children whose location of birth is the same as their district of residence.

Table B.8: Effect of interventions on children's height

Study	Type	Age at baseline	Exposure	Country	HAZ	Height (in cm)
Leroy et al. (2008)	CCT	< 6 mons	2 years	Mexico	0.41**	1.5**
Behrman and Hoddinott (2005)	CCT	12-36 mons	1 year	Mexico		1.016**
Attanasio et al. (2005)	CCT	< 24 mons	1 year	Colombia	0.161*	
Penny et al. (2005)	N	newborns	6 months	Peru	0.272***	0.714**
Lutter et al. (2007)	N	12-14 mons	11 months	Ecuador		0.66

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. CCT: conditional cash transfer program; N: nutrition intervention; HAZ: height-for-age z-score. This table compares the effect size of different interventions on the height of children less than 5 years.

Table B.9: Main results without standardizing temperature variables

	HAZ			Stunted		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
No. of heat-wave days during pregnancy	-0.001 (0.002)	-0.002 (0.002)	-0.002 (0.002)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)
Panel B						
Avg temperature during pregnancy(°C)	-0.007 (0.031)	-0.138* (0.078)	-0.131* (0.077)	-0.003 (0.010)	0.050* (0.026)	0.049* (0.026)
Observations	5602	5602	5590	5602	5602	5590
District-birthmonth FE	No	Yes	Yes	No	Yes	Yes
District-season FE	Yes	No	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	No	Yes	No	No	Yes
Sample	Full	Full	Full	Full	Full	Full

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Table reports results of equation (2.1). Temperature variables are in terms of days in panel A, and °C in panel B. Coefficients represent causal effect of one additional heat-wave day, or increase of °C during the pregnancy on the outcomes. Standard errors clustered at the district level. All regressions control for rainfall during the pregnancy, the child's age and sex. Additional covariates include number of boys and girls under the age of 5, the household head's years of education, and the household's access to water and electricity at the time of the survey.

Table B.10: Exploring the agricultural productivity channel (controlling for NDVI)

	HAZ		Stunted	
	(1)	(2)	(3)	(4)
Panel A				
Heat-wave days	-0.031 (0.064)	-0.019 (0.085)	0.014 (0.018)	0.011 (0.024)
Low NDVI	-0.131 (0.083)	-0.178* (0.096)	0.049* (0.025)	0.059** (0.028)
Heat-wave days $\times$ Low NDVI	0.050 (0.074)	0.040 (0.094)	-0.020 (0.019)	-0.032 (0.024)
Panel B				
Avg. temp	-0.034 (0.060)	-0.221 (0.146)	0.005 (0.020)	0.063 (0.039)
Low NDVI	-0.132 (0.084)	-0.173* (0.099)	0.049* (0.026)	0.059** (0.029)
Avg. temp $\times$ Low NDVI	0.031 (0.060)	0.058 (0.073)	0.002 (0.017)	-0.010 (0.019)
Observations	4029	4029	4029	4029
District-birthmonth FE	No	Yes	No	Yes
District-season FE	Yes	No	Yes	No
Year FE	Yes	Yes	Yes	Yes

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Temperature variables are standardized to mean 0 and unit standard deviation. Standard errors clustered at the district level. All regressions control for rainfall during the pregnancy, the child's age and sex. *LowNDVI* is a binary variable equal to 1 if the average NDVI around a household's location during the pregnancy was below the median in the sample. I obtain NDVI through IFPRI, derived from MODIS/Terra Vegetation Indices Monthly L3 Global 0.05 Deg CMG V006. This product has a resolution of approximately 5 km that allows computation of NDVI at the location of a household, using their GPS location. IFPRI provided data on the monthly mean NDVI in a household's vicinity from 2000 to 2019, which allowed me to compute a measure for the average NDVI around a household during the pregnancy.

Table B.11: Main results with Conley standard errors

	HAZ			Stunted		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Heat-wave days during pregnancy	-0.024 (0.030)	-0.037 (0.039)	-0.031 (0.041)	0.001 (0.011)	0.002 (0.016)	0.000 (0.016)
Panel B						
Avg temperature during pregnancy	-0.010 (0.044)	-0.197** (0.083)	-0.186** (0.082)	-0.005 (0.010)	0.072** (0.031)	0.069** (0.031)
Observations	5602	5602	5590	5602	5602	5590
District-birthmonth FE	No	Yes	Yes	No	Yes	Yes
District-season FE	Yes	No	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	No	Yes	No	No	Yes
Sample	Full	Full	Full	Full	Full	Full

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Table reports results of equation (2.1). Temperature variables are standardized to mean 0 and unit standard deviation. Standard errors are adjusted for spatial correlation within 50km of villages in the sample. All regressions control for rainfall during the pregnancy, the child's age in months and sex. Additional covariates include number of boys and girls under the age of 5, the household head's years of education, and the household's access to water and electricity at the time of the survey.

B.2 Additional tables for Chapter 3

Table B.12: Balance test for school characteristics at baseline (2003)

	Non-Program districts		Program districts		Difference
	# schools (1)	Mean / SD (2)	# schools (3)	Mean / SD (4)	
Located in rural area	6870	0.786 [0.410]	4480	0.867 [0.340]	-0.081***
Total enrollment in grades 6–10	6870	249.770 [351.925]	4480	165.796 [267.701]	83.974***
Total enrollment in grades 6–10, girls	6870	105.793 [264.912]	4480	56.680 [178.519]	49.113***
Total enrollment in grades 6–10, boys	6870	143.977 [289.285]	4480	109.116 [227.469]	34.861***
Total teachers	6870	14.229 [10.479]	4480	12.695 [9.021]	1.534***
Student to teacher ratio	6870	0.040 [0.037]	4480	0.045 [0.029]	-0.006***
Number of classes	6860	9.119 [3.954]	4475	9.191 [3.632]	-0.072
Toilet == 1	6860	0.641 [0.480]	4476	0.645 [0.479]	-0.004
Boundary wall == 1	6858	0.711 [0.453]	4477	0.625 [0.484]	0.086***
Electricity == 1	6836	0.831 [0.375]	4463	0.754 [0.431]	0.077***
Drinking water == 1	6862	0.954 [0.209]	4478	0.948 [0.222]	0.006

Note: Testing balance in characteristics of all government secondary schools in program and non-program districts in 2003. The value displayed for t-tests in column (5) are the differences in the means across the groups. ***, **, and * indicate significance at the 1, 5, and 10 percent level. Data from the Punjab government's Annual School Census for 2003.

Table B.13: Validity check of window length for 30% false cutoff

Window length (1)	Minimum balance test p-value (2)	Nonprogram schools n (3)	Program schools n (4)
Girls			
0.7	0.557	71	82
2.7	0.758	178	82
4.7	0.114	178	352
6.7	0.045	178	991
8.7	0.021	178	1494
10.7	0.015	218	1912
12.7	0.017	218	1961
14.7	0.015	218	2388
16.7	0.023	218	2688
18.7	0.024	218	3000
Boys			
0.7	0.000	87	143
2.7	0.000	240	143
4.7	0.401	240	489
6.7	0.363	240	1122
8.7	0.190	240	1662
10.7	0.686	311	2129
12.7	0.583	311	2182
14.7	0.342	311	2697
16.7	0.267	311	2940
18.7	0.269	311	3287

Note: The table presents balance test results using preprogram values for two covariates: number of teachers and number of classes. These results differ from those in Table 3.3 because they are based on a false program assignment using a 30% cutoff. Column 1 presents the window length. Column 2 shows the minimum p-value across covariates. Columns 3 and 4 show school counts in nonprogram and program districts. Balance tests used the `rdrandinf` package by Cattaneo, Titiunik, and Vazquez-Bare (2016). Own estimates based on data from the Punjab government's Annual School Census for 2003.

Bibliography

- [1] Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. 2023. "When Should You Adjust Standard Errors for Clustering?" *The Quarterly Journal of Economics* 138 (1): 1–35. <https://doi.org/10.1093/qje/qjac038>.
- [2] Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham. 2020. "The Light and the Heat: Productivity Co-Benefits of Energy-Saving Technology." *The Review of Economics and Statistics* 102 (4): 779–92.
- [3] Afridi, Farzana, Kanika Mahajan, and Nikita Sangwan. 2022. "The Gendered Effects of Droughts: Production Shocks and Labor Response in Agriculture." *Labour Economics* 78 (October): 102227. <https://doi.org/10.1016/j.labeco.2022.102227>.
- [4] Ahmed, Akhter. 2013. "Bangladesh Integrated Household Survey (BIHS) 2011-2012." Harvard Dataverse. <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/OR6MHT>.
- [5] Ahmed, Akhter U., Michelle Adato, Ayse Kudat, Daniel Gilligan, Terence Roopnaraine, and Rafik Colasan. 2007. "Impact Evaluation of the Conditional Cash Transfer Program in Turkey: Final Report." International Food Policy Research Institute (IFPRI), Washington, DC; and AGRIN, Ankara, Turkey.
- [6] Alam, Andaleeb, Javier E. Baez, and Ximena V. Del Carpio. 2011. "Does Cash for School Influence Young Women's Behavior in the Longer Term? Evidence from Pakistan." Policy Research Working Paper 5669, World Bank, Washington, DC.
- [7] Almond, Douglas, Janet Currie, and Valentina Duque. 2018. "Childhood Circumstances and Adult Outcomes: Act II." *Journal of Economic Literature* 56 (4): 1360–1446. <https://doi.org/10.1257/jel.20171164>.
- [8] Almond, Douglas, Lena Edlund, and Mårten Palme. 2009. "Chernobyl's Subclinical Legacy: Prenatal Exposure to Radioactive Fallout and School

- Outcomes in Sweden*.” *The Quarterly Journal of Economics* 124 (4): 1729–72. <https://doi.org/10.1162/qjec.2009.124.4.1729>.
- [9] Almond, Douglas, and Bhashkar Mazumder. 2011. “Health Capital and the Prenatal Environment: The Effect of Ramadan Observance during Pregnancy.” *American Economic Journal: Applied Economics* 3 (4): 56–85. <https://doi.org/10.1257/app.3.4.56>.
- [10] Aminjonov, Ulugbek. 2025. Intra-household Resource Allocation after Rainfall Shocks. Unpublished manuscript.
- [11] Andalón, Mabel, João Pedro Azevedo, Carlos Rodríguez-Castelán, Viviane Sanfelice, and Daniel Valderrama-González. 2016. “Weather Shocks and Health at Birth in Colombia.” *World Development* 82 (June): 69–82. <https://doi.org/10.1016/j.worlddev.2016.01.015>.
- [12] Anttila-Hughes, J. K., and S. M. Hsiang. 2013. “Destruction, Disinvestment, and Death: Economic and Human Losses following Environmental Disaster.” Working paper, Goldman School of Public Policy, University of California, Berkeley.
- [13] Attanasio, Orazio, Luis Carlos Gómez, Patricia Heredia, and Marcos Vera-Hernández. 2005. “The Short-Term Impact of a Conditional Cash Subsidy on Child Health and Nutrition in Colombia,” Institute of Fiscal Studies: London, UK.
- [14] Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. “Using Weather Data and Climate Model Output in Economic Analyses of Climate Change.” *Review of Environmental Economics and Policy* 7 (2): 181–98. <https://doi.org/10.1093/reep/ret016>.
- [15] Bangladesh Climate Change Trust and UN Women Bangladesh. 2024. Bangladesh Climate Change and Gender Action Plan 2023.
- [16] Baker, Rachel E., and Jesse Anttila-Hughes. 2020. “Characterizing the Contribution of High Temperatures to Child Undernourishment in Sub-Saharan Africa.” *Scientific Reports* 10 (1): 18796. <https://doi.org/10.1038/s41598-020-74942-9>.
- [17] Baqui, A. H., H. E. Rosen, A. C. C. Lee, J. A. Applegate, S. El Arifeen, S. M. Rahman, N. Begum, R. Shah, G. L. Darmstadt, and R. E. Black. 2013. “Preterm Birth and Neonatal Mortality in a Rural Bangladeshi Cohort: Implications for Health Programs.” *Journal of Perinatology: Official Journal of the California Perinatal Association* 33 (12): 977–81. <https://doi.org/10.1038/jp.2013.91>.
- [18] Barker, D J. 1990. “The Fetal and Infant Origins of Adult Disease.” *BMJ: British Medical Journal* 301 (6761): 1111.

- [19] Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph 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): 105–59. <https://doi.org/10.1086/684582>.
- [20] Barreca, Alan, Olivier Deschenes, and Melanie Guldi. 2018. "Maybe Next Month? Temperature Shocks and Dynamic Adjustments in Birth Rates." *Demography* 55 (4): 1269–93. <https://doi.org/10.1007/s13524-018-0690-7>.
- [21] Behrman, Jere R., and John Hoddinott. 2005. "Programme Evaluation with Unobserved Heterogeneity and Selective Implementation: The Mexican PROGRESA Impact on Child Nutrition." *Oxford Bulletin of Economics and Statistics* 67 (4): 547–69. <https://doi.org/10.1111/j.1468-0084.2005.00131.x>.
- [22] Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen. 2015. "Turning a Shove into a Nudge? A 'Labeled Cash Transfer' for Education." *American Economic Journal: Economic Policy* 7: 86–125.
- [23] Bouguen, Adrien, Yue Huang, Michael Kremer, and Edward Miguel. 2018. "Using RCTs to Estimate Long-Run Impacts in Development Economics." NBER Working Paper 25356, National Bureau of Economic Research, Cambridge, MA.
- [24] Bratti, Massimiliano and Frimpong, Prince Boakye and Russo, Simone, Prenatal Exposure to Heat Waves and Child Health in Sub-Saharan Africa. IZA Discussion Paper No. 14424, Available at SSRN: <https://ssrn.com/abstract=3865449>.
- [25] Brown, Caitlin, Rossella Calvi, and Jacob Penglase. 2021. "Sharing the Pie: An Analysis of Undernutrition and Individual Consumption in Bangladesh." *Journal of Public Economics* 200 (August): 104460. <https://doi.org/10.1016/j.jpubeco.2021.104460>.
- [26] Browning, M., P.-A. Chiappori, and A. Lewbel. 2013. "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power." *The Review of Economic Studies* 80 (4): 1267–1303. <https://doi.org/10.1093/restud/rdt019>.
- [27] Browning, M., P. A. Chiappori, and Y. Weiss. 2014. *Economics of the Family*. Cambridge: Cambridge Univ. Press.
- [28] Buckles, Kasey S., and Daniel M. Hungerman. 2013. "Season of Birth and Later Outcomes: Old Questions, New Answers." *Review of Economics and Statistics* 95 (3): 711–24. https://doi.org/10.1162/REST_a.00314.
- [29] Calvi, Rossella. 2020. "Why Are Older Women Missing in India? The Age Profile of Bargaining Power and Poverty." *Journal of Political Economy*.

- [30] Calvi, Rossella, Jacob Penglase, Denni Tommasi, and Alexander Wolf. 2023. “The More the Poorer? Resource Sharing and Scale Economies in Large Families.” *Journal of Development Economics* 160: 102986.
- [31] Cattaneo, Matias D., and Gonzalo Vazquez-Bare. 2016. “The Choice of Neighborhood in Regression Discontinuity Designs.” *Observational Studies* 2: 134–46.
- [32] Cattaneo, Matias D., Nicolas Idrobo, and Rocío Titiunik. 2020. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.
- [33] ———. 2024. *A Practical Introduction to Regression Discontinuity Designs: Extensions of Elements in Quantitative and Computational Methods for the Social Sciences*. Cambridge: Cambridge University Press.
- [34] Cattaneo, Matias D., Rocio Titiunik, and Gonzalo Vazquez-Bare. 2016. “Inference in Regression Discontinuity Designs under Local Randomization.” *Stata Journal* 16 (2): 331–67.
- [35] ———. 2017. “Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality.” *Journal of Policy Analysis and Management* 36 (3): 643–81.
- [36] Chaudhury, Nazmul, and Dilip Parajuli. 2010. “Conditional Cash Transfers and Female Schooling: The Impact of the Female School Stipend Program on Public School Enrollments in Punjab, Pakistan.” *Applied Economics* 42: 3565–83.
- [37] Cheema, Iftikhar, Martina G. Asia, Simon Hunt, Sarah Javeed, Tanya Lone, and Sean O’Leary. 2016. “Benazir Income Support Program: Evaluation of Waseela-e-Taleem Conditional Cash Transfer.” Oxford Policy Management, London.
- [38] Chiappori, Pierre-André. 1988. “Rational Household Labor Supply.” *Econometrica* 56 (1): 63–90. <https://doi.org/10.2307/1911842>.
- [39] ———. 1992. “Collective Labor Supply and Welfare.” *Journal of Political Economy* 100 (3): 437. <https://doi.org/10.1086/261825>.
- [40] Christian, Paul, Eashani Kandpal, Nethra Palaniswamy, and Vijayendra Rao. 2018. “Safety Nets and Natural Disaster Mitigation: Evidence from Cyclone Phailin in Odisha.” Policy Research Working Paper;No. 8375.
- [41] Currie, Janet, and Maya Rossin-Slater. 2013. “Weathering the Storm: Hurricanes and Birth Outcomes.” *Journal of Health Economics* 32 (3): 487–503. <https://doi.org/10.1016/j.jhealeco.2013.01.004>.

- [42] Das, Jishnu, Stefan Dercon, James Habyarimana, Pramila Krishnan, Karthik Muralidharan, and Venkatesh Sundararaman. 2013. “School Inputs, Household Substitution, and Test Scores.” *American Economic Journal: Applied Economics* 5: 29–57.
- [43] Deschênes, Olivier, and Michael Greenstone. 2011. “Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US.” *American Economic Journal: Applied Economics* 3 (4): 152–85. <https://doi.org/10.1257/app.3.4.152>.
- [44] Deschênes, Olivier, Michael Greenstone, and Jonathan Guryan. 2009. “Climate Change and Birth Weight.” *American Economic Review* 99 (2): 211–17. <https://doi.org/10.1257/aer.99.2.211>.
- [45] Dessy, Sylvain, Luca Tiberti, Marco Tiberti, and David Zoundi. 2021. Polygyny and Farm Households’ Resilience to Climate Shocks. Policy Research Working Papers. The World Bank. <https://doi.org/10.1596/1813-9450-9663>.
- [46] Dhaliwal, Iqbal, Esther Duflo, Rachel Glennerster, and Caitlin Tulloch. 2013. “Comparative Cost-Effectiveness Analysis to Inform Policy in Developing Countries: A General Framework with Applications to Education.” In *Education Policy in Developing Countries*, edited by Paul Glewwe. Chicago: University of Chicago Press.
- [47] Doan, Miki Khanh, Ruth Hill, Stephane Hallegatte, Paul Corral, Ben Brunckhorst, Minh Nguyen, Samuel Freije-Rodriguez, and Esther Naikal. 2023. “Counting People Exposed to, Vulnerable to, or at High Risk From Climate Shocks — A Methodology (English).” Policy Research Working Paper ; no. WPS 10619 Washington, D.C. : World Bank Group.
- [48] Doss, C. 2013. “Intrahousehold Bargaining and Resource Allocation in Developing Countries.” *The World Bank Research Observer* 28 (1): 52–78. <https://doi.org/10.1093/wbro/lkt001>.
- [49] Dunbar, Geoffrey R, Arthur Lewbel, and Krishna Pendakur. 2013. “Children’s Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi.” *American Economic Review* 103 (1): 438–71. <https://doi.org/10.1257/aer.103.1.438>.
- [50] ———. 2019. “Identification of Random Resource Shares in Collective Households Without Preference Similarity Restrictions.” *Journal of Business & Economic Statistics* 39 (2): 402–21.
- [51] Farooqi, Aijaz, Bruno Hägglöf, Gunnar Sedin, Leif Gothefors, and
- [52] Farooqi, Aijaz, Bruno Hägglöf, Gunnar Sedin, Leif Gothefors, and Fredrik Serenius. 2006. “Growth in 10- to 12-Year-Old Children Born at 23 to 25 Weeks’ Gestation in the

- 1990s: A Swedish National Prospective Follow-up Study.” *Pediatrics* 118 (5): e1452-1465. <https://doi.org/10.1542/peds.2006-1069>.
- [53] Filmer, Deon, and Norbert Schady. 2008. “Getting Girls into School: Evidence from a Scholarship Program in Cambodia.” *Economic Development and Cultural Change* 56 (3): 581–617.
- [54] Fishman, Ram, Paul Carrillo, and Jason Russ. 2019. “Long-Term Impacts of Exposure to High Temperatures on Human Capital and Economic Productivity.” *Journal of Environmental Economics and Management* 93 (January): 221–38. <https://doi.org/10.1016/j.jeem.2018.10.001>.
- [55] Fiszbein, Ariel, Norbert Schady, Francisco H. G. Ferreira, Margaret Grosh, Niall Keleher, Pedro Olinto, and Emmanuel Skoufias. 2009. *Conditional Cash Transfers: Reducing Present and Future Poverty*. World Bank Policy Research Report. Washington, DC: World Bank.
- [56] Fruttero, Anna, Daniel Halim, Chiara Broccolini, Bernardo Coelho, Horace Gninafon, and Noël Muller. 2023. *Gendered Impacts of Climate Change: Evidence from Weather Shocks*. Policy Research Working Papers. The World Bank. <https://doi.org/10.1596/1813-9450-10442>.
- [57] Garcia, Sandra, and Juan E. Saavedra. 2017. “Educational Impacts and Cost-Effectiveness of Conditional Cash Transfer Programs in Developing Countries: A Meta-Analysis.” *Review of Educational Research* 87 (5): 921–65.
- [58] Garg, Teevrat, Maulik Jagnani, and Vis Taraz. 2020. “Temperature and Human Capital in India.” *Journal of the Association of Environmental and Resource Economists* 7 (6): 1113–50. <https://doi.org/10.1086/710066>.
- [59] Gelman, Andrew, and Guido Imbens. 2014. “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs.” NBER Working Paper 20405, National Bureau of Economic Research, Cambridge, MA.
- [60] Giannelli, Gianna Claudia, and Eugenia Canessa. 2021. “Women’s Employment and Natural Shocks” IZA Discussion Paper No. 14055.
- [61] ———. 2022. “After the Flood: Migration and Remittances as Coping Strategies of Rural Bangladeshi Households.” *Economic Development and Cultural Change* 70 (3): 1159–95. <https://doi.org/10.1086/713939>.
- [62] Glewwe, Paul, and Karthik Muralidharan. 2016. “Improving Education Outcomes in Developing Countries: Evidence, Knowledge Gaps, and Policy Implications.” In *Handbook*

of the Economics of Education, vol. 5, edited by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, 653–743. Amsterdam: Elsevier.

- [63] Goldenberg, Robert L, and Jennifer F Culhane. 2007. “Low Birth Weight in the United States.” *The American Journal of Clinical Nutrition* 85 (2): 584S-590S. <https://doi.org/10.1093/ajcn/85.2.584S>.
- [64] Graff Zivin, Joshua, Solomon M. Hsiang, and Matthew Neidell. 2018. “Temperature and Human Capital in the Short and Long Run.” *Journal of the Association of Environmental and Resource Economists* 5 (1): 77–105. <https://doi.org/10.1086/694177>.
- [65] Grantham-McGregor, Sally, Yin Bun Cheung, Santiago Cueto, Paul Glewwe, Linda Richter, and Barbara Strupp. 2007. “Developmental Potential in the First 5 Years for Children in Developing Countries.” *The Lancet* 369 (9555): 60–70. [https://doi.org/10.1016/S0140-6736\(07\)60032-4](https://doi.org/10.1016/S0140-6736(07)60032-4).
- [66] Ha, Sandie, Yeyi Zhu, Danping Liu, Seth Sherman, and Pauline Mendola. 2017. “Ambient Temperature and Air Quality in Relation to Small for Gestational Age and Term Low Birthweight.” *Environmental Research* 155 (May): 394–400. <https://doi.org/10.1016/j.envres.2017.02.021>.
- [67] Hallegatte, Stéphane, Adrien Vogt-Schilb, Julie Rozenberg, Mook Bangalore, and Chloé Beaudet. 2020. “From Poverty to Disaster and Back: A Review of the Literature.” *Economics of Disasters and Climate Change* 4 (1): 223–47. <https://doi.org/10.1007/s41885-020-00060-5>.
- [68] Hallegatte,Stephane; Vogt-Schilb,Adrien Camille; Bangalore,Mook; Rozenberg,Julie. 2016. *Unbreakable : building the resilience of the poor in the face of natural disasters*. Climate Change and Development Washington, D.C. : World Bank Group.
- [69] Hasan, Amer. 2010a. “Gender-Targeted Conditional Cash Transfers: Enrollment, Spillover Effects and Instructional Quality.” Policy Research Working Paper 5257, World Bank, Washington, DC.
- [70] ———. 2010b. “Time Allocation in Rural Households: The Indirect Effects of Conditional Cash Transfer Programs.” Policy Research Working Paper 5256, World Bank, Washington, DC.
- [71] Hernandez-de-Benito, Maria. 2024. “This is a Man’s World: Crime and Intra-Household Resource Allocation.” Available at SSRN: <https://ssrn.com/abstract=4712389>.

- [72] Hoddinott, John. 2006. “Shocks and Their Consequences across and within Households in Rural Zimbabwe.” *Journal of Development Studies* 42 (2): 301–21. <https://doi.org/10.1080/00220380500405501>.
- [73] Hoddinott, John, and Lawrence Haddad. 1995. “Does Female Income Share Influence Household Expenditures? Evidence from Côte D’Ivoire.” *Oxford Bulletin of Economics and Statistics* 57 (1): 77–96.
- [74] Hersbach, Hans, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, et al. 2020. “The ERA5 Global Reanalysis.” *Quarterly Journal of the Royal Meteorological Society* 146 (730): 1999–2049. <https://doi.org/10.1002/qj.3803>.
- [75] Hu, Zihan, and Teng Li. 2019. “Too Hot to Handle: The Effects of High Temperatures during Pregnancy on Adult Welfare Outcomes.” *Journal of Environmental Economics and Management* 94 (March): 236–53. <https://doi.org/10.1016/j.jeem.2019.01.006>.
- [76] Imbens, Guido, and Karthik Kalyanaraman. 2012. “Optimal Bandwidth Choice for the Regression Discontinuity Estimator.” *Review of Economic Studies* 79 (3): 933–59.
- [77] Innovative Development Consultants. 2005. “Third Party Validation for Stipend to Girl Students: Final Report.” Innovative Development Consultants, Lahore.
- [78] International Food Policy Research Institute (IFPRI). 2016. “Bangladesh Integrated Household Survey (BIHS) 2015”, <https://doi.org/10.7910/DVN/BXSYEL>, Harvard Dataverse, V4, UNF:6:K4Bn/u8+Ja93lQy8F5d/fQ== [fileUNF]
- [79] International Food Policy Research Institute (IFPRI). 2020. “Bangladesh Integrated Household Survey (BIHS) 2018-2019”, <https://doi.org/10.7910/DVN/NXKLZJ>, Harvard Dataverse, V2, UNF:6:UrJdeN5clEtNcJ0u+2ZmOQ== [fileUNF]
- [80] Isen, Adam, Maya Rossin-Slater, and Reed Walker. 2017. “Relationship between Season of Birth, Temperature Exposure, and Later Life Wellbeing.” *Proceedings of the National Academy of Sciences* 114 (51): 13447–52. <https://doi.org/10.1073/pnas.1702436114>.
- [81] Jacoby, Hanan G., and Ghazala Mansuri. 2015. “Crossing Boundaries: How Social Hierarchy Impedes Economic Mobility.” *Journal of Economic Behavior Organization* 117: 135–54.
- [82] J-PAL (Abdul Latif Jameel Poverty Action Lab). 2019a. “Increasing Student Enrollment and Attendance: Impacts by Gender.” J-PAL Policy Insights. Last modified February 2019. <https://doi.org/10.31485/pi.2262.2018>.

- [83] ———. 2019b. “Reducing Costs to Increase School Participation.” J-PAL Policy Insights. Last modified February 2019. <https://doi.org/10.31485/pi.2264.2018>.
- [84] ———. n.d. “J-PAL Costing Guidelines.” J-PAL, Cambridge, MA. <https://www.povertyactionlab.org/sites/default/files/resources/J-PAL%20Costing%20Guidelines%202016.05.pdf>. Accessed September 2019.
- [85] Karim, Azreen, and Ilan Noy. 2016. “Poverty and Natural Disasters: A Regression Meta-Analysis.” *Review of Economics and Institutions* 7 (2): 26. <https://doi.org/10.5202/rei.v7i2.222>.
- [86] Khandker, Shahidur, Mark Pitt, and Nobuhiko Fuwa. 2003. “Subsidy to Promote Girls’ Education: The Female Stipend Program in Bangladesh.” Paper 23688, Munich Personal RePEc Archive.
- [87] Kim, Jiyeon, Ajin Lee, and Maya Rossin-Slater. 2019. “What to Expect When It Gets Hotter: The Impacts of Prenatal Exposure to Extreme Heat on Maternal and Infant Health.” IZA Discussion Paper No. 12685, Available at SSRN: <https://ssrn.com/abstract=3475791>.
- [88] Kremer, Michael, Gautam Rao, and Frank Schilbach. 2019. “Behavioral Development Economics.” In *Handbook of Behavioral Economics*, edited by Douglas Bernheim, Stefano DellaVigna, and David Laibson. Amsterdam: Elsevier.
- [89] Lechene, Valérie, Krishna Pendakur, and Alex Wolf. 2022. “Ordinary Least Squares Estimation of the Intrahousehold Distribution of Expenditure.” *Journal of Political Economy* 130 (3): 681–731. <https://doi.org/10.1086/717892>.
- [90] Leroy, Jef L., Armando García-Guerra, Raquel García, Clara Dominguez, Juan Rivera, and Lynnette M. Neufeld. 2008. “The Oportunidades Program Increases the Linear Growth of Children Enrolled at Young Ages in Urban Mexico.” *The Journal of Nutrition* 138 (4): 793–98. <https://doi.org/10.1093/jn/138.4.793>.
- [91] Leroy, Jef L., Marie Ruel, Jean-Pierre Habicht, and Edward A. Frongillo. 2015. “Using Height-for-Age Differences (HAD) Instead of Height-for-Age z-Scores (HAZ) for the Meaningful Measurement of Population-Level Catch-up in Linear Growth in Children Less than 5 Years of Age.” *BMC Pediatrics* 15 (1): 145. <https://doi.org/10.1186/s12887-015-0458-9>.
- [92] Lewbel, Arthur, and Krishna Pendakur. 2008. “Estimation of Collective Household Models with Engel Curves.” *Journal of Econometrics* 147 (2): 350–58. <https://doi.org/10.1016/j.jeconom.2008.09.012>.

- [93] ———. 2024. “Estimating a Model of Inefficient Cooperation and Consumption in Collective Households.” *Review of Economics of the Household* 22 (3): 865–907.
- [94] Lutter, Chessa K., Alicia Rodríguez, Guillermo Fuenmayor, Luz Avila, Fernando Sempertegui, and Jessica Escobar. 2008. “Growth and Micronutrient Status in Children Receiving a Fortified Complementary Food.” *The Journal of Nutrition* 138 (2): 379–88. <https://doi.org/10.1093/jn/138.2.379>.
- [95] Maccini, Sharon, and Dean Yang. 2009. “Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall.” *American Economic Review* 99 (3): 1006–26. <https://doi.org/10.1257/aer.99.3.1006>.
- [96] Menon, Martina, Krishna Pendakur, and Federico Perali. 2012. “On the Expenditure-Dependence of Children’s Resource Shares.” *Economics Letters* 117 (3): 739–42. <https://doi.org/10.1016/j.econlet.2012.08.012>.
- [97] Memon, Falak Shad. 2020. “Climate Change and Violence Against Women: Study of A Flood-Affected Population in The Rural Area of Sindh, Pakistan.” *Pakistan Journal of Women’s Studies: Alam-e-Niswan* 27 (1): 65–85. <https://doi.org/10.46521/pjws.027.01.0039>.
- [98] Miguel, Edward. 2005. “Poverty and Witch Killing.” *Review of Economic Studies* 72 (4): 1153–72. <https://doi.org/10.1111/0034-6527.00365>.
- [99] Ministry of Environment, Forest and Climate Change. 2022. *National Adaptation Plan of Bangladesh (2023-2050)*. Dhaka: Ministry of Environment, Forest and Climate Change, Government of the People’s Republic of Bangladesh.
- [100] Mitchell, Tom, Thomas Tanner, and Kattie Lussier. 2007. *We Know What We Need: South Asian Women Speak Out On Climate Change Adaptation*. Institute of Development Studies, University of Sussex, UK.
- [101] Molina Millán, Teresa; Barham, Tania; Macours, Karen; Maluccio, John A.; Stampini, Marco. 2019. *Long-Term Impacts of Conditional Cash Transfers: Review of the Evidence*. World Bank Research Observer. © Published by Oxford University Press on behalf of the World Bank. <http://hdl.handle.net/10986/34347> License: CC BY-NC-ND 3.0 IGO.
- [102] Mora, Camilo, Bénédicte Dousset, Iain R. Caldwell, Farrah E. Powell, Rollan C. Geronimo, Coral R. Bielecki, Chelsie W. W. Counsell, et al. 2017. “Global Risk of Deadly Heat.” *Nature Climate Change* 7 (7): 501–6. <https://doi.org/10.1038/nclimate3322>.

- [103] Mueller, V., and A. Quisumbing. 2011. "How Resilient are Labor Markets to Natural Disasters? The Case of the 1998 Bangladesh Flood." *Journal of Development Studies* 47, no. 12:1954–71.
- [104] Musaddiq, Tareena, and Farah Said. 2023. "Educate the girls: Long run effects of secondary schooling for girls in Pakistan." *World Development* 161: 106115.
- [105] Nellesmann, C., Verma, R., and Hislop, L. (eds). 2011. *Women at the frontline of climate change: Gender risks and hopes. A Rapid Response Assessment*. United Nations Environment Programme, GRID-Arendal.
- [106] Nguyen, Quynh T., and Dhushyanth Raju. 2015. "Private School Participation in Pakistan." *Lahore Journal of Economics* 20: 1–46.
- [107] Nissan, Hannah, Katrin Burkart, Erin Coughlan de Perez, Maarten Van Aalst, and Simon Mason. 2017. "Defining and Predicting Heat Waves in Bangladesh." *Journal of Applied Meteorology and Climatology* 56 (10): 2653–70. <https://doi.org/10.1175/JAMC-D-17-0035.1>.
- [108] Noy, Ilan, Cuong Nguyen, and Pooja Patel. 2021. "Floods and Spillovers: Households after the 2011 Great Flood in Thailand." *Economic Development and Cultural Change*, January. <https://doi.org/10.1086/703098>.
- [109] Pakistan Bureau of Statistics. 2016. *PSLM 2014–15: Pakistan Social and Living Standards Measurement Survey (2014–15): National/Provincial/District*. Islamabad: Pakistan Bureau of Statistics.
- [110] ———. 2017a. "Tehsil Wise Provisional Results of Census—2017." Pakistan Bureau of Statistics, Islamabad.
- [111] ———. 2017b. "District Wise Provisional Results of Census—2017." Pakistan Bureau of Statistics, Islamabad.
- [112] Penny, Mary E, Hilary M Creed-Kanashiro, Rebecca C Robert, M Rocio Narro, Laura E Caulfield, and Robert E Black. 2005. "Effectiveness of an Educational Intervention Delivered through the Health Services to Improve Nutrition in Young Children: A Cluster-Randomised Controlled Trial." *The Lancet* 365 (9474): 1863–72. [https://doi.org/10.1016/S0140-6736\(05\)66426-4](https://doi.org/10.1016/S0140-6736(05)66426-4).
- [113] PERI (Punjab Economic Research Institute). 2017. *Punjab Economic Report 2017*. Lahore: Punjab Planning and Development Department.

- [114] Poaponsakorn, N., Pitsom Meethom, and Kamphol Pantakua. 2015. "The Impact of the 2011 Floods, and Flood Management on Thai Households." In *Resilience and Recovery in Asian Disasters*, ed. D. Aldrich, S. Oum, and Y. Sawada, 75–104. Tokyo: Springer.
- [115] Population Council. 2007. "Evaluation of the Punjab Female School Stipend Program." Population Council, Islamabad.
- [116] Randell, Heather, Clark Gray, and Kathryn Grace. 2020. "Stunted from the Start: Early Life Weather Conditions and Child Undernutrition in Ethiopia." *Social Science Medicine* 261 (September): 113234. <https://doi.org/10.1016/j.socscimed.2020.113234>.
- [117] Rezwana, Nahid, and Rachel Pain. 2021. "Gender-based Violence before, during, and after Cyclones: Slow Violence and Layered Disasters." *Disasters* 45 (4): 741–61. <https://doi.org/10.1111/disa.12441>.
- [118] Seneviratne, S.I., X. Zhang, M. Adnan, W. Badi, C. Dereczynski, A. Di Luca, S. Ghosh, I. Iskandar, J. Kossin, S. Lewis, F. Otto, I. Pinto, M. Satoh, S.M. Vicente-Serrano, M. Wehner, and B. Zhou, 2021: Weather and Climate Extreme Events in a Changing Climate. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 1513–1766, doi: 10.1017/9781009157896.013.
- [119] Skoufias, Emmanuel, and Katja Vinha. 2012. "Climate Variability and Child Height in Rural Mexico." *Economics Human Biology* 10 (1): 54–73. <https://doi.org/10.1016/j.ehb.2011.06.001>.
- [120] Sokullu, Senay, and Christine Valente. 2022. "Individual Consumption in Collective Households: Identification Using Repeated Observations with an Application to PROGRESA." *Journal of Applied Econometrics* 37 (2): 286–304. <https://doi.org/10.1002/jae.2875>.
- [121] Somanathan, E., Rohini Somanathan, Anant Sudarshan, and Meenu Tewari. 2021. "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing." *Journal of Political Economy* 129 (6): 1797–1827. <https://doi.org/10.1086/713733>.
- [122] Stein, Aryeh D., Fernando C. Barros, Santosh K. Bhargava, Wei Hao, Bernardo L. Horta, Nanette Lee, Christopher W. Kuzawa, et al. 2013. "Birth Status, Child Growth, and Adult Outcomes in Low- and Middle-Income Countries." *The Journal of Pediatrics* 163 (6): 1740–1746.e4. <https://doi.org/10.1016/j.jpeds.2013.08.012>.

- [123] Strand, Linn B., Adrian G. Barnett, and Shilu Tong. 2011. "The Influence of Season and Ambient Temperature on Birth Outcomes: A Review of the Epidemiological Literature." *Environmental Research* 111 (3): 451–62. <https://doi.org/10.1016/j.envres.2011.01.023>.
- [124] Tellman, B., J. A. Sullivan, C. Kuhn, A. J. Kettner, C. S. Doyle, G. R. Brakenridge, T. A. Erickson, and D. A. Slayback. 2021. "Satellite Imaging Reveals Increased Proportion of Population Exposed to Floods." *Nature* 596 (7870): 80–86. <https://doi.org/10.1038/s41586-021-03695-w>.
- [125] Tommasi, Denni. 2019. "Control of Resources, Bargaining Power and the Demand of Food: Evidence from PROGRESA." *Journal of Economic Behavior & Organization* 161 (May): 265–86. <https://doi.org/10.1016/j.jebo.2019.04.008>.
- [126] U.S. Geological Survey. "How Are Floods Predicted?" Accessed September 26, 2024. <https://www.usgs.gov/faqs/how-are-floods-predicted>.
- [127] Wilde, Joshua, Bénédicte H. Apouey, and Toni Jung. 2017. "The Effect of Ambient Temperature Shocks during Conception and Early Pregnancy on Later Life Outcomes." *European Economic Review* 97 (August): 87–107. <https://doi.org/10.1016/j.eurocorev.2017.05.003>.
- [128] World Bank. 2019. *World Development Indicators Databank*. Washington, DC: World Bank. <https://databank.worldbank.org/source/world-development-indicators>. Accessed August 2019.
- [129] World Bank. 2021. World Development Indicators. <https://databank.worldbank.org/source/world-development-indicators>.
- [130] World Health Organization. 2006. Length/Height-for-Age, Weight-for-Age, Weight-for-Length, Weight-for-Height and Body Mass Index-for-Age; Methods and Development. Edited by Mercedes de Onis. WHO Child Growth Standards. Geneva: WHO Press.
- [131] Yuan, Ke-Hai, and Wai Chan. 2008. "Structural Equation Modeling with near Singular Covariance Matrices." *Computational Statistics & Data Analysis* 52 (10): 4842–58. <https://doi.org/10.1016/j.csda.2008.03.030>.