

ESSAYS ON ECONOMIC DEVELOPMENT AND CLIMATE CHANGE

by

BENJAMIN A. FITCH-FLEISCHMANN

A DISSERTATION

Presented to the Department of Economics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2015

DISSERTATION APPROVAL PAGE

Student: Benjamin A. Fitch-Fleischmann

Title: Essays on Economic Development and Climate Change

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

Trudy Cameron	Co-Chair
Glen Waddell	Co-Chair
Alfredo Burlando	Core Member
Kristin Yarris	Institutional Representative

and

Scott L. Pratt	Dean of the Graduate School
----------------	-----------------------------

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2015

© 2015 Benjamin A. Fitch-Fleischmann

DISSERTATION ABSTRACT

Benjamin A. Fitch-Fleischmann

Doctor of Philosophy

Department of Economics

June 2015

Title: Essays on Economic Development and Climate Change

The first essay considers the relative effectiveness of government and non-governmental organizations (NGOs) as channels to allocate resources. I use a catastrophic climate-related shock—Hurricane Mitch—to examine the political economy of these channels of aid distribution at the micro level. I combine extensive data on aid received by Nicaraguan households with data on municipal election outcomes and an exogenous, precipitation-based measure of hurricane impact. I find that the hurricane had long-lasting effects on the aid received by households from both NGOs and the government. In the short term, however, the government did not provide aid according to the objective measure of hurricane damage but instead provided aid along political lines.

The second essay presents estimates of a relationship between extreme hot temperatures during gestation and a child's subsequent physical well-being in a sample of children in Peru, thus extending existing evidence constructed from U.S. data. Estimates are constructed using high-resolution gridded climate data and geo-coded household surveys. The results suggest that a period of extreme heat (a month whose average temperature is more than 2σ above the local average) in the period 1 to 3 months before birth is associated with lower weight at birth and a reduction in height (measured 1 to 59 months after birth) that cannot be fully explained by birth weight. There is no evidence of differential maternal investment, as measured by duration of breastfeeding, according to a child's exposure to extreme heat during gestation.

The third essay asks whether improved treatment of HIV/AIDS in Africa can be achieved simply by paying health workers to do more. I present estimates of the impact of financial incentives paid to individual workers at public health facilities in Mozambique. The results suggest that piece-rate incentives increased the delivery of five out of fourteen health services

for which treatment effects can be identified, with estimated increases ranging from 34 to 157 percent, depending on the particular service. I find no evidence of a corresponding decrease in the delivery of services that are not financially incentivized, suggesting that there is no “crowding out” of intrinsic motivation.

CURRICULUM VITAE

NAME OF AUTHOR: Benjamin A. Fitch-Fleischmann

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, OR
University of Montana, Missoula, MT
Claremont McKenna College, Claremont, CA

DEGREES AWARDED:

Doctor of Philosophy, Economics, 2015, University of Oregon
Master of Science, Economics, 2013, University of Oregon
Master of Arts, Economics, 2009, University of Montana
Bachelor of Arts, Economics and Government, 2005, Claremont McKenna College

AREAS OF SPECIAL INTEREST:

Environmental Economics, Economic Development, Econometrics

GRANTS, AWARDS AND HONORS:

Graduate Teaching Fellowship, University of Oregon, 2010-2015
Donald and Darel Stein Graduate Student Teaching Award, University of Oregon, 2014
Raymond F. Mikesell Award, University of Oregon, 2014
Graduate Teaching Fellow Teaching Award, University of Oregon, 2014
Kleinsorge Research Fellowship, University of Oregon, 2014
Dan Kimble First Year Teaching Award, University of Oregon, 2011

ACKNOWLEDGEMENTS

Many people have offered invaluable support and guidance throughout my time in graduate school. I am especially grateful for the hours and effort dedicated to me by the members of my dissertation committee. Trudy, Glen, and Alfredo have each been remarkably generous with their knowledge, time, and enthusiasm. My classmates and the participants in the weekly microeconomics workshop also provided insightful feedback that improved many parts of my research and, more importantly, offered camaraderie and a source of laughter over the last five years. I am particularly grateful for the support I have received from my family and friends, who have kept me going, kept me sane, and expressed far more interest in my research than I ever could have imagined.

TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION	1
II. THE POLITICAL ECONOMY OF AID RESOURCES AFTER CLIMATE CATASTROPHES: EVIDENCE FROM NICARAGUA	3
Introduction	3
Background	6
Data on Hurricanes, Aid, and Political Representation	11
Geographic Patterns in the Impact of Hurricanes in Nicaragua: 1960-2010	15
Empirical Strategy	17
Results	19
Conclusions	23
III. EXTREME HEAT DURING GESTATION AND EARLY-LIFE HEALTH: EVIDENCE FROM PERU	27
Introduction	27
Background	30
Data	31
Empirical Strategy	34
Results	36
Conclusions	39
IV. TEAM INCENTIVES AND HEALTH WORKER PRODUCTIVITY: EVIDENCE FROM MOZAMBIQUE	41
Introduction	41
Background	42

Chapter	Page
The Incentive Program and Associated Data	44
Empirical Strategy	46
Results	48
Conclusions	51
APPENDIX: FIGURES AND TABLES	53
REFERENCES CITED	92

LIST OF FIGURES

Figure	Page
1. Timeline of Hurricanes, Elections, and Data Collection	53
2. Precipitation Data Construction	54
3. Distribution of Hurricane Impacts in Nicaragua: 1960-2010.	55
4. Timeline of Hurricanes Affecting Nicaragua: 1960-2010.	56
5. Geographic Variation in Hurricane Impact (Rainfall Ratio).	57
6. Average Hurricane Impact (Rainfall Ratio) by Municipality: 1960-2010	58
7. Average Hurricane Impact (Rainfall Ratio): 1960-2010.	59
8. Average Height	60
9. Seasonality of Births in Peru	60
10. Average Monthly Temperatures in Peru: 1960-2010	61
11. Locations of Respondents and Weather Data	61
12. Height: Estimated Coefficient of Birthweight, by Age bin (3m)	62
13. Boys Height: Estimated Coefficient of Birthweight, by Age bin (3m)	62
14. Girls Height: Estimated Coefficient of Birthweight, by Age bin (3m)	63
15. Robustness of Key Estimates to HOT Threshold	64
16. Individual Service Trends: HIV Tests	65
17. Individual Service Trends: Newly Initiated on ART	66
18. Individual Service Trends: Longterm Care	67
19. Individual Service Trends: Birth and Babies	68
20. Individual Service Trends: Non-paid Services	69
21. Trends: Total Value of Services	69
22. Data Reporting Trends	70

LIST OF TABLES

Table	Page
1. Descriptive Statistics: Households Surveyed Before and After Hurricane Mitch	71
2. Municipal Hurricane Incidence (Hurricane Rainfall Ratio): 1960-2010	72
3. Household Attrition and Hurricane Exposure for LSMS 1998 Households	73
4. The Effect of Hurricane Mitch on Government Aid Allocations: 0-3 Years After	74
5. Determinants of INEC Hurricane Mitch Designation	75
6. The Effect of Hurricane Mitch on Aid Allocations by NGOs: 0-3 Years After	76
7. The Effect of Hurricane Mitch on Aid Allocations: 3-7 Years After	77
8. The Effect of Hurricane Mitch on Aid Allocations: 4-0 Years Before	78
9. Children (0-59 months old) in Surveyed Households	79
10. Descriptive Statistics: Children in Peru DHS Households	80
11. Effect of Extreme Heat During Gestation on Infant Mortality	81
12. Determinants of Exposure to Extreme Heat During Gestation	82
13. Height (cm) as a Function of Birth Weight (100 g)	82
14. Birth Weight (100 g) as a Function of Extreme Heat During Gestation	83
15. Birth Weight (100 g) and Heat: Wealth Heterogeneity	83
16. Height (cm) and Extreme Heat During Gestation	84
17. Height (cm) and Extreme Heat During Gestation: Boys by Wealth	85
18. Maternal Investment in Boys: Breastfeeding Duration (months)	86
19. Maternal Investment in Girls: Breastfeeding Duration (months)	86
20. Health Services and Prices	87
21. Descriptive Statistics: Facility Characteristics	87
22. Descriptive Statistics: Average Count of Services Delivered, Pre-Treatment	88
23. Paid Services Pre-Treatment Trend Differences and Treatment Effects	89
24. Non-Paid Services Pre-Treatment Trend Differences and Treatment Effects	90
25. Total Value (dollars) of Service Mix	90
26. Heterogeneity of Treatment Effects by Facility Capacity Characteristics	91

CHAPTER I

INTRODUCTION

There is growing concern that human-produced greenhouse gas emissions will alter global temperatures and other climatic conditions. These changes are predicted to lead to increased human exposure both to extreme temperatures and to extreme weather events, such as tropical storms and floods. This dissertation investigates questions about the harm these changes may inflict in developing countries, and the ways that institutions deliver assistance in response.

To study these topics, I take a microeconomic approach and combine a variety of data sources that allow me to associate individual- and household-level economic and health outcomes in developing countries with various climate conditions. This allows me to extend the existing literature by estimating, in data from developing countries, the relationships between climate conditions that are predicted to occur with greater frequency as a consequence of climate change and economic and health outcomes of interest, such as the birth weight of infants or the assistance received by households from governments and non-governmental organizations (NGOs).

In Chapter II, I examine the response to natural disasters associated with climate change by entities that deliver development aid. These responses are considered in the context of Hurricane Mitch and Nicaragua, which is a setting particularly suited to the study of demographic effects of the extreme storms associated with climate change because these storms are predicted to increase in strength and are likely to affect places that are similar to Nicaragua in terms of development levels and climate. I examine differences in the distribution of aid between the Nicaraguan government and NGOs. My results suggest that NGOs are more responsive to the disaster, in the sense that they allocate more aid to households in areas that were hit more severely by the storm, while aid allocations made by the Nicaraguan government were influenced by local political affiliations, but not by hurricane severity until several years after the storm.

In Chapter III, I estimate the effects of exposure to extreme heat during gestation on an infant's birth weight and subsequent physical development, captured by measurements of height collected up to 59 months after birth. These estimates are constructed using a sample of children from Peru and serve to extend the existing evidence to a developing-country context. I find that temperature extremes during gestation are associated with both reduced weight at birth and

lower height later on. Much, but not all of this negative effect on height can be explained by the child's birth weight, which suggests that birth weight is not a complete measure of the health consequences associated with temperature shocks experienced during gestation.

In Chapter IV, I consider the potential for direct financial incentives for health workers to increase the delivery of health services in Mozambique. Many health systems in developing countries suffer from low productivity, and low motivation, among workers. This chapter examines the effectiveness of a performance-based funding mechanism implemented at a set of public health facilities in two provinces in Mozambique. The funding mechanism provided facilities with a piece rate subsidy for performing a variety of services, with half of the funding distributed to the facility workers. I use a difference-in-differences approach to estimate the effect of this payment system on the quantity of services delivered, using a set of health facilities in two other Mozambican provinces as a control group. My results suggest that the program saw modest improvements in about one-third of monitored services.

CHAPTER II

THE POLITICAL ECONOMY OF AID RESOURCES AFTER CLIMATE CATASTROPHES: EVIDENCE FROM NICARAGUA

Introduction

This research presents evidence regarding the effects of a catastrophic storm on the assistance provided to households in and near the area affected, with a specific focus on the source of assistance and political influences. The analysis is motivated by two distinct but related issues in economic development. First, I consider whether there are important differences in the allocations of aid made by the government relative to those made by non-governmental organizations (NGOs), and the results are directly relevant to decisions about how to transmit post-disaster aid from international entities to disaster victims. This is also informative for development aid more generally, as I analyze how these aid channels differ along two important dimensions: their vulnerability to political influence, and the evolution over time of their response to an objective measure of need (i.e., the impact of a hurricane). Using a catastrophic shock to study the political economy of aid is similar in spirit to Besley and Case (1995), who use natural disasters to study the effects of term limits on government responsiveness.

Second, the potential scale of the damage that extreme storms and other natural disasters can cause, and have caused in the past, is often large relative to the size of the domestic economies affected (Freeman et al., 2003). Thus, spreading risk is difficult and the international community is frequently called on to provide aid. However, empirical evidence concerning the micro-level distribution of this aid is surprisingly scarce. Existing studies of the topic are limited to relatively small events (e.g., Takasaki, 2011; Francken et al., 2012), analyze data that reflects only a short period (less than one year) after the disaster (e.g., Morris and Wodon, 2003; Takasaki, 2011; Francken et al., 2012), or lack measures of disaster severity (e.g., Aldrich, 2010; Nose, 2014). The evidence presented in this analysis uses a decade's worth of data on household-level assistance provided by NGOs and the government in combination with local election outcomes and a meteorologic measure capturing geographic variation in the severity of a catastrophic storm. The results suggest that the economic legacy of such an event extends well

beyond one year, and the estimates reveal large differences in who receives assistance from NGOs relative to the government.

The context considered—Nicaragua and Hurricane Mitch—is particularly relevant because climate scientists predict that higher temperatures, in both the atmosphere and ocean surface waters, will increase the likelihood of extremely severe tropical cyclones¹ (Scheraga et al., 2003; Solomon et al., 2007; Pachauri and Reisinger, 2007; Knutson et al., 2010). Specific predictions for the coming century include “substantial increases in the frequency of the most intense cyclones” and “increases of the order of twenty percent in the precipitation rate within 100km of the storm center” (Knutson et al., 2010). Furthermore, the regions of the world historically affected by these storms include many poorer, less-developed nations, and populations in these regions are increasingly concentrating in areas that are vulnerable to climate-related disasters (Freeman et al., 2003; Knutson et al., 2010).

Hurricane Mitch was the deadliest Atlantic hurricane in over 200 years, causing estimated damages in excess of \$5 billion (1998 USD) and more than five times as many deaths as Hurricane Katrina in the U.S. (McCown et al., 1999; Knabb et al., 2005). After Mitch struck Central America in late October of 1998, an immediate donation of supplies was airlifted from Mexico to Nicaragua, while many other countries provided additional supplies and financial donations amounting to hundreds of millions of dollars.² However, efforts to deliver this assistance to the Nicaraguan victims were caught in a political crossfire. The Sandinistas, the major opposition party, accused the ruling coalition of giving aid only to their own supporters. Nicaraguan President Arnoldo Alemán, referred to by some as “El que llegó y se fué,”³ fired back and accused the Sandinistas of greedily distorting the facts to capture the country’s resources. By other accounts, the President’s reaction was indeed callous; he tried to levy a tax on incoming aid and suggested that homeless survivors should be sent to work in the coffee harvest to psychologically cope with their situation (Schenk, 1999). Even the Catholic Church was drawn into the crisis,

¹The term “tropical cyclone” refers to the broad class of weather phenomena that includes Atlantic hurricanes, typhoons (as hurricanes are called in the west Pacific), and less severe (measured by windspeed) tropical storms.

²The largest donors included Spain (\$105 million), Sweden (\$100 million, pledged over three years), and the U.S. (\$80 million) (McCown et al., 1999).

³“He who came and left,” in reference to how quickly he departed the devastated areas without offering help (Olson et al., 2001).

servicing as a state-sponsored vessel to funnel aid—away from the opposition, if the Sandinistas are to be believed (Olson et al., 2001).

While the details of the present study are particular to Nicaragua, similar claims of politically manipulated disaster aid allocations have emerged elsewhere. For example, reports of politicians campaigning with resources intended for the victims surfaced only weeks after Typhoon Haiyan, the strongest storm ever recorded at landfall, struck the Philippines in 2013 (Fischetti, 2013).⁴ Whether longer-term, post-Haiyan reconstruction efforts are influenced by political affiliations remains to be seen.

Following catastrophic storms, where does the aid flow? The primary challenge to retrieving estimates of the causal effect of a natural disaster on aid allocations is the limited availability of appropriate data, and an important feature of the setting studied here is the unusual wealth of available data. While previous studies are limited to data on aid allocations during the twelve months after a storm, the data available for Nicaragua reflect three different time-periods corresponding approximately to the four years before the storm, the three years immediately following it, and the four years after that. These data thus allow the first estimates, to my knowledge, of the longer-term effects of a catastrophe on household-level aid allocations.

A second challenge to the identification of the causal effects of a hurricane on aid allocations, left largely unaddressed in the existing literature, is the possibility that a storm's impact might be correlated with some unobserved factor that determines aid allocations. A simple mechanism by which this might arise is if the impact of hurricanes is correlated over time. The impact of prior storms could affect previous aid allocations, or economic development, which could then affect later allocations of aid. I present evidence that the historical impact of hurricanes in Nicaragua follows no systematic pattern and is therefore unlikely to be correlated with unobservable characteristics that might drive the estimated relationships between hurricane impact and aid. The existence of data covering multiple time periods also allows me to verify that the estimates are not merely reflective of pre-storm patterns in the distribution of aid.

A third concern is that measures of disaster damage may be endogenous to aid allocations. Indications of a disaster's impact that are made by the government are subject to political manipulation, while measures of disaster damage constructed from survey reports could

⁴See *Are Philippine Politicians Using Typhoon Aid to Their Advantage?*, CNN, Nov 22, 2013, at <http://www.cnn.com/2013/11/22/world/asia/philippines-politicians-typhoon-aid-advantage-irpt/>

potentially be misrepresented by individuals attempting to attract aid. I eliminate the potential for these forces to influence the measurement of hurricane impact by constructing an exogenous measure based on independent scientific data on precipitation.

This analysis yields two important findings that, in addition to the methodological improvements described above, expand the evidence on the effects of disasters on aid allocations. First, while the existing literature finds short term effects in the twelve months after a disaster, I find that Mitch had *long-lasting* effects, extending to the period three to seven years after the storm, on aid allocations by both NGOs and the Nicaraguan government. Second, regarding aid allocations in the short-term (zero to three years after the storm), I find that aid from the government was heavily influenced by political affiliations. The evidence strongly suggests that, in the parts of Nicaragua that experienced less than the average hurricane impact, households in areas controlled by the opposition party were about 26 percent less likely to receive assistance from the government than households in areas controlled by the ruling party. Furthermore, governmental assistance during this time was *not* provided according to the objective precipitation-based measure of the hurricane’s impact. On the other hand, there is no evidence of political influence on assistance provided by NGOs, with NGOs more likely to assist households that experienced larger impacts in both the short and long term. These results stand in stark contrast with some of the evidence from national-level studies of disaster aid allocation, such as Becerra et al. (2012), who “[do] not find evidence that political considerations or strategic behavior on the part of donors determine the size of post-disaster aid surges.”

Background

Existing Evidence on Disasters and the Allocation of Aid

Empirical evidence concerning the impact of natural disasters on the distribution of aid at the micro level, and even the community level, is surprisingly rare.⁵ Existing studies are limited to data on aid allocations made within twelve months of the disaster and lack any pre-disaster data for comparison. As one would expect, studies that regress some measure of aid on a measure of disaster severity find a positive correlation (Morris and Wodon, 2003; Takasaki, 2011;

⁵A search for “disaster aid” on EconLit reveals only three studies that regress some measure of aid received by households or individuals on a measure of disaster damage (Morris and Wodon, 2003; Takasaki, 2011; Francken et al., 2012), and two others that study post-disaster aid allocations but do not include a measure of disaster damage (Aldrich, 2010; Nose, 2014).

Francken et al., 2012), though it is difficult to compare magnitudes across measurements (which include various indicators of damage to dwellings and assets, and scales of damage reported by the government) and disasters of different types and severities.

In the analysis most similar to the work described in this paper, Francken et al. (2012) consider the distribution of relief aid separately by NGOs and the government in the eight months after Cyclone Gafilo struck Madagascar in 2004. As in other existing studies, the absence of data on patterns of aid allocations prior to the disaster makes it difficult to know if their estimates reflect changes in resource allocations that are due to the disaster, or whether they simply reflect pre-existing patterns amplified by a surge in the availability of aid following disasters. Nonetheless, Francken et al. (2012) find evidence that suggests political manipulation: among the areas that the government declared were affected by Cyclone Gafilo, aid from the government was more likely to go to communities with greater support for the president. On the other hand, NGO aid was not influenced by presidential support, which perhaps suggests that government aid was more vulnerable to political manipulation.

Three other studies consider the effects on short-term post-disaster aid allocations due to different dimensions of political influence. Takasaki (2011) reports that after a cyclone in Fiji, village elites received assistance before other groups, though they did not receive greater total amounts of aid. Studying aid allocations after the 2004 tsunami in the Indian Ocean, Nose (2014) finds that stronger social ties among fishermen in Indonesia were associated with a higher probability of receiving aid. Aldrich (2010) finds that villages comprised of more members from lower castes in India were less likely to receive aid, and that pre-tsunami wealth (proxied by home-ownership) and aid were positively correlated. Francken et al. (2012) also find that wealth and aid were positively correlated in the Madagascar communities that the government indicated were most affected by Cyclone Gafilo, though Morris and Wodon (2003) find no relationship between pre-hurricane wealth and post-hurricane aid in Honduras during the period six to nine months after Hurricane Mitch.

While the immediate effects of natural disasters can be large, the damage to infrastructure and productive capacity may also have long-term effects. Recent work by Anttila-Hughes and Hsiang (2012) in the Philippines finds that, beyond the immediate destruction, hurricanes do indeed cause longer-term harmful effects on a variety of outcomes. They find that losses to

unearned income and excess infant mortality in the year following a hurricane exceed immediate damages and deaths by a magnitude of fifteen to one. Whether these longer-term damages translate into longer-term effects on aid allocations is a question that the present paper helps to answer.

A much larger literature studies the political economy of disaster aid at the national level. Like the evidence regarding foreign aid in general (e.g., Alesina and Dollar, 2000), some of the evidence suggests that country-to-country flows of post-disaster aid are subject to political influence and strategic considerations. However, the evidence is mixed. Common cultural ties, such as a shared language or colonial history, increase the probability that one country will aid another in the wake of a disaster (Eisenstein and Strömberg, 2007). In a study specific to disaster aid provided by the U.S. government, Drury et al. (2005) find that foreign policy and domestic factors are the “overriding determinant.” However, in their study of 196 countries over 39 years, Becerra et al. (2012) find no evidence of political or strategic influences on the size of country-level post-disaster aid allocations. They also point out that official post-disaster aid typically amounts to only three percent of the estimated cost of damages, even for very large disasters that receive substantial attention from the media. When considering a larger set of post-disaster sources of aid, however, such as foreign lending, foreign direct investment, and remittances from migrants, Yang (2008) finds that total inflows approach eighty percent of total damages by four years after the event.

Studies on the determinants of aid allocations by NGOs, just like the evidence on government aid, are largely limited to national-level estimates. Peter Nunnenkamp, Axel Dreher, and coauthors use data on Official Development Assistance (ODA) from several European countries to consider the so-called “article of faith” (Tendler, 1982) that NGOs are closer to the people they serve and therefore better at targeting aid according to need (Koch et al., 2009; Nunnenkamp et al., 2009; Dreher et al., 2010; Nunnenkamp and Öhler, 2011). The evidence in these studies consistently finds that NGOs are very similar to national governments and, at a national level, “replicate the location choices” (Koch et al., 2009) of governmental destinations of ODA. They also find that NGOs tend to follow each other and cluster their activities in the same countries. Evidence from the U.S. is similar, where NGO aid is found to “mirror” ODA allocations from the U.S. government (Keck, 2014). Contrary to the national-level similarities

revealed by these other studies, I find clear differences in the allocations of aid made by NGOs and the government at the household-level.

Nicaraguan Politics

Nicaragua is the second poorest country in the western hemisphere. GDP per capita was only barely above \$1,000 (2014 USD) at the time of Hurricane Mitch, and the country had a literacy rate of only 60 percent among its 4.45 million inhabitants (ECLAC, 1999; CIA, US, 2013; World Bank, 2014). It is about the size of New York state, bordered by Honduras to the north and Costa Rica to the south, and spans the Central American isthmus.

The recent political history in Nicaragua has been contentious. The primary left wing political party, the *Frente Sandinista de Liberación Nacional*, or Sandinistas, grew out of opposition to the military dictatorship that was in control of Nicaragua in the 1960s. A devastating earthquake in 1972, combined with general unrest and dissatisfaction with the concentration and control of wealth by the regime in power, led to an outpouring of support for the Sandinistas and a subsequent uprising (Black, 1981). The Sandinistas took control of national politics from 1979 to 1990. In the 1980s, the U.S. government, in apparent fear of the Sandinista’s “pro-Cuban” orientation, infamously directed funds towards revolutionary troops—the “Contras”—in the hopes of installing a government in Nicaragua more acceptable to the U.S. administration of the time.⁶ The 1980s were characterized by violent fighting between the Contras and the Sandinistas, until a truce was signed in 1989.

From 1990 until 2006, the national elections were won by the center-right coalition led by the *Partido Liberal Constitucionalista* (PLC). At the time of Hurricane Mitch, the PLC controlled the presidency and national assembly, although 51 *alcaldes*, the locally elected municipal leaders, were affiliated with the Sandinistas (the remaining 92 were affiliated with the PLC). Despite corruption charges in 2000 against some of the party’s leaders, PLC candidate Enrique Bolaños won the presidential election in 2001 and the PLC retained a majority of the seats in the National Assembly. In 2006, the Sandinistas regained control of both the presidency, with the election of Daniel Ortega, and the National Assembly. The contentious nature of politics in Nicaragua has

⁶The U.S. was subsequently accused of violating Nicaraguan sovereignty and ordered by the International Court of Justice to pay \$12 billion to Nicaragua in compensation (Morrison, 1987).

continued, and the transparency and legitimacy of elections have been called into question by international observers, who were banned from monitoring the 2008 elections.

Hurricane Mitch

Nicaragua is in an area susceptible to hurricanes—fourteen passed within 200 miles of its borders between 1960 and 2010—but Hurricane Mitch was unique in the extent of the damage it caused. Mitch made landfall in Honduras on October 26, 1998, just north of the border with Nicaragua.⁷ With maximum sustained winds of 180 miles per hour, the hurricane moved inland and, over five days, dropped as much as fifty inches of rain in some parts of Nicaragua. Mitch caused an estimated 11,000 total deaths (including 3,800 in Nicaragua), vastly more than the deaths caused by other storms that have affected the region.⁸ The total damages in Nicaragua from Hurricane Mitch have been estimated at \$1 billion to \$1.3 billion, with twenty percent of the population left without habitable dwellings, and 1500 miles of roads destroyed along with 300 schools, 90 health clinics, and one-third of agricultural crops (World Bank, 2001a).

This paper is the first to study the political economy of the aid response to Hurricane Mitch in Nicaragua, but other economic and social consequences of Mitch have received attention in the academic literature. Premand (2008) finds a limited and short-term negative economic impact from damage due to Mitch, but no discernible effect on economic growth. Van den Berg (2010) presents evidence that, despite its heavy damage to agriculture, the hurricane did not induce substantial numbers of people to change their strategies for generating income. Jakobsen (2012) finds that Mitch had a significant negative effect on the ownership of durable goods and assets, that the poorest households were affected disproportionately, and that there is “strong suggestive evidence of a geographical poverty trap within the shock-affected areas of the country.”

However, of these analyses, all but Premand (2008) measure hurricane exposure using a governmental designation of which areas were affected. This designation is derived from the decision by Nicaragua’s Instituto Nacional de Estadística y Censos (INEC) to survey households in the areas INEC determined were affected by the hurricane, but the process by which this

⁷This analysis focuses on Nicaragua because the data available is much more extensive than that for Honduras, though the eye of the storm never crossed the border into Nicaragua.

⁸Other hurricanes that have affected Nicaragua since 1985 (when fatality records begin) include Joan (1988, with an estimated 150 fatalities) Gert (1993, 11 fatalities), Cesar (1996, 42 fatalities), Beta (2005, 6 fatalities), Felix (2007, 130 fatalities), and Ida (2009, 0 fatalities).

designation was determined is not transparent.⁹ Given the claims by the opposition party and many in the media (described in Olson et al., 2001) that the government's response was influenced by political affiliations, it is questionable whether this designation accurately reflects the hurricane damage. Following Premand (2008), I use independent precipitation data to construct a measure of hurricane impact that is free of political influences and consider this measure, along with the INEC designation, as alternative measures of the effect of the hurricane on each municipality.

Data on Hurricanes, Aid, and Political Representation

The data for this analysis come from three sources. I use weather data from the British Atmospheric Data Center to construct an exogenous measure of the spatial variation in the extent of hurricane's impact, and survey data from the World Bank's Living Standards Measurement Study to measure allocations of aid in Nicaragua.

To capture the political relationships between the national government and local populations, I use the political party affiliation at the time of Hurricane Mitch of the elected *alcaldes*, the municipal-level leadership position in Nicaragua. These data come from Nicaragua's Consejo Supremo Electoral.¹⁰ In 1998, 51 *alcaldes* were Sandinistas and the remaining 92 were members of the ruling coalition led by the PLC.

Data on Aid Allocations and Household Characteristics

For information on the allocation of aid in Nicaragua, I use the World Bank's Living Standards Measurement Study (LSMS) conducted by Nicaragua's Instituto Nacional de Estadística y Censos (INEC). The LSMS is a nationally representative household survey developed by the World Bank that has been conducted in dozens of countries since the mid-1980s.¹¹ The survey collects extensive information on household economic activity, education, and demographics. Respondents are identified geographically by their municipality (the secondary

⁹The criteria for this determination are not described anywhere in detail (see Premand, 2008). The only description of the process of which I am aware comes from the World Bank (2001a): "Households were selected for inclusion in the post-Mitch survey strictly on the basis that they were located in areas that were: (a) affected by the hurricane; and (b) included in the original 1998 LSMS."

¹⁰Nicaraguan election results can be accessed via <http://www.dgapp-cse.gob.ni/>

¹¹For a detailed description of the LSMS surveys, see Grosh and Glewwe (1995) or the World Bank's LSMS website at <http://iresearch.worldbank.org/lms/lmsurveyFinder.htm>.

administrative unit in Nicaragua, similar to counties in the US) and I use this information to match the household-level information with the precipitation and political data.

In Nicaragua, the LSMS survey was administered three times: Wave I in 1998, Wave II in 2001, and Wave III in 2005. An important feature of these data sets is the timing, relative to Hurricane Mitch, of the periods reflected in the information on aid allocations.¹² Figure 1 presents a timeline of the relevant events (see the Appendix for all figures). The information collected in LSMS Wave I, conducted several months before Hurricane Mitch, reflects aid allocated during 1994 to 1998 (pre-Mitch). The information collected in LSMS Waves II and III reflects aid allocations during 1999 to 2001 and 2001 to 2005, respectively. I refer to these periods as “before”, “short term”, and “long term”.

The information on aid allocations in the LSMS reflects the household head’s answer to the question: “During the period from to , has any member of this household benefitted from ...[X]... ?” The types of activities (X) for which the question is asked include the construction or improvement of roads, health clinics, and schools; the installation or repair of latrines, sewers, and electricity; the provision of health information, nutritional education, job training, and legal assistance; and direct donations of food or medicine. A subsequent question asks about the entity responsible and allows me to determine whether the aid was provided by the government or NGOs.

The answer to this question provides a binary indicator of whether any member of the household benefitted from the particular aid activity in question. I combine the answers across the different aid activities to create a variable for each household indicating whether any member benefitted from aid provided by NGOs, and a separate variable indicating aid provided by the government.¹³ These are the two main outcomes of interest analyzed in this paper. They do not measure the intensive margin of aid, but this approach is the same as that used in other studies, such as Francken et al. (2012) and Morris and Wodon (2003), who point out that it is difficult to

¹²There was also a partial wave of the Nicaragua LSMS conducted in 1999, in the wake of the hurricane, but it was incomplete. Attempts to survey households in the hurricane-affected areas (as determined by INEC) were much less successful in the 1999 wave than in the 2001 wave, and no attempts were made to survey households in areas that were not affected by the hurricane. Thus, the 2001 wave provides the closest-in-time post-hurricane measure of aid across all levels of hurricane exposure.

¹³It could be informative to analyze the aid activities separately. However, the various sub-categories change across the survey years and consistency across surveys requires the combination of many categories.

build a reliable measure of the intensity of aid allocations even when more detailed information is available.¹⁴

The selection of the sample is discussed in greater detail in Section 2.5, but I note here that I focus on households that were surveyed both before and after Hurricane Mitch (i.e., in LSMS Waves I and II). There is the possibility that systematic patterns in attrition could influence the results, however I analyze patterns in attrition and test for this possibility (in Section 2.5) and find no such evidence.

Descriptive Statistics are presented in Table 1 (see the Appendix for all tables). The household characteristics used in the analysis as control variables in the regressions include the education and gender of the household head, household size, per capita household consumption, and indicators of urban locality, access to electricity, and ownership of a television. I also include as a control the dummy variable *INECDesignation* which indicates that a household was located in an area that the Nicaraguan government designated as “affected” by the hurricane. As mentioned before, it is not clear how this determination was made. I explore the possibility that this designation was politically influenced and discuss the results in the context of the other estimates in Section 2.6.

Table 1 also includes the “Rainfall Ratio” for Hurricane Mitch (*MRR*). This ratio is the main measure of hurricane impact used in the analysis, and its construction is described in detail next (in Section 2.3). Simply put, *MRR* is the ratio of rainfall during the hurricane period to average rainfall, measured at the municipality level. For example, $MRR_m = 1.75$ means that municipality *m* experienced 75 percent more rainfall during the hurricane period than its historical average.

Measuring Hurricane Impact with Precipitation Data

To measure the impact of the hurricane, I use precipitation data from version 3.21 of the Climatic Research Unit (CRU) time-series data set constructed by the British Atmospheric Data Center at the University of East Anglia, UK.¹⁵ While other forces during a hurricane,

¹⁴Aid is often not given as cash but as food and clothes or other goods and assistance, or via the construction of public goods, but it is not straightforward to measure the degree to which a public good, such as a re-constructed bridge, benefits one household differently than another, and households differ in their absorptive capacity to make use of clothes, medicine, and other private goods.

¹⁵The data are publicly available at <http://badc.nerc.ac.uk/home/index.html>

such as strong winds, can also cause damage, reports of damage from Hurricane Mitch and other hurricanes that have affected Nicaragua point most frequently to precipitation and the subsequent flooding and landslides as the main cause (Hellin et al., 1999).

The CRU precipitation data are constructed using information from over 4,000 weather stations worldwide. They are reported at a monthly frequency on a 0.5 degree latitude by 0.5 degree longitude grid, which is about 50km-square at Nicaragua's latitude. From each grid-month, I construct a measure of precipitation for each municipality in Nicaragua by interpolating a value for each 0.02 degree cell, using the four nearest grid observations weighted inversely by distance.¹⁶ I then take the average across all 0.02 degree cells contained within a municipality's borders. This is a similar approach to that taken by others when associating weather data with survey data that reports the geographic location of respondents by administrative area (e.g., Strobl, 2012).

Figure 2 depicts the three steps in this process: Panel A presents Nicaragua's municipal boundaries overlaid with the original 0.5 degree precipitation data grid; Panel B shows a color-scaled example of the subsequent distance-weighted estimates for each 0.02 degree cell; and Panel C shows a color-scaled map of the municipal averages for an arbitrarily chosen sample month (here, July 1998).

It is important to recognize that areas differ in their capacity to absorb precipitation and their propensity for flooding or landslides, conditional on a given level of precipitation. For example, a heavy rainfall of, say, four inches in one month is more likely to cause flooding and damage in an area that normally gets one inch of monthly rainfall compared to an area that regularly gets five inches. To incorporate this heterogeneity in resilience into the measure of a given hurricane's impact, I calculate the historical average precipitation for each municipality during each calendar month for the twenty-five years prior to each hurricane. I exclude hurricane months from this calculation so that this average reflects typical, non-hurricane precipitation levels.

I then take the impact of hurricane h in municipality m to be the ratio of precipitation during the month of the hurricane to the average precipitation (over the previous twenty-five

¹⁶These calculations were done using ESRI's ArcGIS software, v10.1.

years) that the municipality experienced during the same calendar month n in which hurricane h occurred.¹⁷ I refer to this as the hurricane Rainfall Ratio:

$$\text{Rainfall Ratio}_{mh} = \text{Precip}_{mh} / \overline{\text{Precip}_{mn}}$$

where

$$\begin{aligned} \text{Precip}_{mh} &= \text{precipitation in municipality } m \text{ during the month of hurricane } h \\ \overline{\text{Precip}_{mn}} &= \text{Average precipitation in municipality } m \text{ in the calendar month } n \text{ during} \\ &\quad \text{which hurricane } h \text{ occurred, excluding hurricane months, over the prior 25} \\ &\quad \text{years} \end{aligned}$$

The Rainfall Ratio during Hurricane Mitch (MRR) ranges from a low of 1.59 to a high of 2.05. The average is 1.75, which means that each municipality experienced, on average, 75 percent more precipitation during the month of Hurricane Mitch (October, 1998) than it typically experienced in the month of October over the previous 25 years.

Geographic Patterns in the Impact of Hurricanes in Nicaragua: 1960-2010

This section presents a summary of the geographic variation in hurricanes across Nicaragua since 1960 (when hurricane tracking data for this region first became available). An important feature of the evidence is that the impact of hurricanes across Nicaragua is reasonably uniform and, therefore, estimates of the influence of Hurricane Mitch on the allocation of aid are unlikely to be biased by unobservable factors that might be correlated with hurricane risk, or driven by hurricane risk in general rather than the specific impact of Hurricane Mitch. In the results, I also control for the impact of Hurricane Cesar, the most recent hurricane prior to Mitch.

The hurricane Rainfall Ratio (RR) introduced in Section 2.3 reflects the variation in hurricane intensity geographically within Nicaragua during a single hurricane, and across different hurricanes. I construct this measure for every municipality for each of the fourteen months during which a hurricane passed within 200 miles of Nicaragua between 1960 and 2010.¹⁸ Figure 3 shows the distribution of this measure across all of these hurricanes.

¹⁷There are no hurricanes in the sample whose impact spanned across multiple months.

¹⁸Hurricane strength tropical cyclones are typically about 300 miles in diameter, though with considerable variation. The 200-mile cutoff represents a natural break point in proximity to Nicaragua among all hurricanes that occurred in the greater Central American region during the period studied. This set of hurricanes was determined using the Weather Underground's Hurricane Archive, which can be accessed at <http://www.wunderground.com/hurricane/hurrarchive.asp>

The average RR , across all hurricanes and all municipalities, is 1.26. This means that during a month in which a hurricane passes within 200 miles of Nicaragua, the average municipal rainfall is 26 percent higher than its non-hurricane seasonal average. (For Hurricane Mitch, the RR ranges from 1.59 to 2.05.) Figure 4 shows the timeline of hurricanes affecting Nicaragua. Each vertical scatter-plot depicts the distribution (across municipalities) of the RR values during each hurricane.

To see the impact of hurricanes leading up to Mitch, which might have lagged effects on aid flows, Figure 5 contains maps that depict the RR for Mitch and the three previous hurricanes. Each of these maps is on its own scale, which highlights the geographic (cross-sectional) variation in the impact of each hurricane. The impact from Hurricane Mitch was largest in the north-central region of Nicaragua. During the three hurricanes before Mitch, the largest impacts were experienced in the south-west during Cesar (1996), the north-west during Gert (1993), and the north-east during Joan (1988).

To get a sense of overall hurricane risk, rather than recent impacts, Figure 6 shows the point estimate, for each municipality, of the mean impact across all hurricanes that affected Nicaragua from 1960-2010, along with 95 percent confidence intervals. The majority of the estimates fall within a tight band around the national average, and no municipality has a mean impact that differs from the national average at a significance level of five percent. The variation of the mean impact across municipalities can be seen geographically in the top map in Figure 7. This figure also includes maps of the impacts during Mitch and the three prior hurricanes, all on the same scale, to provide context to the map of the average municipal impacts (this also reveals the time-series variation across these hurricanes).

In addition to considering the intensive margin of hurricane impact, it is useful to consider the extensive margin. Table 2 shows the municipal incidence of hurricanes using a higher ($RR > 1.75$) and a lower ($RR > 1.50$) threshold to indicate whether a municipality was affected. Between 1960 and 2010, nearly all municipalities—138 out of 143—experienced between 2 and 4 hurricanes as measured by the lower threshold. Using the higher threshold, most municipalities (98 out of 143) experienced one hurricane, and no municipality experienced more than two.

Taken altogether, the evidence presented shows that the impact of hurricanes, along both the intensive and extensive margins, is remarkably similar across Nicaraguan municipalities.

Empirical Strategy

The basic idea, to identify the effects of a catastrophic storm on the micro-level distribution of aid resources, is to (a) estimate the relationship between the impact of the storm and the subsequent distribution of aid and (b) establish convincing evidence that such estimates are unlikely to be driven by factors coincidentally related to the impact of the storm, including prior storms or historical patterns in storm risk. The extent of the data on aid allocations available for Nicaragua, which spans from four years before to seven years after Hurricane Mitch, allows for this.

I model household h 's receipt of aid separately from each source—either government or NGO—and for each period: four-to-zero years before the hurricane, the short term (zero-to-three years after the hurricane), and long term (three-to-seven years after). The basic models for the receipt of aid, contributed to by households h in municipalities m , are captured in Equation 2.1,

$$AID_{hm} = \alpha + \delta_1 MRR_m + \delta_2 Sandinista_m + \delta_3 (MRR_m \times Sandinista_m) + \beta X_h + \epsilon_{hm} \quad (2.1)$$

where I also control for household characteristics, X_h , that may influence the probability that a household receives aid. I present estimates from linear probability models, though the estimates from binary logit models are qualitatively similar to the presented results (and are available upon request).

The main explanatory variables of interest are the municipality-level precipitation-based measure of hurricane impact (MRR_m) and the political affiliation of the leadership in municipality m (captured by the indicator, $Sandinista_m$). Any differential effects will therefore be captured by the interaction of these two covariates ($MRR_m \times Sandinista_m$). These measures vary at the municipal level, so to allow for the possible correlation of unobserved factors between households in the same community, the standard errors are clustered at the municipality level. (There are 125 municipalities represented in the LSMS data set.)

Challenges to Identification

There are two obvious challenges to identifying the causal effect of a hurricane on aid allocations. First, in order to infer from $\hat{\delta}_1$ the causal effect of Hurricane Mitch on aid receipt,

MRR is assumed to vary independently of unobserved factors that themselves influence the distribution of aid. One way the estimates might be biased by a correlation of the hurricane's impact with an omitted variable is if the impact of hurricanes is correlated over time. Past hurricane impacts could affect subsequent aid allocations through their effects on prior aid allocations, or on other factors that affect aid, such as levels of economic development.

However, the evidence presented in Section 2.4 suggests that this is unlikely to be the case—the historical impact of hurricanes across Nicaragua is reasonably uniform. This increases our confidence that regression estimates of the impact of Hurricane Mitch on aid allocations are unlikely to suffer from bias due to omitted variables related to hurricane risk. In the results, I also control for the 1996 impact of Hurricane Cesar, the most recent hurricane prior to Mitch. Hurricane Cesar was significantly smaller than Mitch and not nearly as destructive, though the extent to which it may have affected aid flows is an empirical question. Including Hurricane Cesar as a control leaves the estimates of interest largely unchanged while their precision increases.

Another concern regarding the potential for unobserved variables to influence the results is that the impact of Hurricane Mitch might be related coincidentally to some other, unobserved, factor that influences aid. To speak to this, and in the spirit of a falsification test, I use data on aid allocations during the period before Hurricane Mitch to verify that the estimates I present are not simply reflective of pre-hurricane patterns of aid distribution.

Second, in conditioning the variation in AID_h on controls I am separately capturing potential confounding influences so as to leave the effects of the key variables identified. In so doing, it becomes important that these controls not be outcomes of the hurricane themselves—that is, caused by MRR_m . As it is possible that household characteristics measured *after* the hurricane could both (a) be affected by the hurricane and (b) affect the probability that a household receives aid, I use pre-hurricane measures of controls, collected in the 1998 wave of the LSMS.¹⁹

Focusing on this sample of households that is surveyed both before and after the hurricane introduces the possibility that systematic patterns in attrition from the survey could affect the results. The rate of attrition between LSMS Waves I and II is 26 percent, which is not unusually

¹⁹ In the language of Angrist and Pischke (2008), the post-hurricane measures of these variables are “bad controls” because they were not pre-determined at the time of the hurricane. For example, the hurricane's destruction of agricultural lands could cause an individual to pursue more education, which might in turn increase the probability that he receives aid. If there are positive effects from both MRR and education on the probability of receiving aid, the inclusion of the post-hurricane measure of education in a regression will wrongly attribute part of the effect of the hurricane to education.

high.²⁰ Nonetheless, I conduct a formal test of whether there are systematic differences in household attrition across levels of hurricane impact, estimating a linear-probability model,

$$Attrit_{hm} = \alpha + \delta MRR_m + \beta X_{1998,h} + \gamma(MRR_m \times X_{1998,h}) + \epsilon_{hm} \quad (2.2)$$

where $Attrit_h$ equals one if 1998 household h was not resurveyed in 2001, and is equal to zero otherwise. Estimated standard errors again allow for clustering at the municipal level.

Table 3 presents estimates of the parameters in Equation 2.2. In the simplest specification, the probability of attrition is increasing in hurricane exposure (Column 1), though this relationship is not robust to the inclusion of household control variables, as shown in Column 2. There are some systematic relationships between attrition and household characteristics (Column 2), however there is no systematic difference in attrition across levels of hurricane exposure and any of the key variables of interest (aid from the government, aid from NGOs, or Sandinista representation), or any of the other independent variables for that matter, as indicated by the absence of any statistically significant coefficients on the interaction terms in Column 3. A test of the joint significance of the variable indicating receipt of aid from the government and its interaction with MRR fails to reject the null hypothesis of no relationship ($p = 0.67$), as does a similar test for NGO aid ($p = 0.32$). An F-test of the joint significance of all of the interaction terms also fails to reject the null hypothesis that all terms are zero ($p = 0.25$). Thus, there is no evidence of systematic attrition that would bias estimates of the impact of the hurricane or Sandinista representation (or their interaction) on aid allocations from either source.

Results

I present the estimated effects of Hurricane Mitch on aid allocations separately for each source of aid (government or NGOs) and each period (short or long term). This reveals the evolution of the response to the hurricane for each source of aid. The estimates from the period before the hurricane serve to verify the identification strategy.

²⁰This is close to the attrition rates in the Panel Study of Income Dynamics in the U.S. (see Fitzgerald et al., 1998), and comparable to attrition in other panel surveys in developing countries (see Deaton, 1997). Of the 4,020 households surveyed in 1998 with complete data, 74 percent were surveyed again in 2001. Of these, 89 percent were re-surveyed in 2005.

The main results showing the determinants of aid allocations by the Nicaraguan government in the first three years following Hurricane Mitch are shown in Table 4. The estimated coefficient on MRR in Column 1 is not statistically different from zero, suggesting that the government *did not* allocate more aid to households that experienced greater hurricane damage as measured by precipitation. This estimate remains statistically insignificant across subsequent specifications, and though it is positive, the estimated relationship is weaker than in the period before the hurricane. That the propensity for households in hurricane-affected areas to receive aid is less closely related to hurricane strength after the hurricane than before, however, may be washed out by the across-the-board increase in the level of government aid provided to households (59 percent of all households received some form of government aid in the period before the hurricane, whereas 62 percent received government aid in the period after the hurricane). Column 2 of Table 4 suggests that households in communities that INEC designated as “affected” by the hurricane (as revealed by INEC’s decisions to collect information on household conditions in these communities after the hurricane) were 18 percent more likely to receive aid from the government ($0.122/0.62=0.181$). This result is robust to the inclusion of MRR, as seen in Column 3.

It is difficult to say whether the INEC designation was politically manipulated or is truly reflective of hurricane damage that is not captured by the precipitation-based measure. To explore this, I model the INEC designation as a function of MRR, household characteristics, the Sandinista dummy, and a measure capturing the average slope of the land in each municipality, which might further increase the probability of flooding and landslides.²¹ These estimates, presented in Table 5, are largely ambiguous in revealing how the designation was made. The designation is unrelated to the political affiliation of the municipal leaders, only weakly correlated with MRR and not related to the topographical slope of land in the municipality.

The specification presented in Column 4 of Table 4 includes the variable *Sandinista*, which indicates whether the political leadership of a household’s municipality is affiliated with the main opposition party to the ruling coalition. There is no affect of this affiliation on average (Column 4), but the estimates in Columns 5 and 6 reveal a significant relationship between the interaction of the hurricane impact and municipal political affiliations. The relationship is striking: consistent with their complaints at the time, many Sandinista areas were much less likely

²¹Slope is measured using the standard deviation of land elevation within each municipality.

to receive government aid. This effect is concentrated among the municipalities that experienced a smaller impact during the hurricane. Among households in the least-affected municipalities (MRR=1.59), those represented by the Sandinistas were about 26 percent (all percents relative to the mean) less likely to receive governmental aid than those represented by the ruling coalition (using point estimates from Column 5, the difference in these areas in the propensity to receive aid associated with Sandinista representation is $(-1.157+0.625 \times 1.59)/0.62 = -0.263$). Column 6 includes the estimated interactions between *Sandinista* and dummy variables for MRR quartile (using the 3rd highest quartile as the reference case). These point estimates again suggest that in the municipalities that experienced less than the median impact, households in Sandinista-led municipalities were around 22 to 25 percent *less likely* to receive aid from the government than households in areas represented by the ruling coalition. There is no evidence that political affiliations affected the government's allocations of aid in municipalities in the top quartile of hurricane impact, which might indicate that the ruling coalition felt that political manipulation of aid allocations in these areas would be inappropriate (or, from a more cynical perspective, more noticeable).

Estimates of the parameters from models of the household receipt of assistance from NGOs in the three years following the hurricane are presented in Table 6. These estimates show that, during this period, the NGO response to Hurricane Mitch was different from the government's response in two fundamental ways. First, aid allocations by NGOs were *not related* to political affiliations. Second, unlike government aid, there is evidence that NGOs responded to the physical impact of the hurricane when allocating aid. The probability that a household received aid from an NGO was significantly higher in municipalities that were hit harder by the hurricane as measured by MRR. Using the point estimate (0.371) of the impact of MRR on NGO aid from Column 1, the results suggest that households in the hardest hit areas were 74 percent more likely to receive NGO aid than those that experienced the average impact.²² This effect is large on its own, and perhaps appears larger in the context of the non-response by the government. However, the percentage of households receiving NGO aid was the same in the period following the hurricane as in the period before (in contrast with the 3 percentage point increase in households receiving aid from the government). Like the government, NGOs were more likely

²²The calculation is $[(2.05-1.75) \times 0.371]/0.15 = 0.742$

to allocate aid to households in areas that received INEC's designation, and the coefficient is essentially identical to that estimated for government aid.

The last wave of the LSMS, conducted in 2005, sheds light on the longer term evolution of the effect of the hurricane on aid allocations. Table 7 presents the estimates of the models of longer-term aid allocations, from both NGOs and the government, three to seven years after the hurricane (during 2001 to 2005), as reported by the 2,645 remaining households from the main sample. The estimates suggest that by this time, the Nicaraguan government was allocating more aid to areas that experienced more damage during Hurricane Mitch, with households in the most affected areas about 17 percent more likely to receive aid than households that experienced the average impact.²³ Furthermore, the influence of political affiliations on government aid allocations seen in the earlier period are now gone. The long-term behavior of NGOs is similar to their shorter-term behavior, in that they continue to allocate significantly more aid to areas that were hit harder by the hurricane.

It is perhaps not surprising that Hurricane Mitch has had such a long-term effect on aid allocations in Nicaragua, given the extensive damage it caused. During the early 2000s, among all countries in the western hemisphere, only Haiti had lower per-capita GDP than Nicaragua. Repairing the 1500 miles of roads and 300 schools that were destroyed requires substantial effort, especially in the poor, rural parts of Nicaragua that were hit hardest by the hurricane.

Falsification Test

A straightforward way to test the identification strategy and to verify that the estimates are not driven by a relationship between unobservable patterns in the usual allocation of aid and the impact of Hurricane Mitch can be carried out by conducting the same analysis on the allocation of aid before the hurricane, which was reported in the 1998 wave of the LSMS. These aid allocations should be unrelated to the future impact of the hurricane. Thus, testing whether there are any relationships between pre-hurricane aid allocations and the eventual impact of the hurricane serves as a falsification test for the identification strategy.

Table 8 presents estimates of the parameters for models of aid allocation before the hurricane using the same sample of households (the non-attriters) and the same independent

²³Using the point estimate from Column 1, the calculation is $[(2.05-1.75)*0.448]/0.78=0.172$

variables used in the post-hurricane models. None of the estimated coefficients on the hurricane impact variables have a statistically significant relationship with the allocation of aid by the government or NGOs, with one exception: there is a marginally significant positive association between the household receipt of aid from the government and the future INEC Designation. This is consistent with the concern of an endogenous relationship between this designation, and points to the importance of using a measure of hurricane exposure not chosen by the government. The lack of any significant relationship between MRR, the precipitation-based measure of hurricane exposure, and assistance received by households prior to the hurricane supports the causal interpretation of the estimated effects of MRR on aid allocations discussed in the previous section. Lastly, there is no pre-hurricane relationship between political affiliations and the allocation of aid by the government (or by NGOs).

Conclusions

This paper investigates the political economy of development aid allocations at the micro-level, using Nicaragua as an example. The fortuitous timing of the collection of extensive data on aid received by Nicaraguan households, relative to the catastrophic impact of Hurricane Mitch, allows for an examination of how non-governmental organizations (NGOs) differ from the domestic government in their response to an arguably objective measure of need (the impact of a hurricane). The extent of the available data also allows for important methodological improvements on existing studies of disaster aid and the first estimates, to the best of my knowledge, of the long term impacts of a natural disaster on micro-level development resource allocations. Furthermore, despite predictions that the intensity of tropical cyclones will increase substantially over the next century (Scheraga et al., 2003; Solomon et al., 2007; Pachauri and Reisinger, 2007; Knutson et al., 2010), there is limited empirical evidence on the micro-level aid response to these storms specifically, and to natural disasters in general. This study contributes to filling that gap. The setting studied is also relevant because Nicaragua has similarly low levels of development to many other areas prone to extreme storms, and Hurricane Mitch, while perhaps an outlier in terms of storm severity at the time it occurred, is arguably representative of the extremely powerful hurricanes that are predicted to occur with greater frequency (see Knutson et al., 2010).

To estimate the effects of a natural disaster on aid allocations requires appropriately timed data collection, and we are fortunate to have such an opportunity in Hurricane Mitch's arrival in Nicaragua. I use survey data containing information on aid received by Nicaraguan households that corresponds to three separate blocks of time: zero to four years before the hurricane, zero to three years after (the "short term"), and again three to seven years after (the "long term"). The measures of aid included in the survey are broad and do not permit a detailed analysis by aid type. However, they reflect a household's access to basic resources that are important to economic development, including the construction or repair of roads and schools, improvements in access to electricity and drinking water, and direct provisions of food and medicine. With these data, I estimate the effect of the storm on short- and long-term aid allocations and then test for differences across political affiliations, measured through municipal election outcomes.

In addition to the use of extensive data on aid allocations, this research improves on existing studies that estimate the effects of natural disasters on aid allocations in two other important ways. First, in order to alleviate concerns that the choice of disaster impact designations might be influenced by factors unrelated to damage, such as political connections, I use precipitation data recorded at a relatively fine geographic scale to construct an exogenous measure of the hurricane's impact. While it is possible that official reports of disaster damage could reflect strategic efforts to influence the allocation of aid, these precipitation data should be free from any such manipulation.

Second, I carefully consider the potential for the estimates to be biased by unobserved factors that might influence the distribution of aid. It is possible that patterns in hurricane risk, as opposed to the incidence of any particular storm, could determine aid allocations either directly or indirectly through a correlation with unobserved characteristics (such as prior aid allocations) that might also systematically affect the distribution of aid. I conduct a detailed examination of geographic variations in the historical impact of hurricanes across Nicaraguan municipalities. This assessment reveals that no parts of Nicaragua are particularly more or less likely to be struck by hurricanes and therefore suggests that hurricane impact, at least in Nicaragua, is essentially independent of geographic or population characteristics that might influence the distribution of aid. The data on pre-storm aid allocations also allow me to consider the possibility that the impact of the hurricane was coincidentally related to existing patterns in aid distribution. I

find no evidence of a such a relationship, which, along with the use of exogenous rainfall data, supports the causal interpretation of the estimates I present.

The analysis of aid allocations made by the domestic government separately from those made by NGOs proves to be informative. Following the catastrophe, short-term allocations of aid by NGOs were made according to the precipitation-based measure of the hurricane's impact, with the most severely affected households about 74 percent more likely to receive aid from NGOs than those that experienced the average impact. Furthermore, I find no evidence that NGO aid was influenced by political affiliations. These results stand in contrast to the allocation of aid by the government, which also responded to the hurricane in the short term, but not as one would hope. Among areas that experienced less than the average hurricane impact, the government was significantly *less likely*, by about 26 percent, to provide aid to households in areas led by the major opposition party (relative to similarly affected households in areas controlled by the ruling coalition). While aid was diverted away from these opposition-controlled areas, the government *did not allocate more aid* to the areas that were more heavily affected by the hurricane, as measured objectively by rainfall. There was an increase in the proportion of households receiving aid from the government across the sample period (increasing from 59 percent in the earliest period to 78 percent in the latest), which may have been influenced by the hurricane and the international aid it attracted. However there is no evidence that the distribution of aid from the government was related to hurricane severity, as captured through the precipitation-based measure used in this analysis, in the three years following the hurricane.

I also find that Hurricane Mitch had long-lasting impacts—extending at least three years after the storm—on the allocation of aid in Nicaragua. As in the short term, NGOs were more likely to allocate aid in the long term to households that experienced a more severe impact during the hurricane. Unlike in the short term, the Nicaraguan government actually did allocate more aid in the long term to households that experienced a stronger impact from the hurricane. During this time, households that experienced the largest hurricane impact were about 17 percent more likely to receive aid from the government than those that experienced the average impact. That a hurricane can have such a long-term affect on the allocation of aid within a country is perhaps surprising. However, Mitch was the deadliest Atlantic hurricane in over 200 years and caused significant destruction across large parts of Nicaragua. Repairing such extensive damage to

roads, bridges, agricultural lands, and other infrastructure would be challenging even in wealthy countries, and is likely to be more difficult in a developing country such as Nicaragua.

Given the large differences in the responsiveness of NGOs and the government to the catastrophe, and their differential sensitivities to the political affiliations of the municipalities affected, it is tempting to interpret these results as somewhat disheartening. On the one hand, the political maneuvering reflected in the post-hurricane governmental allocation of aid, and the lack of a response to the physical measure of hurricane damage, suggest that aid might be transmitted more effectively through non-governmental channels. On the other hand, the results may be specific to a context that is characterized by low levels of economic development and contentious politics. As such, any application to other contexts should proceed with caution. That said, post-disaster reports in the media, such as those in the Philippines following Typhoon Haiyan in 2013, often claim that disaster recovery efforts are manipulated for political gain, and the analysis presented here is consistent with these sentiments.

CHAPTER III

EXTREME HEAT DURING GESTATION AND EARLY-LIFE HEALTH: EVIDENCE FROM PERU

Introduction

There is a growing concern that human-produced greenhouse gas emissions will alter global temperatures and other climatic conditions. Understanding the full costs of these emissions, and designing policies to help people adapt to the associated changes in the climate, requires estimates of the health costs associated with changing climate conditions. As a consequence of these greenhouse gas emissions, extreme temperatures are predicted to occur with increasing frequency, and studies that provide estimates of the direct effects of extreme temperatures on economic and health outcomes have been increasing steadily in number.¹ In this paper, I present estimates of the relationships between extreme heat during gestation, a child's weight at birth, and the individual's subsequent height during early childhood (0 to 59 months old) for a sample of children in Peru. I then explore whether a child's birth weight fully captures the negative effects of extreme heat during gestation, and whether mothers invest differently in their children based on their birth weight and exposure to extreme heat during gestation.

Most existing research concerning the effects of temperature extremes on fertility, birth, and later-life outcomes analyzes data from wealthy countries. Understanding the effects of temperature extremes on human health during the early stages of life is particularly relevant for developing countries, however, where low levels of basic infrastructure will amplify the challenges of adaptation (Mendelsohn et al., 2006). Furthermore, most climate change models predict disproportionately large increases in temperature for the developing parts of the world.

A growing body of research suggests that maternal stress during gestation can reduce the birth weight of the in utero child (e.g., Almond and Mazumder, 2011; Almond and Currie, 2011; Brown, 2014). A related literature suggests that physical characteristics of a young child or infant, such as birth weight, can predict outcomes later in life with generally large and positive effects

¹A variety of the potential consequences for human health resulting from extreme temperatures, and other extreme weather events, have received increasing attention in the academic literature and popular press. Common outcomes that have been studied include fertility (e.g., Barreca et al., 2015), infant mortality (e.g., Kudamatsu et al., 2012), birth weight (e.g., Deschênes et al., 2009), later-life educational and labor force outcomes (e.g., Wilde et al., 2014), and mortality (e.g., Burgess et al., 2011; Deschênes and Greenstone, 2011; Barreca, 2012; Barreca et al., 2013).

on income, health, educational attainment, and other outcomes (Black et al., 2007; Almond and Currie, 2011; Currie and Almond, 2011). The influence of child and family characteristics (measured before a child enters school) in predicting future outcomes is “striking” relative to more typically considered factors, such as years of education (Currie and Almond, 2011).

Whether the effects of a shock during these early stages of development can be overcome via parental investment, or other mitigating activities, is important to our understanding of whether the harmful effects of temperature shocks can be mitigated after the fact, or can only be avoided by limiting exposure from the beginning. This research addresses whether extreme hot temperatures during gestation may have negative consequences for the child’s early physical development after birth, and whether these effects can be fully explained by the child’s weight at birth. If temperature shocks have effects on a child’s development beyond those reflected by birth weight, then estimates of the effects of temperature extremes on birth weight will not convey the complete developmental costs of exposure to temperature extremes during gestation.

Although human capital production functions are extremely difficult to estimate (because of their data requirements), the notion of such a function provides a useful framework for understanding the potential relationships between in utero environmental shocks and subsequent outcomes. By outlining the relationships between maternal stress during gestation, health at birth, and health in early childhood, the framework adapted in this paper provides an idea of the extent to which the effects of shocks during gestation may be revealed only after a lag and, thus, not fully reflected simply in birth weights. As in much of the related literature, this paper focuses on temperature shocks during gestation because there is evidence that the time during which a child is in utero is a “critical period” for development, while there is not much evidence of such critical periods after birth (Currie and Almond, 2011).

Data constraints have limited the scope of the existing research on these questions to wealthier countries with more complete data. This paper extends the evidence to a developing-country context by combining high-resolution gridded climate data on monthly temperatures with current height and retrospective birth weight measures collected for children between 0 and 59 months of age in households surveyed in the Demographic and Health Surveys (DHS) conducted in Peru since 2000. Peru serves as a useful setting for the study of these relationships, both because of the availability of data from several waves of the DHS surveys and because the

considerable heterogeneity across micro-climates within Peru provides a useful source of variation in temperature extremes.

The DHS surveys are representative of the population in Peru at the time of survey, but the incomplete availability of all of the measurements required for this analysis may result in a systematically different sample than one representative of the population in Peru and/or the universe of all births in Peru. Exclusion from the sample via infant mortality is one obvious selection mechanism, while a second is the omission of reliable measures of birth weight. I focus on the sample of children who have hospital cards available at the time of survey. This sampling decision trades increased reliability of the birth weight measurements against the possibility that birth in a hospital (and thus the presence of a hospital card at the time of survey) may yield a sample that is not fully representative of the general population of births in Peru.

Within this sample, I focus on mapping the translation of health at birth into health later on, and then estimate how exposure to extreme heat during gestation affects this translation. It is likely that there are unobserved factors that affect both a child's birth weight and her subsequent height, thus the estimated relationship between birth weight and height should not be interpreted as causal. Rather, this relationship is estimated first without, and then including, exposure to extreme heat during gestation in order to see if there is a relationship between in utero shocks and physical development after birth. The four main measurements of interest considered in the analysis reflect a natural chronology: (1) the infant is exposed in utero to the temperature conditions prevailing in the mother's location, (2) the infant is born and weighed, (3) the infant survives to the time of survey, and (4) the child's height is measured. The goal of the analysis is to determine how (1) affects (2) directly, and how (1) affects the relationship between (2) and (4).

The results suggest that a period of extreme heat during the three months preceding birth is associated with a lower birth weight among male infants, while there are no detectable effects on females. I also find that such a period of heat corresponds to reduced height later in life in boys. The magnitude of this effect on height is reduced by roughly one-third when birth weight is included as a control, and becomes only marginally significant, suggesting that much but not all of the developmental harm caused by extreme heat during gestation is reflected in birth weight. It is difficult to say whether these effects are working through biologic or economic channels. Potential mechanisms are discussed in Section 3.5.

Background

Physical Development During Early Childhood

The establishment of standards for human physical development and growth rates (referred to as height or weight “velocity” in the medical literature) has long been of interest to the medical community (Tanner et al., 1966; Tanner and Whitehouse, 1976). Early analyses typically relied on data from British populations, but in the last decade the World Health Organization (WHO) has established global standards for children less than five years of age using data on 8,440 “healthy breastfed infants and young children from diverse ethnic backgrounds and cultural settings” (World Health Organization, 2009). The WHO standards were explicitly designed to reflect the healthy development of individuals likely to achieve their “full genetic growth potential.”²

These studies show that typical development during the first five years of life proceeds with an initial period of rapid growth, up to about the first year of age, followed by slower, near-linear growth in the subsequent four years. Figure 8 depicts average height by age for the sample considered here and reveals a similar pattern. Growth velocities vary substantially across children, however, and over time for a given child, and “catch-up” growth is expected following periods of low growth velocity due to illness or other environmental shocks (World Health Organization, 2009). Given this within-child variability in growth rates, and the potential for catch-up growth, the remarkable predictive power of birth weight is all the more impressive. It remains unclear, however, whether shocks or interventions that affect birth weight have only temporary effects or if they actually translate into different outcomes later on. This paper attempts to answer this question by exploring the relationship between a significant temperature shock in utero and health later in life, controlling for an individual’s birth weight.

Extreme Heat During Gestation

A growing body of evidence suggests that maternal environmental stress during gestation—such as exposure to violent conflict (Brown, 2014), earthquakes (Torche, 2011), or extreme heat (Deschênes et al., 2009; Cil and Cameron, 2014)—can have negative consequences for in utero infants, including reduced birth weights and shorter gestational lengths. Using data from the U.S.,

²As of 2011, 125 countries had adopted the WHO standards (de Onis et al., 2012).

both Deschênes et al. (2009) and Cil and Cameron (2014) find that the negative effect on birth weight from extreme heat during gestation is largest during the second and third trimesters of pregnancy.

There is also evidence that temperature extremes can affect fertility. Barreca et al. (2015) find that high temperatures in the U.S. reduced birth rates eight to ten months later, with this initial decline followed by a subsequent increase in birth rates eleven to thirteen months after the temperature extreme.

Given the evidence that birth weight predicts outcomes later in life, it is natural to ask whether shocks during gestation affect later outcomes, or whether they cause only cause temporary disturbances to physical development. As Deschênes et al. (2009) summarize, an “important question left unanswered...is whether these changes in birth weight are related to welfare losses, for example, through changes in health, human capital, income, or neonatal mortality. If there is indeed a relationship between in utero exposure to high temperature and welfare measures, then these impacts should become part of the calculations of the benefits of greenhouse gas reductions.”

Wilde et al. (2014), using data from sub-Saharan Africa, present evidence that heat waves can in fact have long-term effects on important outcomes, but their evidence suggests that these effects are *positive*: they find that heat waves around the time of conception correspond to higher educational attainment and literacy, fewer disabilities, and lower child mortality. They argue that this perhaps counterintuitive finding is most likely due to fetal selection. Contrary to the fetal origins hypothesis, and also contrary to some of the evidence presented here, Wilde et al. (2014) find no effects of hot temperatures during gestation on health or education later in life.

Data

Two data sources are used in this analysis. Extreme heat is measured using monthly observations of average local temperatures reported by the British Atmospheric Data Center. Children’s heights, weights at birth, and the demographic characteristics of their mothers come from the geo-coded MEASURE Demographic and Health Surveys (DHS) conducted in Peru.

The DHS, funded by the U.S. Agency for International Development, includes extensive information on household demographics as well as birth histories for women aged 15-49 in

surveyed households, and measurements of height and weight for the children in the household. These surveys were conducted in Peru in different waves: once in 2000 and continuously from 2004 to 2009. Table 9 displays the number of children living in surveyed households during each year. For this analysis, I use all observations for children that have complete height, birth weight, and location information (location is reported using the geographic coordinates of the center of the survey cluster). The total numbers of births for all women included in the Peru DHS surveys are depicted by month in Figure 9. There are no strong seasonal patterns in fertility in this sample of over 30,000 Peruvian women.

The temperature data come from Version 3.21 of the Climatic Research Unit (CRU) time-series data set constructed by the British Atmospheric Data Center at the University of East Anglia, UK.³ The CRU data are constructed using information from over 4,000 weather stations worldwide. The data are reported at a monthly frequency on a 0.5 degree latitude by 0.5 degree longitude grid, which is about a 55 km square at Peru's latitude. Average monthly temperatures for the entire country, over the period from 1960 to 2010, are presented in Figure 10. While Peru is known for having particularly diverse micro-climates, the average monthly temperature for the country remains fairly steady between 70 and 73 degrees F throughout the year.

I assign temperatures to DHS respondents using the CRU grid point nearest to the center of each household's survey cluster. Figure 11 depicts the geographic locations of the centers of each survey cluster included in the Peru DHS data, overlaid with the grid-points at which the CRU temperature data are reported. I measure exposure to extreme hot temperatures during gestation by using an indicator for whether the average monthly temperature at grid point g was greater than two standard deviations above that grid point's historical average, calculated from 1960 to 2010, during any of the three months comprising each trimester of an infant's gestation.⁴

In Table 10, I present summary statistics for the distributions of birth weights and heights across children, along with demographic characteristics of their households and the measures of exposure to extreme heat during each of the three trimesters of gestation. The first set of statistics (left panel) in Table 10 shows the means, standard deviations, and sample sizes for each of these variables for all children included in the DHS. The second set of statistics (right panel) is

³The data are publicly available at <http://badc.nerc.ac.uk/home/index.html>

⁴Because the exact date of conception is unknown, the trimesters are approximated using three-month blocks of time going back in time from the date of birth.

restricted to the sample of children that have no missing information among the key variables and can thus be included in the analysis in this paper (these variables are birth weight, height, and exposure to heat during gestation (the last of which is missing only if location information for the child's household is missing)).

Sample Selection

As mentioned, there are several important ways in which the sample analyzed here may not be fully representative of all births in Peru during the period studied, and any inference from the estimates constructed with this sample must thus be accompanied by that important caveat. The first way by which the sample considered here may differ from the set of all live births in Peru is that children included in this analysis have survived until the time of the DHS survey. If mortality is inversely related to physical size, as much of the evidence suggests (e.g., Black et al., 2007), then selection into the sample in this way will likely be positively correlated with birth weight and height. If extreme heat during gestation can have sufficiently harmful effects to increase infant mortality, then this sample may exclude the individuals that would reflect the most harmful effects of extreme heat during gestation. For the sample of all births to DHS mothers that occurred within five years leading up to the survey date, Table 11 presents estimates of a regression of an indicator of whether that child survived on demographic characteristics of the mother and exposure to heat during gestation. There appears to be no increase in infant mortality due to heat exposure during gestation in this sample of births.

Second, while the vast majority of births in the U.S. and other wealthy countries occur in hospitals, or with health professionals in attendance, this is not the case in many developing countries. Thus, measures of birth weight are often not available. In this paper, I focus on the sample of children who had health records available at the time their household was surveyed and thus can be matched to a reliable measure of their birth weight. This restricts the sample to children who were both (a) born in a health facility, and (b) whose household has the child's health records available to show to the enumerator. Perhaps due to this restriction, the children included in the sample for this analysis (relative to those surveyed in the DHS but excluded from the analysis because of incomplete information) have better educated mothers, are more likely to live in an urban environment, and come from wealthier households. The children are also younger,

which may be due to the fact that possession by the household of the child’s hospital card is required for inclusion in the sample, and hospital cards may increasingly tend to become lost over time.

Third, the parents of just under 3.5 percent of the children living in the surveyed households refused height and weight measurements for their children, and some children were not present in the home at the time of the DHS survey (though repeated visits were made). The data contain height measurements for 88 percent of all eligible children (i.e., younger than 60 months of age) living in surveyed households.

Fourth, location information is not available for all households and thus their survey information cannot be matched with the temperature data. The geographic coordinates are missing for a small fraction (less than one percent) of respondents, while eight percent of otherwise eligible children have changed location (determined using information on each mother’s tenure at her current residence) between the date of the survey and one year before the child’s month of birth (thus spanning the child’s time in utero).

Empirical Strategy

To determine whether an infant’s birth weight can be expected fully to capture the health effects of exposure to extreme hot temperatures during gestation, I proceed in three steps. First, I estimate the relationship in this sample between birth weight and height at 0 to 59 months of age. As mentioned, this relationship is not intended to be interpreted causally. Rather, it serves to provide an understanding of the typical relationship between these measures within this sample. Second, I present evidence that the effects on birth weight of extreme heat during gestation are similar in this sample of children to other estimates in the literature. Third, I estimate the effect of extreme heat during gestation directly on a child’s height, comparing models with and without birth weight included as a control variable to see what, if any, incremental explanatory power can be attributed to exposure to extreme heat during gestation beyond that captured by birth weight.

To estimate the effect of periods of extreme heat during gestation on a child’s birth weight, I estimate the following equation:

$$BirthWeight_{ig} = \alpha + \sum_{t=1}^3 \beta_t HOT_{ig}^t + \mu_{YOB} + \gamma_{MOB} + \delta_g + \eta_{BirthOrder} + \theta_{SurveyYear} + \epsilon_{ig} \quad (3.1)$$

The variables that measure exposure to extreme heat during gestation, HOT_{ig}^t , are indicators that take the value of 1 if, during any of the three months in trimester t of child i 's gestation, the average monthly temperature (measured at the nearest grid point g to child i 's household) was more than two standard deviations above that grid point's average, where the average and standard deviation of temperature are calculated for each grid point over the period from 1960 to 2010.

There are several potential reasons why extreme high temperatures during gestation or the first year of life may be correlated with physical health at birth and in early childhood, even if no causal relationship exists. For example, periods of high heat are associated with seasons, and there is evidence that month of birth may be correlated with subsequent outcomes (e.g., Buckles and Hungerman, 2013). It may also be the case that populations in areas subject to extreme heat, or extreme variations in temperature, are systematically different than populations in other areas. For example, coastal areas often have milder climates and higher levels of economic development.

To identify the effect of exposure to extreme heat using random weather fluctuations within a given location, I include locational fixed effects (δ_g) where g indexes the points on the half-degree latitude by half-degree longitude grid at which the monthly temperature data are reported. I also include fixed effects for year of birth, month of birth, birth order, and survey year. Standard errors are adjusted for clustering within location (g). If $\hat{\beta}_3$ (for example) is significantly different from zero, then exposure to extreme heat during approximately the third trimester of gestation is correlated with birth weight.

Before estimating the effect of extreme heat during gestation on subsequent height, I model the relationship between a child's weight at birth and his/her subsequent height. To reiterate, the medical literature shows that typical childhood physical development includes an initial period of rapid growth in height, through about twelve months of age, followed by a decrease in the growth rate and a near-linear progression through the fifth year of age (Tanner et al., 1966; Tanner and Whitehouse, 1976; World Health Organization, 2009). I include both age and age-squared in the models of height to capture this relationship. More specifically, I estimate:

$$\begin{aligned}
 Height_{ig} = & \alpha + \beta_1 Age_i + \beta_2 Age_i^2 + \mu_{YOB} + \gamma_{MOB} + \delta_g + \eta_{BirthOrder} + \theta_{SurveyYear} \\
 & + BirthWeight_i \times (\beta_3 + \beta_4 Age_i + \beta_5 Age_i^2) + \epsilon_{ig}
 \end{aligned} \tag{3.2}$$

for child i born at location g . I again include fixed effects for location as well as year of birth, month of birth, birth order, and survey year. The interactions of birth weight with the age variables are included to determine whether a child's weight at birth influences the intercept, or the trajectory of their subsequent growth, or both. I also consider models with dummy variables for age intervals to model the growth path more flexibly. Standard errors are adjusted to allow for clustering at the location level (g). Finally, to reveal the extent to which exposure to extreme heat during gestation may affect physical development in ways that are not reflected in a child's birth weight, I add HOT_{ig}^t as an additional independent variable in Equation 3.2.

Considerable foresight could perhaps allow potential mothers in a particular location to influence their exposure to extreme heat during gestation. To see whether variation in exposure to extreme heat during pregnancy within a given location is related to observable characteristics of mothers, Table 12 presents estimates of the relationship between mother's demographic characteristics and the probability that her child(ren) in the sample were exposed to extreme heat during each trimester of pregnancy. There is no difference in exposure to extreme heat during pregnancy according to a mother's age, literacy, or urban/rural location. Wealthy mothers are slightly less likely to be exposed to extreme heat during the second trimester, though the point estimates are small and only marginally significant. Furthermore, none of the mother characteristics are predictive of heat exposure during the third trimester, which the results suggest is the period during which extreme heat has the strongest effects on both birth weight and later height.

Results

In Table 13, I present estimates of how height in early childhood differs according to a child's weight at birth. The estimated coefficient for birth weight in a simple model that does not include any measure of age (column 1) suggests that 100 grams of additional birth weight is associated with approximately 0.2 centimeters of additional height across all ages (0 to 59 months) in this sample.

As expected, the age of a child is a strong predictor of a child's height. Age-squared is also included in the models to capture the nonlinear relationship between height and age shown in the medical literature. Interestingly, the estimated intercept shift in the time-path of height that is

associated with birth weight is similar when modeling age as a quadratic or more flexibly with age dummies (column 2). It is striking that the estimated coefficients on the interaction terms between birth weight and each of the age variables are statistically zero (column 4). This suggests that the effect of a child’s birth weight on his/her later height is essentially constant throughout early childhood development (as opposed to altering the slope of the growth trajectory).

Figure 12 plots the estimated coefficient of birth weight (in the model of later height) estimated on samples stratified by three-month age bins. The estimates are remarkably constant throughout a child’s development. Figures 13 and 14 repeat this plot of birth weight coefficients separately for boys and girls. The nature of this relationship—that differences in birth weight produce an intercept shift and nothing more—is in fact similar to the effect of low birth weight on cognitive development through the school career (Figlio et al., 2014).

The estimates in Table 14 show that the effect of exposure to extreme heat during gestation in this sample of children is similar to that found elsewhere in the literature (i.e., Deschênes et al., 2009; Cil and Cameron, 2014). (Table 14 presents results from samples stratified by wealth quintile, while Table 15 includes estimates from models with an indicator of wealth (measured as the top two wealth quintiles) and an interaction of this indicator with the heat exposure variables.) Exposure to an average monthly temperature that exceeds the local average by more than two standard deviations has harmful effects, concentrated in the third trimester, on a child’s birth weight. The estimates suggest that a period of extreme heat in the three months before birth is associated with a decrease in birth weight of just under 90 grams. Interestingly, this effect is apparent for boys (column 2) but not for girls (column 3). As a check on the robustness of the key estimates, Figure 15 shows how the point estimates vary when the threshold for defining HOT_{ig}^t is reduced from two standard deviations above the mean. As the threshold is lowered, the point estimates move slightly towards zero, though they are in general quite stable.

The estimates presented in Table 16 extend Model 3 from Table 13 further by adding variables that capture exposure to extreme heat during gestation as controls. Column 2 in Table 16 shows that the harmful effects of extreme heat during gestation are reflected in a child’s later height, as well as in his/her birth weight, as was evidenced in Table 14. However, when birth weight is also included as an explanatory variable, the magnitude of the coefficient on “Hot third trimester” (which reveals the relationship between extreme heat and height) decreases from

-0.70 to -0.52 and is no longer statistically distinguishable from zero at conventional levels of significance. Columns 4-9 in Table 16 repeat this analysis separately for boys and girls and the results reveal a marginally significant negative relationship between extreme heat during the third trimester later height among boys, even when controlling for birthweight. I repeat the analysis on boys while stratifying the sample according to their wealth, and the results, presented in Table 17, show that the harmful effects are driven by boys in the lower three wealth quintiles.

It may be tempting to interpret the loss of statistical significance on the estimated relationship between extreme heat during gestation and height in the full sample that comes with the inclusion of birth weight as a control variable as evidence that birth weight could serve as a sufficient statistic for health at birth. However, the subsequent results also show a large, statistically significant relationship between extreme heat during gestation and subsequent height, controlling for birth weight, among the less wealthy boys.

Mechanisms

Though the results do not conclusively indicate an across-the-board relationship between extreme heat during gestation and subsequent height after controlling for birth weight, the results hint at the possibility of this among boys. An effect like this could work through either biologic or economic channels. For example, a heat wave may cause crop failures that serve as a negative income shock to the child's family, or may directly affect a child's nutritional intake after birth, thus affecting the child's height. Alternatively, extreme heat might affect the public sector, either reducing the supply or increasing the demand for public goods in a way that might affect a child's development through, for example, access to vaccinations or health services. Unfortunately, the data considered in this analysis do not include information on economic conditions at the time of birth, though such data would provide a fertile ground for future research.

It may also be the case that exposure to extreme heat during gestation has physical effects on the infant that are not reflected in birth weight. It is not possible to test this directly without measurements of other dimensions of physical health, although the incorporation of data that allow for an examination of the potential economic mechanism could provide indirect evidence regarding this possibility by ruling out other explanations.

One mechanism through which the harmful effects of extreme heat during gestation may be reflected in birth weight (but not in height) is if mothers respond to low birth weight with increased investment in the child's health. To explore this, I model the duration that a child is breastfed (in months) as a function of his/her birth weight and exposure to heat during gestation. However, these estimates, presented in Tables 18 (for boys) and Tables 19 (for girls), reveal no relationship between heat exposure and duration of breastfeeding. The results suggest that, if anything, mothers are more likely to breastfeed girls longer if they have *higher* birth weights.

Conclusions

I combine (a) data on birth weight (recorded on hospital cards) for a sample of Peruvian children who were subsequently surveyed and measured for height 0 to 59 months after birth, with (b) data on local monthly average temperatures. I find negative effects of extreme heat during gestation on birth weights and subsequent (0 to 59 months later) heights for the boys in the sample, but no apparent effect among the girls. Among boys, the estimates suggest that exposure to a month with average temperature more than two standard deviations above the local average during the third trimester is associated with a large reduction in birth weight (about 140 grams) and reduced height (about 0.9 cm) in early childhood. The reduction in height is similar in magnitude across all ages, suggesting that the harmful effects of exposure to heat during gestation are seen immediately after birth and remain over time.

The magnitude of the estimated effect of extreme heat during gestation on boys' subsequent heights remains marginally significant when controlling for (endogenous) birthweight, though the effect is reduced by about thirty percent. This suggests that extreme heat during gestation may have harmful effects on physical development in early childhood that are not reflected solely in a child's weight at birth. Differential parental investment in a child's development is one mechanism through which these harmful effects may be reduced. However, I find no evidence of differential patterns in breastfeeding according to a child's birth weight or according to exposure to extreme heat during gestation, with the exception that birth weight and breastfeeding duration are positively correlated for girls.

This analysis is based on a sample of children that may be expected to have greater health than a representative sample would have, because these children were from wealthier families with

more educated parents, were born in hospitals, survived from birth to the time at which their household was surveyed, and lived in a household that retained possession of the child's birth documentation until the time of survey. If the harmful effects of maternal stress during gestation are less likely to be experienced by healthy infants, or among offspring to parents with greater wealth and access to hospital care, then these estimates may best be interpreted as lower bounds. Similar analyses among other populations of children and births in developing countries will be useful to determine whether these estimates are robust.

CHAPTER IV

TEAM INCENTIVES AND HEALTH WORKER PRODUCTIVITY: EVIDENCE FROM MOZAMBIQUE

Introduction

Financial incentives connected to a worker's productivity have been shown to increase output (e.g. Lazear (2000); Paarsch and Shearer (2001); Prendergast (1999); Shearer (2004)). However, such financial motivation may not be as effective as simple intuition might suggest. There is evidence that payment linked to productivity can lead to lower output than fixed pay (Gneezy and Rustichini, 2000), can reduce cooperation and helping among coworkers (Fehr and Gächter, 2002; Drago and Garvey, 1998), and can potentially crowd out intrinsic motivation (Frey and Jegen, 2001).¹ For tasks with important social spillovers, such as the delivery of health care considered in this study, Ashraf et al. (2012) find that financial incentives do not crowd out intrinsic motivation, but they also find that financial incentives, whether large or small, are much less effective than non-financial incentives.

Anecdotally, the Mozambican public health system has suffered from very low motivation among workers – many are observed to work slowly and inefficiently, and to leave work early. In this context, financial incentives tied to worker activity may be a promising way to increase productivity and perhaps encourage workers to innovate and improve the system of health care delivery. However, workers have no previous experience with entrepreneurial input in facility operations because the health system in Mozambique is highly centralized and health facility managers have very little say in the labor and capital decisions at their facilities. All staff, equipment, and medicine allocation decisions are made by the national Ministry of Health with little external input. In this paper, I estimate the impact of piece-rate financial incentives on the performance of public health facilities in Mozambique.

¹Additional theories of how financial motivation might decrease productivity include the signaling of social norms (Sliwka, 2007), reduced pleasure from performing the task (Deci and Porac, 1978) and decreased perceptions of an agent's own ability (Benabou and Tirole, 2003). There is also evidence that learning about another's intentions can affect one's decisions even when it does not affect material consequences (Charness and Rabin, 2002; Falk et al., 2008), and, further, that agents may reduce effort in response to a principal's controlling actions (Falk and Kosfeld, 2006).

Background

Much of the literature on piece-rate incentives considers the classic principal-agent situation in the context of a single task. Holmstrom and Milgrom (1991) analyze a scenario more relevant to this context in which there are multiple tasks (or a single task with multiple dimensions) whose completion benefits the principal. With multiple tasks, an incentive to complete a particular task can be provided either by rewarding the completion of that task, or by reducing rewards to competing tasks and consequently lowering the opportunity cost of time devoted to the original task. Thus, paying a piece rate for the delivery of one health service affects the incentives for all other tasks a worker may perform.

Holmstrom and Milgrom (1991) find that, if possible, each task should be the responsibility of only one agent and tasks should be grouped by ease of measurement. That is, tasks that are easy to measure should be grouped into one job performed by one worker, and harder-to-measure tasks should be grouped into another job performed by a different worker. Not surprisingly, piece-rate incentives work better for the job composed of more measurable tasks. However, health services are typically delivered by combined effort from multiple workers, and it is rarely feasible to make each task the sole responsibility of one worker.

In the context studied here, the quality of service is an important aspect of health care, and it is not possible to separate this difficult-to-measure component into a distinct job. Accordingly, the financial incentive scheme implemented in Mozambique is based on the number of services delivered, with payments adjusted by a quality index designed to capture some aspects of health care that are more difficult to measure. However, the quality index is collected infrequently and concerns that workers may substitute quantity for quality are inevitable when quality is not completely monitored.

A further issue arises when piece-rate incentives are paid for team production activities because payments based on group outputs can create incentives for free riding.² Prendergast (1999) refers to this as the “1/N problem.” By sharing the benefits of an individual’s effort among N workers, the benefits of the incentive pay are not fully internalized by each worker when choosing an effort level. Newhouse (1973) observes that in group medical practices, the number

²Group incentives can also affect the composition of teams, as explored in Bandiera et al. (2012). In the context studied here, this is unlikely to happen because jobs are assigned centrally and workers have no input concerning their assignments.

of hours that physicians work decreases as the fraction of revenue shared with others increases.³ Despite incentives to free ride, however, payment schemes based on group effort are common (e.g. employee stock ownership plans) and there is evidence that these schemes increase output. Lavy (2002) shows that group financial incentives for teachers caused significant improvements in student outcomes. Knez and Simester (2001) argue that free riding can be overcome by mutual monitoring of effort among workers, and Kandel and Lazear (1992) provide an analysis of the ways that peer pressure can overcome incentives to free ride.

Despite free riding, group incentives may still be preferred because of difficulties in measuring individual effort as well as potentially harmful individual incentives that might reduce cooperation or cause competition for inputs. Boning et al. (2007) find that in complex production processes the adoption of a team approach provides workers with a valuable opportunity to solve problems and increase productivity. Furthermore, if productivity depends on group effort then worker absence has higher costs, and group incentive pay in this context has been found to reduce worker absenteeism (Heywood et al., 2008). Given the significant absenteeism of teachers and medical practitioners in developing countries (Chaudhury et al., 2006), group incentive pay may be a particularly useful management strategy for health systems in these places.

Incentive schemes to increase and improve the delivery of health care have been implemented in many health systems. Some schemes have created perverse incentives resulting in workers gaming the system (see, e.g., Gravelle et al. (2010)), while other schemes have had no effect (e.g. Rosenthal et al. (2005), and Nicholas et al. (2011)). However, in the scenario most similar to that studied here, Basinga et al. (2011) find that a group incentive scheme in Rwandan health facilities increased delivery of the services that had the highest payment rates and required the lowest worker effort. This result suggests clearly strategic behavior by health facility workers and fits nicely with intuitive economic predictions of profit and utility maximization. However, it would be troubling if increases in the delivery of paid services were achieved only by diverting resources away from other important dimensions of health care. In this paper I also estimate the effects of incentive pay for some services on the delivery of other services that are not financially incentivized.

³Newhouse's observation comes from comparisons across firms and therefore does not account for obvious concerns about the endogeneity of sharing rules. Gaynor and Pauly use risk aversion to exogenously identify variation in practice size.

The Incentive Program and Associated Data

To increase the delivery of health care services to the Mozambican population, a non-governmental organization (NGO), in conjunction with the Mozambique Ministry of Health (MOH), has implemented a performance-based financing system for a subset of public health facilities in Mozambique. This NGO has historically supported health facilities in Mozambique by compensating them for documented qualifying expenses up to a pre-determined limit. In 2011, the organization began supporting a group of clinics (the “treatment” facilities) in Mozambique with payments based on the quantity of services delivered at the facility, rather than facility expenditures. The NGO pays these facilities a set price for the delivery of certain services. Prices vary by the type of service, ranging from \$0.13 for each child fully vaccinated to \$10.00 for each Human Immunodeficiency Virus (HIV) positive youth patient receiving a full year of anti-retroviral therapy (ART).

Each quarter, the total value of the facility’s services is deposited into a bank account controlled by the facility manager. Half of the total incentive funds earned by a facility are allocated to non-labor spending at the facilities’ discretion and the facility reports on these expenditures to the NGO. The remaining funds are distributed as salary bonuses to the facility workers. The funds are distributed among all staff proportionally by job classification and seniority, and the average bonus for an individual worker is around 10 or 15 percent of their quarterly salary. The result is a piece-rate incentive scheme in which an individual worker’s payment is determined by the collective output of all workers. The NGO continues to support another set of facilities (the “comparison” facilities) with quarterly cost-reimbursement funding approximately equivalent to the average quarterly incentive payments received by Treatment facilities.⁴ The cost-reimbursement funding at these facilities is restricted to non-labor uses.

Health services at public facilities in Mozambique are provided free of charge to patients, so ‘price’ in this context refers only to the per-unit service fee that the NGO pays to facilities as part of the incentive pay program. Though the monetary cost of health services for patients is zero, there may be substantial travel costs for some. There are also likely to be other demand-side

⁴The NGO sets a quarterly limit on the amount of reimbursement funds they will pay to each facility, but cannot control the expenses that each facility submits for reimbursement and so it is difficult to maintain precise equality of payments across treatment and comparison facilities. However, in almost no cases has a facility submitted expenses equaling or exceeding the limit, suggesting that the constraints on NGO-provided cost-reimbursement funding are not limiting productivity.

factors that influence the ability of health care providers to respond to the financial incentives introduced by the program. To the extent that these characteristics have similar trends across treatment and comparison facilities, they can be accounted for when deriving estimates of the program's impact with facility fixed effects. However, if these trends are not similar across treatment and comparison facilities, or if the program impacts the demand for health care, incorporating information on the catchment populations of each facility through demographic and health surveys would be a promising avenue for future research.

Mozambique and HIV/AIDS

Mozambique has the eighth highest prevalence of HIV in the world and the seventh highest number of deaths due to HIV and Acquired Immune Deficiency Syndrome (AIDS) (CIA, US, 2013). A large portion of health facility activity in Mozambique involves the detection and treatment of HIV/AIDS. Additionally, the NGO has a particular focus on the prevention and treatment of HIV/AIDS and the majority of services for which treatment facilities receive incentive pay involve HIV/AIDS. Importantly, the treatment facilities are paid a per-unit fee for many, but not all, of the different types of services they administer. This strategy of payment for the delivery of only some services could cause a substitution of resources away from other important services or dimensions of health care.

Data

For 64 health facilities over nine quarters, data were collected on the number of times each of twenty-one different health services was delivered. Twenty-four of the facilities received incentive payments based on the number of services they delivered and the remaining 48 facilities received cost-reimbursement funding. The data cover the period from 2010Q1 to 2012Q1 and the funding program began in 2011Q1. This panel of service data is combined with a cross-sectional census of facility characteristics collected before the initiation of the program. Table 20 lists the prices and descriptions of the individual services, which are grouped into five service-type categories: (1) Testing for HIV, (2) Initiation of Anti-Retroviral Therapy, (3) Long-term Care and Treatment of HIV/AIDS, (4) Family Planning, Pregnancy, and Births, and (5) Non-paid services.

Due to a change in protocol for the prevention of mother-to-child transmission (PMTCT) of HIV/AIDS, in the subsequent analysis I have summed the raw data for pregnant women started on ART (newART_preg) and ARV to create the variable PMTCT. Previously, it was thought that ARV could be administered to HIV+ women during pregnancy and breastfeeding to prevent transmission of the virus to the child. Starting in 2010, however, the protocol shifted to put these women on ART as it became clear that ARV was not as successful as expected in preventing transmission.⁵ Facilities still use ARV for PMTCT, but a patient will receive only ART or ARV, not both, so totaling them provides a more accurate representation of the number of women treated to prevent transmission of the virus to their infants.

Empirical Strategy

The incentive pay system was explained to the provincial health directors of the four provinces in Mozambique where the NGO operates and they were asked to submit a request if they wished to participate. The directors of Cabo Delgado and Maputo provinces said that they were interested in joining, but believed the MOH was not in favor of the program, so they declined to participate. The directors of Gaza and Nampula provinces opted into the program. Twenty-four of the public health facilities in Gaza and Nampula began receiving incentive pay in January, 2011, and eighteen facilities received a delayed roll-out of the program (the data used in this paper are from before these additional facilities began the incentive pay program).⁶

Statistical identification of a relationship between the size of the piece rates and their impacts on output is limited by a lack of data on the cost or effort required to provide each service. Furthermore, the piece rates were endogenously assigned based on which services the NGO felt were more likely to respond and more important to incentivize. In Section 4.5, I present an indirect approach to detect any potential heterogeneity of treatment effects due to variation in the sizes of incentives.

The treatment and comparison facilities are similar across all but two of the staffing and facility characteristics collected in the facility census. Summary statistics are presented in

⁵To clarify, ARV is the name of the drug administered to patients undergoing ART, but ARV also refers to the administration of a shorter course of the drug particularly for PMTCT.

⁶The facilities receiving delayed implementation of the incentive pay program were previously supported by a different NGO. The delay was to allow time for the formulation of contractual agreements necessary for their inclusion under the implementing NGO's support. They began receiving incentive pay in the second quarter of 2012.

Table 21. The two groups have similar numbers of health workers and support staff. Comparison facilities are larger in terms of bed capacity, on average, though the difference is not statistically significant. Treatment facilities are more likely to have an electrical generator on site, and they are slightly less likely (but not statistically significantly so) to have electricity from the power grid. These two sources of electrical power are negatively correlated, and every facility in the sample has at least one or the other. Treatment facilities are much more likely to have a lab on site, which may affect a facility’s ability to process HIV tests. However, there is a regularly used system in place for facilities without labs to send test kits to the nearest facility with a lab for processing.

The average quarterly performance of each service, prior to treatment, is presented in Table 22. Despite the similarities in capacity and staffing between the treatment and comparison facilities, the comparison facilities perform significantly greater numbers of many of the services. Figures 16 through 22 depict the trends in quantities of services delivered over time separately for treatment and comparison facilities. Due to this difference in levels, a simple difference estimator would not identify treatment effects. However, there is variation in exposure to incentive pay across treatment and comparison facilities, as well as over time at treatment facilities. This allows for difference-in-differences estimation, conditional on common trends prior to treatment, of the effect of the incentive pay mechanism on the quantity of services delivered. I include facility fixed effects to control for all time-invariant differences in facilities and patient populations, and quarterly fixed effects to control for seasonality.⁷

More formally, let Y_{ft}^s be the number of times the service s is performed in facility f during time period t ($t = 1, \dots, 9$). Let TG_f be an indicator for whether facility f is in the Treatment Group ($TG_f = 1$). By restricting the sample to the pre-treatment period, δ_2 in Equation 4.1 measures the difference (between treatment and comparison facilities) in pre-treatment trends for service s . Facility fixed effects are represented by α_f , and seasonality is captured via quarterly fixed effects (ρ_q for $q = 1, \dots, 4$).

$$Y_{ft}^s = \alpha_f + \rho_q + \delta_1 t + \delta_2 (TG_f * t) + \epsilon_{ft} \quad \text{for } s = 1, \dots, 20 \text{ and } t < 5 \quad (4.1)$$

⁷Treatment and comparison facilities are quite far apart, so there are unlikely to be any issues stemming from patients self-selecting into treatment facilities.

Letting TP_t indicate whether the observation is collected during the Treatment Period ($TP_t = 1$), the standard difference-in-differences model is given by Equation 4.2. For those services for which the identifying assumption of common pre-treatment trends in treatment and comparison facilities (i.e. $\delta_2 = 0$) cannot be rejected, β_3 in Equation 4.2 measures the impact of the incentive pay program on the output, Y_{ft}^s , of service s .

$$Y_{ft}^s = \alpha_f + \rho_q + \beta_1 TP_t + \beta_2 TG_f + \beta_3 (TG_f * TP_t) + \epsilon_{ft} \quad \text{for } s = 1, \dots, 20 \quad (4.2)$$

Results

Table 23 presents the estimated pre-treatment trend differences (δ_2 from Equation 4.1) and, for the services with common trends across treatment and comparison facilities prior to the initiation of the incentive program, the estimated treatment effects (β_3 from Equation 4.2). Although some of the data are higher-order counts, most are low-order, and I use a Poisson specification in all models to ease comparison between estimates for different services.⁸ According to the NGO, it is likely that every service is performed at least once per quarter at each facility. However, for the individual services, around ten percent of the counts are recorded as zeros, perhaps raising concerns that some missing data were reported as zeros. To see if the results are possibly driven by improved record keeping alone, I include an alternative set of results (models 2, 4, and 6) excluding all zeros from the sample. In the interpretation, I focus on the results from model 3, which are estimated on the largest sample, including observations from all periods and those reported as zeros.

The estimated difference-in-difference coefficients (models 3 to 6) can be interpreted as approximately the percentage change in the count of reported services delivered that is attributable to the incentive pay program.⁹ In the Appendix, I provide graphs of the trends over time for each of the services. The last panel of results (Table 23, columns 5 and 6) presents the treatment effect estimates from a sample that excludes the two quarters immediately following the initiation of the program to allow for a period of learning in Treatment facilities.

⁸Results from a continuous log-normal specification are qualitatively similar to those presented here.

⁹The interpretation of the coefficients as percentage impacts is most appropriate for those small in magnitude. As the estimates increase in magnitude, they are better interpreted as a conservative (lower bound) estimate of the percentage impact. To convert the presented estimates precisely into percentage terms, the formula is $(e^\beta - 1) * 100$.

The trends prior to treatment differ at conventional levels of statistical significance only for the number of HIV tests given and the number of women treated for pre-eclampsia. Of the 14 paid services for which I cannot reject the assumption of common trends prior to treatment, the results suggest a statistically significant positive impact from the incentive pay program on five: (1) the number of pregnant women completing four ante-natal care visits, (2) the number of family planning consultations, (3) the treatment of HIV+ patients for TB, and the initiation of ART for (4) adults and (5) TB patients. The initiation of children on ART shows a marginally significant ($\alpha = 0.10$) positive impact. The longterm services that require effort over a full year of treatment show no response to the incentive pay program. However, the number of HIV+ patients treated for Tuberculosis, which can require multiple patient visits depending on the severity of the infection, did increase. Curiously, the number of postpartum consultations decreased by about 28 percent with the implementation of the incentive pay program. Staff members of the NGO have reported that they believe this may be a result of the increase in ante-natal visits, which could be a substitute (from the perspective of patients and/or clinic workers) to postpartum consultations.

The estimates for the procedures for which workers did not receive incentive pay are presented in Table 24. The results suggest that the incentive pay program had no effect on the delivery of services for which workers received no financial incentives. The NGO's goal in implementing incentive pay was to increase the total productivity of health workers, not to create a system that causes workers to respond more to financial incentives than to patient needs. These results suggest that the increases identified for the delivery of paid services are not coming at the cost of the non-paid activities for which there are data. Increases in productivity could potentially come at the cost of the quality of care, but there are no reliable data on service quality available to analyze in conjunction with the incentive pay program.

Incentive Size Effects

The endogenous assignment of the piece rates makes it difficult to determine a relationship between incentive size and worker effort. However, the incentive payment is shared among all N workers at a facility and there is variation in N , so there is consequently variation across facilities in the size of incentives received by workers for the delivery of a given service. To see if the impact of the program varies with the size of the incentives, I allow the treatment effect

to vary with a variety of facility characteristics, including the number of workers at a facility. To allow for potentially heterogeneous treatment effects, an additional interaction term can be added to Equation 4.2. In Equation 4.3 below, γX captures systematic variation in the estimated treatment effect, where X is a vector of facility characteristics and γ is a vector of parameters to be estimated.

$$Y_{ft} = \alpha_f + \rho_q + \beta_1 TP_t + \beta_2 TG_f + (\beta_3 + \gamma X)(TG_f * TP_t) + \epsilon_{ft} \quad (4.3)$$

I estimate Equation 4.3 using the total value of a facility's service mix as the dependent variable, where total value is calculated for all facilities and quarters using the piece rates from the incentive pay program. Let P be the 1 x 21 vector of piece rates and Q_{ft} be the 1 x 21 vector of the counts of services delivered at facility f in quarter t . Then PQ'_{ft} is the total value, in terms of incentive pay, of facility f 's quarter t performance (though only treatment facilities in the treatment period are actually paid this as part of the program). The average total service values for treatment and comparison facilities for the pre- and post-treatment periods are presented in Table 25.

The estimated parameters of Equation 4.3, with facility characteristics included in X , are presented in Table 26.¹⁰ If the response to incentive pay is affected by the size of the incentives, we might expect to find a negative relationship between the number of workers (among whom the bonus pay is shared) and the treatment effect. Or, it may be that skilled health workers can increase the delivery of services but support staff are limited in their ability to do so, and we might expect a negative sign on the ratio of support staff to health workers. The results suggest, however, that whether staff size is measured as total workers, only skilled workers, or as the ratio of staff-to-skilled workers, there is no effect on the response to incentive pay. These results suggest that there is no relationship between incentive size and worker response.

Table 26 also presents the estimates of potentially heterogeneous treatment effects in the physical capacity of facilities. These are captured by the log of the number of beds, and an indicator of whether the facility is a hospital (as opposed to a smaller clinic). I also include estimates of how the treatment effect varies for facilities with a generator or laboratory on site,

¹⁰The difference in trends for total service value between treatment and comparison facilities is not statistically significantly different (P=0.74).

since these were the observable characteristics by which Treatment and Comparison facilities differed. None of these characteristics have a statistically significant impact on how facilities respond to the incentive pay program.

Conclusions

This paper adds to a growing body of evidence on the ability of financial incentives for health workers to increase the delivery of health services. Many health systems in developing countries are struggling to provide adequate care for their people, and low wages and under-motivation among workers are thought to be major contributors to these struggles.

This paper analyzes the impact of a non-randomly assigned system that gives workers bonus payments based on the number of services delivered at their facility. There are many perverse outcomes that may potentially arise from paying health workers based on the numbers and types of services they perform, yet I find no evidence of such behavior. The financial incentives in this scheme did not increase the delivery of every type of service, but there were large increases, ranging from 30 to over 150 percent, for about one-third of the services. In particular, there were notable increases in the number of new patients put on anti-retroviral therapy for HIV/AIDS.

Importantly, the increase in the delivery of these services does not appear to have come at the cost of a reduction in the delivery of other services. With the data used in this study it is impossible to determine whether overall productivity increased, or whether the increases observed are due to a substitution of effort away from un-monitored activities. However, it is heartening that there were no observed decreases in the non-paid services, though one of the paid services did decrease and potential explanations for this should be explored. Finally, although I use an indirect method, I find no evidence that these health facility workers respond to variation in the size of incentives. This is counter to basic economic theory, though the findings may also be a result of limited power to detect such an effect in these data.

This paper presents a basic analysis of the effects of an incentive pay scheme implemented in Mozambique on the delivery of eighteen different health services. Further research on potentially heterogeneous treatment effects by facility or patient population characteristics could reveal much more about the mechanisms by which incentive pay impacts worker activity. Data

from this program could potentially be combined with data from demographic and health surveys to incorporate information on the demand for health care. Another important issue, left largely unaddressed due to a lack of data, concerns the potential impact that quantity incentives have on the quality of care. Though a quality index was collected as part of the program analyzed here, the components of the index changed during the course of data collection. Future research using a survey of health facility managers and workers, or data collected directly from patients, would be very useful for an examination of the impacts of incentive pay on quality of care.


APPENDIX
FIGURES AND TABLES

FIGURE 1. Timeline of Hurricanes, Elections, and Data Collection

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005
Hurricanes			Cesar		Mitch							Beta
LSMS collected					I			II				III
Aid Period	[.....“before”.....][.....“short-term”.....][.....“long-term”.....]											
Households					4,020			2,960				2,645
Municipal Elections			x				x				x	
National Elections			x					x				


Notes: Hurricane Mitch occurred in October of 1998, after the collection of data in LSMS Wave I. Hurricane Beta occurred in October of 2005, after the collection of data in LSMS Wave III.

FIGURE 2. Precipitation Data Construction




Notes: The top left map depicts the 0.5 degree latitude by 0.5 degree longitude grid for which monthly precipitation data are available from the British Atmospheric Data Center's Climate Research Unit. These data are used to interpolate a measure of rainfall for each 0.02 degree cell from the four nearest observations, weighted inversely by distance (top right). This measure is then averaged within each municipality to create a measure of rainfall at the municipal level (depicted lower left for a sample month (here, July 1998)).

FIGURE 3. Distribution of Hurricane Impacts in Nicaragua: 1960-2010.




Notes: Construction of the Hurricane Rainfall Ratio is described in detail Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years. There are 2,002 total observations (143 municipalities times 14 hurricanes). The mass of this distribution that is less than one reflects the fact that some municipalities may experience below average rainfall at the same time that other municipalities are affected by a hurricane

FIGURE 4. Timeline of Hurricanes Affecting Nicaragua: 1960-2010.




Notes: Each point represents one municipality during a hurricane. Construction of the Hurricane Rainfall Ratio is described in detail Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years.

FIGURE 5. Geographic Variation in Hurricane Impact (Rainfall Ratio).



Notes: The maps are individually scaled to highlight the geographic variation in the impact of each hurricane. Construction of the Rainfall Ratio is described in detail Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years.

FIGURE 6. Average Hurricane Impact (Rainfall Ratio) by Municipality: 1960-2010






Notes: There are 14 hurricanes reflected in the data, and there are 143 municipalities in Nicaragua. Error bars represent 95 percent confidence intervals. Construction of the Rainfall Ratio is described in detail Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years.


FIGURE 7. Average Hurricane Impact (Rainfall Ratio): 1960-2010.

Average Hurricane Rainfall Ratio


Legend




Mitch 1998



Cesar 1996



Gert 1993



Joan 1998

Note: All maps use the same color scale. The small maps are provided to give context to the large map of the average impacts. Construction of the Rainfall Ratio is described in detail Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years.

FIGURE 8. Average Height




FIGURE 9. Seasonality of Births in Peru




FIGURE 10. Average Monthly Temperatures in Peru: 1960-2010




FIGURE 11. Locations of Respondents and Weather Data






FIGURE 12. Height: Estimated Coefficient of Birthweight, by Age bin (3m)




Notes: The graphic presents point estimates and 95 percent confidence intervals from regressions of height (centimeters) on birthweight (100 grams) using successive samples of children falling within three-month age bins.

FIGURE 13. Boys Height: Estimated Coefficient of Birthweight, by Age bin (3m)



Notes: The graphic presents point estimates and 95 percent confidence intervals from regressions of height (centimeters) on birthweight (100 grams) using successive samples of children falling within three-month age bins.

FIGURE 14. Girls Height: Estimated Coefficient of Birthweight, by Age bin (3m)




Notes: The graphic presents point estimates and 95 percent confidence intervals from regressions of height (centimeters) on birthweight (100 grams) using successive samples of children falling within three-month age bins.


FIGURE 15. Robustness of Key Estimates to HOT Threshold

From Table 14

(Column 1)




(Column 2)




From Table 16

(Column 5)




(Column 6)




From Table 17

(Column 9)



(Column 10)



Trends in Services and Data Reporting

The following figures depict the number of each service delivered each quarter, averaged separately across the Treatment and Comparison facilities. The time period t (quarter) is on the horizontal axis and the number of services is on the vertical axis. The vertical line indicates the initiation of the incentive pay program at the start of Quarter 5 (January, 2011) at the Treatment facilities.

FIGURE 16. Individual Service Trends: HIV Tests




FIGURE 17. Individual Service Trends: Newly Initiated on ART




FIGURE 18. Individual Service Trends: Longterm Care




FIGURE 19. Individual Service Trends: Birth and Babies




FIGURE 20. Individual Service Trends: Non-paid Services




FIGURE 21. Trends: Total Value of Services





FIGURE 22. Data Reporting Trends

Paid Service Data



Non-paid Service Data




TABLE 1. Descriptive Statistics: Households Surveyed Before and After Hurricane Mitch

Variable	Mean or Proportion	Std. Dev.	Min.	Max.
1998 (Pre-Mitch) Household Characteristics				
Household size	5.70	2.76	1	20
HH head has \geq elementary education	0.32	0.47	0	1
HH head is female	0.28	0.45	0	1
Lives in urban locality	0.55	0.50	0	1
HH has electricity	0.64	0.48	0	1
HH owns a television	0.52	0.50	0	1
Consumption (annual, US\$100 per person)	6.67	8.21	0.60	168.26
Locally represented by Sandinistas	0.29	0.45	0	1
Household Received Aid				
Before Mitch: -4 - 0 yrs	0.69	0.47	0	1
from Nicaraguan government	0.59	0.49	0	1
from NGOs	0.15	0.36	0	1
Short Term after Mitch: 0 - 3 yrs	0.72	0.45	0	1
from Nicaraguan government	0.62	0.48	0	1
from NGOs	0.15	0.36	0	1
Long Term after Mitch: 3 - 7 yrs	0.83	0.38	0	1
from Nicaraguan government	0.78	0.41	0	1
from NGOs	0.12	0.32	0	1
Hurricane Mitch Impact				
Mitch (1998) Rainfall Ratio (MRR)	1.75	0.12	1.59	2.05
INEC Hurricane Mitch Designation	0.15	0.36	0	1

Notes: N=2,960 for all variables except the measures of aid in the Long Term period (where N=2,645). Selection of the sample and analysis of attrition are discussed in detail in Section 3.4. Construction of the Rainfall Ratio is described in detail in Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years. INEC Hurricane Area Designation indicates that the municipality was deemed "affected" by the hurricane by the Nicaraguan government (INEC).

TABLE 2. Municipal Hurricane Incidence (Hurricane Rainfall Ratio): 1960-2010

Frequency of Hurricanes (X)	Number of Municipalities Experiencing X Hurricanes with Impact:	
	HRR>1.75	HRR>1.50
0	26	0
1	98	3
2	19	50
3	0	76
4	0	12
5	0	2
6+	0	0
Total	143	143

Note: There were 14 hurricanes during the period considered. Construction of the Rainfall Ratio is described in detail Section 2.3. It is equal to the municipal rainfall during the hurricane month divided by the municipality's average, non-hurricane, rainfall for the same month over the prior 25 years.

TABLE 3. Household Attrition and Hurricane Exposure for LSMS 1998 Households

	Attrition between 1998-2001			Attrition between 2001-2005		
	(1)	(2)	(3)	(4)	(5)	(6)
Hurricane Mitch Rainfall Ratio (HRR)	0.160** (0.069)	0.047 (0.069)	0.061 (0.353)	-0.048 (0.050)	-0.064 (0.051)	-0.147 (0.230)
Household size		-0.020*** (0.003)	-0.026 (0.041)		-0.010*** (0.002)	-0.020 (0.028)
HH head has \geq elementary education		0.057*** (0.017)	0.047 (0.292)		0.009 (0.015)	-0.094 (0.217)
HH head is female		-0.070*** (0.015)	0.217 (0.236)		-0.020* (0.012)	0.058 (0.180)
Lives in urban locality		0.035* (0.020)	-0.410 (0.266)		0.039** (0.020)	-0.318 (0.223)
Has electricity		-0.010 (0.021)	0.069 (0.346)		-0.017 (0.023)	-0.002 (0.276)
Owens Television		-0.111*** (0.022)	-0.225 (0.339)		-0.040** (0.017)	-0.019 (0.250)
Received aid from government		-0.061*** (0.017)	0.192 (0.259)		-0.026** (0.011)	-0.299* (0.179)
Received aid from NGO		-0.075*** (0.018)	0.374 (0.282)		-0.028* (0.015)	0.075 (0.162)
Opposition		-0.052** (0.022)	-0.158 (0.311)		-0.034* (0.017)	0.004 (0.220)
log(Consumption per capita)		-0.023* (0.013)	0.014 (0.203)		0.006 (0.012)	0.157 (0.156)
HRR \times Household size			0.003 (0.024)			0.006 (0.016)
HRR \times HH head has \geq elementary education			0.006 (0.169)			0.060 (0.127)
HRR \times HH head is female			-0.164 (0.136)			-0.045 (0.102)
HRR \times Lives in urban locality			0.254 (0.154)			0.204 (0.125)
HRR \times Has electricity			-0.046 (0.195)			-0.009 (0.155)
HRR \times Owens Television			0.065 (0.196)			-0.013 (0.145)
HRR \times Received aid from government			-0.144 (0.148)			0.156 (0.102)
HRR \times Received aid from NGO			-0.253 (0.163)			-0.058 (0.091)
HRR \times Opposition			0.060 (0.176)			-0.022 (0.122)
HRR \times log(Consumption per capita)			-0.021 (0.115)			-0.086 (0.089)
Households	4020	4020	4020	2960	2960	2960
R-squared	0.00	0.05	0.06	0.00	0.02	0.03
Mean of Dep Var	0.26	0.26	0.26	0.08	0.08	0.08
F test: Interaction terms jointly 0; Prob> F			0.25			0.21

Notes: Coefficients represent estimates from linear probability models. Estimated standard errors are in parentheses, adjusted for any clustering at the municipal level. Mitch Rainfall Ratio ranges from 1.59 to 2.05 and represents the rainfall experienced during the hurricane as a percentage of the average non-storm rainfall, adjusted for seasonality. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 4. The Effect of Hurricane Mitch on Government Aid Allocations: 0-3 Years After

	(1)	(2)	(3)	(4)	(5)	(6)
Mitch Rainfall Ratio (MRR) 1998	0.241 (0.178)		0.167 (0.182)	0.253 (0.201)	-0.233 (0.241)	-0.063 (0.235)
INEC Hurricane Area Designation		0.122** (0.047)	0.115** (0.048)	0.114** (0.048)	0.118** (0.047)	0.091* (0.049)
Sandinista				-0.055 (0.051)	-1.157*** (0.442)	0.012 (0.059)
Sandinista \times MRR					0.625** (0.247)	
Sandinista \times =1(MRR lowest quartile)						-0.158** (0.065)
Sandinista \times =1(MRR second quartile)						-0.139* (0.073)
Sandinista \times =1(MRR top quartile)						-0.019 (0.062)
Cesar Rainfall Ratio (CRR) 1996	0.055 (0.143)	-0.051 (0.105)	0.021 (0.151)	0.133 (0.195)	0.125 (0.184)	0.148 (0.185)
log(Consumption per capita)	-0.033 (0.022)	-0.034 (0.022)	-0.031 (0.021)	-0.031 (0.021)	-0.030 (0.021)	-0.031* (0.018)
Household (HH) size	0.010*** (0.004)	0.010*** (0.004)	0.011*** (0.004)	0.011*** (0.004)	0.011*** (0.004)	0.011*** (0.004)
HH head has \geq elementary education	0.025 (0.024)	0.027 (0.024)	0.027 (0.024)	0.027 (0.023)	0.028 (0.023)	0.032 (0.022)
HH head is female	0.011 (0.020)	0.016 (0.021)	0.016 (0.021)	0.015 (0.021)	0.017 (0.020)	0.017 (0.019)
Lives in urban locality	-0.035 (0.034)	-0.023 (0.034)	-0.022 (0.034)	-0.020 (0.034)	-0.023 (0.033)	-0.021 (0.033)
Has electricity	0.029 (0.036)	0.030 (0.037)	0.030 (0.037)	0.032 (0.037)	0.030 (0.037)	0.034 (0.036)
Owens a television	0.026 (0.027)	0.030 (0.027)	0.029 (0.027)	0.030 (0.027)	0.027 (0.027)	0.026 (0.027)
Constant	0.084 (0.494)	0.644*** (0.169)	0.233 (0.517)	-0.057 (0.606)	0.819 (0.617)	0.484 (0.617)
Households	2960	2960	2960	2960	2960	2960
R-squared	0.01	0.02	0.02	0.02	0.02	0.03
Mean of Dep Var	0.62	0.62	0.62	0.62	0.62	0.62

Notes: Coefficients represent estimates from linear probability models. Estimated standard errors are in parentheses, adjusted for any clustering at the municipal level. Mitch Rainfall Ratio ranges from 1.59 to 2.05 and represents the rainfall experienced during the hurricane as a percentage of the average non-storm rainfall, adjusted for seasonality. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 5. Determinants of INEC Hurricane Mitch Designation

	(1)	(2)
Mitch Rainfall Ratio (MRR) 1998		0.886*
		(0.497)
Slope (sd of municipal elevation)	-0.002	0.005
	(0.002)	(0.005)
Slope * MRR	0.002	-0.003
	(0.001)	(0.003)
log(Consumption per capita)	-0.008	-0.009
	(0.013)	(0.013)
HH head has \geq elementary education	-0.014	-0.015
	(0.014)	(0.014)
Sandinista	0.031	0.039
	(0.056)	(0.055)
Received Govt aid	0.034*	0.036*
	(0.020)	(0.020)
Received NGO aid	0.006	0.011
	(0.030)	(0.029)
Lives in urban locality	-0.128***	-0.132***
	(0.037)	(0.037)
Owens a television	-0.032	-0.027
	(0.024)	(0.023)
Constant	0.177***	-1.335
	(0.058)	(0.829)
Households	2960	2960
R-squared	0.07	0.08

Notes: Coefficients represent estimates from linear probability models. Estimated standard errors are in parentheses, adjusted for any clustering at the municipal level. Mitch Rainfall Ratio ranges from 1.59 to 2.05 and represents the rainfall experienced during the hurricane as a percentage of the average non-storm rainfall, adjusted for seasonality. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 6. The Effect of Hurricane Mitch on Aid Allocations by NGOs: 0-3 Years After

	(1)	(2)	(3)	(4)	(5)	(6)
Mitch Rainfall Ratio (MRR) 1998	0.371*** (0.140)		0.297** (0.141)	0.324** (0.162)	0.480* (0.286)	0.383 (0.243)
INEC Hurricane Area Designation		0.127*** (0.037)	0.115*** (0.038)	0.115*** (0.038)	0.113*** (0.038)	0.105*** (0.037)
Sandinista				-0.017 (0.036)	0.339 (0.428)	0.013 (0.043)
Sandinista \times MRR					-0.202 (0.252)	
Sandinista \times =1(MRR lowest quartile)						-0.029 (0.046)
Sandinista \times =1(MRR second quartile)						-0.036 (0.035)
Sandinista \times =1(MRR top quartile)						-0.054 (0.051)
Cesar Rainfall Ratio (CRR) 1996	0.211** (0.094)	0.049 (0.051)	0.178** (0.083)	0.212* (0.119)	0.215* (0.116)	0.220* (0.120)
log(Consumption per capita)	-0.015 (0.012)	-0.018 (0.011)	-0.013 (0.011)	-0.013 (0.011)	-0.013 (0.011)	-0.013 (0.012)
Household (HH) size	0.005* (0.003)	0.005* (0.003)	0.005** (0.003)	0.005** (0.003)	0.005** (0.003)	0.005** (0.003)
HH head has \geq elementary education	0.029* (0.016)	0.031** (0.015)	0.031** (0.015)	0.031** (0.015)	0.030* (0.015)	0.033** (0.015)
HH head is female	0.006 (0.014)	0.011 (0.013)	0.010 (0.014)	0.010 (0.014)	0.010 (0.014)	0.011 (0.014)
Lives in urban locality	-0.096*** (0.028)	-0.084*** (0.027)	-0.082*** (0.027)	-0.082*** (0.028)	-0.080*** (0.028)	-0.080*** (0.029)
Has electricity	-0.046* (0.027)	-0.046 (0.027)	-0.045* (0.027)	-0.045* (0.027)	-0.044 (0.027)	-0.043 (0.027)
Owns a television	-0.031* (0.017)	-0.027 (0.017)	-0.028 (0.017)	-0.028 (0.017)	-0.027 (0.017)	-0.028 (0.017)
Constant	-0.754** (0.360)	0.127 (0.083)	-0.606* (0.347)	-0.696 (0.428)	-0.979 (0.621)	-0.811 (0.551)
Households	2960	2960	2960	2960	2960	2960
R-squared	0.06	0.07	0.07	0.07	0.07	0.07
Mean of Dep Var	0.15	0.15	0.15	0.15	0.15	0.15

Notes: Coefficients represent estimates from linear probability models. Estimated standard errors are in parentheses, adjusted for any clustering at the municipal level. Mitch Rainfall Ratio ranges from 1.59 to 2.05 and represents the rainfall experienced during the hurricane as a percentage of the average non-storm rainfall, adjusted for seasonality. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 7. The Effect of Hurricane Mitch on Aid Allocations: 3-7 Years After

	(1)	(2)	(3)
Panel A: Aid from the Nicaraguan Government			
Mitch Rainfall Ratio (HRR) 1998	0.448*** (0.146)		0.463*** (0.155)
INEC Hurricane Area Designation		-0.005 (0.034)	-0.023 (0.035)
Controls	yes	yes	yes
Households	2645	2645	2645
R-squared	0.04	0.04	0.04
Mean of Dep Var	0.78	0.78	0.78
Panel B: Aid from Non-Governmental Organizations			
Mitch Rainfall Ratio (HRR) 1998	0.510*** (0.123)		0.502*** (0.128)
INEC Hurricane Area Designation		0.033 (0.027)	0.013 (0.026)
Controls	yes	yes	yes
Households	2645	2645	2645
R-squared	0.06	0.05	0.06
Mean of Dep Var	0.12	0.12	0.12

Notes: Coefficients represent estimates from linear probability models. Estimated standard errors are in parentheses, adjusted for any clustering at the municipal level. Mitch Rainfall Ratio ranges from 1.59 to 2.05 and represents the rainfall experienced during the hurricane as a percentage of the average non-storm rainfall, adjusted for seasonality. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 8. The Effect of Hurricane Mitch on Aid Allocations: 4-0 Years Before

	(1)	(2)	(3)
Panel A: Aid from the Nicaraguan Government			
Mitch Rainfall Ratio (HRR) 1998	0.344 (0.258)		0.304 (0.263)
INEC Hurricane Area Designation		0.074* (0.040)	0.061 (0.041)
Controls	yes	yes	yes
Households	2960	2960	2960
R-squared	0.04	0.04	0.04
Mean of Dep Var	0.59	0.59	0.59
Panel B: Aid from Non-Governmental Organizations			
Mitch Rainfall Ratio (HRR) 1998	0.067 (0.192)		0.060 (0.193)
INEC Hurricane Area Designation		0.014 (0.034)	0.012 (0.034)
Controls	yes	yes	yes
Households	2960	2960	2960
R-squared	0.03	0.03	0.03
Mean of Dep Var	0.15	0.15	0.15

Notes: Coefficients represent estimates from linear probability models. Estimated standard errors are in parentheses, adjusted for any clustering at the municipal level. Mitch Rainfall Ratio ranges from 1.59 to 2.05 and represents the rainfall experienced during the hurricane as a percentage of the average non-storm rainfall, adjusted for seasonality. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 9. Children (0-59 months old) in Surveyed Households

Survey Year	Total		With Information on			With Complete Information
	Height	Birth Weight	Location at Survey	Location during Gestation		
2000	13,697	11,795	1,906	13,620	11,762	1,615
2003	161	0	51	161	138	0
2004	2,376	0	590	2,370	1,990	0
2005	2,631	2,321	881	2,618	2,301	750
2006	2,779	0	822	2,779	2,406	0
2007	2,696	2,401	903	2,687	2,356	803
2008	6,546	5,773	2,316	2,533	2,243	740
2009	10,289	9,406	3,532	10,277	8,937	3,086
Totals	41,175	31,696	11,001	37,045	32,133	6,994

TABLE 10. Descriptive Statistics: Children in Peru DHS Households

Variable	All Children (0-59m) in Surveyed Households			Children With Hospital Cards and Complete [†] Information		
	Mean	Std. Dev.	N	Mean	Std. Dev.	N
Birthweight (grams)	3238	497	11,001	3234	494	6,994
Height (cm)	83.9	13.8	31,696	81.5***	14.1	6,994
Age (months)	30.2	17.3	41,175	25.6***	16.8	6,994
Months Breastfed	14.6	8.7	41,175	13.9***	8.4	6,994
Household wealth quintile	2.6	1.3	27,478	2.8***	1.3	5,379
Mother literate	0.79	0.40	41,175	0.87***	0.33	6,994
Urban	0.52	0.50	39,668	0.64***	0.48	6,994
Has electricity	0.63	0.48	41,175	0.80***	0.40	6,994
Not moved since conception	0.87	0.34	41,175	1.00***	0.00	6,994
Hot 1st Trimester	0.05	0.21	37,045	0.03***	0.18	6,994
Hot 2nd Trimester	0.05	0.21	37,045	0.03***	0.17	6,994
Hot 3rd Trimester	0.05	0.21	37,045	0.03***	0.17	6,994

Stars indicate a statistically significant difference between the selected and non-selected samples at the 0.10 (*), 0.05 (**) and 0.01 (***) levels.

[†]These observations have no missing values for the key variables in the analysis (birth weight, height, and location (to determine temperature)).

TABLE 11. Effect of Extreme Heat During Gestation on Infant Mortality

	All (1)	Boys (2)	Girls (3)
Hot 1st Trimester	-0.007 (0.01)	-0.019*** (0.01)	0.006 (0.01)
Hot 2nd Trimester	-0.002 (0.01)	0.001 (0.01)	-0.004 (0.01)
Hot 3rd Trimester	-0.006 (0.01)	-0.004 (0.01)	-0.009 (0.01)
Children	37045	18872	18173
R-squared	0.02	0.03	0.03
Mean of dep var	0.03	0.04	0.03
Survey Year FEs	yes	yes	yes
Month of Birth FEs	yes	yes	yes
Year of Birth FEs	yes	yes	yes
Birth Order FEs	yes	yes	yes
Location FEs	yes	yes	yes

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 12. Determinants of Exposure to Extreme Heat During Gestation

	Hot TR1	Hot TR1	Hot TR2	Hot TR2	Hot TR3	Hot TR3
Mother's age (yrs)	-0.000 (0.00)	-0.000 (0.00)	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)
Mother literate	-0.007 (0.01)	-0.006 (0.01)	0.002 (0.00)	0.006 (0.00)	-0.006 (0.01)	-0.001 (0.01)
Urban	-0.006 (0.01)	-0.000 (0.01)	0.003 (0.01)	-0.003 (0.01)	0.000 (0.01)	-0.002 (0.00)
Has electricity	-0.013** (0.01)	-0.009 (0.01)	-0.002 (0.01)	0.002 (0.01)	-0.003 (0.01)	-0.004 (0.01)
Wealth quintile	0.002 (0.00)	0.005* (0.00)	-0.004* (0.00)	-0.006* (0.00)	0.001 (0.00)	0.000 (0.00)
Children	5379	5379	5379	5379	5379	5379
R-squared	0.00	0.12	0.00	0.10	0.00	0.14
Mean of dep var	.03	.03	.03	.03	.03	.03
Survey Year FEs		yes		yes		yes
Month of Birth FEs		yes		yes		yes
Year of Birth FEs		yes		yes		yes
Birth Order FEs		yes		yes		yes
Location FEs		yes		yes		yes

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. 1,615 children from the main sample are missing information about their household wealth quintile and are excluded from these regressions. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 13. Height (cm) as a Function of Birth Weight (100 g)

	(1)	(2)	(3)	(4)
Birth weight (100g)	0.220*** (0.02)	0.209*** (0.01)	0.204*** (0.01)	0.213*** (0.03)
Age (months)			1.323*** (0.02)	1.397*** (0.08)
Age ²			-0.009*** (0.00)	-0.011*** (0.00)
Birth weight × Age				-0.002 (0.00)
Birth weight × Age ²				0.000 (0.00)
Children	6994	6994	6994	6994
R-squared	0.78	0.94	0.94	0.94
Birth weight Source	card	card	card	card
Age FEs		yes		
Survey Year FEs	yes	yes	yes	yes
Month of Birth FEs	yes	yes	yes	yes
Year of Birth FEs	yes	yes	yes	yes
Birth Order FEs	yes	yes	yes	yes
Location FEs	yes	yes	yes	yes
Location	known	known	known	known

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 14. Birth Weight (100 g) as a Function of Extreme Heat During Gestation

	All	Boys			Girls		
	(1)	All (2)	Top 2 Wealth Quintiles (3)	Bottom 3 Quintiles (4)	All (5)	Top 2 Wealth Quintiles (6)	Bottom 3 Quintiles (7)
Hot 1st Trimester	-0.48 (0.37)	-0.32 (0.55)	-1.28 (0.85)	-0.09 (0.66)	-0.44 (0.54)	0.31 (1.79)	-0.64 (0.61)
Hot 2nd Trimester	0.37 (0.38)	0.08 (0.61)	0.56 (1.96)	-0.08 (0.69)	0.80 (0.50)	0.06 (1.70)	0.80 (0.51)
Hot 3rd Trimester	-0.88** (0.42)	-1.38** (0.63)	-1.88 (1.20)	-0.89 (0.69)	-0.47 (0.55)	-0.27 (1.53)	-0.61 (0.65)
Hot 1-3m post-birth	0.07 (0.26)	0.09 (0.34)	0.27 (0.89)	0.03 (0.41)	-0.13 (0.36)	0.56 (0.94)	-0.07 (0.36)
Children	6994	3584	832	2752	3410	802	2608
R-squared	0.10	0.14	0.19	0.16	0.13	0.23	0.15
Mean of Dep. Var.	3234.1	3282.7	3366.8	3257.3	3183.0	3289.7	3150.2
Birth Weight Source	card	card	card	card	card	card	card
Survey Year FEs	yes	yes	yes	yes	yes	yes	yes
Month of Birth FEs	yes	yes	yes	yes	yes	yes	yes
Year of Birth FEs	yes	yes	yes	yes	yes	yes	yes
Birth Order FEs	yes	yes	yes	yes	yes	yes	yes
Location FEs	yes	yes	yes	yes	yes	yes	yes
Location	known	known	known	known	known	known	known

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 15. Birth Weight (100 g) and Heat: Wealth Heterogeneity

	All (1)	Boys (2)	Girls (3)
Hot 1st Trimester	-0.47 (0.42)	-0.19 (0.63)	-0.57 (0.60)
Hot 2nd Trimester	0.39 (0.36)	0.02 (0.65)	1.01* (0.52)
Hot 3rd Trimester	-0.67 (0.46)	-0.94 (0.68)	-0.62 (0.65)
Hot 1-3m post-birth	0.08 (0.26)	0.13 (0.34)	-0.13 (0.36)
=1(Wealthy)	0.66*** (0.20)	0.59** (0.28)	0.73*** (0.23)
=1(Wealthy) × Hot 1st Trimester	-0.29 (1.04)	-0.88 (1.26)	0.49 (1.88)
=1(Wealthy) × Hot 2nd Trimester	-0.18 (1.19)	0.52 (1.73)	-1.82 (1.58)
=1(Wealthy) × Hot 3rd Trimester	-1.37 (1.14)	-2.43 (1.55)	0.65 (1.69)
Children	6994	3584	3410
R-squared	0.10	0.14	0.14
Mean of Dep. Var.	3234.1	3282.7	3183.0
Birth weight Source	card	card	card
Survey Year FEs	yes	yes	yes
Month of Birth FEs	yes	yes	yes
Year of Birth FEs	yes	yes	yes
Birth Order FEs	yes	yes	yes
Location FEs	yes	yes	yes
Location	known	known	known

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 16. Height (cm) and Extreme Heat During Gestation

	All (1)	All (2)	All (3)	Boys (4)	Boys (5)	Boys (6)	Girls (7)	Girls (8)	Girls (9)
Birth weight (100g)	0.20*** (0.01)		0.20*** (0.01)	0.20*** (0.01)		0.20*** (0.01)	0.19*** (0.01)		0.19*** (0.01)
Hot 1st Trimester		0.05 (0.35)	0.15 (0.34)		0.27 (0.45)	0.33 (0.45)		-0.01 (0.45)	0.07 (0.45)
Hot 2nd Trimester		0.29 (0.31)	0.22 (0.33)		-0.17 (0.40)	-0.20 (0.43)		0.90* (0.50)	0.75 (0.48)
Hot 3rd Trimester		-0.70** (0.32)	-0.52 (0.32)		-0.92*** (0.35)	-0.65* (0.35)		-0.50 (0.54)	-0.41 (0.55)
Age (months)	1.32*** (0.02)	1.32*** (0.02)	1.32*** (0.02)	1.35*** (0.03)	1.35*** (0.03)	1.35*** (0.03)	1.30*** (0.03)	1.29*** (0.03)	1.30*** (0.03)
Age ²	-0.01*** (0.00)								
Children	6994	6994	6994	3584	3584	3584	3410	3410	3410
R-squared	0.94	0.93	0.94	0.94	0.94	0.94	0.94	0.93	0.94
Mean of Dep. Var.	81.5	81.5	81.5	82.1	82.1	82.1	80.8	80.8	80.8
Survey Year FEs	yes								
Month of Birth FEs	yes								
Year of Birth FEs	yes								
Birth Order FEs	yes								
Location FEs	yes								
Location	known								

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 17. Height (cm) and Extreme Heat During Gestation: Boys by Wealth

	All Boys			Top 2 wealth quintiles			Bottom 3 wealth quintiles			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Birth weight (100g)	0.20*** (0.01)		0.20*** (0.01)	0.19*** (0.01)	0.16*** (0.03)		0.16*** (0.03)	0.20*** (0.01)		0.20*** (0.01)
Hot 1st Trimester		0.27 (0.45)	0.33 (0.45)	0.14 (0.49)		-0.16 (1.11)	0.11 (1.09)		0.06 (0.49)	0.08 (0.49)
Hot 2nd Trimester		-0.17 (0.40)	-0.20 (0.43)	-0.52 (0.49)		1.29 (0.84)	1.15 (0.82)		-0.62 (0.47)	-0.61 (0.49)
Hot 3rd Trimester		-0.92*** (0.35)	-0.65* (0.35)	-0.94** (0.38)		0.28 (0.58)	0.57 (0.46)		-1.10*** (0.41)	-0.91** (0.38)
Age (months)	1.35*** (0.03)	1.35*** (0.03)	1.35*** (0.03)	1.34*** (0.03)	1.38*** (0.05)	1.39*** (0.05)	1.38*** (0.05)	1.32*** (0.03)	1.32*** (0.03)	1.32*** (0.03)
Age ²	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
=1(Wealthy)				1.50*** (0.18)						
=1(Wealthy) × Hot 1st Trimester				0.24 (1.05)						
=1(Wealthy) × Hot 2nd Trimester				1.63** (0.78)						
=1(Wealthy) × Hot 3rd Trimester				1.05 (0.67)						
Children	3584	3584	3584	3584	832	832	832	2752	2752	2752
R-squared	0.94	0.94	0.94	0.94	0.96	0.95	0.96	0.94	0.94	0.94
Mean of Dep. Var.	82.1	82.1	82.1	82.1	84.4	84.4	84.4	81.4	81.4	81.4
Survey Year FEs	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Month of Birth FEs	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Year of Birth FEs	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Birth Order FEs	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Location FEs	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Location	known	known	known	known	known	known	known	known	known	known

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 18. Maternal Investment in Boys: Breastfeeding Duration (months)

	(1)	(2)	Top 2 Wealth Quintiles (3)	Bottom 3 Wealth Quintiles (4)
Birth weight (100g)	0.035 (0.04)	0.037 (0.04)	-0.044 (0.08)	0.052 (0.05)
Hot 3rd Trimester	0.180 (0.80)	0.928 (0.82)	-4.666 (3.96)	0.678 (0.87)
Age (months)	0.500*** (0.09)	0.515*** (0.08)	0.459** (0.20)	0.527*** (0.10)
Age ²	-0.006*** (0.00)	-0.006*** (0.00)	-0.006** (0.00)	-0.007*** (0.00)
=1(Wealthy)		-2.054*** (0.50)		
=1(Wealthy) × Hot 3rd Trimester		-4.960 (3.09)		
Children	2551	2551	602	1949
R-squared	0.19	0.20	0.21	0.23
Birth weight Source	card	card	card	card
Survey Year FEs	yes	yes	yes	yes
Month of Birth FEs	yes	yes	yes	yes
Year of Birth FEs	yes	yes	yes	yes
Birth Order FEs	yes	yes	yes	yes
Location FEs	yes	yes	yes	yes
Location	known	known	known	known

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 19. Maternal Investment in Girls: Breastfeeding Duration (months)

	(1)	(2)	Top 2 Wealth Quintiles (3)	Bottom 3 Wealth Quintiles (4)
Birth weight (100g)	0.079** (0.03)	0.089*** (0.03)	0.124 (0.09)	0.072** (0.04)
Hot 3rd Trimester	0.707 (1.16)	0.613 (1.16)	0.510 (3.85)	0.415 (1.21)
Age (months)	0.415*** (0.09)	0.414*** (0.09)	0.015 (0.15)	0.593*** (0.12)
Age ²	-0.005*** (0.00)	-0.004*** (0.00)	0.001 (0.00)	-0.007*** (0.00)
=1(Wealthy)		-2.112*** (0.53)		
=1(Wealthy) × Hot 3rd Trimester		2.144 (3.80)		
Children	2411	2411	553	1858
R-squared	0.19	0.19	0.23	0.22
Birth weight Source	card	card	card	card
Survey Year FEs	yes	yes	yes	yes
Month of Birth FEs	yes	yes	yes	yes
Year of Birth FEs	yes	yes	yes	yes
Birth Order FEs	yes	yes	yes	yes
Location FEs	yes	yes	yes	yes
Location	known	known	known	known

Notes: Standard errors, adjusted for clustering at the climate grid point level, are in parentheses. *, **, and *** denote significance at the 0.10, 0.05, and 0.01 levels, respectively.

TABLE 20. Health Services and Prices

Variable Name	Description	Price
Testing for HIV		
test_babies	Babies born to HIV+ women tested for HIV between 9-12 months after birth	\$4.00
test_other	Total number of HIV tests performed excluding test_babies	\$0.50
Initiation of Anti-Retroviral Therapy (ART)		
newART_pregnant	HIV+ pregnant women initiating anti-retroviral therapy (ART)	\$6.00
ARV	HIV+ mothers given ART for prevention of mother-to-child transmission	\$4.00
newART_baby	HIV+ children 0-24 months old initiating ART	\$6.00
newART_child	HIV+ children 2-14 years old initiating ART	\$5.00
newART_adult	HIV+ adults (excluding pregnant women) initiating ART	\$3.00
newART_TB	HIV+ adults co-infected with Tuberculosis initiating ART	\$3.00
Longterm Care and Treatment of HIV/AIDS		
TBtreat_HIV+	HIV+ patients treated for Tuberculosis	\$1.50
alive0014	HIV+ patients, 14 yrs or younger, alive and on treatment for 12 months	\$10.00
alive15up	HIV+ patients, 15 yrs or older, alive and on treatment for 12 months	\$5.00
Patients recovered	Former (lost to follow-up) HIV patients recovered and back on treatment	\$5.00
Family Planning, Pregnancy, and Births		
family planning	Family planning consultations where a contraception method is provided	\$0.15
preg_4ANC	Pregnant women completing 4 ante-natal visits	\$0.22
deliveries	Institutional births	\$1.00
children vaccinated	Children fully vaccinated by 9 months of age	\$0.13
post-partum cons	Post-partum consultations with new mothers between 3-28 days after birth	\$0.18
Non-paid Services		
triage	Patients who present for treatment of any acute illness or injury	\$0.00
well-child cons	Well-child consultations	\$0.00
preg_Malaria	Pregnant women treated with malarial prophylaxis	\$0.00
pre-eclampsia	Pregnant women treated for pre-eclampsia	\$0.00

TABLE 21. Descriptive Statistics: Facility Characteristics

Variable	Treatment	Comparison	P-value
	Mean (1)	Mean (2)	(1)=(2) (3)
Health workers	54.05 (54.28)	58.39 (33.80)	0.70
Other staff	16.60 (30.88)	14.93 (18.41)	0.79
Total workers	70.65 (81.49)	73.32 (36.65)	0.85
Beds	52.80 (37.67)	66.30 (66.42)	0.40
Access to water on site	0.900 (0.308)	0.795 (0.408)	0.31
Electricity	0.700 (0.470)	0.795 (0.408)	0.41
Generator on site	0.650 (0.489)	0.341 (0.479)	0.02
Refrigerator on site	0.950 (0.224)	0.932 (0.255)	0.78
Lab on site	0.700 (0.470)	0.295 (0.462)	0.001
Observations ^a	20	44	

Standard deviations are in parentheses.

^a Facility census information is missing for 8 facilities.

TABLE 22. Descriptive Statistics: Average Count of Services Delivered, Pre-Treatment

Variable	Treatment		Comparison		P-value (1)=(2) (3)
	Mean (1)	N [†]	Mean (2)	N [†]	
Paid Services					
Babies tested for HIV	2.6 (6.2)	91	12.4 (17.7)	126	0.00
All others tested for HIV	1009.7 (1229.7)	32	851.0 (605.9)	72	0.38
PMTCT	25.7 (29.9)	92	43.2 (47.9)	163	0.00
Babies started on ART	3.2 (4.3)	84	3.8 (4.6)	154	0.31
Children (2-14) started on ART	3.2 (4.4)	84	5.3 (7.8)	151	0.03
Non-pregnant Adults started on ART	50.2 (73.8)	84	80.4 (108.8)	139	0.03
Tuberculosis patients started on ART	3.2 (6.1)	80	6.2 (9.3)	135	0.01
HIV+ patients treated for Tuberculosis	8.5 (32.5)	74	8.1 (13.6)	134	0.90
Youth (0-14) patients alive and on treatment for 12 months	5.0 (9.4)	75	5.2 (6.1)	120	0.83
Adult (15+) patients alive and on treatment for 12 months	34.7 (53.2)	79	57.3 (68.0)	104	0.02
Former ART patients recovered and back on treatment	5.7 (15.9)	80	10.3 (11.9)	90	0.03
Pregnant women completing 4 ante-natal care visits	7.3 (10.4)	24	65.8 (94.1)	42	0.00
Institutional deliveries (births)	346.4 (276.0)	96	344.6 (331.2)	187	0.96
Children completely vaccinated	258.4 (213.9)	96	353.7 (212.2)	173	0.00
Family planning consultations	271.0 (297.0)	95	367.7 (320.8)	160	0.02
Post-partum consultations	380.8 (330.4)	96	371.7 (210.6)	112	0.81
Non-paid Services					
Triage visits	11426.6 (7084.2)	96	14339.8 (9314.8)	187	0.01
Well-child consultations	1112.5 (2235.2)	92	3945.8 (4283.4)	183	0.00
Pregnant women given malarial prophylaxis	147.6 (227.8)	61	167.7 (160.0)	124	0.49
Pregnant women treated for pre-eclampsia	6.7 (5.6)	23	12.6 (15.9)	87	0.08

Notes: Standard deviations are in parentheses.

[†]A service with no missing data in Treatment facilities will have $N = 96$ (24 facilities X 4 quarters). Complete data in Comparison facilities is $N = 192$ (48 facilities X 4 quarters).

TABLE 23. Paid Services Pre-Treatment Trend Differences and Treatment Effects

Dependent Variable	Trend Difference (Pre-Treatment)		Diff-in-Diff Point Estimate			
	Full Sample (1)	Zeros Excluded (2)	All Periods		Periods 5 & 6 Excluded	
			Full Sample (3)	Zeros Excluded (4)	Full Sample (5)	Zeros Excluded (6)
Paid Services						
test_baby	-0.24** (0.11)	-0.30** (0.07)	1.31** † (0.19)	0.94** † (0.14)	1.34** † (0.19)	0.95** † (0.15)
test_other	0.16** (0.06)	0.16** (0.06)	-0.26 † (0.17)	-0.25 † (0.17)	-0.24 † (0.19)	-0.24 † (0.19)
PMTCT	0.07 (0.05)	0.07 (0.05)	0.13 (0.23)	0.13 (0.23)	0.11 (0.21)	0.11 (0.21)
newART baby	-0.04 (0.08)	-0.07 (0.07)	0.13 (0.12)	0.01 (0.10)	0.21* (0.19)	0.08 (0.11)
newART child	0.05 (0.09)	0.04 (0.08)	0.29* (0.16)	0.25* (0.15)	0.44** (0.19)	0.39** (0.17)
newART adult	0.06 (0.05)	0.06 (0.50)	0.34** (0.15)	0.35** (0.15)	0.40** (0.16)	0.41** (0.16)
newART TB	0.08 (0.12)	0.02 (0.09)	0.82** (0.37)	0.70* (0.37)	1.06** (0.40)	0.92** (0.39)
TBtreat HIV+	0.67* (0.38)	0.78* (0.40)	0.92** (0.42)	0.73* (0.37)	1.03** (0.46)	0.80** (0.41)
alive 0-14	0.07 (0.14)	0.06 (0.14)	-0.08 (0.22)	-0.06 (0.22)	-0.25 (0.18)	-0.29* (0.16)
alive 15up	-0.01 (0.05)	-0.01 (0.05)	0.00 (0.11)	0.00 (0.11)	0.14 (0.10)	0.14 (0.10)
patients recovered	-0.48* (0.26)	-0.37* (0.20)	0.49 (0.40)	0.18 (0.35)	0.84** (0.42)	0.48 (0.35)
preg 4ANC	0.08 (0.27)	0.09 (0.28)	1.57** (0.55)	1.36** (0.65)	1.99** (0.57)	1.53** (0.68)
deliveries	-0.05* (0.02)	-0.05* (0.02)	0.01 (0.06)	0.01 (0.06)	0.03 (0.06)	0.03 (0.06)
children vaccinated	-0.06 (0.05)	-0.08 (0.05)	-0.12 (0.12)	-0.15 (0.13)	-0.08 (0.12)	-0.11 (0.12)
family planning	0.08 (0.08)	0.08 (0.08)	0.37** (0.16)	0.38** (0.16)	0.45** (0.18)	0.45** (0.18)
post-partum cons	-0.03 (0.05)	-0.03 (0.05)	-0.28** (0.10)	-0.25** (0.10)	-0.27** (0.14)	-0.27** (0.14)

Notes: All models include facility and quarter (for seasonality) fixed effects. Standard errors clustered at the facility level are in parentheses. Stars denote significance at the 10%(*) or 5%(**) levels.

† These point estimates should not be interpreted as treatment effects because these services violate the assumption of common trends before treatment.

TABLE 24. Non-Paid Services Pre-Treatment Trend Differences and Treatment Effects

Dependent Variable	Trend Difference (Pre-Treatment)		Diff-in-Diff Point Estimate	
	Full Sample	Zeros Excluded	Full Sample	Zeros Excluded
	(1)	(2)	(3)	(4)
triage	0.00 (0.04)	0.00 (0.04)	0.00 (0.09)	0.00 (0.09)
well-child cons	0.09 (0.07)	0.09 (0.07)	0.00 (0.14)	0.00 (0.14)
preg malaria	-0.19 (0.14)	-0.12 (0.13)	-0.20 (0.21)	-0.17 (0.22)
pre-eclampsia	0.23** (0.09)	0.18** (0.07)	-0.18 † (0.25)	-0.21 † (0.25)

Notes: All models include facility and quarter (for seasonality) fixed effects. Standard errors clustered at the facility level are in parentheses. Stars denote significance at the 10%(*) or 5%(**) levels.

† These point estimates should not be interpreted as treatment effects because these services violate the assumption of common trends before treatment.

TABLE 25. Total Value (dollars) of Service Mix

	Treatment Facilities	Comparison Facilities	Difference
Pre	1160.2 (1259.3)	1278.1 (1117.2)	-107.8
Post	1963.3 (1456.6)	1778.0 (1611.3)	185.3
Difference in Differences			293.1 (145.8)

Notes: Standard deviations are in parentheses.

TABLE 26. Heterogeneity of Treatment Effects by Facility Capacity Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
TGxTP	293.09** (145.8)	82.88 (576.1)	253.05 (554.5)	357.46** (154.0)	289.65* (158.8)	-258.30 (607.3)	298.65* (166.1)	260.33 (204.6)	-1037.21* (609.0)
x ln(health workers)		71.00 (150.9)							288.25 (208.9)
x ln(total workers)			21.38 (131.1)						-145.51 (116.0)
x Staff-to-skilled-worker Ratio				-40.45 (37.2)					
x Hospital					13.78 (223.7)				-279.39 (300.6)
x ln(beds)						163.46 (162.3)			250.50 (199.0)
x Generator							-26.08 (185.8)		-176.12 (163.1)
x Laboratory								52.42 (213.8)	136.11 (146.9)
Observations	648	630	630	630	648	621	612	648	594
R ²	0.189	0.178	0.177	0.178	0.189	0.186	0.165	0.189	0.175

Notes: All models include facility and quarter (for seasonality) fixed effects. Standard errors clustered at the facility level are in parentheses. Stars denote significance at the 10%(*) or 5%(**) levels.

REFERENCES CITED

- Ahrend, R. (2002). *Press Freedom, Human Capital and Corruption*. Delta.
- Albala-Bertrand, J.-M. (1993). Natural disaster situations and growth: A macroeconomic model for sudden disaster impacts. *World Development*, 21(9):1417–1434.
- Aldrich, D. P. (2010). Separate and unequal: Post-tsunami aid distribution in southern India. *Social Science Quarterly*, 91(5):1369–1389.
- Alesina, A. and Dollar, D. (2000). Who gives foreign aid to whom and why? *Journal of Economic Growth*, 5(1):33–63.
- Almond, D. and Currie, J. (2011). Killing me softly: The fetal origins hypothesis. *Journal of Economic Perspectives*, 25(3):153.
- Almond, D. and Mazumder, B. (2011). Health capital and the prenatal environment: The effect of Ramadan observance during pregnancy. *American Economic Journal: Applied Economics*, 3(4):56–85.
- Anbarci, N., Escaleras, M., and Register, C. A. (2005). Earthquake fatalities: The interaction of nature and political economy. *Journal of Public Economics*, 89(9):1907–1933.
- Angrist, J. D. and Pischke, J.-S. (2008). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Anttila-Hughes, J. K. and Hsiang, S. M. (2012). Destruction, disinvestment, and death: Economic and human losses following environmental disaster. *Unpublished Manuscript*.
- Arnové, R. F. (1981). The Nicaraguan national literacy crusade of 1980. *Comparative Education Review*, 25(2):244–260.
- Ashraf, N., Bandiera, O., and Jack, K. (2012). No margin, no mission? A field experiment on incentives for pro-social tasks. *CEPR Discussion Paper No. DP8834*.
- Baez, J. E. and Santos, I. V. (2007). Childrens vulnerability to weather shocks: A natural disaster as a natural experiment. *Social Science Research Network, New York*.
- Bandiera, O., Barankay, I., and Rasul, I. (2012). Team incentives: Evidence from a firm level experiment. *CEPR Discussion Paper No. DP8776*.
- Bardhan, P. and Mookherjee, D. (2000). Capture and governance at local and national levels. *The American Economic Review*, 90(2):135–139.
- Barker, D., Bull, A., Osmond, C., and Simmonds, S. (1990). Fetal and placental size and risk of hypertension in adult life. *BMJ*, 301(6746):259–262.
- Barker, D. J. (1995). Fetal origins of coronary heart disease. *BMJ*, 311(6998):171–174.
- Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2013). Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the 20th century. Technical report, National Bureau of Economic Research.
- Barreca, A., Deschenes, O., and Guldi, M. (2015). Maybe next month? The dynamic effects of ambient temperature on fertility. *Unpublished manuscript*.

- Barreca, A. I. (2012). Climate change, humidity, and mortality in the united states. *Journal of Environmental Economics and Management*, 63(1):19–34.
- Basinga, P., Gertler, P., Binagwaho, A., Soucat, A., Sturdy, J., and Vermeersch, C. (2011). Effect on maternal and child health services in Rwanda of payment to primary health-care providers for performance: An impact evaluation. *The Lancet*, 377(9775):1421–1428.
- BBC News (2013). Typhoon Haiyan death toll rises over 5,000. <http://www.bbc.co.uk/news/world-asia-25051606>. Accessed: 2014-09-30.
- Becerra, O., Cavallo, E. A., and Noy, I. (2012). *Foreign Aid in the Aftermath of Large Natural Disasters*. Inter-American Development Bank, Research Department.
- Benabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *The Review of Economic Studies*, 70(3):489–520.
- Besley, T. and Burgess, R. (2002). The political economy of government responsiveness: Theory and evidence from India. *The Quarterly Journal of Economics*, 117(4):1415–1451.
- Besley, T. and Case, A. (1995). Does electoral accountability affect economic policy choices? Evidence from gubernatorial term limits. *The Quarterly Journal of Economics*, 110(3):769–798.
- Besley, T. and Prat, A. (2006). Handcuffs for the grabbing hand? Media capture and government accountability. *The American Economic Review*, 96(3):720–736.
- Biderman, J. (1983). The development of capitalism in Nicaragua: A political economic history. *Latin American Perspectives*, 10(1):7–32.
- Black, G. (1981). *Triumph of the People: The Sandinista Revolution in Nicaragua*. Zed Press London.
- Black, S. E., Devereux, P. J., and Salvanes, K. G. (2007). From the cradle to the labor market? The effect of birth weight on adult outcomes. *The Quarterly Journal of Economics*, 122(1):409–439.
- Blankespoor, B., Dasgupta, S., Laplante, B., and Wheeler, D. (2010). The economics of adaptation to extreme weather events in developing countries. *Center for Global Development Working Paper 199*.
- Boning, B., Ichniowski, C., and Shaw, K. (2007). Opportunity counts: Teams and the effectiveness of production incentives. *Journal of Labor Economics*, 25(4):613–650.
- Brown, R. (2014). The Mexican drug war and early-life health: The impact of violent crime on birth outcomes. *Unpublished manuscript*.
- Brunetti, A. and Weder, B. (2003). A free press is bad news for corruption. *Journal of Public Economics*, 87(7):1801–1824.
- Buckles, K. S. and Hungerman, D. M. (2013). Season of birth and later outcomes: Old questions, new answers. *Review of Economics and Statistics*, 95(3):711–724.
- Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2011). Weather and death in India. *Unpublished manuscript*.
- Burnside, C. and Dollar, D. (2000). Aid, policies, and growth. *The American Economic Review*, 90(4):847–868.

- Cavallo, E., Galiani, S., Noy, I., and Pantano, J. (2013). Catastrophic natural disasters and economic growth. *Review of Economics and Statistics*, 95(5):1549–1561.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869.
- Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., and Rogers, F. (2006). Missing in action: Teacher and health worker absence in developing countries. *The Journal of Economic Perspectives*, 20(1):91–116.
- CIA, US (2013). The World Factbook.
<https://www.cia.gov/library/publications/the-world-factbook/index.html>.
 Accessed: 2013-011-12.
- Cil, G. and Cameron, T. A. (2014). Heat waves: Effects on fertility decisions and pregnancy outcomes. *University of Oregon Working Paper*.
- Cohen, C. and Werker, E. D. (2008). The political economy of ‘natural’ disasters. *Journal of Conflict Resolution*, 52(6):795–819.
- Currie, J. (2011). Inequality at birth: Some causes and consequences. *The American Economic Review*, 101(3):1–22.
- Currie, J. and Almond, D. (2011). Human capital development before age five. *Handbook of Labor Economics* 4, pages 1315–1486.
- Dasgupta, A. and Beard, V. A. (2007). Community driven development, collective action and elite capture in Indonesia. *Development and Change*, 38(2):229–249.
- Davie, R., Butler, N., and Goldstein, H. (1972). From birth to seven: The second report of the national child development study 1958 cohort. *London Longmans*.
- de Onis, M., Onyango, A., Borghi, E., Siyam, A., Blössner, M., and Lutter, C. (2012). Worldwide implementation of the who child growth standards. *Public Health Nutrition*, 15(09):1603–1610.
- Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. World Bank Publications.
- Deci, E. and Porac, J. (1978). *The Hidden Costs of Reward: New Perspectives on the Psychology of Human Motivation*. Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Deryugina, T. (2011). The dynamic effects of hurricanes in the US: The role of non-disaster transfer payments. *MIT Center for Energy and Environmental Policy Research Working Paper*.
- Deschênes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, 3(4):152–185.
- Deschênes, O., Greenstone, M., and Guryan, J. (2009). Climate change and birth weight. *The American Economic Review Papers and Proceedings*, 99(2):211–217.
- Drago, R. and Garvey, G. (1998). Incentives for helping on the job: Theory and evidence. *Journal of Labor Economics*, 16(1):1–25.

- Dreher, A., Mölders, F., and Nunnenkamp, P. (2010). Aid delivery through non-governmental organisations: Does the aid channel matter for the targeting of Swedish aid? *The World Economy*, 33(2):147–176.
- Dreher, A., Nunnenkamp, P., Öhler, H., and Weisser, J. (2012). Financial dependence and aid allocation by Swiss NGOs: A panel tobit analysis. *Economic Development and Cultural Change*, 60(4):829–867.
- Drury, A. C., Olson, R. S., and Belle, D. A. V. (2005). The politics of humanitarian aid: US foreign disaster assistance, 1964–1995. *Journal of Politics*, 67(2):454–473.
- ECLAC (1999). Nicaragua: Assessment of the damage caused by Hurricane Mitch, 1998. *United Nations Economic Commission for Latin America and the Caribbean (ECLAC)*.
- Eisensee, T. and Strömberg, D. (2007). News droughts, news floods, and US disaster relief. *The Quarterly Journal of Economics*, 122(2):693–728.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness: Intentions matter. *Games and Economic Behavior*, 62(1):287–303.
- Falk, A. and Kosfeld, M. (2006). The hidden costs of control. *The American Economic Review*, 96(5):1611–1630.
- Fehr, E. and Gächter, S. (2002). Do incentive contracts undermine voluntary cooperation? *Zurich IEER Working Paper*.
- Fehr, E. and Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868.
- Ferreira, S., Hamilton, K., and Vincent, J. R. (2013). Does development reduce fatalities from natural disasters? New evidence for floods. *Environment and Development Economics*, 18(06):649–679.
- Figlio, D., Guryan, J., Karbownik, K., and Roth, J. (2014). The effects of poor neonatal health on children’s cognitive development. *The American Economic Review*, 104(12):3921–3955.
- Fischetti, M. (2013). Was Typhoon Haiyan a record storm? <http://blogs.scientificamerican.com/observations/2013/11/12/was-typhoon-haiyan-a-record-storm/>. Accessed: 2014-09-30.
- Fitzgerald, J., Gottschalk, P., and Moffitt, R. A. (1998). An analysis of sample attrition in panel data: The Michigan Panel Study of Income Dynamics. *Journal of Human Resources*, 33(2):251–299.
- Francken, N., Minten, B., and Swinnen, J. F. (2012). The political economy of relief aid allocation: Evidence from Madagascar. *World Development*, 40(3):486–500.
- Freeman, P., Keen, M., and Mani, M. (2003). Dealing with increased risk of natural disasters: Challenges and options. *IMF Working Paper*.
- Frey, B. and Jegen, R. (2001). Motivation crowding theory. *Journal of Economic Surveys*, 15(5):589–611.
- Fritzen, S. A. (2007). Can the design of community-driven development reduce the risk of elite capture? Evidence from Indonesia. *World Development*, 35(8):1359–1375.
- Gaynor, M. and Pauly, M. (1990). Compensation and productive efficiency in partnerships: Evidence from medical groups practice. *Journal of Political Economy*, pages 544–573.

- Gneezy, U. and Rustichini, A. (2000). Pay enough or don't pay at all. *The Quarterly Journal of Economics*, 115(3):791–810.
- Gravelle, H., Sutton, M., and Ma, A. (2010). Doctor behaviour under a pay for performance contract: Treating, cheating and case finding? *The Economic Journal*, 120(542):F129–F156.
- Grosh, M. E. and Glewwe, P. (1995). *A Guide to Living Standards Measurement: Study Surveys and Their Data Sets*. The World Bank.
- Guiney, J. L. and Lawrence, M. B. (1999). Preliminary report: Hurricane Mitch: 22 October–05 November 1998. *United States National Oceanic and Atmospheric Administration's National Hurricane Center*.
- Hanemann, U. (2006). Nicaragua's literacy campaign. *Paper commissioned for the EFA Global Monitoring Report*, pages 1–13.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.
- Hellin, J., Haigh, M., and Marks, F. (1999). Rainfall characteristics of Hurricane Mitch. *Nature*, 399(6734):316–316.
- Heywood, J., Jirjahn, U., and Wei, X. (2008). Teamwork, monitoring and absence. *Journal of Economic Behavior & Organization*, 68(3):676–690.
- Holmstrom, B. and Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization*, 7(special issue):24–52.
- Horwitz, J. R. and Nichols, A. (2009). Hospital ownership and medical services: Market mix, spillover effects, and nonprofit objectives. *Journal of Health Economics*, 28(5):924–937.
- Iannotti, L. L., Zavaleta, N., Huasaquiche, C., Leon, Z., and Caulfield, L. E. (2012). Early growth velocities and weight gain plasticity improve linear growth in Peruvian infants. *Maternal & Child Nutrition*, 11(1):127–137.
- Jakobsen, K. T. (2012). In the eye of the storm: The welfare impacts of a hurricane. *World Development*, 40(12):2578–2589.
- Jayne, T. S., Strauss, J., Yamano, T., and Molla, D. (2001). Giving to the poor? Targeting of food aid in rural Ethiopia. *World Development*, 29(5):887–910.
- Jayne, T. S., Strauss, J., Yamano, T., and Molla, D. (2002). Targeting of food aid in rural Ethiopia: Chronic need or inertia? *Journal of Development Economics*, 68(2):247–288.
- Kahn, M. E. (2005). The death toll from natural disasters: The role of income, geography, and institutions. *Review of Economics and Statistics*, 87(2):271–284.
- Kandel, E. and Lazear, E. (1992). Peer pressure and partnerships. *Journal of Political Economy*, 100(4):801–817.
- Keck, M. (2014). Comparing the determinants of US-funded NGO aid versus US Official Development Aid. *VOLUNTAS: International Journal of Voluntary and Nonprofit Organizations*, pages 1–23, published online, August 05, 2014.
- Kellenberg, D. K. and Mobarak, A. M. (2008). Does rising income increase or decrease damage risk from natural disasters? *Journal of Urban Economics*, 63(3):788–802.

- Knabb, R. D., Rhome, J. R., and Brown, D. P. (2005). Tropical cyclone report: Hurricane Katrina. Technical report, United States National Oceanic and Atmospheric Administration's National Hurricane Center.
- Knez, M. and Simester, D. (2001). Firm-wide incentives and mutual monitoring at Continental Airlines. *Journal of Labor Economics*, 19(4):743–772.
- Knutson, T. R., McBride, J. L., Chan, J., Emanuel, K., Holland, G., Landsea, C., Held, I., Kossin, J. P., Srivastava, A., and Sugi, M. (2010). Tropical cyclones and climate change. *Nature Geoscience*, 3(3):157–163.
- Koch, D.-J., Dreher, A., Nunnenkamp, P., and Thiele, R. (2009). Keeping a low profile: What determines the allocation of aid by non-governmental organizations? *World Development*, 37(5):902–918.
- Kudamatsu, M., Persson, T., and Strömberg, D. (2012). Weather and infant mortality in Africa. *CEPR Discussion Paper No. DP9222*.
- Lavy, V. (2002). Evaluating the effect of teachers group performance incentives on pupil achievement. *Journal of Political Economy*, 110(6):1286–1317.
- Lazear, E. (2000). Performance pay and productivity. *The American Economic Review*, 90(5):1346–1361.
- McCown, S., Graumann, A., Ross, T., and Lott, N. (1999). Mitch: The deadliest Atlantic hurricane since 1780. *National Climatic Data Center, Asheville, NC*.
- Mendelsohn, R., Dinar, A., and Williams, L. (2006). The distributional impact of climate change on rich and poor countries. *Environment and Development Economics*, 11(02):159–178.
- Morris, S. S. and Wodon, Q. (2003). The allocation of natural disaster relief funds: Hurricane Mitch in Honduras. *World Development*, 31(7):1279–1289.
- Morrison, F. L. (1987). Legal issues in the Nicaragua opinion. *The American Journal of International Law*, 81(1):160–166.
- NDRRMC (2013). Sitrep no. 46 effects of Typhoon Yolanda (Haiyan). Technical report, Republic of the Philippines.
- Newhouse, J. (1973). The economics of group practice. *Journal of Human Resources*, 8(1):37–56.
- Newhouse, J. P. (1970). Toward a theory of nonprofit institutions: An economic model of a hospital. *The American Economic Review*, 60(1):64–74.
- Nicholas, L., Dimick, J., and Iwashyna, T. (2011). Do hospitals alter patient care effort allocations under pay-for-performance? *Health Services Research*, 46(1p1):61–81.
- Nose, M. (2014). Micro responses to disaster relief aid: Design problems for aid efficacy. *Economic Development and Cultural Change*, 62(4):727–767.
- Noy, I. (2009). The macroeconomic consequences of disasters. *Journal of Development Economics*, 88(2):221–231.
- Nunnenkamp, P. and Öhler, H. (2011). Aid allocation through various official and private channels: Need, merit, and self-interest as motives of German donors. *World Development*, 39(3):308–323.

- Nunnenkamp, P., Weingarh, J., and Weisser, J. (2009). Is NGO aid not so different after all? Comparing the allocation of Swiss aid by private and official donors. *European Journal of Political Economy*, 25(4):422–438.
- Olson, R. S., Alvarez, R., Baird, B., Estrada, A., Gawronski, V., and Sarmiento, J. P. (2001). The storms of 98: Hurricanes Georges and Mitch, impacts, institutional response, and disaster politics in three countries, Special Publication 38. *University of Colorado Natural Hazards Research and Applications Information Center*.
- Owens, T., Hoddinott, J., and Kinsey, B. (2003). Ex-ante actions and ex-post public responses to drought shocks: Evidence and simulations from Zimbabwe. *World Development*, 31(7):1239–1255.
- OXFAM (2013). Oxfam reaches nearly 250,000 typhoon survivors in the Philippines. <http://www.oxfamamerica.org/articles/oxfam-reaches-nearly-250-000-typhoon-survivors-in-the-philippines>. Accessed: 2014-12-12.
- Paarsch, H. and Shearer, B. (2001). Piece rates, fixed wages, and incentive effects: Statistical evidence from payroll records. *International Economic Review*, 41(1):59–92.
- Pachauri, R. K. and Reisinger, A. (2007). Climate Change 2007: Synthesis Report. Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. *Intergovernmental Panel on Climate Change*.
- Parsons, T. J., Power, C., and Manor, O. (2001). Fetal and early life growth and body mass index from birth to early adulthood in 1958 British Cohort: Longitudinal study. *British Medical Journal*, 323(7325):1331–1335.
- Pelling, M., Maskrey, A., Ruiz, P., and Hall, L. (2004). Reducing disaster risk: A challenge for development. *United Nations Development Bank, Bureau for Crisis Prevention and Recovery, New York*.
- Philippines, Republic of the (2013). Sitrep no. 46 effects of Typhoon Yolanda (Haiyan).
- Platteau, J.-P. (2004). Monitoring elite capture in community-driven development. *Development and Change*, 35(2):223–246.
- Plümper, T. and Neumayer, E. (2009). Famine mortality, rational political inactivity, and international food aid. *World Development*, 37(1):50–61.
- Premand, P. (2008). Hurricane Mitch and consumption growth of Nicaraguan agricultural households. Technical report, University of Oxford, Department of Economics.
- Prendergast, C. (1999). The provision of incentives in firms. *Journal of Economic Literature*, 37(1):7–63.
- Rocha, J. L. and Christopolos, I. (2001). Disaster mitigation and preparedness on the Nicaraguan post-Mitch agenda. *Disasters*, 25(3):240–250.
- Rosenthal, M., Frank, R., Li, Z., and Epstein, A. (2005). Early experience with pay-for-performance. *Journal of the American Medical Association*, 294(14):1788–1793.
- Sadowski, N. C. and Sutter, D. (2005). Hurricane fatalities and hurricane damages: Are safer hurricanes more damaging? *Southern Economic Journal*, 72(2):422–432.
- Schenk, A. (1999). Hurricane Mitch and disaster relief: The politics of catastrophe. <http://www.solidarity-us.org/site/node/1752>. Accessed: 2014-09-29.

- Scheraga, J., Ebi, K. L., Furlow, J., and Moreno, A. (2003). From Science to Policy: Developing Responses to Climate Change. In *Climate Change and Human Health: Risks and Responses*, pages 237–266. OMS.
- Shearer, B. (2004). Piece rates, fixed wages and incentives: Evidence from a field experiment. *Review of Economic Studies*, 71(2):513–534.
- Sliwka, D. (2007). Trust as a signal of a social norm and the hidden costs of incentive schemes. *The American Economic Review*, 97(3):999–1012.
- Smith, V. K., Carbone, J. C., Pope, J. C., Hallstrom, D. G., and Darden, M. E. (2006). Adjusting to natural disasters. *Journal of Risk and Uncertainty*, 33(1-2):37–54.
- Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, B., Tignor, M., and Miller, H. (2007). Climate Change 2007: Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. *Intergovernmental Panel on Climate Change*.
- Stocker, T., Qin, D., Plattner, G., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P. (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
- Strobl, E. (2012). The economic growth impact of natural disasters in developing countries: Evidence from hurricane strikes in the Central American and Caribbean Regions. *Journal of Development Economics*, 97(1):130–141.
- Strömberg, D. (2001). Mass media and public policy. *European Economic Review*, 45(4):652–663.
- Strömberg, D. (2007). Natural disasters, economic development, and humanitarian aid. *The Journal of Economic Perspectives*, 21(3):199–222.
- Takasaki, Y. (2011). Do local elites capture natural disaster reconstruction funds? *Journal of Development Studies*, 47(9):1281–1298.
- Tanner, J. and Whitehouse, R. (1976). Clinical longitudinal standards for height, weight, height velocity, weight velocity, and stages of puberty. *Archives of Disease in Childhood*, 51(3):170–179.
- Tanner, J. M., Whitehouse, R., and Takaishi, M. (1966). Standards from birth to maturity for height, weight, height velocity, and weight velocity: British children. *Archives of Disease in Childhood*, 41(219):454.
- Tendler, J. (1982). *Turning Private Voluntary Organizations Into Development Agencies: Questions for Evaluation*. Washington DC USAID Bureau of Food and Voluntary Assistance Office of Program and Management Support.
- Torche, F. (2011). The effect of maternal stress on birth outcomes: Exploiting a natural experiment. *Demography*, 48(4):1473–1491.
- Van den Berg, M. (2010). Household income strategies and natural disasters: Dynamic livelihoods in rural Nicaragua. *Ecological Economics*, 69(3):592–602.
- Wilde, J., Apouey, B., and Jung, T. (2014). Heat waves at conception and later life outcomes. *Unpublished Manuscript*.
- World Bank (2001a). Nicaragua Poverty Assessment. *The World Bank, Washinton, DC*.

- World Bank (2001b). World Development Report 2000/2001: Attacking Poverty. *The World Bank, Washinton, DC.*
- World Bank (2014). World Development Indicators. *The World Bank, Washinton, DC.*
- World Bank (2015). The World Bank Data Catalog.
<http://data.worldbank.org/country/peru>. Accessed: 2015-03-23.
- World Health Organization (2006). *WHO Multicentre Growth Reference Study Group. WHO Child Growth Standards: Length/Height-for-age, Weight-for-age, Weight-for-length, Weight-for-height and Body Mass Index-for-age: Methods and Development.* Geneva: World Health Organization.
- World Health Organization (2009). *WHO Child Growth Standards: Growth Velocity Based on Weight, Length, and Head Circumference - Methods and Development.* Geneva: World Health Organization.
- Yang, D. (2008). Coping with disaster: The impact of hurricanes on international financial flows, 1970-2002. *The BE Journal of Economic Analysis & Policy*, 8(1). ISSN (Online) 1935-1682.