Essays in Development, Environmental, and Health Economics

Luís Ricardo Maertens Odria

TESI DOCTORAL UPF / ANY 2017

DIRECTOR DE LA TESI

Antonio Ciccone
Departament d'Economia i Empresa

A mis padres y a Michelle

Acknowledgements

Curiosity and hopes of leaving a positive footprint made me gravitate towards research. I have many to thank for guiding my curiosity and educating my thinking. I have many more to thank for their encouragement, friendship, and love.

I am deeply grateful to Antonio Ciccone, from whom I have learned how to think critically about economic phenomena and to conduct rigorous research. His guidance, honesty, and support have been invaluable. I am also especially grateful to Alessandro Tarozzi, his comments and suggestions have significantly improved my work, while his encouragement has made the work that much more enjoyable. I also gratefully acknowledge comments and suggestions from the applied microeconomics faculty at UPF, especially from Libertad González, Gianmarco León, and Christian Fons-Rosen.

Francesco Amodio and Bruno Caprettini, colleagues and friends, have been immensely generous with me. I thank them both for all the time spent discussing my research, and for sharing their research with me. Their encouragement and support are also greatly appreciated.

I thank Cristina Bellés, Kizkitza Biguri, Miguel De Jesus, Fernando Fernández, Helena Hernández, Niklas Heusch, Francisco Queirós, María Lombardi, Marco Nieddu, Lorenzo Pandolfi, Thomas Stoerk, Darren Tang, Ana Varatojo Dos Santos, Tomás Williams, and Dijana Zejcirovic for their comments, support, and friendship. They have made Barcelona a familiar place to me.

I thank my parents for their love, support, and patience. Writing this thesis required attention to details, effort, discipline, and tenacity. I first learned these things from them.

I thank Michelle Benard for her help proof-reading my work. She has also made me a kinder person, filled my life with beauty and happiness, and shown me unconditional support.

Abstract

This thesis is composed of three independent essays. In the first chapter, I analyze the effect of a biofuel-favorable policy in the U.S. on fetal health. I show that the policy led to an expansion in the production of corn, a pesticide-intensive crop, and to increased risk of fetal conditions previously associated with exposure to corn pesticides. In the second chapter, I examine the role of agricultural productivity as a mechanism linking rainfall shocks to civil wars in African countries. I show that rainfall over agricultural land and during the growing season has a hump-shaped relationship with agricultural output, which is mirrored by a U-shaped relationship with civil war risk. In the third chapter, I examine the effect of various selling schemes for testing tubewell water for arsenic on test uptake and, conditional on adverse news, on health-protective behavior. I find that uptake is increased by fees that depend on test results, and that social networks and public information can promote health-protective behavior.

Resumen

Esta tesis consta de tres ensayos independientes. En el primer capítulo, analizo el efecto de una ley estadounidense que favorece la producción de biocombustibles sobre la salud fetal. Demuestro que la ley aumentó la producción de maíz, un cultivo con altos requerimientos de pesticidas, y el riesgo de enfermedades fetales asociadas con la exposición a pesticidas. En el segundo capítulo, estudio cómo la productividad agrícola puede mediar la relación entre *shocks* de lluvia y guerras en países africanos. Midiendo el nivel de lluvia sobre el territorio agrícola y durante la fase de crecimiento, demuestro que éste tiene una relación en forma de U-invertida con la producción agrícola, y una relación en forma de U con la incidencia de guerras civiles. En el tercer capítulo, estudio el efecto de diversas modalidades de venta de pruebas de arsénico para agua de pozo sobre la demanda por las mismas y, para familias que reciben noticias adversas, sobre su comportamiento para evitar el arsénico. Encuentro que la demanda aumenta cuando el precio a pagar depende de los resultados de la prueba, y que las redes sociales e información pública pueden promover medidas para evitar el agua contaminada.

Preface

This thesis brings together the results from three self-contained essays spanning the areas of development, environmental, and health economics. While the three essays are empirical in nature, they reflect dialog with economic theory, either testing it or leveraging on it for causal identification.

The introduction of the 2005 Renewable Fuel Standard in the United States increased the demand for corn ethanol and led to heterogeneous increases in the production of corn, a pesticide-intensive crop. In the first chapter, I estimate the effect of the resulting policy-induced increases in corn production on the incidence of two fetal conditions previously associated with exposure to corn pesticides and on the incidence of perinatal death in the U.S. Corn Belt. To better understand the heterogeneity in cross-county corn expansion, I present a model of corn production tailored to the U.S. context, where corn is regularly rotated with soy, and develop a new empirical measure for potential for corn expansion following the Renewable Fuel Standard. By combining the introduction of the Renewable Fuel Standard, county-level variation in potential for corn expansion, seasonal variation in corn pesticide applications during the growing year, and variation in fetal month of conception, I find that the policy-induced increase in corn production had a positive and significant effect on the incidence of abdominal wall defects, on being born small-for-gestational age, and on perinatal death.

News reports and policy makers frequently link African civil conflicts and wars to agricultural crises caused by droughts. However, the empirical evidence on the relationship between rainfall and civil conflict or war in Africa remains inconclusive. This evidence comes from studies linking the presence of civil war or conflict in a country to annual rainfall over a country's entire territory. In the second chapter, I argue that to better understand whether rainfall shocks affect the risk of civil war and conflict through agricultural productivity, it is useful to first examine the effect of rainfall shocks on agricultural output. Following recent work in agricultural economics, I relate the agricultural output of African countries to rainfall over agricultural land during growing seasons and allow for a hump-shaped effect of rainfall. This yields a robust, hump-shaped relationship between rainfall and agricultural output. Hence increases in rainfall raise agricultural output at low levels and decrease agricultural output at high levels. If rainfall

affects civil war and conflict through its effect on agricultural productivity, the effect of rainfall on the risk of civil war and conflict should therefore be U-shaped. I find this to be the case.

Lack of reliable information on environmental risk is often a key constraint limiting the extent of risk-avoiding behavior in developing countries. Such considerations are salient in Bangladesh, where naturally-occurring low-dose arsenic is frequently present in tubewell water, a primary source of drinking water for millions of households. Because arsenic contamination varies considerably across space, even within very narrow areas, the provision of information on arsenic contamination has been shown to be effective at allowing households relying on unsafe water to switch to safer and nearby sources. However, the safety status of millions of tubewells remains unknown, and there is no well-established market for tests. In the third chapter of my thesis, co-authored with Kazi Matin Ahmen, Alessandro Tarozzi, and Alexander van Geen, we describe the results from a randomized controlled trial in Sonargaon, Bangladesh, where tests were sold under different conditions. We find that sales were increased neither by "nudges" in the form of visible metal placards indicating safety status, nor by offers that attempted to promote the sharing of safe water through informal agreements, but contracts requiring payment only in case of "good news" more than doubled demand. Conditional on learning about the unsafe status of one's tubewell water, informal agreements, visible placards, and fees-for-good-news (but only at lower prices) nearly doubled the fraction of households which stopped drinking water from contaminated wells.

Contents

List of Figures

List of Tables

Chapter 1

BIOFUELING POOR FETAL HEALTH?

1.1 Introduction

Corn is the most widely produced crop in the United States and since 1971 it has also been the top pesticide-user (Fernandez-Cornejo et al., 2014). Nationally representative studies from the National Water-Quality Assessment Program have found atrazine, a pesticide used predominantly in corn production, to be the single most frequently detected pesticide in streams and wells (DeSimone et al., 2009; Gilliom et al., 2006). Epidemiological studies have documented an association between exposure to atrazine (and related pesticides in the triazine family) and two fetal conditions, gastroschisis, an abdominal wall defect (AWD) whereby the abdominal wall fails to close, and being born small-for-gestational age (SGA), a measure of intrauterine growth retardation (IUGR).[1] This has led to concerns that corn production, with its associated pesticide use, may have negative fetal health externalities (Duhigg, 2009).

Following the U.S. Congress' 2005 enactment of the so-called Renewable Fuel Standard (RFS), which aimed to reduce greenhouse gas emission, U.S. corn acreage grew rapidly to its largest level over the past 30 years. A main driver of

[1] Empirical studies of the effect of triazine pesticides on SGA and AWD include Chevrier et al. (2011); Mattix et al. (2007); Munger et al. (1997); Ochoa-Acuña et al. (2009); Shaw et al. (2014); Villanueva et al. (2005); Waller et al. (2010). Below, I discuss this evidence further.

this increase in corn acreage was the surge in the production of corn ethanol, a biofuel that is a cost effective way to comply with the RFS's mandate of blending gasoline with increasing amounts of biofuels (Schnepf and Yacobucci, 2013). I combine the introduction of the RFS biofuel mandate in 2005 with cross-county variation in a novel measure of potential for corn expansion to provide instrumental variables estimates of the effect of corn production —and its associated pesticide use— on the incidence of AWD and SGA. Because perinatal mortality has been linked to both AWD and SGA, and is of direct policy interest, I further provide estimates of the effect of corn production on perinatal death. Intuitively, my empirical analysis compares changes in fetal health outcomes, before and after the introduction of the RFS, across counties whose pre-shock potential for corn expansion led to differential expansions of corn production. Because triazine pesticides are pre-emergence pesticides, which are mainly applied during the planting season, I further exploit the variation in exposure to pre-emergence corn pesticides induced by the interaction of the month of conception and planting times. These estimates are then used to quantify the fetal health externalities associated with total policy-induced increases in corn acreage.

I find that RFS-induced increases in corn acreage significantly increased the risk of abdominal wall defects, being born small-for-gestational age, and perinatal (within one hour of birth) death, for births exposed to the planting season during critical gestational periods. My estimates imply that the RFS increased the risk of AWD by between 92 and 94 percent for births exposed at conception, the risk of SGA by between 1.4 and 2.2 percent for births exposed during their third gestational trimester, and the risk of perinatal death by between 20.5 and 26 percent for births exposed during their last two gestational trimesters. After the RFS, I estimate that the average county in the bottom quartile of the distribution of potential for corn expansion increased annual corn acreage by about 390, while the average county in the top quartile increased annual corn acreage by about 8,530. The instrumental variables estimates of the effect of corn production imply that the greater mean annual corn expansion of 8,140 acres experienced by the counties in the top quartile, relative to those in the bottom, translated into a larger increase in the risk of AWD of over 10 cases per 10,000 births, a larger increase in the risk SGA of over 34 cases per 10,000 births, and a larger increase in the risk perinatal

death of over 5 cases per 10,000 births, for births exposed to the planting season during critical gestational periods. These estimates control for unobserved factors that may translate into permanently higher risk of these conditions across counties, seasons, and years (among others). I show that my results are not driven by changes in seasonal labor supply during the planting season, permanent shocks (e.g., income shock) associated with higher corn prices and increased corn production, or maternal selection. Further, I provide suggestive evidence of fetal selection in utero, implying that my estimates are likely lower bounds of the true effects. This set of results suggest a causal link between increases in corn production, induced by the RFS, and negative fetal health externalities.

The costs associated with these negative health externalities are likely to be large. In the U.S., aggregate hospital charges associated with the medical care of newborns with intrauterine growth retardation have been estimated to be $1.07 billion in 1989 (Almond et al., 2005); the corresponding number for gastroschisis was about $220 million in 2003 (CDC, 2007). Moreover, a body of work in health economics has documented the adverse effects of IUGR on outcomes measured soon after birth e.g., APGAR score, infant mortality risk, and later in life e.g., academic achievement, educational attainment, height, IQ, wages, and welfare usage (Almond et al., 2005; Bharadwaj et al., 2010; Black et al., 2007; Behrman and Rosenzweig, 2004; Oreopoulos et al., 2008; Royer, 2009).[2] While the study of the potential effects of gastroschisis has received less attention, research in epidemiology indicates that it is associated with increased infant mortality risk, pediatric intestinal failure, and bowel transplantation (CDC, 2007; Chabra and Gleason, 2005; Shah et al., 2012).

In the U.S., a two-year corn-soy rotation scheme accounted for the majority of the acreage dedicated to both corn and soy between 1996 and 2010 (Wallander, 2013). The likely reason is that the rotation of corn and soy leads to a yield premium for corn, as soy fixes nitrogen in the soil, improves soil structure, and can reduce corn-pests populations (Padgitt, 1994; Roth, 1996). To better understand which places had a greater potential for expanding corn acreage following the RFS —necessary for my empirical work— I build a two-period model of corn and soy

[2]This literature uses within twin-pair variation in measures of IUGR to control for family and genetic differences.

production in the U.S. that identifies the county-specific characteristics that predict corn expansion following a corn price increase. Profit-maximizing farmers take into account the relative price of corn with respect to soy, the relative suitability of the soil, and complementarities from crop rotation when deciding whether to allocate land to corn or soy monoculture, or a corn-soy rotation scheme. In my model, complementarities from corn-soy rotation imply that in equilibrium, land with intermediate relative suitability levels rotate corn and soy while the land most suitable for corn or soy is used for monocultivation of the respective crop. The model predicts that counties with a larger potential for corn expansion are those with larger soy acreage before the corn price increase, and where land used for soy has a higher relative corn suitability.

This paper brings together a new empirical measure of potential for corn expansion with detailed and non-public birth and linked birth-infant death records from the United States National Center for Health Statistics for the 2001-2011 period. County-level potential for corn expansion following the passage of the RFS in 2005 is measured as the total pre-determined soy acreage weighted by the land's relative corn suitability. I construct this measure by combining detailed geographical data on agricultural land use from the National Agricultural Statistics Service's Crop Data Layer with data on land suitability from the Food and Agriculture Organization's Global Agro-Ecological Zones dataset. My instrumental variables strategy relies on the (testable) assumption that counties with greater potential for corn expansion are more likely to increase corn production after the introduction of the RFS. My estimates support this assumption and indicate that the RFS can account for an annual increase in corn operations of nearly 2 million acres in the U.S. Corn Belt, which is 49 percent of the total increase in mean corn acreage between the 2001-2005 and 2006-2011 periods.

At the time of the enactment of the RFS in 2005, the potential agricultural, environmental, and health externalities of the RFS's biofuels mandate were not well understood. Recently, policymakers —including Congress and the Environmental Protection Agency— have conducted studies aimed at evaluating these effects (Mittal, 2010; Sissine, 2010). These studies raise concerns over the increased pesticide use associated with the expansion of corn production, especially the effects on air quality due to pesticide volatilization and the effects on water quality due to

pesticide run-off leakage into streams and filtration into groundwater. However, the Environmental Protection Agency concludes that further research is needed to determine whether the increased pesticide use, associated with the expansion of corn production, has the potential to negatively affect human health (Sissine, 2010). To the best of my knowledge, this is the first paper to provide evidence on the negative health effects associated with the introduction of the RFS.

Previous epidemiological research has documented a more than two-fold increase in the national incidence of gastroschisis between 1995 and 2012 (Jones, 2016) and a 30 percent increase in the national incidence of SGA between 2002 and 2011 (Ewing et al., 2016).[3] These trends do not match the evolution of known and prominent risk factors for gastroschisis and SGA (young maternal age and maternal smoking, respectively) and have become a public health puzzle (Donahue et al., 2010; Jones, 2016). Between the 2001-2005 and 2006-2011 periods, the incidence of AWD among all births in the U.S. Corn Belt increased by 0.8 per 10,000 births, and the incidence of SGA increased by 46 cases per 10,000 births. My estimates of the effect of policy-induced increases in corn production on AWD and SGA suggest that the RFS can account for over 60 and 8 percent of the respective increases, thus shedding light on the causes of the upward trends in these fetal conditions.

My work relates to a literature in economics aimed at studying the effects of early-life exposure to pollution on fetal and infant health. The reason for the focus on the health of fetuses and infants is twofold. First, unlike adults, they have a short window of time in which they can be exposed to pollutants, which facilitates the measurement of lifetime exposure to pollution and exposure at different developmental stages. Second, early-life health status has been linked to future human capital accumulation (Currie, 2009, 2011; Currie et al., 2014). This literature has further focused on the effects of air pollution. In two seminal papers, Chay and Greenstone (2003a,b) use the variation in reductions of total suspended particulates across U.S. counties, induced by the passage of the Clean Air Act Amendments and the 1981-1982 recession, and find that these reductions signif-

[3]The study by Ewing et al. (2016) only analyzes the prevalence of SGA among single-birth, term (over 37 weeks of gestation) newborns. However, I find similar trends among all single-birth newborns (see Figure 1.2).

icantly decreased infant mortality. Other studies exploiting changes in exposure to various air pollutants (e.g., CO, NO_2, O_3, SO_2) find significant adverse effects on birth weight, fetal loss, infant mortality, prematurity, among other health outcomes (see Currie et al. 2014 for a review).[4]

The study of the effects of other forms of pollution on fetal and infant health has received less attention. A notable exception is the work by Currie et al. (2013), who find a positive and significant relationship between drinking water contamination during pregnancy and the incidence of low birth weight. More closely related to my work, Brainerd and Menon (2014) study the effect of fertilizer agrochemical concentration in the water on fetal health in India. By combining the time and geographical variation in state-level rice and wheat acreage with the seasonal variation in fertilizer applications, they find that exposure to fertilizers at the time of conception significantly increases infant and neonatal mortality. Their estimates rely on the assumption that changes in crop acreage are unrelated to changes in fetal health outcomes other than through changes in fertilizer use. I contribute to this literature by analyzing the effects of corn production —and its associated pre-emergence pesticide use— on fetal health. My empirical strategy exploits the combination of a policy change, the plausibly exogenous geographic variation in potential for corn expansion, and the seasonal variation in corn pesticide applications during the growing year to provide estimates of the effect of corn production on the incidence of AWD, SGA, and perinatal death.

This paper further relates to a literature in epidemiology assessing the effect of exposure to triazine pesticides on AWD and SGA.[5] Mattix et al. (2007) find a positive and significant correlation between atrazine concentration in surface water and the incidence of AWD by month of conception.[6] A study by Waller et al. (2010) finds a significantly higher incidence of gastroschisis among new-

[4]These studies include Agarwal et al. (2010); Arceo et al. (2016); Coneus and Spiess (2012); Currie et al. (2005, 2009); Currie and Schmieder (2009); Currie and Walker (2011); Jayachandran (2009); Knittel et al. (2016); Luechinger (2014); Sanders and Stoecker (2015); Sneeringer (2009).

[5]Corn pesticides have also been associated with limb birth defects, e.g., adactyly polydactyly, syndactyly, club foot (Winchester et al., 2009; Ochoa-Acuña and Carbajo, 2009). I do not expand on the evidence documenting this association, nor do I investigate this relationship empirically, as U.S. birth records stopped recording limb birth defects systematically after the year 2003.

[6]Mattix et al. (2007) study jointly the incidence of gastroschisis and omphalocele, both types of AWD.

borns to mothers living close to high atrazine concentration sites (within 50km). Shaw et al. (2014) relate the incidence of gastroschisis to maternal exposure to triazine pesticides, where mothers are considered exposed if a positive amount of pesticide was applied within 500m of her residence, between one month before conception and two months after. They find a significantly positive association between exposure to triazine pesticides and gastroschisis, but significance is lost when controlling for maternal characteristics.

Munger et al. (1997) find a significantly higher incidence of SGA in communities served by water sources with high triazine pesticide concentration. Villanueva et al. (2005) find that atrazine concentration in drinking water between the months of June and September —months of peak atrazine concentration in the water— is positively and significantly associated with the incidence of SGA for births whose third trimester of gestation coincided partly of fully with the June-September time period. Similarly, Ochoa-Acuña et al. (2009) find a significant and positive association between atrazine concentration in drinking water during the last trimester of gestation, and during the entire pregnancy, and the incidence of SGA. Chevrier et al. (2011) find a positive and significant association between the presence of atrazine in maternal urine and SGA and having a small head circumference, while a negative and significant association is found with birth weight.

These epidemiological studies analyze the geographic or seasonal correlation between fetal health outcomes and triazine pesticide exposure, and are therefore unable to account for unobserved differences that may translate into permanently greater risk of poor fetal health in certain localities or seasons. Unlike these studies, my empirical strategy allows me to exploit plausibly exogenous variation in corn production, while controlling for unobservables that translate into permanently greater risk of adverse fetal health across locations, seasons, and years. Further, it considers a larger population, both in terms of years studied and geographical space. My analysis thus provides supporting evidence of a causal association between corn production —and its associated pesticide use— and AWD, SGA, and perinatal mortality.

The remainder of the paper is structured as follows. Section 1.2 provides some information on the Renewable Fuel Standard's biofuel mandate and its effect on corn production. Section 1.3 provides some background knowledge on abdominal

wall defects and being born small-for-gestational age, and their relationship with seasonal exposure to pre-emergence corn pesticides. Section 1.4 presents a model of corn production and suggests a measure that predicts corn acreage expansion. Section 1.5 describes the data sources used in the paper. Section 1.6 delineates the empirical strategy, with results being presented in Section 1.7. A series of robustness checks are conducted in Section 1.8. Conclusions are presented in Section 1.9.

1.2 The Renewable Fuel Standard and its time series correspondence with corn production and fetal health

The Renewable Fuel Standard (RFS) is one of the most ambitious federal policies promoting the use of renewable fuels in U.S. history.[7] Its stated goals are to improve environmental quality, foster rural development, and secure energy independence.

Enacted through the Energy Policy Act of 2005, the RFS mandated the blending of a minimum volume of biofuels in the national gasoline supply, starting at 4 billion gallons (Bgal) in 2006 and rising to 7.5 Bgal by 2012. While the original version of the RFS placed little restrictions on the type of biofuel to be blended with gasoline, costs and technological considerations made corn ethanol the biofuel of choice for compliance (Schnepf and Yacobucci, 2013). The RFS was later expanded through the Energy Independence and Security Act of 2007, dividing renewable fuels into four nested categories and mandating specific volumetric requirements for each. These categories group biofuels based on their estimated reduction of lifecycle greenhouse gas emissions (GHG) —with respect to the gasoline or diesel fuel they replace— and on the biomass feedstock used in their production.[8] In particular, the law differentiates between conventional biofu-

[7]This section draws extensively from Schnepf and Yacobucci (2013) and Stock (2015). Duffield et al. (2008) and Kovarik (1998) provide comprehensive reviews of the U.S. ethanol policy history.

[8]The expanded RFS also places restrictions on the types of lands from which feedstock can come from, excluding explicitly virgin agricultural land cleared or cultivated after December 19,

els (e.g., corn ethanol), which provide an estimated reduction in GHG emissions of between 20 and 50 percent, and advanced biofuels, which provide an estimated reduction of over 50 percent. The latter category further encompasses the cellulosic and biomass-based diesel categories. Because of the nested classification of biofuels, those belonging to narrower categories (e.g., cellulosic biofuel) can be used to comply with requirements from broader categories (e.g., advanced or conventional biofuels) at specified rates that reflect each biofuel's energy content.[9]

For the year 2008, the RFS required 9 Bgal of conventional biofuel (e.g. corn ethanol) with no requirements for other renewable fuels with higher GHG emissions reductions. For 2009, the mandate required 11.1 Bgal of renewable fuels, with at least 0.6 Bgal coming from advanced biofuels —which excludes corn ethanol. The required amount of conventional biofuel mandated by the RFS increases every year until 2015, after which it is capped at 15 Bgal until 2022, while the required amount of advanced biofuels is mandated to increase progressively through 2022.

The Environmental Protection Agency (EPA) is in charge of regulating and overseeing the implementation of the RFS. To this purpose, EPA first calculates the mandated use of renewable fuels as a percentage of the projected total U.S. transportation fuel use, for each of the four renewable fuel categories. Fuel blenders and exporters are required to include renewable fuels in these same proportions in their yearly supply of fuel. The RFS also contemplates a waiver program, whereby EPA can waive in part or in total the yearly cellulosic and/or biomass-based diesel requirements. Since 2010, EPA has repeatedly waived the cellulosic mandate due to supply constraints; for example, in 2010 and 2011 there was no reported commercial production of cellulosic fuel.[10] In those same years, corn ethanol represented approximately 99 and 98 percent of all the renewable fuels blended with transportation fuels, respectively.[11]

The mandated large increase in biofuel usage has been followed by the largest increase in corn prices and land used for corn production in the past 30 years. As

2007 (when the Energy Independence and Security Act was enacted).

[9]The nested structure of the four biofuel categories contemplated in the RFS is depicted in appendix Figure A1.1.

[10]See `http:www.epa.gov/otaq/fuels/rfsdata/`

[11]See `http://www.eia.gov/dnav/pet/pet_pnp_inpt_dc_nus_mbbl_m.htm`

can be seen in Figure 1.1, the average corn price and total acres devoted to corn production in the U.S. increased by over 68 and 9 percent, respectively, in the five years following the passage the of the RFS compared to the preceding five years. The large increase in corn acreage raises concerns over potential adverse health effects associated with increased pesticide use. For instance, in their regulatory impact analysis of the RFS, the Environmental Protection Agency concludes that further research is needed to determine whether increased pesticide use, associated with increased corn production, has the potential to negatively affect human health (Sissine, 2010).

Figure 1.2 depicts the evolution of the relative incidence rates (with 2001-2005 averages as bases) of abdominal wall defects (AWD) and being born small-for-gestational age (SGA), alongside the time-series of total acres of corn planted in the U.S. Corn Belt. These two fetal conditions have been previously associated with exposure to triazine pesticides, prominently used in corn production.[12] As is clear from the figure, the sharp increase in corn production, following the passage of the RFS, is mirrored by a sharp increase in the incidence of AWD and a more subtle increase in the incidence of SGA. The upward trends in the incidence of AWD —gastroschisis in particular— and SGA have been documented elsewhere and do not correspond with the trends of known risk factors.[13] The mean incidence rates of these two conditions increased by about 19 and 5 percent, respectively, in the six years following the passage of the RFS compared to the preceding five years. The correspondence among the timing of the introduction of the RFS, the increase in corn acreage, and the increases in the incidence rates of AWD and SGA warrant further investigation of the relationship between policy-induced increases in corn acreage and adverse fetal health.

[12] See epidemiological literature reviewed above.
[13] See Ewing et al. (2016) for SGA and Jones (2016) for gastroschisis.

1.3 Abdominal wall defects and small-for-gestational age: background and seasonal patterns

In this section I provide some background knowledge on abdominal wall defects (AWD) and small-for-gestational age (SGA) and their relationship with seasonal exposure to pre-emergence corn pesticides.

1.3.1 Abdominal wall defects and small-for-gestational age

Gastroschisis is a congenital defect where the anterior abdominal wall fails to close (usually to the right of the umbilicus) and results in herniation of the abdominal content into the amniotic sac. Gastroschisis can be detected by week 10 of gestation, when regular embryo development results in closure of the abdominal wall (Chabra and Gleason, 2005; Kastenberg and Dutta, 2013). While no definitive genetic or environmental cause for gastroschisis has been identified, its increasing prevalence rates across different populations suggests a role for environmental insults (Chabra and Gleason, 2005). Lubinsky (2012) has put forth a hypothesis linking estrogen related thrombosis —clotting of the blood— to gastroschisis.[14] Because atrazine is a known endocrine and, in particular, estrogen disruptor, Lubinsky (2012, p. 810) concludes "A link with one such chemical, atrazine, is suggestive, but with hundreds of such estrogen "mimics" in the environment, effects may involve different substances either alone or combined."

Gastroschisis and omphalocele are the most common AWD and have been historically treated as a single entity. In the U.S. no birth record differentiated between the two anomalies until 2003, when a revision of the birth record standard started being implemented in some states.[15] In my data, because of the slow phase-in of the new standard, I cannot systematically distinguish between cases of gastroschisis and omphalocele, and thus study the joint incidence of these two types of AWD. Historically, the incidence of omphalocele was twice that of gastroschisis, however, while the former has remained stable over time, the latter has

[14]The hypothesis suggests that palmitic acid by-products of thrombosis would attach to (many) proteins, affecting cell signaling prior to closure of the abdominal wall.

[15]Importantly, these two conditions represent distinct pathologies, and while omphalocele is associated other structural anomalies, gastroschisis is not (Kastenberg and Dutta, 2013).

been trending upwards (Kirby et al., 2013; Marshall et al., 2015).[16]

Intrauterine growth retardation (IUGR) is a marker of an unhealthy pregnancy and, as reviewed above, has been associated with a host of negative outcomes including infant death. While there is still a vivid debate on how to best measure IUGR, SGA has been the measure most widely used by medical doctors and researchers and it is defined as being in the bottom 10 percent (or some other percentile) of a weight-by-gestational age, sex-specific, and usually time-invariant distribution.[17] While epidemiological studies have associated SGA risk with exposure to pre-emergence corn pesticides, the pathophysiological mechanism linking these variables is still not fully understood.

1.3.2 Seasonal exposure to pre-emergence corn pesticides and seasonal variation in fetal health

Pre-emergence corn pesticides are applied during the corn planting season (which in the U.S. varies between March and May), with exact planting times changing geographically to reflect different climate conditions and soil characteristics. This seasonality in pesticide application times has led to a corresponding seasonality in biomarkers of pesticide exposure in farmers. For instance, Bakke et al. (2009) find that Iowa farmers have significantly higher increases in urinary levels of corn pesticides 2,4-D, acetochlor, and atrazine during the planting season (relative to the pre-planting season) than non-farmer controls.[18] Further, another Iowa study by Curwin et al. (2007) finds that urinary levels of atrazine for fathers, mothers, and children are higher for farming households around the time of atrazine application with respect to non-farming households.

Nevertheless, little is known about the relevance of different pathways through which seasonal exposure may occur. Farmers and professional pesticide applicators can be affected by dermal exposure and inhalation while mixing and applying

[16] These estimates are based on medical charts, which allow for the independent assessment of the presence of gastroschisis and omphalocele.

[17] Zhang et al. (2010) review various measures used to gauge IUGR.

[18] This seasonality can be explained by the fact that several pesticides do not tend to bioaccumulate in the humans. For instance, toxicological studies have documented that atrazine has a half-life of one day, i.e., it takes the human body about one day to eliminate half of the atrazine it absorbed after exposure (Gilman et al., 1998).

pesticides. Further, there is evidence of pesticides being brought into their houses, possibly through traces in shoes and clothes (Curwin et al., 2005). More generally, people living in proximity to agricultural fields might experience seasonal exposure to pesticides through inhalation or water consumption. While there is little evidence of the former mechanism, likely due to the lack of data, atrazine concentration in U.S. streams has been found to be correlated with the timing and intensity of pesticide applications (Gilliom et al., 2006). Further, data from the Environmental Protection Agency's Atrazine Monitoring Program shows a corresponding seasonal pattern in atrazine concentration in finished drinking water (EPA, 2016). The degree of exposure is likely to be larger for the 15 percent of the U.S population that draws drinking water from domestic wells (Hutson, 2004), as these have been found to regularly be contaminated by pesticides (DeSimone et al., 2009) and are not subject to water quality regulation.

The closure of the abdominal wall by week 10 of gestation implies the existence of a critical period, around conception and up to week 10, in which environmental insults (e.g., pre-emergence corn pesticides) have the potential to increase AWD risk. SGA status, however can reflect nutritional deprivation, or maternal behaviors (e.g., smoking) throughout the pregnancy. Because no clear pathophysiological mechanism linking corn pesticides to SGA has been established, exposure to corn pesticides around conception and throughout the gestational period could potentially lead to growth retardation in utero. While maternal and paternal *seasonal* exposure to pesticides could potentially have a *permanent* (year-round) effect on fetal health, the epidemiological literature has found a corresponding seasonal pattern in the incidence of AWD (Mattix et al., 2007; Shaw et al., 2014) and SGA (Ochoa-Acuña et al., 2009; Villanueva et al., 2005) during critical gestational periods.[19] Moreover, the studies of Mattix et al. (2007) and Shaw et al. (2014) relate AWD to periconceptional exposure to triazine pesticides, which is consistent with the decreased importance of late exposure to pesticides due to the early closure of the abdominal wall.

[19]Paternal exposure to corn pesticides could potentially affect fetal health, independently of maternal exposure, through sperm quality. Swan et al. (2003) find that biomarkers of exposure to alachlor, atrazine, diazinon, and metolachlor pesticides (all used in corn production) are associated with poor semen quality in fertile men.

1.4 A model of corn and soy production

In this section I present a two-period model of corn and soy production that predicts a county's potential for corn expansion following a positive demand shock. Farmers are price-takers and they make decisions over plots of lands with heterogeneous yields for corn and soy. Guided by the model, I propose an empirical measure of potential for corn expansion following the introduction of the Renewable Fuel Standard (RFS).

1.4.1 Setup

A typical county, denoted by c, is inhabited by a farmer and is made out of a continuum of plots of land of mass L_c. Each plot j in the county is characterized by a potential yield for corn and soy production, $\pi_{j,c}^{corn}$ and $\pi_{j,c}^{soy}$, respectively. These are potential in the sense that they characterize yields under normal weather conditions and input utilization, and capture farmers' information sets at the time planting decisions are made. There are two periods ($t = 1, 2$), and each plot of land can be used for corn monoculture (corn in both periods), soy monoculture (soy in both periods), or rotation (either corn-soy or soy-corn). The prices of corn (p^{corn}) and soy (p^{soy}) are the same in both periods and are know to farmers.

The expected revenue from corn monoculture, in plot j in county c under normal weather and input utilization conditions, is $2p^{corn}\pi_{j,c}^{corn}$. Analogously, the expected revenue from soy monoculture in the same plot is $2p^{soy}\pi_{j,c}^{soy}$. The expected revenue from rotation is given by $(1+\delta)p^{corn}\pi_{j,c}^{corn} + p^{soy}\pi_{j,c}^{soy}$, $0 < \delta < 1$. Where δ captures the complementarity in yields that arises from rotating corn and soy, with soy providing a yield boost for corn (Padgitt, 1994; Roth, 1996).[20]

1.4.2 Crop choice decision

Denote by P the relative price of corn (p^{corn}/p^{soy}). Then, a farmer's decision on whether to engage in corn monoculture, soy monoculture, or rotation in plot

[20]There is no time discount factor. Further, the model abstracts from the heterogeneity in input costs for corn and soy production across different plots. Input costs are not modeled because of lack of data at the sub-county level. There would be no loss of generality in the results from this section if corn and soy input costs were the same at every plot.

j can be characterized as a partition of the space of relative suitability for soy (π^{soy}/π^{corn}). Plot j will be used for

- corn monoculture iff $\pi^{soy}_{j,c}/\pi^{corn}_{j,c} < P(1-\delta)$
- rotation iff $P(1-\delta) \leq \pi^{soy}_{j,c}/\pi^{corn}_{j,c} < P(1+\delta)$
- soy monoculture iff $\pi^{soy}_{j,c}/\pi^{corn}_{j,c} \geq P(1+\delta)$.

The conditions imply that plots with intermediate relative suitability levels rotate corn and soy, capitalizing on the yield complementarity brought about by rotation (δ), while the land most suitable for corn and soy is used for monoculti-vation of the respective crop.

Denote the cumulative distribution of π^{soy}/π^{corn} in county c by $F_c(x)L_c$, where $F_c(+\infty) = 1$. The amount of land that will be planted with corn in county c on average, over the two periods, can be written as

$$\begin{aligned} L^{corn}_c &= \{F_c(P(1-\delta)) + [F_c(P(1+\delta)) - F_c(P(1-\delta))]/2\}L_c, \\ &= [F_c(P(1-\delta)) + F_c(P(1+\delta))]L_c/2, \end{aligned} \tag{1.1}$$

which is the sum of the amount of land used for corn monoculture and half of the land that is used for rotation.

Assume further that $F_c(x)$ is differentiable, then the derivative of the amount of land used for corn in a typical year, with respect to an increase in the relative price of corn is given by

$$\frac{\partial L^{corn}_c}{\partial P} = [F'_c(P_0(1-\delta))(1-\delta) + F'_c(P_0(1+\delta))(1+\delta)]L_c/2, \tag{1.2}$$

where P_0 is the initial relative price of corn. The comparative statics of this expression hinge on three county-level characteristics. A county will see a greater expansion in its use of land for corn: (1) the larger is the density of plots rotating corn and soy at the switching threshold between corn monoculture and rotation (high $F'_c(P_0(1-\delta))$); (2) the larger is the density of plots dedicated to soy mono-culture at the switching threshold between rotation and soy monoculture (high $F'_c(P_0(1+\delta))$); and (3) the larger is the agricultural land (L_c). The intuition for the

first point is as follows, given a relative price increase, a county will increase corn production more —by moving away from rotation and into corn monoculture— the more plots used for rotation exist with the highest relative corn suitability among all plots not initially used for corn monocultivation. The intuition for the second point is analogous, a county will increase corn production more —by moving away from soy monoculture and into rotation— the more plots used for soy monocultivation exist with the highest relative corn suitability among all plots initially used for soy monocultvation.[21]

Figures 1.3a-d illustrate these points. Panels a and b depict two counties, A and B, respectively, of equal size that vary in their distribution of relative soy suitability. Plots in county A are on average relatively less soy suitable, i.e., more corn suitable, than plots in county B. Additionally, plots in county B are, on average, more heterogeneous in their relative soy suitability than those in A. At the initial relative corn price of $P_0 = 1$, the amount of land used for corn monoculture, soy monoculture and rotation, is depicted as the areas colored in yellow, green, and brown, respectively. Panels c and d depict in red the additional land that moves away from rotation into corn monoculture and away from soy monoculture into rotation as a result of an increase in the relative price of corn. The expansion in the area devoted to corn is larger in county B because it initially had more plots of land used for soy, both in the form of rotation and monocultivation, at the switching thresholds for corn monocultivation and rotation, respectively. Intuitively, a corn price increase is more conducive to corn expansion in county B because it initially had more physical space for corn to expand into, the land used for soy, and because this land was adequately suitable for corn. Lastly, the third point above indicates that, everything else equal, larger counties will have a larger capacity for corn expansion.

[21]The model formalizes the predictions of several agricultural economists, around the passage of the RFS, indicating that increased acreage for corn would come from land previously used for soy (Babcock and Hennessy, 2006; Hart, 2006; Westcott, 2007). The model adds to these predictions that a relative corn price increase is more likely to be conducive to the expansion of corn, into land used for soy, the larger is the relative corn suitability of that land.

1.4.3 Empirical measure of potential for corn expansion

The five years following the introduction of the RFS saw a 59 percent increase in the mean relative price of corn with respect to soy, compared to the preceding five years (USDA, 2016c). Constructing an empirical measure of potential for corn expansion following the introduction of the RFS is useful because it allows me to identify the counties that were in the capacity to react to the policy-induced increase in the relative price of corn by increasing corn production.

I lack the data to implement the exact formula in equation 1.2 for the increase in corn acreage following a positive corn price shock associated with the introduction of the RFS. As a result, I will use a proxy for the county-level increase in corn acreage following the introduction of the RFS that captures two main features of the formula in equation 1.2 and can be calculated with the available data. This proxy is pre-shock soy acreage weighted by the land's relative corn suitability, and is given by

$$PCE_c = \sum_{j \in J_c} \overline{soy}_{j,c}^{pre} \left(\frac{\pi_{j,c}^{corn}}{\pi_{j,c}^{soy}} \right), \tag{1.3}$$

where J_c is the set of plots of land j in county c, $\overline{soy}_{j,c}^{pre}$ is the share of years that plot j was used for soy between 2001 and 2005 (before the introduction of the RFS), and $\pi_{j,c}^{corn}$ and $\pi_{j,c}^{soy}$ are plot j's potential yields for corn and soy production, respectively. Hence, in line with the model, my empirical measure of potential for corn expansion is increasing in initial amount of land used for soy and in the relative corn suitability of that land.

1.5 Data

This section describes the main variables and sources used in this paper.

Fetal health data come from the U.S. National Center for Health Statistics' natality, linked birth-infant death, and fetal death records. I accessed a non-public version of these data that allows me to identify a mother's county of residence. These records provide rich information on fetal characteristics at birth or stillbirth as well as on maternal characteristics. Throughout this paper, I focus on health

outcomes from single births.

The outcomes studied in this paper are: abdominal wall defects (AWD), being born small-for-gestational age (SGA), and perinatal death. Natality records include an indicator variable for the presence of gastroschisis or omphalocele, which are types of abdominal wall defects, but do not allow me to systematically differentiate between the two conditions. A live birth is coded as having an AWD if its birth record indicates the presence of either condition. These records also have data on birth weight, length of the gestational period (in weeks), and sex that I use to construct an indicator variable for SGA. A live birth is coded as being born SGA if it lies in the bottom decile of a gender-by-gestational age, time-invariant weight distribution. I use weight distributions for single births in the U.S. compiled by Olsen et al. (2010) and defined for males and females and for every gestation length (in weeks) between 22 and 44 weeks. The SGA indicator variable is not defined for gestation lengths of less than 22 or more than 44 weeks.

I also use linked birth-infant death records to construct an indicator variable of perinatal death, i.e., whether a newborn died within an hour of birth. The linked birth-infant death records used are linked by year, meaning that it is only possible to tell if live birth from a given year died in that same year. These records are suitable for analyzing deaths that happen shortly after birth, but unsuitable for analyzing deaths that happen later in life, e.g., infant mortality (within one year of birth).

The fetal death record provides a registry of occurrences of fetal loss for fetuses that reach the age of 20 weeks. Abortions are not included. I combine the natality and fetal death records to construct an indicator variable for fetal death. I use this indicator variable to analyze fetal selection in utero. Unlike the birth and linked birth-infant death records, there is evidence of substantial underreporting (MacDorman et al., 2012) in the fetal death registry. I expand on this issue in the robustness checks section.

Maternal characteristics in these records include: marital status, age, educational attainment, ethnicity, tobacco use, number of previous live births, and the presence of diabetes, chronic hypertension, pregnancy-associated hypertension, eclampsia, among others.

Potential for corn expansion is measured by combining detailed geographical data on agricultural land use and land suitability for corn and soy.

The data on agricultural land use come from the United States Department of Agriculture, National Agricultural Statistics Service's (NASS) Cropland Data Layer (CDL). NASS uses satellite imagery coming from the Deimos-1, UK-DMC 2, Landsat TM/EMT+, and the Indian Remote Sensing Advanced Wide Field sensors, along with in-house and commercial classification software, to produce high-resolution, geo-referenced, yearly data on agricultural land use for various crops in the U.S. (USDA, 2016b). The resolution of these data varies over time, CDL data for the periods 2001-2005 and 2010-2011 has a ground resolution of 30m by 30m, while data for the 2006-2009 period comes at a 56m by 56m ground resolution. These satellite data have a classification accuracy of about 95 percent for the crops considered in this paper (corn and soy). Classification accuracy of the CDL is assessed using ground truth data coming from NASS's June Acreage Survey, which includes data from over 10 thousand area segments, across the U.S.; farmers are asked to report the acreage, by crop, that has been planted, that they intend to plant, and that they intend to harvest.[22] Land use data prior to 2005 was only systematically collected for all counties in 6 states —Illinois, Indiana, Iowa, Mississippi, Nebraska, and North Dakota— and a subset of counties in eastern Arkansas and Missouri. This results in a sample of 588 counties in the U.S. Corn Belt.

Data on potential corn and soy yields come from the Food and Agriculture Organization's (FAO) Global Agro-Ecological Zones dataset (Fischer et al., 2012). These potential yields measure the agronomically possible upper limit production for a given crop and are calculated using crop growth models that incorporate data on historical weather conditions (e.g., precipitation, temperature, sunshine), soil characteristics (e.g., soil nutrient availability, soil nutrient retention capacity), irrigation availability, and input levels or management practices (e.g., fertilizer use, pesticide use, mechanization). In particular, I use an index for potential corn and soy yields, the so-called crop suitability index, that consider rain-fed agriculture practices under high input use, i.e., full mechanization, use of high-yielding

[22]For more information on the CDL data and its accuracy refer to USDA (2016a).

seeds, and optimal fertilizer and pesticide use.[23] The crop suitability index, both for corn and soy, is available world-wide at a resolution of 5 arc min by 5 arc min (approximately 9.3km by 9.3km at the Equator).

To construct my measure of potential for corn expansion I first calculate the average number of years that every unit of land in the CDL data was used for soy between 2001 and 2005 ($\overline{soy}_{i,c}^{pre}$). Because the resolutions of the land use and land suitability datasets are different, I work with the coarser resolution of the FAO-GAEZ dataset. Then, for any given county, I select all the cells from the coarser data that lie partially or totally within the county's borders. For each one of these cells I calculate the number of acres, within the county, that were used for soy in the typical year between 2001 and 2005. Average soy acreage at the cell-by-county level is depicted in Figure 1.4a. I then aggregate to the county level by adding cell-level soy acreage from all selected cells, weighting these values by each cell's relative corn suitability. Corn and soy suitability indices for my Corn Belt sample are illustrated in Figures 1.4b-c. My county level measure of potential for corn expansion (PCE), depicted in Figure 1.4d.

Corn acreage data is taken from USDA-NASS's Agricultural Survey estimates (USDA, 2016c), whenever possible. While satellite corn acreage data is available for all years between 2001 and 2011 in my Corn Belt sample, survey estimates incorporate both satellite and ground-based data, and are considered by the USDA-NASS to be the most reliable source of crop acreage data.

Usual planting season as well as harvest season times for corn come from USDA-NASS's Agricultural Handbook (USDA, 2010). I rely on the dates of the start and end of the "most active" period within each season. Because pre-emergence corn pesticides are usually applied between one month before planting and at planting (Hartzler and Owen, 2005), I adjust planting season start and end dates 15 days backwards. Throughout, I refer to this adjusted period as the planting season. Planting and harvesting times vary by state, and I consider any given calendar month to be part of the planting or harvest season in a given state, as long as it includes at least 10 days from the usual planting or harvest season, respectively. These data are used to determine fetal exposure to the planting and harvest seasons. I do this by using data on the month of conception from birth records,

[23] The index further takes weather conditions from 1961 to 1990 as the baseline.

data on planting season months in the mother's state of residence, and assuming a gestational length of 37 weeks.[24] A fetus is defined as exposed to the planting season at a given gestational period if at least one month of that gestational period overlaps with the planting season in the mother's state of residence.

Table 1.1, panel A, provides some basic descriptive statistics for my sample of counties in U.S. Corn Belt for the 2001-2005 and 2006-2011 periods. Panel B further provides summary statistics on fetal conditions and maternal characteristics by the time of fetal/maternal exposure to the planting season.

1.6 Empirical strategy

This section describes the econometric methods used to estimate the effect of corn production —and its associated pre-emergence pesticide use— on the incidence of AWD, SGA, and perinatal mortality. The analysis exploits two main sources of variation: (1) the heterogeneity in policy-induced changes in corn production across counties with differential potential for corn expansion, and (2) the timing of the application of corn pesticides.

The structural equation to be estimated relates a fetal health outcome Y (i.e., the risk of AWD, SGA, or perinatal death) to the acreage of corn planted in the maternal county of residence, the sex of the newborn and a set of maternal characteristics (X), as well as month of conception-, county-, and state-by-year-fixed effects (δ_m, δ_c, and $\delta_{s,t}$, respectively), county-specific linear trends (ρ_c), and an error term ϵ.[25] The equation is given by

$$Y_{i,c,s,m,t} = \beta corn_{c,s,t} + \theta X_{i,c,t} + \delta_m + \delta_c + \delta_{s,t} + \rho_c t + \epsilon_{i,c,s,m,t}, \qquad (1.4)$$

[24] The month of conception is approximated using birth records' data on maternal last month of normal menstruation, whenever available. For the cases where this this data is missing, I use birth records' clinical estimate of the length of gestation (in weeks) and data on the month of birth to infer the month of conception.

[25] Maternal characteristics included in equation 1.4 are indicator variables for: marital status, age range ([11,15), [16, 20), [21, 25), ..., [41, 45), [46, 53]), educational attainment, ethnicity, tobacco use, number of previous live births, diabetes, chronic hypertension, pregnancy-associated hypertension, and eclampsia. I also include indicator variables for missing values for all controls.

where subscripts i, c, s, m, and t index birth, county, state, month of conception, and year, respectively. The coefficient of interest is β, which captures the effect of corn production and its associated pesticide use on fetal health. Because pre-emergence pesticides are applied during the planting season, the relationship is assumed to hold for births that were exposed to the planting season during a critical gestational period. In the case of AWD, that critical period is the conception month. Because the relevant timing of exposure is less clear for SGA and perinatal death, I explore the effect of corn production on the risk of these conditions for births exposed to the planting season at conception and during the three gestational trimesters, separately.[26]

OLS estimation of the effect of corn production on fetal health using equation 1.4 controls for unobserved factors that may translate into permanently higher risk of poor fetal health across months of conceptions and counties, as well as for time shocks affecting all counties within a state. It further controls for unobserved county-level characteristics that might be trending over time and that could be correlated with fetal health. Nevertheless, this approach would be biased if there exist unobserved shocks that affect both changes in corn production and fetal health. For instance, farmers might choose to expand corn production as a response to worsening out-of-farm labor opportunities. Indeed, in the appendix Table A1.1 I show that changes in corn production are negatively related to changes in total employment and employment in the services sector. If these unfavorable conditions reduce likelihood of pre-natal care or worsen maternal nutritional intake, they would likely lead to an upward bias of β. If, however, these unfavorable conditions reduce disposable income and the likelihood of maternal drinking or smoking, they could lead to a downward bias of β. Further, OLS estimates of β would be biased if corn acreage data were measured with error.[27]

[26] A fetus is defined as exposed to the planting season at a given gestational period if at least one month of that gestational period overlaps with the planting season in the mother's state of residence.

[27] Measurement error in county-level corn acreage estimates by the USDA-NASS is likely to be substantial due to the following reasons. First, until 2010 these estimates were based on survey data representative only at the state level (and complimented with administrative data). Second, county-level corn acreage estimates must add to previously estimated corn acreage data at the agricultural statistical district and state levels, facilitating the propagation of measurement error down to lower spatial scales. Third, human error is a likely source of measurement error, as estimates for corn acreage are determined by the Agricultural Statistics Board and "To date, no

Because of the possibility of unobserved shocks that affect both changes in corn production and fetal health and of measurement error in the corn acreage data, I will also consider instrumental-variables estimation of equation 1.4. The instrument will be the introduction of the RFS in 2005 interacted with the new measure of potential for corn expansion following the introduction of the RFS. The introduction of the RFS in 2005 generated a demand shock that was accompanied by a sharp increase in the price of corn. This led to an increase in corn production that was larger for counties with greater pre-shock potential for corn expansion.

The first stage equation of the relationship between corn production and the RFS is given by

$$Corn_{c,s,t} = \lambda(PCE_c \cdot Post_{2005t}) + \phi \bar{X}_{i,c,t} + \bar{\delta}_m + \delta_c + \delta_{s,t} + \rho_c t + \eta_{c,s,t}, \quad (1.5)$$

where an upper bar indicates county-year average (e.g., $\bar{\delta}_m$ is the share of births conceived in month m) and η is an error term capturing the unobserved determinants of corn production.[28] OLS estimation of λ provides the difference-in-differences estimate of the effect of the RFS on corn production. Here, county-level potential for corn expansion measures the intensity with which counties could respond to the policy-induced demand shock by expanding their corn operations. Consistent estimation of λ requires counties with differential potential for corn expansion to be on parallel corn acreage trends, conditional on all controls. I note that the inclusion of county-specific linear trends in equation 1.5 lends support to a causal interpretation of my estimates. My analysis provides evidence of a positive and significant effect of the RFS on corn acreage.

Hence, the interaction of a post 2005 dummy variable ($Post_{2005}$) with the measure of potential for corn expansion (PCE) could serve as an instrument if counties' predetermined levels of soy acreage weighted by the land's relative corn suitability are unrelated to changes in fetal health around 2005 other than through their effect on changes in corn acreage (the so-called exclusion restric-

model-based estimate has served as the published official statistic" (Cruze et al., 2016).

[28] The county-year averages are over the same births included in the second stage.

tion).[29] Further, if measurement error in the instrument is orthogonal to that in the corn acreage data, IV estimation of β would yield consistent estimates where OLS suffered from attenuation bias.

This approach should work well for documenting any effects of the combination of pre-emergence pesticides used in corn production on fetal health. However, the approach could not be used to study the specific effect of atrazine or other specific triazine pesticide on fetal health once such data become available. This is because atrazine is regularly combined (and detected in water) with other pre-emergence pesticides, and the policy-induced increases in corn production I exploit in my analysis are unlikely to generate idiosyncratic increases in specific pesticides.

1.7 Results

I begin by showing that the introduction of the RFS was accompanied by increases in corn production in the U.S. Corn Belt, and that these increases were larger in counties with greater potentials for corn expansion. Table 1.2, column 1, estimates the mean county-level increase in corn production between the 2001-2005 and 2006-2011 periods. The average county in the sample increased its annual area planted with corn by over 8,480 acres, which represents about a 10 percent increase with respect to the pre-RFS period.

Table 1.2, columns 2-4, examines the cross-county heterogeneity in the effect of the RFS on corn acreage. Column 2 shows that counties with higher predetermined levels of potential for corn expansion (PCE) increased their corn acreage significantly more between the 2001-2005 and 2006-2011 periods than counties with lower potential for corn expansion. Column 3 includes state-year fixed effects, so as to control for corn production shocks affecting all counties in a given state and year, and the magnitude and significance of the effect of the RFS remains nearly unchanged. To address possible concerns over heterogeneous trends

[29] The Renewable Fuel Standard could have had a direct impact on fetal health through its effect on seasonal labour supply or by inducing a permanent income shock. These issues are considered in Section 1.8. Further, in the appendix Table A1.1 I also show that the measure of potential for corn expansion following the introduction of the RFS is unrelated to changes in total employment and employment in the services sector, which are correlated with changes in corn production.

in corn acreage, column 4 includes county-specific linear trends. I note that this specification has the potential to underestimate the effect of the RFS, if counties do not fully adjust to a higher level of corn production immediately after its introduction. In this specification, I still find a positive and significant effect of the RFS on corn production —albeit smaller than those from models that do not control for linear trends. According to the estimate in column 4, the effect of increasing a county's potential for corn expansion from that of the average county in the bottom quartile of the distribution of potential for corn expansion to that of the average county in the top quartile is to increase annual corn acreage by 8,140. This estimate further implies that the RFS can account for an annual increase in corn operations of nearly 2 million acres, which explains 48.9 percent of the total increase in mean corn acreage between the 2001-2005 and 2006-2011 periods, in the U.S. Corn Belt.

Table 1.3, columns 1 and 2, provides IV and OLS estimates based on equation 1.4 of the effect of corn production on the risk of AWD. As indicated above, this equation is assumed to hold for births that were exposed to the planting season at conception. Because closure of the abdominal wall happens by week 10 in healthy fetuses, increased corn production and pesticide use later in the gestational period (e.g., second or third trimester), should be unrelated to AWD risk if seasonal maternal exposure has transitory effects on fetal health. To test for this, columns 3-4, 5-6, and 7-8, further provide IV and OLS estimates of equation 1.4 for birth that were exposed to the planting season in their first, second, and third gestational trimester, respectively.

The IV estimate of the effect of corn production on AWD risk for births exposed to the planting season during conception in column 1 is positive and significant. The estimate implies that RFS-induced increases in corn production increased the risk of AWD by 4.21 cases per 10,000 births, for births conceived during the planting season.[30] This represents a 93.8 percent increase in the inci-

[30]Estimates of the effect of RFS-induced increases in corn production on the incidence of fetal health outcomes per 10,000 births are calculated as follows:
$\hat{\beta}_{IV}\hat{\lambda}_{OLS} \cdot [(\sum_{t=2006}^{2011} \sum_{c \in CB} N(\tau)_{c,t} \cdot PCE_c)/(\sum_{t=2006}^{2011} \sum_{c \in CB} N(\tau)_{c,t})] \cdot 10000$,
where c and t index counties in my Corn Belt (CB) sample and years, respectively, and $N(\tau)_{c,t}$ refers to the number of births in county c and year t that were exposed to the planting season during the critical gestational period τ.

dence of AWD with respect to that of the 2001-2005 period. The estimate further implies that increasing a county's annual land planted with corn by 8,140 acres increases the incidence of AWD by 10.47 cases per 10,000 births, for births exposed to the planting season at conception. Where the chosen corn expansion is the estimated effect of increasing a county's potential for corn expansion form that of the mean in the bottom quartile to that of the mean in the top quartile of the respective distribution. The bottom part of the table reports first-stage statistics for each IV specification. It can be seen that the F-statistic associated with the hypothesis that the RFS is unrelated to corn acreage has a value of over 60, which alleviates concerns over weak instruments. Nevertheless, following Stock and Andrews (2005) and Chernozhukov and Hansen (2008), throughout the paper I report p-values robust to weak inference and standard errors clustered at the county level.

Columns 3, 5, and 7 present the IV estimates of the effect of corn production on AWD risk for birth that spent at least one month of their first, second, and third trimester in the planting season, respectively. The no-effect estimates found for births exposed during the second and third trimester are consistent with seasonal maternal exposure having transitory (immediate) effects on AWD risk. The fact that the estimate for births exposed to the planting season during the first trimester is smaller than that estimated for births exposed during conception is not surprising, as the first trimester includes the conception period, but also other weeks later in the gestational period.

Column 2, 4, 6, and 8 presents OLS estimates of the effect of corn production on AWD risk for births exposed to the planting season at conception and during their first, second, and third gestational periods, respectively. These point estimate are found to be insignificant. The discrepancy between the IV and OLS estimates, in columns 1 and 2, could be explained by the existence of unobserved determinants of AWD that are correlated with changes in corn production, but not potential for corn expansion. Moreover, the fact that both estimates are positive but that the OLS estimate is smaller could also be explained by attenuation bias in the latter due to classical measurement error in the corn acreage data.

Table 1.4 presents the IV and OLS estimates of the effect of corn production on SGA risk for births exposed to the planting season at conception, and during

their first, second, and third gestational trimesters. I note that while most of the evidence linking pre-emergence corn pesticides to SGA has found significant effects for exposures during the third trimester of gestation, the lack of a known pathophysiological mechanism means that there could potentially be significant effects arising from exposures at other gestational times.

IV estimates of the effect of corn production on the risk of SGA for births exposed to the planting season at conception, during their first, and second gestational trimesters are of a no-effect type (columns 1, 3, and 5). However, consistent with previous findings (Villanueva et al., 2005; Ochoa-Acuña et al., 2009), column 7 shows that there is a positive and significant effect of corn production on SGA risk for births that spent at least one month of their third gestational trimester in the planting season. The estimate implies that RFS-induced increases in corn production increased the incidence of SGA by 14.18 cases per 10,000 births, for births exposed to the planting season during their third gestational trimester. This represents a 1.4 percent increase in the incidence of SGA, with respect to that of the 2001-2005 period. The estimate further implies that increasing a county's annual land planted with corn by 8,140 acres increases the incidence of SGA by 34.22 cases per 10,000 births, for births exposed to the planting season during their third gestational trimester. Just like in the previous table, all OLS estimates are insignificant.

Table 1.5 presents the IV and OLS estimates of the effect of corn production on perinatal risk, for births exposed to the planting season at conception and during their first, second, and third trimester of gestation, respectively. The instrumental variables estimates in columns 1 and 3 indicate no effect of corn production on perinatal death for births exposed to the planting season at conception or during their first gestational trimester. However, there is a positive and significant effect for births exposed during their second and third gestational trimester. The estimates imply that policy-induced increases in corn production increased the risk of perinatal death for births exposed to the planting season during their second or third gestational trimesters by 2.26 cases per 10,000 births. This represents a 26 percent increase in the incidence of perinatal death with respect to the 2001-2005 period.[31] The estimate further implies that increasing a county's annual land

[31]These calculations are based on IV estimates of β for births exposed to the planting season

planted with corn by 8,140 acres increases the incidence of perinatal death by 5.5 cases per 10,000 births, for births exposed to the planting season during their second or third gestational trimesters. The corresponding OLS estimates are smaller and significant, and the OLS estimate for exposure during the third trimester has the opposite sign. The OLS estimate for exposure at conception is of the no-effect type and the OLS estimate for exposure during the first trimester is positive, small, and significant.

All in all, Tables 1.3-1.5 suggest a causal link between policy-induced increases in corn production and the risk of AWD, SGA, and perinatal death.[32] These effects can help explain some puzzling upward trends in the incidence of AWD and SGA. My instrumental variables estimates of β suggest that the RFS can account for about 61.8 percent of the increase in the incidence rate of AWD between the 2001-2005 and 2006-2011 periods, and for about 8.9 percent of the increase in the incidence rate of SGA.

To examine whether measurement error in the corn acreage data is driving the difference between IV and OLS estimates, I consider an alternative OLS estimation strategy using county-level average corn acreage, before and after the introduction of the RFS. Averaging can lead to more precise estimates of corn acreage if yearly changes in observed corn acreage data, within the 2001-2005 and 2006-2011 periods, are mainly driven by measurement error.[33] I then use this measure to provide OLS estimates the effect of changes in average corn acreage, around the introduction of the RFS, on the risk of AWD, SGA, and perinatal death. Table 1.6, columns 1 and 3, shows that increases in average corn acreage have a positive and significant effect on the risk od AWD and perinatal death, for births exposed to the planting season during critical gestational periods. Column 2 shows that the estimated effect of increased average corn acreage on the incidence of SGA is positive but insignificant. Thus, these results indicate that measurement error

during their second or third gestational periods. The estimated effect of an additional 10,000 acres of corn is an increase in the risk of perinatal death of 6.76 cases per 10,000 births (with a standard error of 2.55). This estimate can be directly compared with those of Table 1.5, for births exposed to the planting season during their second or third gestational trimesters, separately.

[32]Reduced form estimates of the relationship between the RFS and these health outcomes can be found in appendix Table A1.2.

[33]The fact that both the price of corn and the relative price of corn with respect to soy remained stable in the 2001-2005 and 2006-2011 periods renders support to this assumption.

in yearly, county-level corn acreage data is a likely source of bias in the previous OLS estimates, which would in turn explain the discrepancy between IV and OLS estimates.

1.8 Robusteness checks

Causal interpretation of the IV estimates requires that the RFS have no effect on fetal health outcomes other than through increased corn production and pesticide use. In this section, I explore possible threats to identification.

1.8.1 Increased seasonal labor supply

Policy-induced increases in corn production could have led to seasonal increases in maternal or paternal labor supply, precisely in the counties with larger potential for corn expansion and in the planting season. These increases in labor supply could transitorily increase household income and lead to a downward bias of the estimate of β, if these positive income shocks increase maternal nutritional intake or the likelihood or prenatal care —which are likely to improve fetal health. If positive transitory income shocks, however, increase maternal risky behaviors like smoking and drinking, the converse would happen. Further, increased maternal (strenuous) work in the fields could potentially affect fetal health negatively, leading to upward bias.

In the sample period, mean agricultural employment during the harvest season is 17 percent larger than during the planting season (BLS, 2016).[34] If the IV estimates were driven by increased seasonal employment, one would find that the RFS increased the risk of AWD, SGA, and perinatal death for births that spent critical gestational times in the harvest season. Table 1.7, columns 1-3, presents reduced form estimates of the effect of the RFS on AWD, SGA, and perinatal death, for births exposed to the harvest season in the respective critical gestational periods. I find no evidence of seasonal employment driving poor fetal health outcomes.

[34] In appendix Table A1.3 I find suggestive evidence showing that the RFS increased agricultural employment both in the planting and harvest seasons, and that the increase is larger in the latter.

1.8.2 The Renewable Fuel Standard as a permanent income shock

The most salient threat to the exclusion restriction is that the RFS could have had a permanent and positive income effect in counties with higher potential for corn expansion, through increased crop value, and that this shock is (at least partly) driving my results. If a permanent income shock were to lead to permanent changes in maternal behavior or characteristics with direct impacts on fetal health, I should find that policy-induced increases in corn production increase the risk of poor fetal health year-round, and not just for births exposed to the planting season during critical gestational periods. The case is made most clearly for AWD, for which there is a well established critical period (conception) in which pre-emergence pesticides have the potential affect the risk of this condition. As is shown in Table 1.3, corn production has no effect on AWD incidence for births exposed to the planting season during their second and third gestational trimesters.

To further alleviate concerns, I implement an alternative IV strategy that is robust to the presence of policy-induced permanent health shocks, and that exploits the differential effect of corn production on fetal health across seasons of high and low (or null) pesticide application. In particular, I compare the changes (before and after the introduction of the RFS) in the differences of the incidence of fetal health outcomes between births that were and were not exposed during a critical gestational period to the planting season, across counties with different policy-induced changes in corn production. I note that unexposed births include both those whose gestational period never overlapped with the planting season and those that were exposed to the planting season but during non-critical gestational periods (e.g., third trimester for AWD). The structural equation to be estimated is given by

$$Y_{i,c,s,m,t} = \beta_2 PS(\tau)_{i,s,m} \cdot corn_{c,s,t} + \gamma X_{i,c,t} + \delta_{m,c} + \delta_{m,s,t} + \delta_{c,t} + \rho_{c,m} t + \nu_{i,c,s,m,t}, \tag{1.6}$$

where $PS(\tau)$ is a binary variable indicating whether birth i, conceived in month m, was exposed during the critical gestational period τ to the planting season of state s. $\delta_{m,c}$, $\delta_{m,s,t}$, and $\delta_{c,t}$ are month of conception-by-county, month

of conception-by-state-by-year, and county-year fixed effects, respectively; $\rho_{c,m}$ are county-by-month of conception linear trends and ν is an error term. This models nearly saturates on the fixed effects from the previous strategy. The coefficient of interest here is β_2, which identifies the heterogeneity in the effect of corn production across seasons, depending on whether the season is characterized by intensive use of pre-emergence corn pesticide. Positive estimates of β_2 would indicate that RFS-induced increases in corn production increase the risk of fetal ailments for births exposed during critical gestational times to the planting season, above and beyond any possible effect it might have had on unexposed births. $\delta_{m,c}$ controls for unobserved characteristics that translate into permanently increased risk of fetal health across different times of the year in every county (e.g., weather, seasonal flu), $\delta_{m,s,t}$ controls for time shocks (at the month of conception-by-year level) common to all counties in a given state, and $\delta_{c,t}$ control for time shocks at the county level (e.g. income shock), and $t_{c,m}$ control for trending unobserved determinants of fetal health common to all births conceived in a given month and county. In this alternative approach, the triple interaction of the indicator variable for exposure to the planting season, the measure of potential for corn expansion, and an indicator variable for the introduction of the RFS ($PS(\tau) \cdot PCE \cdot Post_{2005}$) serves as a candidate instrument for the interaction of the indicator variable for exposure to the planting season and corn acreage ($PS(\tau) \cdot corn$).[35]

IV estimates of β_2 from equation 1.6 for AWD, SGA, and perinatal death risk are presented in Table 1.8, columns 1, 3, and 5. The critical gestational periods for exposure for AWD, SGA, and perinatal death are taken to be conception, the third trimester, and the last two gestational trimesters, respectively. All three estimates are positive and significant, and they imply that RFS-induced increases in corn production significantly increased the risk of AWD, SGA, and perinatal death by 92.1, 2.2, and 20.5 percentage points, respectively, for births exposed during critical gestational periods to the planting season. These results net out any year-round effect that the RFS might have had on fetal health, and isolates the effect of increased corn acreage for births exposed to times of high pesticide use. However, the agreement in the magnitudes of the estimated β and β_2 for AWD and perinatal

[35]The first stage of this IV/2SLS strategy is estimated at the individual-level. A county-level first stage, as in equation 1.5 is not feasible, as $PS(\tau)$ varies within county.

death —and to a lesser extent SGA— suggest a null or negligible permanent effect of the RFS on these fetal conditions. These results provide further evidence of a causal effect of policy-induced increases in corn production and fetal health. OLS estimates (columns 2, 4, and 6) are insignificant. My estimates of β_2 imply that the RFS can account for about 60.7 and 14.1 percent of the increases in the incidence rates of AWD and SGA, respectively.

1.8.3 Fetal selection

The U.S. registry of live births is a selected census of fetuses, those born alive, and as such does not account for fetuses that died in utero. The linked birth-death registry suffers from the same problem. Estimates of the effect of policy-induced increases in corn production on perinatal death are likely lower bounds of the true effect of corn production on fetal death (from conception and up to an hour after birth). In some cases, the bias could be so severe that it could lead to estimates with the opposite sign. For instance, an increase in corn acreage could make fetuses that would have otherwise been born alive (and died soon after) die in utero. The decreased rate of (observed) perinatal death could lead to a negative estimate of the effect of corn production on perinatal death. Under the plausible assumption that fetuses with AWD and growth restrictions are more likely to die in utero than other fetuses, estimates of the effect of policy-induced increases in corn production on AWD and SGA using the birth registry are also lower bounds of the true effect.

I gauge the magnitude of this selection problem in two ways. First, I construct a panel of fetuses using the registry of fetal deaths and of live births, and estimate the effect of corn production on the risk of death in utero, for births exposed to the planting season at different gestational periods (equation 1.4). Fetuses that die before week 20 are considered miscarriages and are not part of the fetal death registry. While this strategy allows me to account for a greater number of fetuses, it is prone to the same problems mentioned above, specially if policy-induced increases in corn production make fetuses that would have otherwise survived week 20 of gestation (and died in utero) die before week 20. Further, these selection issues might remain large, as the vast majority of fetal deaths occur early in the

pregnancy and there is substantial evidence that not all fetal deaths for which reporting is required are actually reported (MacDorman et al., 2012). For these same reasons, Sanders and Stoecker (2015) have suggested the use of changes in sex ratios at birth to gauge the effect of policy-induced changes in fetal insults on fetal death. Their strategy hinges on the reported larger susceptibility of males in utero, relative to females (Kraemer, 2000). Second, following Sanders and Stoecker (2015), I estimate the effect of policy-induced increases in corn production on the probability of being born male (equation 1.4 with an indicator variable for male as the dependent variable). These estimates are lower bounds of the effect of policy-induced increases in corn production on fetal death, as females are likely to be affected too.

Table 1.9, panel A, presents the IV estimates of the effect of policy-induced increases in corn production on fetal death, for fetuses exposed to the planting season at different times in their gestational period (equation 1.4). Columns 1-4 show that corn production does not have a significant effect on fetal death risk for fetuses exposed at conception, or during their first, second, or third gestational trimester, respectively. However, because the fetal death registry only includes deaths happening after week 20 of gestation, these estimates can only account for fetal selection happening late in the gestational period. If policy induced increases in corn production indeed lead to fetal selection in utero, then these estimates indicate that the selection must happen before week 20 of gestation, i.e., at conception or during the first two gestational trimesters.

Table 1.9, panel B, follows Sanders and Stoecker (2015) and estimates the effect of policy-induced increases in corn acreage on the probability of being born male (equation 1.4, not controlling for gender). I find significantly negative and large effects of policy-induced increases in corn production on the probability of being born male for fetuses exposed to the planting season during their first and second gestational trimesters. The estimates suggests that the RFS decreased the probability of being born a male by about 40 cases per 10,000 conceptions, for fetuses exposed during their first or second gestational trimester.[36] I further es-

[36]These calculations are based on IV estimates of β for births exposed to the planting season during their first or second gestational periods. The estimated effect of an additional 10,000 acres of corn is a decrease in the probability of being born a male of 118.78 cases per 10,000 births (with a standard error of 53.74). This estimate can be directly compared with those of Table1.9,

timate the heterogeneous effect of policy-induced increases in corn production, across seasons of high and low (or null) use of pre-emergence corn pesticide, on the probability of being born male (equation 1.6, not controlling for gender). I find that policy-induced increases in corn production decreased the probability of being born male about 55 cases pero 10,000 conceptions, for fetuses exposed during their first or second gestational trimester (regression estimates are presented in appendix Table A1.4). A conservative estimate, asuming that the RFS did not affect females, suggests that the RFS can account for an increase in the incidence of missing males (i.e., predicted number of missing males over total live births of males and females) by between 18.6 and 20.4 cases per 10,000 live births.

1.8.4 Maternal selection

The introduction of the RFS could have changed the composition of mothers across counties and seasons. Estimation of β_2 could, therefore, suffer from selection bias if policy-induced increases in corn production were to have heterogeneous effects on maternal characteristics influencing fetal heath, depending on the time they were exposed to the planting season. Such characteristics include young maternal age and maternal smoking, prominent risk factors for AWD and SGA, respectively. To evaluate whether maternal selection is a potential source of bias in my results, I estimate a reduced-form version of equation 1.6, with $PS(\tau) \cdot PCE \cdot Post_{2005}$ instead of $PS(\tau) \cdot corn$, and with different maternal characteristics as the dependent variable (excluding them one at a time from the set of controls).[37]

Table 1.10, columns 1-4, shows the estimates of the differential effect of my instrument for corn acreage ($PCE \cdot Post_{2005}$) on maternal characteristics depending on whether these mothers were exposed to the planting season at conception, during the first two gestational trimesters, third trimester, and last two gestational trimesters, respectively. I select these exposure periods as these are the ones for

panel B, for births exposed to the planting season during the first or second gestational trimesters, separately.

[37]Regressions with dependent variables constructed from non-binary maternal characteristics exclude all indicator variables for the different values of the respective non-binary variable. For instance, regressions with maternal age under 20 years exclude all maternal age bins from the set of controls.

which I find significant effects of policy-induced increases in corn production on AWD, the probability of being born male, SGA, and perinatal death, respectively. Importantly, I find no evidence of the instrument having heterogeneous effects on the probability of young maternal age (under 20 years) or maternal smoking for any of the exposure periods considered. Further, my instrument does not have significant heterogeneous effects on the probability of being married, being hispanic, or black, for any of the exposure periods considered. However, there is evidence suggesting that my instrument increased the probability that a mother does not have a high school diploma for mothers exposed during their third trimester of gestation. Still, overall Table 1.10 alleviates concerns about maternal selection driving my results.

1.9 Conclusions

The introduction of the Renewable Fuel Standard generated a demand shock for corn that was followed by heterogeneous increases in corn production across counties in the U.S. Corn Belt. To better understand this cross-county heterogeneity, I present a model of corn production tailored to the U.S. context and develop a measure to predict increases in corn acreage following a corn price increase. Using this new, county-level measure of potential for corn expansion I find that the RFS can account for an annual increase in corn acreage that amounts to almost half of its total increase in the U.S. Corn Belt between the 2001-2005 and 2006-2011 periods. Further, I find that policy-induced increases in corn acreage increased the risk of AWD, SGA, and perinatal death for births exposed to the planting season —a time of high pre-emergence pesticide use— during critical gestational periods. My estimates imply that the RFS increased the risk of: (1) AWD by more than 90 percent for births exposed at conception, (2) SGA by around 2 percent for births exposed during their third gestational trimester, and (3) perinatal death by more than 25 percent for births exposed during their last two trimesters. My estimates are robust to controlling for time shocks at the county level and are likely lower bounds of the true effects as I find suggestive evidence of fetal selection in utero.

To the best of my knowledge, this is the first paper to document a causal ef-

fect of corn production on human health, specifically, fetal health. Further, the established link between corn production and AWD and SGA, combined with the policy-induced increases in corn production between the 2001-2005 and 2006-2011 periods, help shed light on the upward trends in the incidence rates of these two conditions —a current a puzzle in the field of epidemiology. While this paper is able to shed light on the health consequences of increased corn production, in the context of the pesticide practices of my sample period, much could be gained from understanding which pesticides —or combination of pesticides— are most pernicious to health. Further research is needed in this regard.

Figures and Tables

Figure 1.1: Corn price and production in the U.S.: 1996-2011

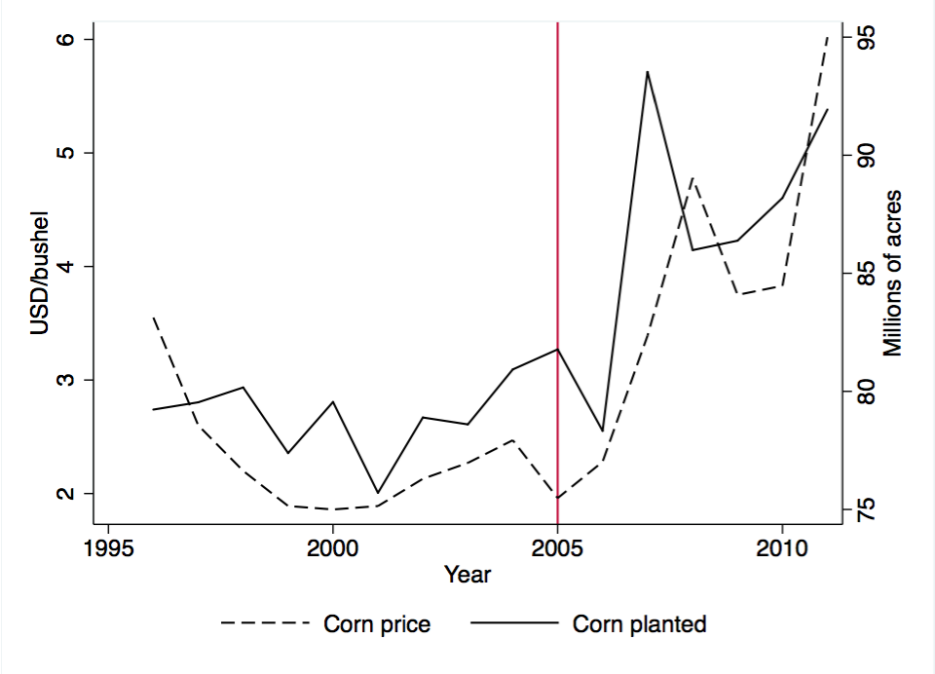

Source: USDA-NASS's Quick Stats. Available at `https://quickstats.nass.usda.gov` (accessed July, 2016).

Figure 1.2: Abdominal wall defects, small-for-gestational age, and corn production trends in the U.S. Corn Belt: 2001-2011

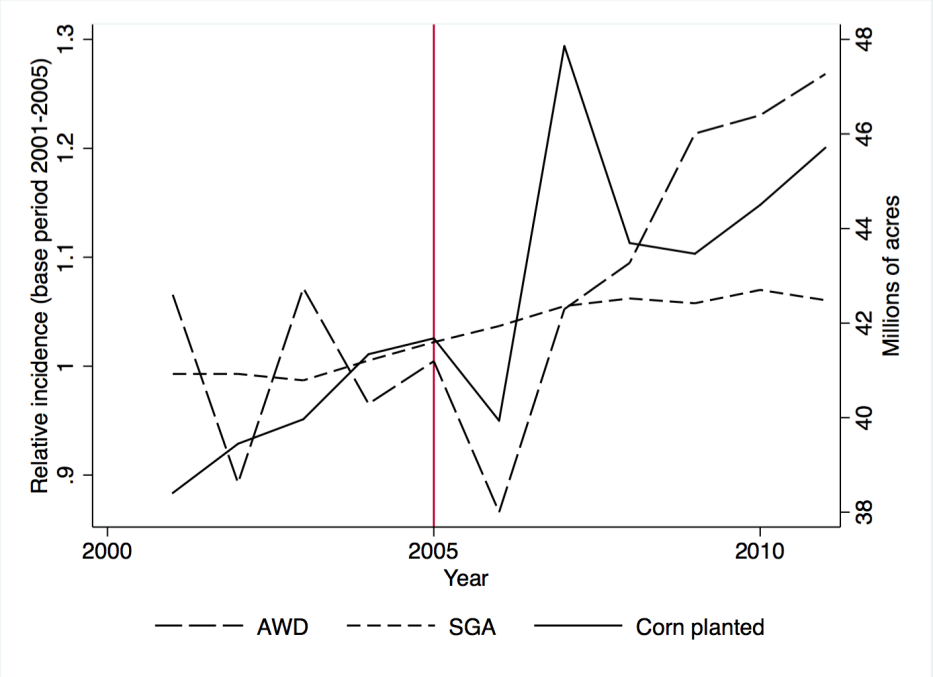

Notes: AWD and SGA time-series are based on all single births from 588 counties in the U.S. Corn Belt, these are all the counties for which there is detailed data on land use before 2005. Corn acreage data come from the same set of counties. The base incidence rates (2001-2005 average) for AWD and SGA are 3.91 per 10,000 births and 10.15 per 100 births, respectively
Sources: Health data: National Center for Health Statistics' natality files. Corn production: USDA-NASS's Quick Stats. Available at `https://quickstats.nass.usda.gov` (accessed July, 2016).

Figure 1.3: Characterization of county-level potential for corn expansion

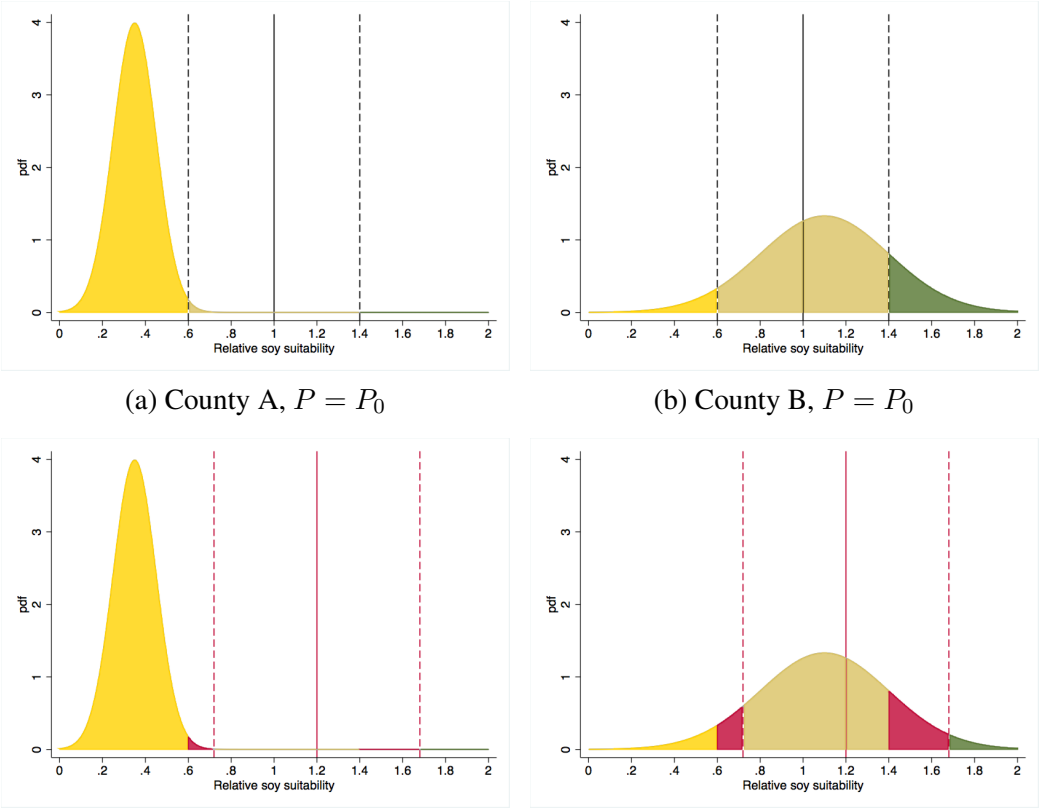

(a) County A, $P = P_0$

(b) County B, $P = P_0$

(c) County A, $P = P'$

(d) County B, $P = P'$

Notes: Author's calculations. Parameters: $\delta = 0.2$, $P_0 = 1$ in solid black line, $P' = 1.2$ in solid red line. Amount of land dedicated to corn monoculture, soy monoculture, and rotation are depicted in yellow, green, and brown, respectively. The increase in the amount of land dedicated to corn due to a price increase from P_0 to P' is depicted in red.

Figure 1.4: Land use, land suitability, and potential for corn expansion following the introduction of the Renewable Fuel Standard

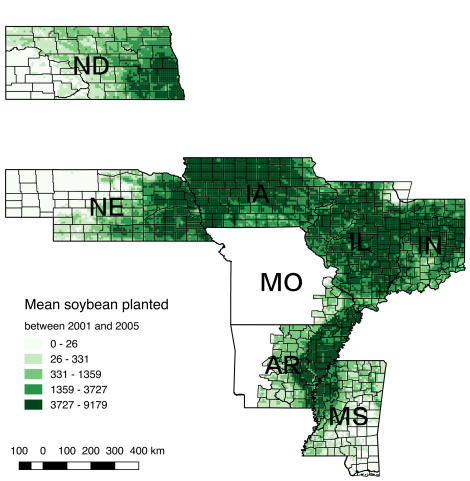

(a) Average soy acres 2001-2005

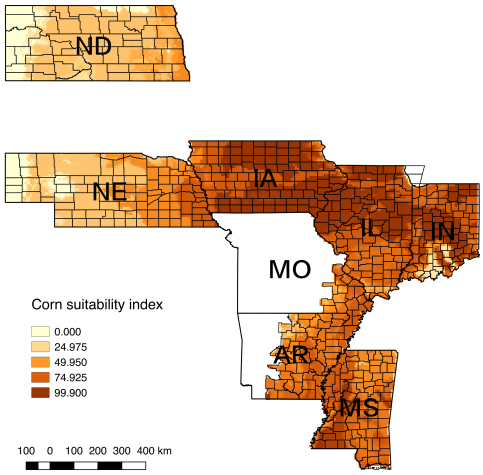

(b) Corn suitability index

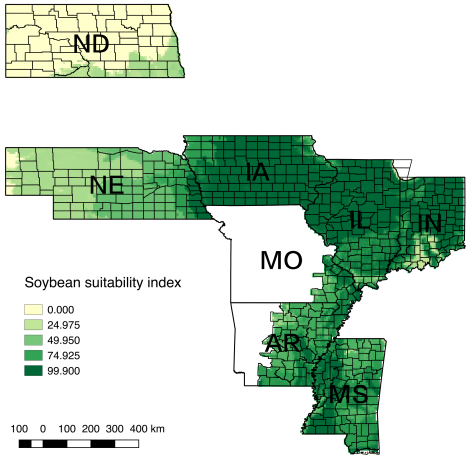

(c) Soy suitability index

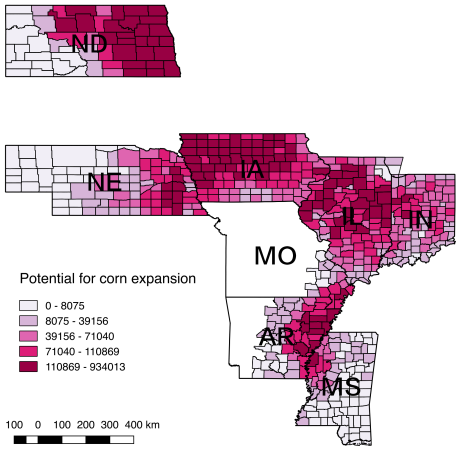

(d) Potential for corn corn expansion

Notes: Author's calculations based on data from USDA-NASS's Crop Data Layer, and FAO-GAEZ's crop suitability index.

Table 1.1: Descriptive statistics for the U.S. Corn Belt between 2001 and 2011

	2001-2005 Period					2006-2011 Period				
Panel A: Descriptive statistics pre- and post-RFS										
Number of counties	523					533				
Number of births	1,836,399					1,956,191				
Mean county-level corn acres	79,635					89,946				
S.D. county-level corn acres	64,736					70,579				
Mean county-level PCE						81,625				
S.D. county-level PCE						89,078				
Panel B: Descriptive statistics pre- and post-RFS by time of exposure to the planting season										
Gestational period	Any (1)	Concep. (2)	1^{st} Trim. (3)	2^{nd} Trim. (4)	3^{rd} Trim. (5)	Any (6)	Concep. (7)	1^{st} Trim. (8)	2^{nd} Trim. (9)	3^{rd} Trim. (10)
Fetal conditions and sex										
AWD*	3.91	4.53	4.40	4.15	3.39	4.67	5.64	4.93	4.60	4.56
SGA*	1015	1064	1062	1025	1003	1062	1111	1104	1084	1053
Perinatal death*	8.71	9.26	9.40	8.92	8.49	7.70	7.84	7.91	8.08	7.62
Fetal death**	60.15	65.18	64.55	58.61	60.05	56.53	60.26	58.65	56.78	57.47
Percent of births male	51.27	51.41	51.25	51.30	51.33	51.20	51.05	51.09	51.17	51.28
Maternal characteristics										
Share under 20 years	15.90	17.41	16.81	15.99	15.89	15.31	16.74	16.14	15.51	15.15
Mean age (years)	27.02	26.67	26.82	26.95	26.97	27.06	26.70	26.87	26.99	27.03
Share smoker	13.32	15.02	14.31	13.73	13.47	4.47	4.70	4.66	4.54	4.39
Share married	63.50	61.30	61.57	63.30	63.86	58.98	56.33	56.94	58.54	59.50
Share no high school diploma	20.52	21.58	21.45	20.41	20.28	18.69	20.26	19.58	18.73	18.39
Share black	15.65	17.47	16.88	15.82	15.42	15.10	17.01	16.50	15.47	14.76
Share hispanic	14.37	11.99	13.61	13.61	13.56	13.37	11.44	12.62	12.63	12.79

Notes: *Incidence rate in cases per 10,000 births. ** Incidence rate in cases per 10,000 fetuses (both those born alive and those that reach week 20 of gestation but die in utero). Summarized data come from the 588 counties for which there is data on PCE and for all county-years for which these counties have data on corn acreage. Gestational period makes reference to the time in which fetuses are exposed to the planting season: conception or any of the three gestational trimesters. "Any" considers all fetuses irrespective of whether they were exposed or not to the planting season.

Table 1.2: The effect of the Renewable Fuel Standard on corn production

	(1)	(2)	(3)	(4)
$PCE \cdot Post_{2005}$		0.093	0.097	0.049
s.e.		(0.007)	(0.008)	(0.006)
F-stat		158.632	141.440	74.193
$Post_{2005}$	8,482.66	879.522		
s.e.	(523.028)	(610.866)		
County FE	Y	Y	Y	Y
State-year FE	N	N	Y	Y
County trends	N	N	N	Y
Adjusted R^2	0.980	0.984	0.990	0.994
Observations	5,464	5,464	5,464	5,464

Notes: The dependent variable is corn planted (in acres). Estimation method is OLS. Standard errors are clustered at the county level.

Table 1.3: The effect of corn production on the incidence of abdominal wall defects by the time of exposure to the planting season

Gest. period	Conception		1^{st} Trimester		2^{nd} Trimester		3^{rd} Trimester	
	IV	OLS	IV	OLS	IV	OLS	IV	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Corn	12.868	0.845	6.521	0.141	-2.329	-0.104	-0.857	-0.618
s.e.	(3.657)	(0.728)	(2.395)	(0.514)	(1.945)	(0.407)	(1.998)	(0.426)
p-val	0.000	0.246	0.001	0.784	0.196	0.799	0.749	0.148
Observations	415,189	415,189	1,006,093	1,006,093	1,081,770	1,081,770	1,088,373	1,088,373
First Stage								
$PCE \cdot Post_{2005}$	0.048		0.050		0.049		0.050	
s.e.	(0.006)		(0.006)		(0.006)		(0.006)	
F-stat	67.651		75.479		72.476		72.600	
Observations	5,349		5,438		5,446		5,439	

Notes: The dependent variable is the incidence of abdominal wall defects, which has been multiplied by 10,000. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. The heading indicates the gestational period in which births were exposed to the planting season. Corn is measured in 10,000 acres and PCE has been divided by 10,000. All standard errors are clustered at the county level. IV/2SLS p-values are robust to weak instruments.

Table 1.4: The effect of corn production on the incidence of being born small-for-gestational age by the time of exposure to the planting season

Gest. period	Conception		1^{st} Trimester		2^{nd} Trimester		3^{rd} Trimester	
	IV	OLS	IV	OLS	IV	OLS	IV	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Corn	-55.314	-6.797	-39.866	-3.727	12.138	0.139	42.037	-5.091
s.e.	(45.721)	(9.383)	(27.555)	(6.451)	(21.313)	(6.166)	(23.258)	(5.440)
p-val	0.297	0.469	0.124	0.564	0.867	0.982	0.053	0.350
Observations	411,618	411,618	997,090	997,090	1,073,359	1,073,359	1,078,692	1,078,692
First Stage								
$PCE \cdot Post_{2005}$	0.048		0.049		0.049		0.049	
s.e.	(0.006)		(0.006)		(0.006)		(0.006)	
F-stat	66.986		75.221		71.753		72.316	
Observations	5,348		5,437		5,445		5,439	

Notes: The dependent variable is the incidence of being born small-for-gestational age, which has been multiplied by 10,000. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. The heading indicates the gestational period in which births were exposed to the planting season. Corn is measured in 10,000 acres and PCE has been divided by 10,000. All standard errors are clustered at the county level. IV/2SLS p-values are robust to weak instruments.

Table 1.5: The effect of corn production on the incidence of perinatal death by the time of exposure to the planting season

Gest. period	Conception		1^{st} Trimester		2^{nd} Trimester		3^{rd} Trimester	
	IV	OLS	IV	OLS	IV	OLS	IV	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Corn	-4.693	1.139	1.216	1.207	8.428	0.855	4.379	-1.015
s.e.	(4.168)	(0.880)	(2.502)	(0.520)	(3.542)	(0.500)	(2.427)	(0.489)
p-val	0.199	0.196	0.572	0.021	0.000	0.088	0.031	0.038
Observations	415,648	415,648	1,010,544	1,010,544	1,083,513	1,083,513	1,089,196	1,089,196
First Stage								
$PCE \cdot Post_{2005}$	0.048		0.049		0.049		0.049	
s.e.	(0.006)		(0.006)		(0.006)		(0.006)	
F-stat	66.242		75.018		73.848		71.714	
Observations	5,347		5,438		5,447		5,438	

Notes: The dependent variable is the incidence of perinatal death, which has been multiplied by 10,000. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. The heading indicates the gestational period in which births were exposed to the planting season. Corn is measured in 10,000 acres and PCE has been divided by 10,000. All standard errors are clustered at the county level. IV/2SLS p-values are robust to weak instruments.

Table 1.6: The effect of changes in average corn acreage, before and after the Renewable Fuel Standard, on fetal health

Dep. Var. Gest. period	AWD Conception (1)	SGA 3^{rd} Trimester (2)	Perinatal death 2^{nd} or 3^{rd} Trimester (3)
$\overline{Corn}$	4.685***	10.212	1.387*
s.e.	(1.183)	(8.258)	(0.790)
p-val	0.000	0.217	0.080
Observations	415,189	1,078,692	2,050,714

Notes: The heading indicates the gestational period in which births were exposed to the planting season. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. $\overline{Corn}$ is mean corn acreage, before and after the introduction of the RFS, and is measured in 10,000 acres. All dependent variables were multiplied by 10,000. Estimation method is OLS. All standard errors are clustered at the county level.

Table 1.7: The effect of the Renewable Fuel Standard on fetal health for births exposed during critical gestatational periods to the harvest season

Dep. Var. Gest. period	AWD Conception (1)	SGA 3^{rd} Trimester (2)	Perinatal death 2^{nd} or 3^{rd} Trimester (3)
$PCE \cdot Post_{2005}$	-0.118	-1.926	0.160
s.e.	(0.162)	(1.445)	(0.102)
p-val	0.466	0.183	0.116
Observations	447,143	991,674	1,939,757

Notes: The heading (Gest. period) indicates the gestational period in which births were exposed to the harvest season. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. All dependent variables were multiplied by 10,000 and PCE has been divided by 10,000. Estimation method is OLS. All standard errors are clustered at the county level.

Table 1.8: The seasonal effect of corn production on the incidence of abdominal wall defect, small-for-gestational age, and perinatal death

Dep. Var.	AWD		SGA		Perinatal death	
Gest. period	Conception		3^{rd} Trimester		2^{nd} or 3^{rd} Trimester	
	IV	OLS	IV	OLS	IV	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
$PS(\tau) \cdot Corn$	14.924	1.152	94.791	-3.506	7.993	-0.574
s.e.	(5.701)	(0.765)	(45.992)	(6.645)	(5.429)	(0.534)
p-val	0.000	0.133	0.005	0.598	0.100	0.283
Observations	3,769,994	3,769,994	3,739,283	3,739,283	3,774,030	3,774,030
First Stage						
$PS(\tau) \cdot Corn \cdot PCE \cdot Post_{2005}$	0.041		0.035		0.033	
s.e.	(0.012)		(0.012)		(0.013)	
F-stat	11.873		8.026		6.742	
Observations	3,769,994		3,739,283		3,774,030	

Notes: The heading (Gest. period) indicates the gestational period (τ) in which births were exposed to the planting season. All regressions include month of conception-by-county, month of conception-by-state-by-year, and county-by-year FEs, as well as month of conception-by-county linear trends. All dependent variables were multiplied by 10,000, corn is measured in 10,000 acres, and PCE has been divided by 10,000. $PS(\tau)$ is a binary variable indicating whether a given birth was exposed to the planting season during its τ gestational period. All standard errors are clustered at the county level. IV/2SLS p-values are robust to weak instruments.

Table 1.9: In utero selection by the time of exposure to the planting season

Gest. period	Conception (1)	1^{st} Trimester (2)	2^{nd} Trimester (3)	3^{rd} Trimester (4)
Panel A: Fetal Death				
Corn	-0.087	-8.359	6.019	19.866
s.e.	(7.877)	(6.227)	(7.712)	(14.054)
p-val	0.392	0.903	0.394	0.162
Observations	418,267	1,016,799	1,089,794	1,095,626
First Stage				
$PCE \cdot Post_{2005}$	0.048	0.049	0.049	0.049
s.e.	(0.006)	(0.006)	(0.006)	(0.006)
F-stat	66.108	75.018	73.848	71.714
Observations	5,348	5,438	5,447	5,438
Panel B: Male				
Corn	-32.454	-171.811	-79.299	12.624
s.e.	(66.785)	(74.934)	(45.019)	(42.983)
p-val	0.892	0.000	0.049	0.887
Observations	415,620	1,006,979	1,082,376	1,088,960
First Stage				
$PCE \cdot Post_{2005}$	0.048	0.049	0.049	0.050
s.e.	(0.006)	(0.006)	(0.006)	(0.006)
F-stat	67.097	74.992	72.213	72.716
Observations	5,349	5,438	5,446	5,439

Notes: The heading indicates the gestational period in which births were exposed to the planting season. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. All dependent variables were multiplied by 10,000 and PCE has been divided by 10,000. Estimation method is IV/2SLS. All standard errors are clustered at the county level and p-values are robust to weak inference.

Table 1.10: Maternal selection

Gest. period	Conception (1)	1^{st} or 2^{nd} Trim. (2)	3^{rd} Trim. (3)	2^{nd} or 3^{rd} Trim. (4)
Panel A: <20 years				
$PS(\tau) \cdot PCE \cdot Post_{2005}$	0.724	-0.058	-1.289	-1.452
(s.e.) p-val	(1.878) 0.700	(1.215) 0.962	(1.694) 0.447	(1.482) 0.327
Observations	3,772,560	3,774,155	3,774,033	3,774,148
Panel B: Smoker				
$PS(\tau) \cdot PCE \cdot Post_{2005}$	-1.784	0.739	0.715	0.429
(s.e.) p-val	(1.953) 0.361	(1.275) 0.562	(1.336) 0.593	(1.258) 0.733
Observations	3,344,555	3,345,974	3,345,861	3,345,967
Panel C: Married				
$PS(\tau) \cdot PCE \cdot Post_{2005}$	-2.754	-2.246	1.408	-1.387
(s.e.) p-val	(2.702) 0.309	(1.711) 0.190	(2.663) 0.597	(3.709) 0.709
Observations	3,772,560	3,774,155	3,774,033	3,774,148
Panel D: No HSD				
$PS(\tau) \cdot PCE \cdot Post_{2005}$	0.837	-0.550	3.398	2.240
(s.e.) p-val	(1.609) 0.603	(1.283) 0.668	(1.435) 0.018	(1.701) 0.188
Observations	3,484,956	3,486,547	3,486,425	3,486,540
Panel E: Black				
$PS(\tau) \cdot PCE \cdot Post_{2005}$	1.019	0.692	-1.424	-1.694
(s.e.) p-val	(1.828) 0.577	(1.147) 0.546	(1.963) 0.469	(1.804) 0.348
Observations	3,763,951	3,765,518	3,765,396	3,765,511
Panel F: Hispanic				
$PS(\tau) \cdot PCE \cdot Post_{2005}$	0.256	1.837	-1.662	-0.665
(s.e.) p-val	(1.505) 0.865	(1.556) 0.238	(1.543) 0.282	(1.022) 0.516
Observations	3,763,951	3,765,518	3,765,396	3,765,511

Notes: The heading indicates the gestational period (τ) in which births were exposed to the planting season. All regressions include month of conception-by-county, month of conception-by-state-by-year, and county-by-year FEs, as well as month of conception-by-county linear trends. All dependent variables were multiplied by 10,000 and PCE has been divided by 10,000. $PS(\tau)$ is a binary variable indicating whether a given birth was exposed to the planting season during its τ gestational period. Estimation method is OLS. All standard errors are clustered at the county level.

1.10 Appendix

Figure A1.1: Biofuels categories contemplated in the Renewable Fuel Standard

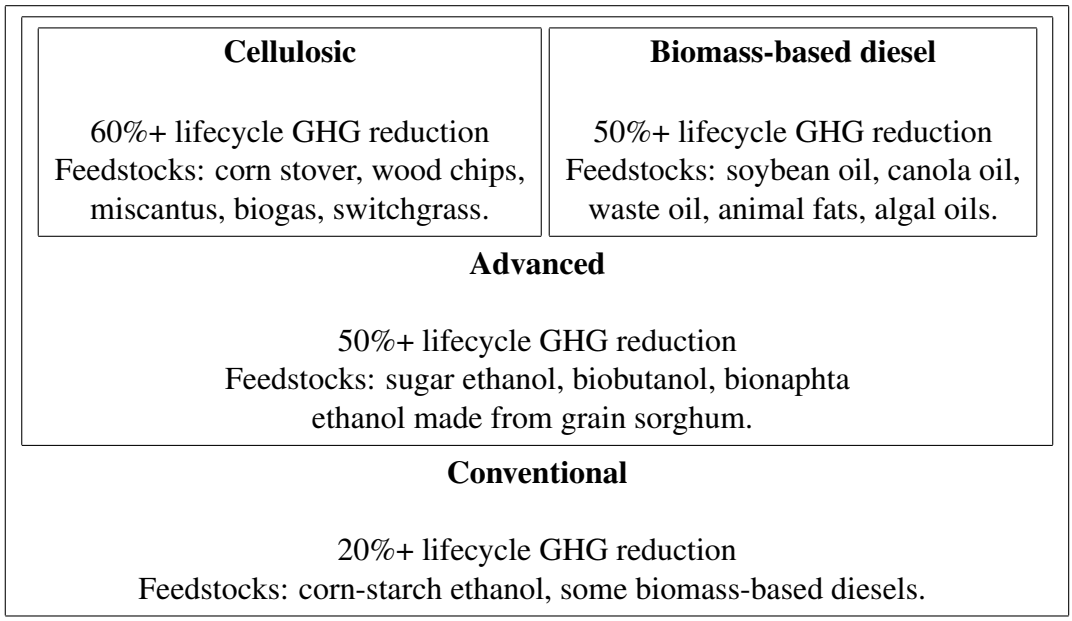

Note: The graph depicts the nested structure of the four biofuel categories contemplated in the Renewable Fuel Standard. It further indicates the requirements for biofuels, in terms of lifecycle greenhouse gas (GHG) emissions reduction and biomass feedstock criteria, to qualify under each category.
Source: Stock (2015), complemented with additonal information from Schnepf and Yacobucci (2013).

Table A1.1: The effect of corn production and the Renewable Fuel Standard on employment by sector

Dep. Var.	Total employment		Employment: Goods		Employment: Services	
	(1)	(2)	(3)	(4)	(5)	(6)
Corn	-117.991		-18.722		-99.281	
s.e.	(70.030)		(21.108)		(51.808)	
p-val	0.093		0.376		0.056	
$PCE \cdot Post_{2005}$		-0.004		-0.001		-0.003
s.e.		(0.004)		(0.001)		(0.003)
p-val		0.360		0.435		0.339
Observations	5,201	5,201	5,193	5,193	5,198	5,198

Notes: The dependent variables measure total mean monthly employment (columns 1 and 2), mean monthly employment in the goods sector (columns 3 and 4), and mean monthly employment in the services sector (columns 5 and 6) at the county level. All regression include county and state-by-year FEs, as well as county specific linear trends. Corn is measured in 10,000 acres and PCE has been divided by 10,000. Estimation method is OLS. All standard errors are clustered at the county level. Employment data come from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages.

Table A1.2: Reduced form estimates of the effect of the Renewable Fuel Standard on fetal health

Gest. period	Conception (1)	1^{st} Trimester (2)	2^{nd} Trimester (3)	3^{rd} Trimester (4)
Panel A: AWD				
$PCE \cdot Post_{2005}$	0.617	0.374	-0.114	-0.032
s.e.	(0.129)	(0.116)	(0.088)	(0.101)
p-val	0.000	0.001	0.196	0.749
Observations	415,189	1,006,093	1,081,770	1,088,373
Panel B: SGA				
$PCE \cdot Post_{2005}$	-2.797	-2.217	0.198	1.934
s.e.	(2.678)	(1.438)	(1.176)	(0.997)
p-val	0.297	0.124	0.867	0.053
Observations	411,618	997,090	1,073,359	1,078,692
Panel C: Perinatal death				
$PCE \cdot Post_{2005}$	-0.250	0.071	0.447	0.255
s.e.	(0.195)	(0.126)	(0.119)	(0.118)
p-val	0.199	0.572	0.000	0.031
Observations	415,648	1,010,544	1,083,513	1,089,196

Notes: The heading indicates the gestational period in which births were exposed to the planting season. All regressions include county, state-by-year, and month of conception FEs, as well as county specific linear trends. All dependent variables were multiplied by 10,000 and PCE has been divided by 10,000. Estimation method is OLS. All standard errors are clustered at the county level.

Table A1.3: Effect of the Renewable Fuel Standard on employment by season

Dep. Var.	Total employment				Agricultural employment			
Season	Planting (1)	Harvest (2)	Planting (3)	Harvest (4)	Planting (5)	Harvest (6)	Planting (7)	Harvest (8)
$PCE \cdot Post_{2005}$	41.586	41.628	-30.390	-21.064	0.352	0.811	-0.260	-0.152
s.e.	(29.081)	(28.628)	(34.310)	(27.668)	(0.209)	(0.312)	(0.150)	(0.264)
p-val	0.153	0.147	0.376	0.447	0.092	0.010	0.083	0.565
County FE	Y	Y	Y	Y	Y	Y	Y	Y
State-year FE	Y	Y	Y	Y	Y	Y	Y	Y
County trends	N	N	Y	Y	N	N	Y	Y
Observations	5,458	5,457	5,458	5,457	2,995	2,855	2,995	2,855

Notes: The dependent variables measure total mean monthly employment (across all sectors) at the county level during the planting season (columns 1 and 3) and the harvest season (columns 2 and 4), and mean monthly agricultural employment at the county level during the planting season (columns 5 and 7) and the harvest season (columns 6 and 8). PCE has been divided by 10,000. Estimation method is OLS. All standard errors are clustered at the county level. Employment data come from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages.

Table A1.4: Heterogeneity in the effect of corn production on the probability of being born male

Gest. period	1^{st} or 2^{nd} Trimester	
	IV	OLS
	(1)	(2)
$PS(\tau) \cdot Corn$	-163.731	1.090
s.e.	(96.335)	(11.346)
p-val	0.015	0.923
Observations	3,774,155	3,774,155
First Stage		
$PS(\tau) \cdot PCE \cdot Post_{2005}$	0.032	
s.e.	(0.012)	
F-stat	6.776	
Observations	3,774,155	

Notes: The heading indicates the gestational period (τ) in which births were exposed to the planting season. All regressions include month of conception-by-county, month of conception-by-state-by-year, and county-by-year FEs, as well as month of conception-by-county linear trends. The dependent variable was multiplied by 10,000, corn is measured in 10,000 acres, and PCE has been divided by 10,000. $PS(\tau)$ is a binary variable indicating whether a given birth was exposed to the planting season during its τ gestational period. All standard errors are clustered at the county level. IV/2SLS p-value is robust to weak instruments.

Chapter 2

ADVERSE RAINFALL SHOCKS AND CIVIL WAR: MYTH OR REALITY?

2.1 Introduction

According to a recent BBC news report, "Ethiopia has suffered periodic droughts and famines that lead to a long civil conflict in the 20^{th} Century" (BBC, 2015). Similar news reports are frequent. Their message —that droughts in Africa lead to food shortages that trigger civil conflict and war— seems plausible, and the idea that adverse rainfall shocks are a cause of civil conflict and war has by now become pervasive in policy circles. For example, U.S. President Barack Obama linked the rise of the terrorist group Boko Haram in Nigeria to droughts and the Secretary General of the United Nations Ban Ki-moon stated that droughts fueled the 1983-2005 civil war in Sudan (Ki-moon, 2007; Obama, 2015).

However, the empirical evidence on the relationship between rainfall and civil conflict or war in Africa is inconclusive. While some empirical studies find that Sub-Saharan African countries are more likely to see civil conflict or war following adverse rainfall shocks, others studies do not find such a link (Miguel et al., 2004; Ciccone, 2011; Miguel and Satyanath, 2011; Couttenier and Soubeyran, 2014). I contribute to this literature with an empirical approach that differs from

previous work in two main ways. Existing studies of the relationship between rainfall and civil conflict or war in African countries link the presence of civil conflict or war in a country to annual rainfall over a country's entire territory. I use new satellite data on African countries' growing seasons and data on their agricultural areas to focus on rainfall over countries' agricultural land during their growing seasons.[1] This should yield a rainfall measure that is more closely related to a country's agricultural output than rainfall during calendar years over a country's entire territory as agricultural output should be unaffected by rainfall before planting, after harvest, or in places where little or nothing is grown. In addition, my approach takes into account the evidence in agricultural economics that the relationship between rainfall and agricultural output is hump-shaped, with rainfall beyond a threshold decreasing output (Guiteras, 2009; Lobell et al., 2011; Schlenker and Roberts, 2009; Schlenker and Lobell, 2010).

My empirical analysis proceeds in three steps. I first combine data on rainfall and agricultural land with new, high-resolution satellite data on growing seasons for 51 African countries to construct a country-level measure of rainfall over agricultural land during growing seasons from 1980 to 2013. I refer to this new measure of country-level rainfall as agricultural rainfall.[2] I then combine the agricultural rainfall data with data on agricultural output to confirm the hump-shaped relationship between rainfall and agricultural output documented in agricultural economics in my data. In the last step I examine the effect of agricultural rainfall on the risk of civil conflict and war in all African countries and the subsample of Sub-Saharan African countries as many previous studies focused on Sub-Saharan Africa. A key feature of my analysis is that I allow for a U-shaped relationship between agricultural rainfall and the risk of civil conflict or war. This permits the relationship between agricultural rainfall and civil conflict or war to mirror the hump-shaped relationship between agricultural rainfall and agricultural output.

My main finding is a robust, U-shaped relationship between agricultural rainfall and the risk of civil war onset and incidence in (Sub-Saharan) African coun-

[1] The Food and Agriculture Organization (FAO) defines growing seasons as time periods when temperature and soil moisture allow for crop growth (Fischer et al., 2012).

[2] Monthly satellite rainfall data are available since 1979. However, because data from two calendar years are needed to measure rainfall during the growing season, agricultural rainfall data are available starting in 1980.

tries. The U-shaped relationship implies that the quantitative effect of rainfall shocks on the risk of civil conflict or war depends on the base level of rainfall.

I find that a negative rainfall shock that takes a country from the 50^{th} to the 25^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war onset and incidence in Africa by 2.2 and 2.3 percentage points, respectively. A positive shock that takes a country from the 50^{th} to the 75^{th} percentile of the distribution of agricultural rainfall decreases the risk of civil war onset and incidence by 0.6 and 0.8 percentage points, respectively. Nevertheless, large enough positive shocks have the opposite effect, increasing civil war onset and incidence risk. Going from the 50^{th} to the 90^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war onset and incidence by 1 and 0.7 percentage points, respectively. I also find a robust, U-shaped relationship between agricultural rainfall and the risk of civil conflict incidence. Moreover, the effect of rainfall on civil war and civil conflict onset and incidence risk is qualitatively the same in Sub-Saharan Africa.

Determining if and when rainfall shocks cause civil conflicts and especially civil wars is important because of the enormous cost of civil conflict and war in terms of human lives and living conditions (Sambanis, 2002). A better understanding of whether civil conflicts and wars might be triggered by rainfall shocks informs policymakers on how the risk of civil conflicts and wars might be diminished. A better understanding of the effect of rainfall shocks on civil conflicts and wars has also become pressing given the consensus that climate change will make extreme rainfall events more likely (IPCC, 2014). Simulations from twelve global circulation models predict increased heavy precipitation in east Africa and the opposite in the southern region of the continent (Seneviratne et al., 2012). These new weather patterns are expected to affect food security in many poor and agricultural countries. In Sub-Saharan Africa, predicted reductions in agricultural yields by the mid-century range between 8% and 22%, depending on the crop (Schlenker and Lobell, 2010).

Related Literature

My work is closely related to empirical studies examining whether Sub-Saharan African countries were more likely to experience civil conflict or war following low-rainfall years. In a seminal study, Miguel et al. (2004) find that Sub-Saharan

African countries experiencing low year-on-year rainfall growth were more likely to see civil conflict and war over the 1981-1999 period. Their civil conflict and war indicators are based on the Uppsala Conflict Data Program and the Peace Research Institute Oslo's (UCDP-PRIO) Armed Conflict Dataset. Civil conflict is defined as "a contested incompatibility that concerns government or territory or both where the use of armed forces between two parties results in at least 25 battle-related deaths. Of these two parties, at least one is the government of a state." (Gleditsch et al., 2002, pp. 168-619). Civil war is defined as a civil conflict with more than a 1000 deaths per year. An attractive feature of the panel-data approach by Miguel et al. (2004) is that it allows controlling for unobservables that translate into permanently greater civil conflict risk in some countries (country fixed effects) or some years (year fixed effects), as well as country-specific trends in conflict risk. Later studies with the same panel-data approach for Sub-Saharan Africa but for longer time periods do not find a statistically significant relationship between rainfall levels or year-on-year rainfall growth on the one hand and civil conflict or war on the other. See Ciccone (2011) and Miguel and Satyanath (2011) for Sub-Saharan African countries over the 1981-2009 period and Couttenier and Soubeyran (2014) for Sub-Saharan African countries over the 1945-2005 period, the latter only considers rainfall in levels.

The rainfall measures used in these empirical studies of the link between rainfall and civil conflict or war, aggregate rainfall during calendar years and over the totality of a country's territory. Recent research in agricultural economics on the relationship between rainfall and agricultural output has taken a different approach. At the local level, Schlenker and Roberts (2009) construct crop-specific measures of rainfall for U.S. counties by aggregating rainfall during the growing season and over the counties' cropland.[3] Schlenker and Lobell (2010) and Lobell et al. (2011) have generalized these crop-specific rainfall measures to the country level for Sub-Saharan Africa and a world panel, respectively. Additionally, this literature has documented the existence of a hump-shaped relationship between rainfall and agricultural output; the evidence comes from India (Guiteras, 2009), the U.S. (Schlenker and Roberts, 2009), Sub-Saharan Africa (Schlenker and Lo-

[3]Guiteras (2009) measures rainfall in Indian districts during the growing season but does not take into account land use.

bell, 2010), and a world panel (Lobell et al., 2011). Following this literature, I measure rainfall during the growing season and over a country's agricultural land. Further, I allow my rain measure and agricultural output to have a hump-shaped relationship and confirm it holds at the country level in my sample.

There is also empirical work examining the link between rainfall and inter-group violent events at the local level. For Africa, between 1960 and 2004, Theisen et al. (2011) find no statistically significant relationship between year-on-year rainfall growth or rainfall anomalies on the one hand and civil war battle locations on the other hand. Their data on battle location is derived from UCDP-PRIO's Armed Conflict Dataset. von Uexkull (2014) uses the UCDP Georeferenced Event Dataset (UCDP-GED) for Sub-Saharan Africa between 1989 and 2008 and finds that sustained drought is more likely to lead to conflict in locations with rainfed agriculture. Harari and La Ferrara (2013) find that negative shocks to the so-called standardized precipitation evapotranspiration index (SPEI) during the growing season increase the risk of inter-group violence incidence in Africa between 1997 and 2011 using UCDP-PRIO's Armed Conflict Location and Event Data Project (ACLED) dataset. This effect is mainly driven by increased battle risk, increased violence against civilians, and increased riot risk. My result on the existence of a significant relationship between civil conflict and war at the country level and rainfall over agricultural areas during the growing season resonate with those of Harari and La Ferrara (2013) at the local level. Outside the African context, my paper relates to a recent study by Crost et al. (2015) on the relationship between seasonal rainfall and inter-group violence in Philippine provinces over the 2001-2009 period. Using military reports, the authors find that more rainfall during the dry season decreases the risk of violent events while more rainfall during the wet season increases the risk of violent events.

There is also a growing theoretical literature in the social sciences that has examined the relationship between income and civil war (Besley and Persson, 2011; Chassang and Padró i Miquel, 2009; Dal Bó and Dal Bó, 2011; Fearon, 2007; Grossman, 1991), highlighting that civil war risk is increasing in the size of the appropriable resources (i.e., the loot) and decreasing in the opportunity cost of participating in civil war (e.g., foregone agricultural income).[4] Empirical tests

[4]A comprehensive review of the theoretical literature on the causes of civil wars is beyond the

of the opportunity cost mechanism, however, need to address the issue that the size of appropriable resources is seldom observable and that it will often be correlated with the opportunity cost of fighting (Chassang and Padró i Miquel, 2009; Fearon, 2007), leading to omitted-variable bias. For instance, consider the decision of an agricultural worker that has to choose whether to work on the fields or, alternatively, become a rebel and fight over the control of the state's resources. A negative and persistent agricultural shock (e.g., soil erosion, long-lasting pest) would reduce the returns to working the land, increasing the likelihood of conflict. However, at the same time, it would reduce the value of the economy —in the present and into the future— and, hence, the incentives to capture the state. A way to test for the opportunity cost mechanism is to look at the effect of effect of transitory rainfall shocks or transitory rainfall-induced income shocks on civil war. Chassang and Padró i Miquel (2009) have developed a model that underscores that while transitory income shocks have a direct impact on the opportunity cost of engaging in war, the effect on the total value of the economy is orders of magnitude smaller. By definition, the transitory shock will quickly dissipate and the size of the economy in the future will go back to its pre-shock value. All in all, theory predicts that following adverse transitory rainfall shocks in agricultural economies one should observe an increased risk of civil war.

The remainder of the paper is structured as follows. Section 2.2 introduces the data and discusses the construction of the agriculture-relevant rainfall measure. Section 2.3 draws from the agricultural economics literature to inform the mapping from rainfall onto agricultural output and civil war. Section 2.4 outlines the empirical strategy and presents the main results. Section 2.5 concludes.

2.2 Data

2.2.1 Agricultural weather and agricultural output

To construct the new country-level measure of agricultural rainfall, I combine raw data on rainfall with data on growing seasons and land use in Africa. I also construct an analogous variable for agricultural temperature. The sources of the

scope of this paper. The reader is referred to Blattman and Miguel (2010) and Sambanis (2002).

data are:

- The precipitation data (in mm) come from the Global Precipitation Climatology Project (GPCP Version 2.2) in a 2.5° latitude by 2.5° longitude global grid. The dataset combines gauge station information with satellite instruments to produce monthly rainfall estimates.[5]

- The growing season data for Africa come from a new data set, on an 8 km by 8 km grid, based on satellite images from the Advanced Very High Resolution Radiometer (AVHRR) sensor (Vrieling et al., 2013). The sensor effectively monitors phenological changes on land surface, and allows for the detection of green-up and senescence of vegetation for every year between 1981 and 2011. Because growing seasons —whether there is just one or two within 12 months— can span more than one calendar year, data from two calendar years are used to determine the start and end of the growing season(s) each year. The dataset reports the average start and end dates of the growing season(s), over the whole sample, for each grid cell.

- The land use data come from the Land Degradation Assessment in Drylands Project (LADA Version 1.1), which indicates whether the area of any cell on a 5 by 5 arc minutes grid (approximately 9km by 9 km at the Equator) was used for agricultural purposes in the year 2000 (Nachtergaele and Petri, 2013).

- The temperature data (in $^\circ$K) come from the National Center for Environmental Prediction and the U.S. Department of Energy (NCEP-DOE R2) in a T62 Gaussian grid.[6] The dataset combines gauge station, marine, aircraft, and satellite data, among other, using a climate model, to produce 6-hour temperature estimates.[7]

[5]Miguel et al. (2004) use this same rainfall dataset. The reader is referred to their paper for an introduction to the data and to Adler et al. (2003) for a technical discussion.

[6]The T62 Gaussian grid is made out of 192 point along each parallel and 94 points along each meridian. Points are equally spaced along the longitude dimension at a distance of 1.875°, and unequally spaced along the latitude dimension at a distance of approximately 1.904° —with the spacing becoming (marginally) smaller as one approaches the poles.

[7]The reader is referred to Kalnay et al. (1996) for an introduction to the data set and Kanamitsu

I use the data to construct country-level rainfall and temperature measures over agricultural land during the growing seasons following the agricultural economics literature (Guiteras, 2009; Schlenker and Roberts, 2009; Schlenker and Lobell, 2010; Lobell et al., 2011). Gridded data are mapped into political maps using country borders from Weidmann et al. (2010). Given that the growing seasons data is at a higher resolution than the weather data, I first construct mean growing season start and end dates for grid cells in Africa that match the resolution of the rainfall and temperature grids. For precipitation, I calculate the total amount of rainfall (in dm) during the growing season in each cell.[8] For temperature, I calculate the fraction of time (i.e., 6-hour readings) during the growing season that every cell was exposed to temperatures in the following temperature bins (in °C): $(-\infty, 0)$, $[0, 3)$, $[3, 6)$, ..., $[36, 39)$, $[39, +\infty)$.[9] Additionally, and for the sake of comparability with previous work that controls for average temperature, I also calculate the mean temperature (in °C) during the growing season. I then aggregate spatially these annual data to the country level. For any given country, I first select all the cells that "touch" the country (i.e., that lie fully or partially within the country's borders). Then, for each of these cells, I calculate the amount of agricultural land from the selected country that lies within the respective cell. Aggregation is done by a averaging the annual weather measurements of these cells, weighting them by their share of the country's agricultural land.

For comparability with previous work in the conflict literature, I also construct rainfall and temperature measures over countries' entire territories and during the calendar year. I term these variables aggregate rainfall and aggregate temperature,

---

et al. (2002) for a description of the latest improvements to the data. This data set also provides 6-hour estimates for precipitation, but I do not use them in this paper because of reliability problems (see Kalnay et al. (1996, p. 448)) —which are not present in the temperature data.

[8] For example, if the average growing season start and end months in a grid-cell are June and September, respectively, for each calendar year, rainfall is aggregated between those months in the same calendar year. When growing seasons span different calendar years (e.g., starts in November and ends in March), for each calendar year, rainfall is aggregated between the start month of the previous calendar year and the end month of the corresponding calendar year.

[9] I do this following the agricultural economics literature that has highlighted the need to exploit high-frequency temporal variation in the study of the effect of temperature on agricultural yields. As Schlenker and Roberts (2009, p. 15594) put it "... similar average temperatures may arise from two very different days, one with little temperature variation and one with wide temperature variation. Holding the average temperature constant, days with more variation will include more exposure to extreme outcomes, which can critically influence yields."

respectively. Aggregate rainfall corresponds exactly to the rainfall measure used by Miguel et al. (2004); the reader is referred to their paper for the details on how this variable is constructed. For temperature, the method for constructing the aggregate data is slightly different in that all cells that touch a country are used in the construction of the aggregate variable, and not just those whose centers lie within a given country. I do this for two reasons: (i) There is always some cell that touches a country, while there is not always a cell whose center lies within a country, thus, my process is discretion-free in the assignment of cells to countries. (ii) For comparability with the agricultural rainfall data.

The agricultural production data —which is at the country level— come from FAO's Statistical Division FAOSTAT. In particular, I measure agricultural production using the crops gross production index (GPI) as it is a quantity index of agricultural production (the base period is 2004-2006).

2.2.2 Civil war and civil conflict

Civil war and civil conflict data come from UCDP-PRIO's Armed Conflict Dataset (Gleditsch et al., 2002), version 4 (see section above). The original dataset codes dyads made out of the government of a state and an armed group that result in at least 25 or 1000 deaths per year for civil conflict and civil war, respectively. I construct a civil war incidence measure, at the country level, by coding a country as experiencing civil war in a given year if and only if it experienced an internal civil war (with or without foreign intervention) with at least one armed group. I, thus, exclude all dyads that involve extrasystemic (colonial) wars and interstate wars. To study the start of civil wars, I construct a civil war onset variable that is unity in period t if there was no civil war in $t-1$ but there was a civil war in t. It takes the value of zero if there was no war at $t-1$ nor at t. The civil war onset variable is not defined if a civil war was ongoing in $t-1$. Civil conflict incidence and onset variables are defined in an analogous way.

Table 2.1 shows that the average African country experienced civil war in 7.28% of the years during the 1981-2013 period and experienced the onset of a civil war in 2.21% of the years, the numbers are 21.54% and 5.75%, respectively, for civil conflict.

2.3 Rainfall and agriculture

Opportunity cost theories of the link between rain and civil war in Africa are based on the premise that rainfall affects agricultural output. I therefore start by investigating the effect of agricultural rain on agricultural output. Following the agricultural economics literature, I use a quadratic specification in agricultural rainfall —which allows for a hump-shaped relationship— to approximate the conditional expectation function (CEF) of agricultural output for African countries. The quadratic specification allows the effect of rainfall increments on output to depend on the base rainfall level. Hence, increased rainfall at low levels can have a positive effect on agricultural output, while the same increment at high rainfall levels (i.e., floods) could have a negative effect. Additionally, I also report results with linear agricultural rainfall. Table 2.2, columns 1-4, presents results from OLS estimations of the following equation

$$y_{c,t} = \beta_1 rain_{c,t} + \beta_2 rain_{c,t}^2 + \gamma temp_{c,t} + \delta_c + \delta_t + t_c + \epsilon_{c,t}, \qquad (2.1)$$

where y is agricultural production, $rain$ is agricultural rainfall, $temp$ is either mean agricultural temperature or a full set of agricultural temperature bins, δ_c are country fixed-effects, δ_t are year fixed-effects, t_c are linear trends, and ϵ is an error term. Subscripts c and t index countries and years, respectively. The vector $[\beta_1 \ \beta_2 \ \gamma]$ of regression coefficients is identified exploiting (exogenous) agricultural weather variation after controlling for country fixed effects, yearly shocks common to all African countries, and country-specific linear trends. Results from column 1, where rainfall enters linearly ($\beta_2 = 0$), indicate that rainfall has no effect on agricultural output. Column 2, follows Guiteras (2009), Lobell et al. (2011), Schlenker and Lobell (2010), and Schlenker and Roberts (2009) in using a quadratic relationship between agricultural output and agricultural rainfall. The results indicate that agricultural rainfall significantly affects agricultural output, both the linear and quadratic terms are significant at the 99% confidence level. At low rainfall levels, increased rain is positive for agricultural output, while the opposite is true at high levels. The high significance of the quadratic term confirms the non-monotonicity of the effect of rainfall on agricultural output and rejects a

linear relationship between these two variables.[10] In my sample, about 19% of the country-year observations lie on the decreasing section of the estimated relationship.

Columns 3 and 4 in Table 2.2 replace mean agricultural temperature with a full set of agricultural temperature bins so as to control for temperature more flexibly. The adjusted R^2 for the quadratic specification with the temperature bins controls (column 4) is larger than the one controlling for mean temperature (column 2). The effect of agricultural rain on agricultural output remains qualitatively and quantitatively the same. In what follows I only report results that flexibly control for temperature and relegate results controlling for average temperature to the appendix.

Figure 2.1 illustrates the hump-shaped relationship between agricultural rain and agricultural output in an augmented component-plus-residuals plot. It depicts the fitted values of agricultural output (as predicted by linear and quadratic agricultural rain) from OLS estimation of equation 2.1 plus the residuals, against agricultural rainfall. In the construction of the augmented component-plus-residuals plot, the estimation of equation 2.1 takes $temp$ to be the full set of agricultural temperature bins.[11]

Columns 5-8 in Table 2.2 estimate columns 1-4 using aggregate weather variables instead of agricultural weather variables. It is worth noting that the adjusted R^2s are always larger in the regressions that use agricultural weather variables. Moreover, in the quadratic specification with temperature bins, in columns 4 and 8, the share of the residual variation in agricultural output —after controlling for country fixed effects, year fixed-effects, and linear trends— that is explained by agricultural weather variables alone (R^2 (p)) is over four times larger than that explained by aggregate weather variables. This is what one would expect if the use of rainfall data from outside the growing season and from places where little or nothing is grown adds measurement error to the aggregate rainfall measure.

[10] In Table OA1 in the online appendix, I also compare the quadratic specification to several other parametric specifications that have been used in the conflict literature. The quadratic specification always has a higher explanatory power in terms of adjusted R^2.

[11] See Ashraf and Galor (2013), Ashraf and Michalopoulos (2015), Duranton et al. (2014), and Liebman et al. (2004) for other applications of augmented component-plus-residuals plots and Mallows (1986) for a general discussion. Standard partial residual plots of agricultural output on linear and quadratic agricultural rainfall terms are presented in Figure A2.1a-b in the Appendix.

As shown in Table OA2 in the online appendix, the hump-shaped relationship between agricultural output and agricultural rainfall holds for Sub-Saharan Africa (SSA) also.

Theory only predicts an unambiguous negative effect of adverse rainfall shocks or rain-induced income shocks, when these are transitory (Chassang and Padró i Miquel, 2009). To test if agricultural rainfall shocks are indeed short-lived, I estimate a modified version of equation 2.1, augmented with once-lagged weather variables. Lagged agricultural rainfall is never significant, whether one controls for mean temperature or a full set of temperature bins (results are presented in Table A2.1 in the appendix).

2.4 Rainfall and civil war

The section above has provided evidence, supporting previous work in agricultural economics, showing that (i) agricultural rainfall is a better predictor of agricultural output than aggregate rainfall, and (ii) transitory agricultural rainfall shocks have non-monotonic effects on agricultural output. Additionally, Table 2.2 showed that a linear specification relating agricultural rainfall to agricultural output masks this relationship. These pieces of evidence beg the question of whether previous inconclusive findings relating rainfall to civil war are due to a true no-effect (i.e., civil war risk being independent of rainfall in a statistical sense) or a combination of mismeasurement and misspecification.[12]

2.4.1 Empirical strategy

To estimate the effect of agricultural rainfall shocks on civil war onset (war), I relate the latter to a linear and quadratic term in agricultural rain, some measure of agricultural temperature —either mean agricultural temperature or a full set of agricultural temperature bins—, country fixed-effects, year fixed-effects, country-specific linear trends, and an error term. This specification (equation 2.2) allows

[12]The non-conclusive findings using aggregate rainfall are also true in my data as I show in appendix Tables A2.2, A2.3, and A2.4.

rainfall and civil war to have a U-shaped relationship, mirroring the hump-shaped relationship between rainfall and agricultural output.

$$war_{c,t} = \beta_1 rain_{c,t} + \beta_2 rain_{c,t}^2 + \gamma temp_{c,t} + \delta_c + \delta_t + t_c + \epsilon_{c,t} \tag{2.2}$$

The coefficients of interest are β_1 and β_2, and these are identified out of the (exogenous) rainfall variation, after controlling for time-invariant country differences, shocks common to all countries in a given year, country-specific linear trends, and temperature. $\beta_1 < 0$ and $\beta_2 > 0$ would be consistent with the opportunity cost mechanism, whereby decreased agricultural production, either due to droughts or excess rain, leads to increased civil war outbreak risk.

To study the effect of agricultural rainfall shocks on civil war incidence, I relate this variable to all the independent variables in equation 2.2. Additionally, and to account for the fact that civil wars tend to be persistent events, I also control for lagged civil war incidence —note that, by construction, civil war onset is not persistent. Again, a negative β_1 and a positive β_2 would be consistent with the theoretical effects of rainfall shocks on civil war.

The effect of agricultural rainfall shocks on civil conflict onset and incidence risk is estimated in an analogous way.

2.4.2 Estimates of the effect of agricultural rainfall shocks on civil war and civil conflict

Column 1, in (panel A) Table 2.3, reports the OLS estimates of the effect of agricultural rainfall shocks on civil war onset risk for Africa (1981-2013) using a quadratic specification. Robust standard errors clustered at the country level are presented in parenthesis. Both the linear and quadratic agricultural rainfall coefficients are significant at the 95% confidence level —with $\beta_1 < 0$ and $\beta_2 > 0$— evidencing a U-shaped relationship between civil war onset risk and agricultural rainfall shocks. Column 2 presents the estimates of a modified version of equation 2.2, where the quadratic agricultural rainfall term has been eliminated. Linear agricultural rain is not significantly related to civil war onset. This result comes as

no surprise, if agricultural rain has non-monotonic effects on agricultural output and it, in turn, affects civil war risk; using a linear specification will mask the link between agricultural rain and both agricultural output (as shown in the section above) and civil war. Column 3 presents estimates from regressions using agricultural rainfall growth rates for comparison with previous work in the conflict literature. Again, the results are of the no-effect type.

Panel a, in Figure 2.2, illustrates the quadratic, U-shaped relationship between civil war onset risk and agricultural rainfall in an augmented component-plus-residuals plot. It depicts the fitted values of civil war onset risk (as predicted by linear and quadratic agricultural rain) from OLS estimation of equation 2.2 plus the residuals, against agricultural rainfall.[13]

Columns 4-6 and 7-9, in (panel A) Table 2.3, replicate the analysis in columns 1-3, but for SSA between 1981-1999 and 1981-2013, respectively. The first of these samples corresponds to the one analyzed by Miguel et al. (2004) and the second one to an updated version of it. While I find no effect of linear agricultural rainfall or agricultural rainfall growth on civil war onset risk, I find that a more flexible specification —the quadratic— uncovers a significant relationship between agricultural rainfall and civil war onset risk in SSA.

The U-shaped relationship implies that the quantitative effect of agricultural rainfall on civil war onset risk depends on the baseline level of rainfall. A negative rainfall shock that takes a country from the 50^{th} to the 25^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war onset in Africa (1981-2013), SSA (1981-1999), and SSA (1981-2013) by 2.2, 3.1, and 2.3 percentage points, respectively. A positive shock that takes a country from the 50^{th} to the 75^{th} percentile of the distribution of agricultural rainfall decreases the risk of the outbreak of civil war by 0.6, 0.1, and 0.8 percentage points, in the respective samples. Nevertheless, large enough positive shocks increase civil war onset risk. Going from the 50^{th} to the 90^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war outbreak by 1, 2.8 and 0.6 percentage points, respectively.

Column 1, in (panel B) Table 2.3, reports the OLS estimates of the effect

[13] Standard partial residual plots of civil war onset risk on linear and quadratic agricultural rainfall terms are presented in Figure A2.2a-b in the Appendix.

of agricultural rainfall shocks on civil war incidence risk for Africa (1981-2013) using a quadratic specification. Again, both the linear and quadratic agricultural rainfall coefficients are significant at the 95% confidence level —with $\beta_1 < 0$ and $\beta_2 > 0$— evidencing a U-shaped relationship between civil war incidence risk and agricultural rainfall shocks. This quadratic relationship is illustrated in in Panel b, in Figure 2.2, by means of an augmented component-plus-residuals plot.[14] Column 2 and column 3 present results from specifications using linear agricultural rainfall ($\beta_2 = 0$) and agricultural rainfall growth. Once again, neither linear agricultural rainfall nor agricultural rainfall growth are significantly related to civil war onset risk. The (quadratic) relationship between agricultural rainfall and civil war onset risk is qualitatively similar in SSA for the 1981-1999 period and the 1981-2013 period.

Quantitatively, a negative rainfall shock that takes a country from the 50^{th} to the 25^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war incidence in Africa (1981-2013), SSA (1981-1999), and SSA (1981-2013) by 2.3, 3.9, and 2.4 percentage points, respectively. A positive shock that takes a country from the 50^{th} to the 75^{th} percentile of the distribution of agricultural rainfall decreases the risk of civil war incidence by 0.8, 1.5, and 0.8 percentage points, in the respective samples. However, large enough positive shocks increase civil war incidence risk. Going from the 50^{th} to the 90^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war incidence by 0.7, 0.3 and 0.7 percentage points, respectively. The quantitative effect of agricultural rainfall on civil war onset and incidence risk, for all three samples and for a larger combination of rainfall shocks, is presented in appendix Table A2.5.

Table 2.4, presents the same analysis as Table 2.3 but for civil conflict. Panel A shows that, unlike the results for civil war onset, agricultural rainfall shocks are not significantly related to the start of conflicts that do not necessarily exceed the 1000 deaths threshold. Panel B presents the estimates of the effect of agricultural rainfall on civil conflict incidence. The results largely mimic those for civil war incidence, indicating that agricultural rainfall shocks not only have non-monotonic effects on the incidence of fully fledged war, but also on the incidence

[14] Standard partial residual plots of civil war incidence risk on linear and quadratic agricultural rainfall terms are presented in Figure A2.2c-d in the Appendix.

of smaller scaled conflicts. Further, linear agricultural rainfall and growth rates are never significantly related to civil conflict onset or incidence risk. The quantitative effect of agricultural rainfall on civil conflict onset and incidence risk, for all three samples and for a larger combination of rainfall shocks, is presented in appendix Table A2.6.

The equivalents of Tables 2.3 and 2.4, controlling for average temperature instead of the flexible set of temperature bins can be found in appendix Tables A2.7 and A2.8. All results remain qualitatively the same.

The estimated hump-shaped relationship between agricultural rainfall and agricultural output implies the existence of a turning point beyond which extra rainfall decreases agricultural output. Similarly, the estimated U-shaped relationships between agricultural rainfall and civil war and civil conflict imply the existence of turning points beyond which extra rainfall increases civil war and civil conflict risk. Importantly, when equations 2.1 and 2.2, for civil war incidence and onset and civil conflict incidence, are estimated in a seemingly unrelated regression (SUR) framework, I cannot reject the joint, null hypothesis that the estimated turning points from civil war and civil conflict regressions are the same as the turning point in the agricultural output regression at the 90% confidence level, for Africa (1981-2013) and SSA (1981-1999).[15]

2.5 Conclusions

Policy makers and the media around the world have associated crop failures caused by droughts to civil war and conflict in African countries. However, empirical work on the effect of adverse rainfall shocks on African civil wars and conflicts has been inconclusive. I argue that to better understand whether rainfall shocks affect the risk of civil war and conflict through agricultural productivity, it is useful to first examine the effect of rainfall shocks on agricultural output. Following recent work in agricultural economics, I relate the agricultural output of African

[15]Standard errors for each regression are clustered at the country level. The p-values for the joint, null hypotheses for Africa (1981-2013), SSA (1981-1999), and SSA (1981-2013) are 0.072, 0.051, and 0.104, respectively. I do not include civil conflict onset estimates in the multiple-hypotheses test because linear and quadratic agricultural rainfall are not significant in any of the samples.

countries to rainfall over agricultural land during the growing seasons and allow for a hump-shaped effect of rainfall. This yields a robust, hump-shaped relationship between rainfall and agricultural output. Hence, increases in rainfall raise agricultural output at low levels and decrease agricultural output at high levels. If rainfall affects civil war and conflict through its effect on agricultural productivity, the effect of rainfall on the risk of civil war and conflict should therefore be U-shaped. I find this to be the case. Hence, increases in rainfall lower the risk of civil war and conflict at low levels and raise the risk of war and conflict at high levels. In particular, I find that a negative rainfall shock that takes a country from the 50^{th} to the 25^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war onset and incidence in Africa by 2.2 and 2.3 percentage points, respectively. A positive shock that takes a country from the 50^{th} to the 75^{th} percentile of the distribution of agricultural rainfall decreases the risk of civil war onset and incidence by 0.6 and 0.8 percentage points, respectively. However, large enough positive shocks have the opposite effect, increasing civil war onset and incidence risk. Going from the 50^{th} to the 90^{th} percentile of the distribution of agricultural rainfall increases the risk of civil war onset and incidence by 1 and 0.7 percentage points, respectively. The effect of rainfall on civil war onset and incidence risk is qualitatively the same for Sub-Saharan African countries. These results resonate with a recent literature that has linked rainfall shocks to other (local) forms of political violence and inform the policy debate on the effects of adverse rainfall shocks, in general, and climate change, in particular.

Figures and Tables

Figure 2.1: Agricultural output and agricultural rainfall

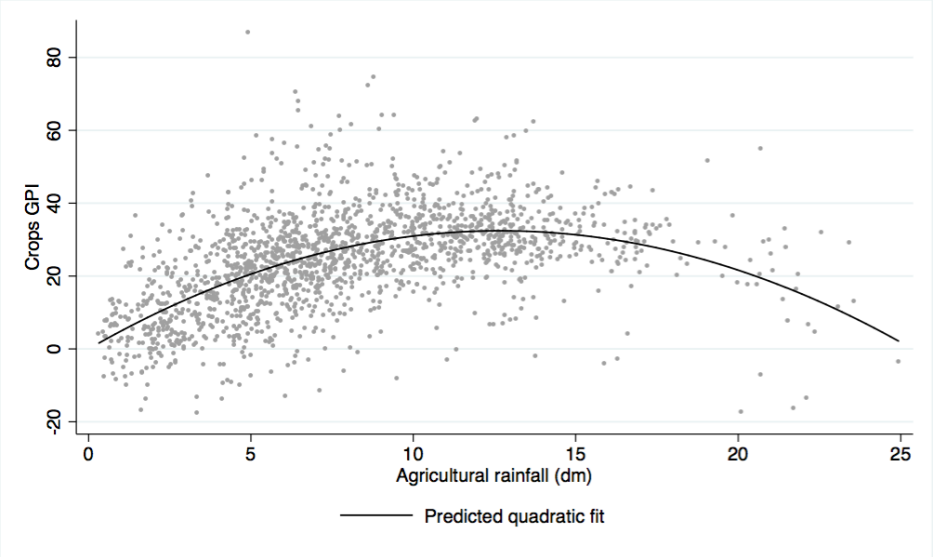

Notes: Author's calculations, see the Data section for details on data sources and variable construction. The graph shows an augmented component-plus-residual plot of the relationship between agricultural output and agricultural rainfall. The underlying regression corresponds to the specification in column 4 of Table 2.2.

Figure 2.2: Civil war onset and incidence risk and agricultural rainfall

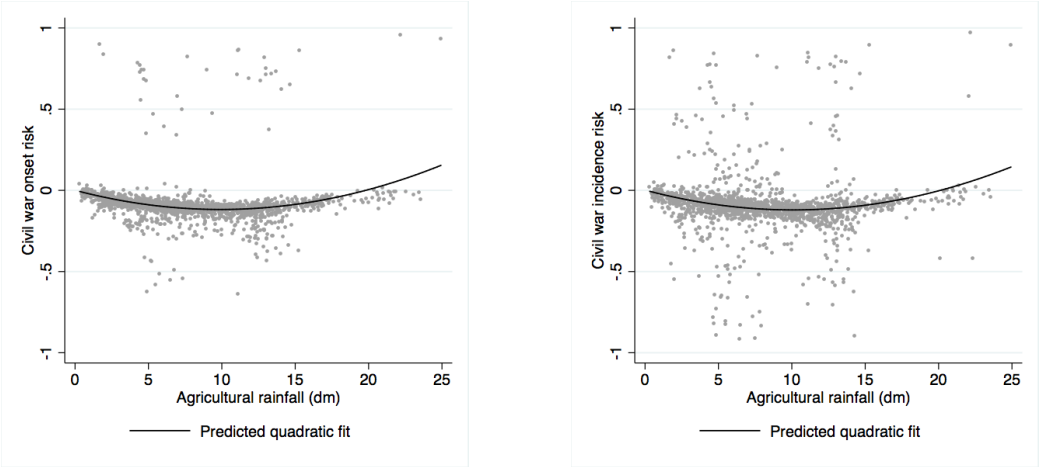

(a) Civil war onset

(b) Civil war incidence

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Panel a (b) shows an augmented component-plus-residual plot of the relationship between civil war onset (incidence) and agricultural rainfall. The underlying regression corresponds to the specification in column 1, panel A (B), of Table 2.3.

Table 2.1: Descriptive statistics

	Obs	Mean	S.D.	Min	Max
Civil war incidence	1,662	0.073	0.260	0	1
Civil conflict incidence	1,662	0.215	0.411	0	1
Civil war onset	1,538	0.022	0.147	0	1
Civil conflict onset	1,305	0.057	0.233	0	1
Crops GPI†	1,650	85.306	28.307	25.190	234.520
Agg. rain (dm)	1,662	9.249	5.893	0.191	26.197
Agg. mean temp (°C)	1,662	23.431	2.740	14.357	28.238
Agri. rain (dm)	1,662	8.036	4.314	0.320	24.929
Agri. mean temp (°C)	1,662	23.133	3.435	14.672	29.900

Notes: Author's calculations, see the Data section for details on data sources and variable construction. The sample is made out of 51 African countries between 1981-2013. Aggregate (Agg.) variables summarize temporally and spatially disaggregated weather data over the entire calendar year and the totality of a country's territory. Agricultural (Agri.) variables summarize information during the growing seasons and over agricultural land. † Data on the crops gross production index (GPI) for Ethiopia between 1981 and 1992 are missing.

Table 2.2: Agricultural production and rainfall in Africa between 1981 and 2013

	Agricultural weather				Aggregate weather			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rain	0.883	5.380	0.796	5.128	0.257	4.417	0.207	4.323
	(0.603)	(1.217)	(0.583)	(1.173)	(0.450)	(0.997)	(0.459)	(1.018)
$rain^2$		-0.214		-0.203		-0.158		-0.155
		(0.047)		(0.046)		(0.032)		(0.031)
Observations	1,650	1,650	1,650	1,650	1,650	1,650	1,650	1,650
Obs. decreasing section (%)	n.a.	19.27	n.a.	19.21	n.a.	24.79	n.a.	25.03
Adjusted R-squared	0.822	0.828	0.826	0.831	0.819	0.824	0.823	0.827
Adjusted R-squared (p)	-0.056	-0.023	-0.033	-0.006	-0.076	-0.049	-0.052	-0.026
Mean temperature	Y	Y	N	N	Y	Y	N	N
Temp. bins	N	N	Y	Y	N	N	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. The dependent variable is the crops gross production index. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Robust standard errors are clustered at the country level and are presented in parentheses. The adjusted R-squared (p) is the adjusted R-squared from regressions where country fixed-effects, year fixed-effects, and country-specific linear time trends have been partialled out from all variables. Obs. decreasing section (%) refers to the percentage of the observations that lie on the decreasing section of the estimated relationship between agricultural output and rainfall. Agricultural variables summarize information during the growing seasons and over agricultural land. Aggregate weather variables summarize temporally and spatially disaggregated weather data over the entire calendar year and the totality of a country's territory.

Table 2.3: The effect of agricultural rainfall on civil war onset and incidence risk

	Africa 1981-2013			SSA 1981-1999			SSA 1981-2013		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Civil war onset									
rain	-0.024	0.002		-0.044	0.003		-0.030	0.002	
	(0.011)	(0.004)		(0.022)	(0.007)		(0.013)	(0.005)	
$rain^2$	0.001			0.002			0.001		
	(0.001)			(0.001)			(0.001)		
rain growth			0.011			0.005			0.014
			(0.008)			(0.017)			(0.011)
Observations	1,538	1,538	1,536	662	662	662	1,228	1,228	1,227
Panel B: Civil war incidence									
rain	-0.024	0.001		-0.046	-0.002		-0.031	0.002	
	(0.012)	(0.005)		(0.024)	(0.008)		(0.015)	(0.006)	
$rain^2$	0.001			0.002			0.001		
	(0.001)			(0.001)			(0.001)		
rain growth			0.014			-0.024			0.011
			(0.010)			(0.022)			(0.013)
lagged dep. variable	0.369	0.370	0.370	0.185	0.186	0.186	0.338	0.341	0.341
	(0.073)	(0.072)	(0.072)	(0.073)	(0.075)	(0.074)	(0.070)	(0.070)	(0.070)
Observations	1,660	1,660	1,660	743	743	743	1,343	1,343	1,343
Temp. bins	Y	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Robust standard errors are clustered at the country level and are presented in parentheses. Agricultural variables summarize information during the growing seasons and over agricultural land. Columns 4-6 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 7-9 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013.

Table 2.4: The effect of agricultural rainfall on civil conflict onset and incidence risk

	Africa 1981-2013			SSA 1981-1999			SSA 1981-2013		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Civil conflict onset									
rain	-0.000	-0.004		-0.002	-0.004		-0.009	-0.007	
	(0.013)	(0.005)		(0.017)	(0.008)		(0.015)	(0.006)	
$rain^2$	-0.000			-0.000			0.000		
	(0.000)			(0.001)			(0.001)		
rain growth			-0.012			0.002			-0.025
			(0.021)			(0.029)			(0.031)
Observations	1,305	1,305	1,303	575	575	575	1,032	1,032	1,031
Panel B: Civil conflict incidence									
rain	-0.032	-0.004		-0.042	-0.013		-0.043	-0.005	
	(0.015)	(0.008)		(0.024)	(0.010)		(0.019)	(0.010)	
$rain^2$	0.001			0.001			0.002		
	(0.001)			(0.001)			(0.001)		
rain growth			-0.028			-0.062			-0.049
			(0.020)			(0.033)			(0.028)
lagged dep. variable	0.386	0.389	0.389	0.076	0.078	0.078	0.373	0.378	0.378
	(0.055)	(0.056)	(0.056)	(0.073)	(0.073)	(0.073)	(0.056)	(0.059)	(0.059)
Observations	1,660	1,660	1,660	743	743	743	1,343	1,343	1,343
Temp. bins	Y	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Robust standard errors are clustered at the country level and are presented in parentheses. Agricultural variables summarize information during the growing seasons and over agricultural land. Columns 4-6 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 7-9 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013.

Appendix

Figure A2.1: Agricultural output and agricultural rainfall

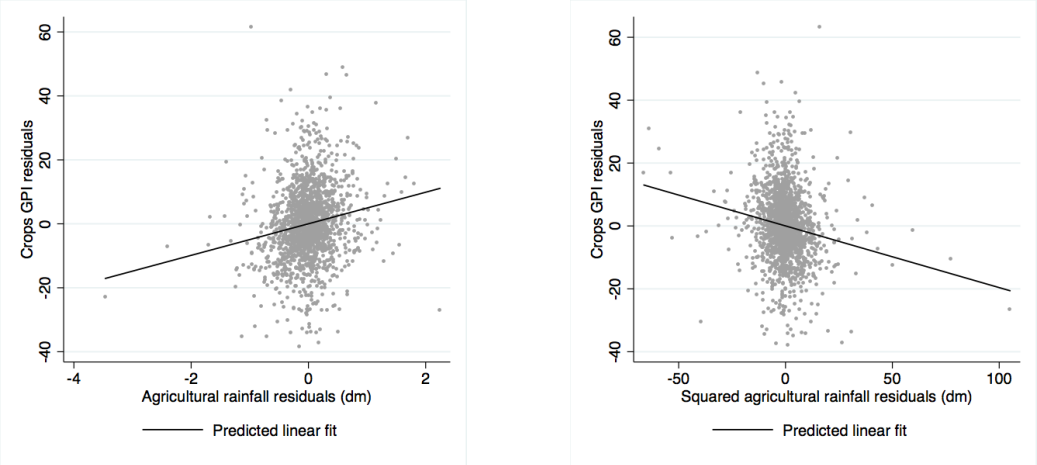

(a) First-order effect

(b) Second-order effect

Notes: Author's calculations, see the Data section for details on data sources and variable construction. This figure illustrates the positive first-order- and negative second-order- partial effects of agricultural rainfall on agricultural output, in Panel a and b, respectively. Panel a (b) plots the residuals of agricultural output from a regression on quadratic (linear) agricultural rainfall, country fixed-effects, year fixed-effects, country-specific linear time trends, and the full set of agricultural temperature bins against the residuals of linear (quadratic) agricultural rainfall on quadratic (linear) agricultural rainfall and same set of controls.

Figure A2.2: Civil war onset and incidence risk and agricultural rainfall

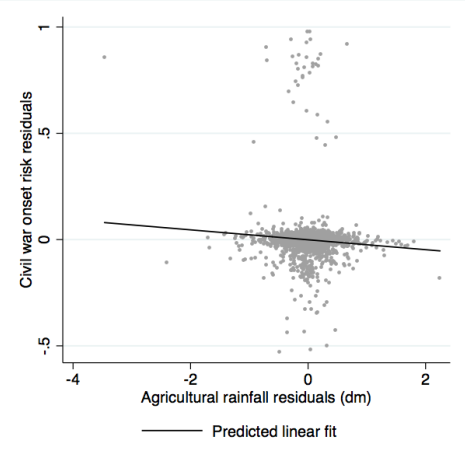

(a) Civil war onset risk: first-order effect

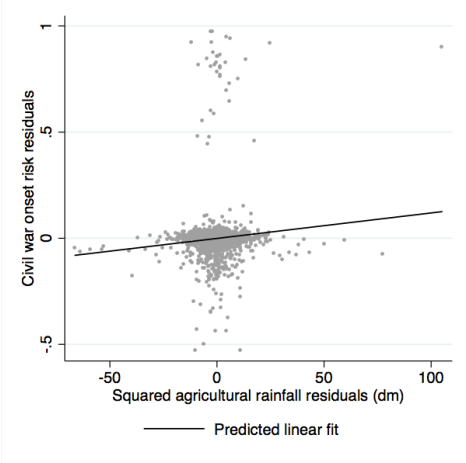

(b) Civil war onset risk: second-order effect

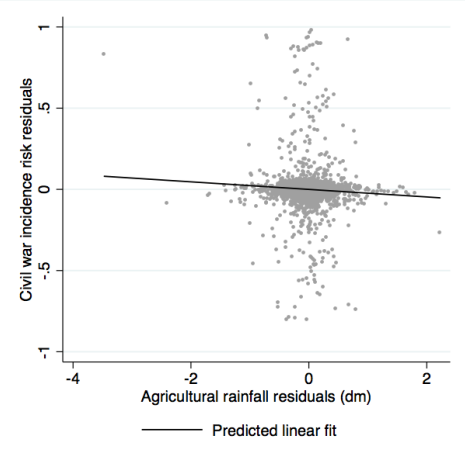

(c) Civil war incidence risk: first-order effect

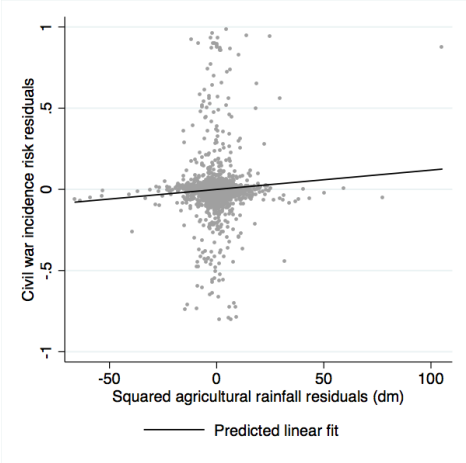

(d) Civil war incidence risk: second-order effect

Notes: Author's calculations, see the Data section for details on data sources and variable construction. This figure illustrates the negative first-order- and positive second-order- partial effects of agricultural rainfall on civil war onset (Panels a and b) and incidence (Panels c and d), respectively. Panel a (b) plots the residuals of civil war onset risk from a regression on quadratic (linear) agricultural rainfall, country fixed-effects, year fixed-effects, country-specific linear time trends, and the full set of agricultural temperature bins against the residuals of linear (quadratic) agricultural rainfall on quadratic (linear) agricultural rainfall and same set of controls. Panel c (d) plots the residuals of civil war incidence risk from a regression on quadratic (linear) agricultural rainfall, lagged civil war incidence, country fixed-effects, year fixed-effects, country-specific linear time trends, and the full set of agricultural temperature bins against the residuals of linear (quadratic) agricultural rainfall on quadratic (linear) agricultural rainfall and same set of controls.

Table A2.1: Agricultural production and rainfall in Africa 1981-2013

	Agricultural weather				Aggregate weather			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rain	0.917	5.173	0.869	5.091	0.272	4.310	0.154	4.089
	(0.540)	(1.172)	(0.520)	(1.093)	(0.393)	(0.965)	(0.393)	(0.969)
$rain^2$		-0.201		-0.195		-0.152		-0.146
		(0.043)		(0.040)		(0.030)		(0.030)
l.rain	-0.438	1.109	-0.556	0.877	-0.110	2.529	-0.080	2.474
	(0.523)	(1.076)	(0.516)	(1.108)	(0.444)	(0.980)	(0.412)	(0.935)
$l.rain^2$		-0.066		-0.059		-0.093		-0.090
		(0.058)		(0.060)		(0.034)		(0.033)
Observations	1,648	1,648	1,648	1,648	1,648	1,648	1,648	1,648
Adjusted R-squared	0.823	0.829	0.827	0.832	0.819	0.826	0.826	0.832
Adjusted R-squared (p)	-0.055	-0.022	-0.031	-0.002	-0.076	-0.039	-0.035	-0.002
Mean temperature & lag	Y	Y	N	N	Y	Y	N	N
Temp. bins & lags	N	N	Y	Y	N	N	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. The dependent variable is the crops gross production index. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Robust standard errors are clustered at the country level and are presented in parentheses. The adjusted R-squared (p) is the adjusted R-squared from regressions where country fixed-effects, year fixed-effects, and country-specific linear time trends have been partialled out from all variables. Agricultural variables summarize information during the growing seasons and over agricultural land. Aggregate weather variables summarize temporally and spatially disaggregated weather data over the entire calendar year and the totality of a country's territory.

Table A2.2: The effect of aggregate rainfall on civil war and civil conflict onset risk

Dep. variable	1981-1999				1981-2013			
	Civil war		Civil conflict		Civil war		Civil conflict	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rain	0.004		-0.002		0.001		-0.001	
	(0.005)		(0.007)		(0.004)		(0.005)	
l.rain	-0.005		-0.009		-0.003		-0.009	
	(0.005)		(0.009)		(0.004)		(0.006)	
rain growth		0.015		0.023		0.001		0.038
		(0.020)		(0.037)		(0.018)		(0.032)
l.rain growth		0.004		-0.011		0.014		-0.048
		(0.028)		(0.061)		(0.013)		(0.035)
Observations	662	661	575	574	1,227	1,226	1,031	1,030
Mean temperature & lag	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Columns 1-4 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 5-8 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013. Robust standard errors are clustered at the country level and are presented in parentheses. Aggregate weather variables summarize temporally and spatially disaggregated weather data over the entire calendar year and the totality of a country's territory.

Table A2.3: The effect of aggregate rainfall on civil war and civil conflict incidence risk (no lagged dependent variable)

Dep. variable	1981-1999				1981-2013			
	Civil war		Civil conflict		Civil war		Civil conflict	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rain	-0.002		0.003		-0.000		0.007	
	(0.006)		(0.009)		(0.006)		(0.010)	
l.rain	-0.007		-0.004		-0.004		-0.005	
	(0.007)		(0.010)		(0.007)		(0.012)	
rain growth		-0.004		-0.011		-0.006		-0.010
		(0.027)		(0.060)		(0.025)		(0.046)
l.rain growth		-0.019		-0.001		0.001		-0.041
		(0.036)		(0.045)		(0.012)		(0.038)
Observations	743	742	743	742	1,343	1,342	1,343	1,342
Mean temperature & lag	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Columns 1-4 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 5-8 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013. Robust standard errors are clustered at the country level and are presented in parentheses. Aggregate weather variables summarize temporally and spatially disaggregated weather data over the entire calendar year and the totality of a country's territory.

Table A2.4: The effect of aggregate rainfall on civil war and civil conflict incidence risk (including lagged dependent variable)

Dep. variable	1981-1999				1981-2013			
	Civil war		Civil conflict		Civil war		Civil conflict	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rain	-0.002		0.003		-0.000		0.005	
	(0.005)		(0.008)		(0.005)		(0.007)	
l.rain	-0.008		-0.004		-0.005		-0.008	
	(0.007)		(0.010)		(0.005)		(0.008)	
rain growth		-0.015		-0.009		-0.011		-0.008
		(0.028)		(0.061)		(0.023)		(0.047)
l.rain growth		-0.026		-0.001		-0.002		-0.049
		(0.040)		(0.046)		(0.015)		(0.035)
l.dep variable	0.187	0.186	0.080	0.080	0.343	0.343	0.378	0.378
	(0.077)	(0.077)	(0.075)	(0.076)	(0.070)	(0.071)	(0.058)	(0.058)
Observations	743	742	743	742	1,343	1,342	1,343	1,342
Mean temperature & lag	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Columns 1-4 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 5-8 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013. Robust standard errors are clustered at the country level and are presented in parentheses. Aggregate weather variables summarize temporally and spatially disaggregated weather data over the entire calendar year and the totality of a country's territory.

Table A2.5: The quantitative effect of agricultural rainfall on civil war onset and incidence risk when moving from/to different percentiles of the distribution of agricultural rainfall

	Africa 1981-2013					SSA 1981-1999					SSA 1981-2013				
Panel A: Onset															
from\to	10	25	50	75	90	10	25	50	75	90	10	25	50	75	90
10	0	-3.18	-5.40	-6.01	-4.41	0	-3.20	-6.29	-6.38	-3.52	0	-2.44	-4.77	-5.54	-4.22
25	3.18	0	-2.22	-2.83	-1.23	3.20	0	-3.09	-3.18	-0.32	2.44	0	-2.33	-3.10	-1.78
50	5.40	2.22	0	-0.61	0.98	6.29	3.09	0	-0.09	2.77	4.77	2.33	0	-0.77	0.55
75	6.01	2.83	0.61	0	1.60	6.38	3.18	0.09	0	2.86	5.54	3.10	0.77	0	1.32
90	4.41	1.23	-0.98	-1.60	0	3.52	0.32	-2.77	-2.86	0	4.22	1.78	-0.55	-1.32	0
Panel B: Incidence															
from\to	10	25	50	75	90	10	25	50	75	90	10	25	50	75	90
10	0	-3.25	-5.55	-6.32	-4.85	0	-3.65	-7.57	-9.03	-7.32	0	-2.56	-5.00	-5.77	-4.35
25	3.25	0	-2.31	-3.07	-1.61	3.65	0	-3.93	-5.38	-3.67	2.56	0	-2.44	-3.21	-1.79
50	5.55	2.31	0	-0.76	0.70	7.57	3.93	0	-1.45	0.25	5.00	2.44	0	-0.77	0.65
75	6.32	3.07	0.76	0	1.46	9.03	5.38	1.45	0	1.71	5.77	3.21	0.77	0	1.43
90	4.85	1.61	-0.70	-1.46	0	7.32	3.67	-0.25	-1.71	0	4.35	1.79	-0.65	-1.43	0

Notes: Author's calculations, see the Data section for details on data sources and variable construction. The table shows the change in civil war onset and incidence risk when when moving from(rows)/to(columns) different percentiles of the distribution of agricultural rainfall (in the respective sample). Results are based on the estimates from the quadratic specifications in Table 2.3. The unconditional probabilties of civil war onset (incidence) risk for Africa 1981-2013, SSA 1981-1999, and SSA 1981-2013 are 2.21% (7.28%), 3.02% (11.04%), and 2.61% (8.48%), respectively.

Table A2.6: The quantitative effect of agricultural rainfall on civil conflict onset and incidence risk when moving from/to different percentiles of the distribution of agricultural rainfall

	Africa 1981-2013					SSA 1981-1999					SSA 1981-2013				
Panel A: Onset															
from\to															
10	0	-0.37	-1.02	-2.40	-3.79	0	-0.44	-1.24	-2.45	-3.43	0	-1.28	-3.07	-5.48	-7.06
25	0.37	0	-0.65	-2.03	-3.42	0.44	0	-0.79	-2.00	-2.99	1.28	0	-1.79	-4.20	-5.79
50	1.02	0.65	0	-1.38	-2.77	1.24	0.79	0	-1.21	-2.20	3.07	1.79	0	-2.41	-4.00
75	2.40	2.03	1.38	0	-1.39	2.45	2.00	1.21	0	-0.99	5.48	4.20	2.41	0	-1.59
90	3.79	3.42	2.77	1.39	0	3.43	2.99	2.20	0.99	0	7.06	5.79	4.00	1.59	0
Panel B: Incidence															
from\to															
10	0	-4.71	-8.55	-11.25	-11.06	0	-4.16	-9.59	-14.40	-16.10	0	-4.09	-8.58	-11.88	-11.94
25	4.71	0	-3.83	-6.54	-6.35	4.16	0	-5.43	-10.25	-11.94	4.09	0	-4.49	-7.78	-7.85
50	8.55	3.83	0	-2.71	-2.52	9.59	5.43	0	-4.81	-6.50	8.58	4.49	0	-3.30	-3.36
75	11.25	6.54	2.71	0	0.19	14.40	10.25	4.81	0	-1.69	11.88	7.78	3.30	0	-0.06
90	11.06	6.35	2.52	-0.19	0	16.10	11.94	6.50	1.69	0	11.94	7.85	3.36	0.06	0

Notes: Author's calculations, see the Data section for details on data sources and variable construction. The table shows the change in civil conflict onset and incidence risk when when moving from(rows)/to(columns) different percentiles of the distribution of agricultural rainfall (in the respective sample). Results are based on the estimates from the quadratic specifications in Table 2.4. The unconditional probabilties of civil conflict onset (incidence) risk for Africa 1981-2013, SSA 1981-1999, and SSA 1981-2013 are 5.75% (21.54%), 6.26% (23.42%), and 6.49% (23.36%), respectively.

Table A2.7: The effect of agricultural rainfall on civil war onset and incidence risk (controlling for mean agricultural temperature)

	Africa 1981-2013			SSA 1981-1999			SSA 1981-2013		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Civil war onset									
rain	-0.023	0.002		-0.036	0.006		-0.028	0.003	
	(0.010)	(0.004)		(0.019)	(0.006)		(0.013)	(0.005)	
$rain^2$	0.001			0.002			0.001		
	(0.001)			(0.001)			(0.001)		
rain growth			0.011			0.020			0.019
			(0.009)			(0.016)			(0.011)
Observations	1,538	1,538	1,536	662	662	662	1,228	1,228	1,227
Panel B: Civil war incidence									
rain	-0.021	0.002		-0.040	0.001		-0.028	0.003	
	(0.011)	(0.004)		(0.022)	(0.007)		(0.014)	(0.006)	
$rain^2$	0.001			0.002			0.001		
	(0.001)			(0.001)			(0.001)		
rain growth			0.016			0.003			0.018
			(0.009)			(0.018)			(0.012)
lagged dep. variable	0.370	0.372	0.372	0.186	0.186	0.185	0.342	0.344	0.344
	(0.073)	(0.073)	(0.072)	(0.075)	(0.077)	(0.077)	(0.071)	(0.070)	(0.070)
Observations	1,660	1,660	1,660	743	743	743	1,343	1,343	1,343
Mean temperature	Y	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Robust standard errors are clustered at the country level and are presented in parentheses. Agricultural variables summarize information during the growing seasons and over agricultural land. Columns 4-6 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 7-9 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013.

Table A2.8: The effect of agricultural rainfall on civil conflict onset and incidence (controlling for mean agricultural temperature)

	Africa 1981-2013			SSA 1981-1999			SSA 1981-2013		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Civil conflict onset									
rain	0.000	-0.004		-0.002	-0.004		-0.007	-0.007	
	(0.012)	(0.005)		(0.017)	(0.007)		(0.015)	(0.006)	
rain^2	-0.000			-0.000			0.000		
	(0.000)			(0.001)			(0.001)		
rain growth			-0.012			0.021			-0.019
			(0.022)			(0.029)			(0.031)
Observations	1,305	1,305	1,303	575	575	575	1,032	1,032	1,031
Panel B: Civil conflict incidence									
rain	-0.031	-0.004		-0.035	-0.010		-0.042	-0.004	
	(0.016)	(0.008)		(0.024)	(0.010)		(0.020)	(0.010)	
rain^2	0.001			0.001			0.002		
	(0.001)			(0.001)			(0.001)		
rain growth			-0.028			-0.041			-0.044
			(0.021)			(0.034)			(0.029)
lagged dep. variable	0.385	0.388	0.388	0.081	0.082	0.083	0.375	0.379	0.380
	(0.055)	(0.057)	(0.057)	(0.075)	(0.076)	(0.076)	(0.056)	(0.059)	(0.059)
Observations	1,660	1,660	1,660	743	743	743	1,343	1,343	1,343
Mean temperature	Y	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Author's calculations, see the Data section for details on data sources and variable construction. Estimation method is OLS. All regressions include country fixed-effects, year fixed-effects, and country-specific linear time trends. Robust standard errors are clustered at the country level and are presented in parentheses. Agricultural variables summarize information during the growing seasons and over agricultural land. Columns 4-6 correspond to the sample used in Miguel et al. (2004) with the only difference being that these regressions treat rainfall for Namibia in 1989 as a mising value –the country obtained its independence only in 1990. Columns 7-9 correspond to the same set of countries as in Miguel et al. (2004), but extend the sample up to 2013.

Chapter 3

KNOWING (YOUR) WELL? AN EVALUATION OF ALTERNATIVE STRATEGIES FOR SELLING WELL-WATER ARSENIC TESTS IN SONARGAON, BANGLADESH

Joint with Kazi Matin Ahmed, Alessandro Tarozzi, and Alexander van Geen

3.1 Introduction

Widespread poor health stands out as one of the most unfortunate and common features of life in less developed countries (LDCs). Several factors contribute to the persistence of the problem, ranging from the poor availability and high cost of good quality health care, to the insufficient investment in prevention, to the frequent reliance on ineffective and sometimes unnecessarily expensive treatments, see Dupas (2012), Dupas and Miguel (2016), and Tarozzi (2016) for recent reviews. Information campaigns on health risks are sometimes seen as an appealing policy tool in environmental and other health policy, because they can be relatively inexpensive to run when compared to other policy options. On the other

hand, Governments in LDCs may lack the resources or the political will to carry out even such relatively simple information campaigns. In addition, we frequently observe that information alone is not sufficient to promote positive changes in behavior.

In this paper, we describe the results of a randomized controlled trial (RCT) carried out in Sonargaon, Bangladesh, to study demand for information on the quality of drinking water, and the adoption of risk-mitigating behavior conditional on the information received. Despite much progress in numerous health indicators (Chowdhury et al. 2013), Bangladesh remains in the midst of an extremely severe health emergency due to the widespread presence of low-dose, naturally occurring arsenic (As) in shallow aquifers, see Ahmed et al. (2006), Johnston et al. (2014), and Pfaff et al. (2017). The problem, due to the particular hydro-geological conditions present in large parts of the country, is compounded by millions of households in rural areas relying on water from privately owned, un-regulated shallow tubewells for drinking. Using nationwide data from 2009, Flanagan et al. (2012) estimated that, in a country of more than 150 million people, about 20 million were likely exposed to As levels above the official Bangladesh standard of 50ppb, while almost one third of the population was likely exposed to levels above the significantly lower threshold of 10ppb adopted by the World Health Organization (WHO).

The most visible health consequences of chronic exposure to As from drinking tubewell water in South Asia, such as cancerous skin lesions and loss of limb, were recognized in the state of West Bengal, India in the mid-1980s (Smith et al. 2000). It has since been shown on the basis of long-term studies in neighboring Bangladesh that As exposure increases mortality due to cardiovascular disease and several forms of cancer, and may inhibit intellectual development in children (Wasserman et al. 2007; Argos et al. 2010; Rahman et al. 2010; Chen et al. 2011). These health effects are accompanied by significant economic impacts. Pitt et al. (2015) recently estimated that a nine percent reduction in household income is associated with each earner in the household exposed to arsenic in Bangladesh, while Flanagan et al. (2012) calculated, in addition, that the estimated arsenic-related mortality rate of 1 in every 18 adult deaths represents an additional economic burden of USD13 billion in lost productivity alone over the next 20 years.

Piped water from regulated and monitored suppliers is likely the best policy answer, but such solution will require immense investment in infrastructure, and because it will likely remain unfeasible for years to come, identifying short-term mitigation strategies remains essential. The consensus view now is that household-level water treatment, dug wells, and rain-water harvesting are not viable alternatives for lowering As exposure because of the cost and logistics of maintaining such systems in rural South Asia (Ahmed et al. 2006; Howard et al. 2006; Sanchez et al. 2016). In contrast, and despite being the main source of As exposure, tubewells also offer an effective way of providing safe drinking water to the rural population of Bangladesh. With the exception of the most severely affected areas of Bangladesh, the spatial distribution of high- and low-As wells is highly mixed, even over small distances. At the same time, whether a well is contaminated with As or not rarely changes over time (van Geen et al. 2007; McArthur et al. 2010). Therefore, exposure among users of As-contaminated wells can often be avoided by switching to a nearby safe well, be it a shallow private well or a deeper —hence usually safer— community well (van Geen et al. 2002; van Geen et al. 2003). Previous studies in Bangladesh have documented switching rates from an unsafe to a safe well after testing between one-third and three-quarters, depending on how much effort was put into persuading them to do so (Opar et al. 2007; Chen et al. 2007; Madajewicz et al. 2007; Bennear et al. 2013; Balasubramanya et al. 2014; George et al. 2012; Inauen et al. 2014; Pfaff et al. 2017).

Well-sharing as an effective risk mitigation strategy, however, relies on knowledge about the safety of both one's own water source and that of neighbors. Between 1999 and 2005, the Bangladesh Department of Public Health Engineering (DPHE), with the assistance of the World Bank, DANIDA, UNICEF, and few other NGOs, conducted a blanket testing campaign that tested close to 5 million wells, and identified them as 'safe' or 'unsafe' —according to the Bangladesh standard of 50ppb— by painting the well spout with green or red paint, respectively.[1] Unfortunately, such a testing campaign, coordinated through

[1] In some villages, testing was conducted using low-precision kits and thus only wells with no detectable arsenic concentration were classified as safe and wells with concentrations above 100ppb were classified as unsafe.

the Bangladesh Arsenic Mitigation and Water Supply Program (BAMWSP), has not been repeated, and may not be expected in the foreseeable future. Meanwhile, millions of new wells have sprouted in the country, and in most cases the well owner does not know the As level of the water. There are a few commercial laboratories in Dhaka with the capability to test wells for As, but few rural households are aware of these services. The cost of well testing is greatly reduced and the logistics are greatly simplified by the use of field kits, which have become increasingly reliable and easy to use (George et al. 2012; van Geen et al. 2014), but even these tests are rarely available, and the willingness to pay for them is not known.

The first objective of this paper was to make progress in understanding the sustainability of a market for As field tests, by offering them at a price of BDT45 (about USD0.60) in 49 randomly selected villages in Sonargaon sub-district. Such amount was close to the price of one kg of rice in Dhaka, and was gauged to be high enough to cover for the salary of the surveyors hired for the project. The results contribute to a growing literature that studies demand for health-protecting technologies in developing countries. Prior research in different countries have documented very low uptake of a variety of such products, ranging from insecticide-treated nets (Cohen and Dupas 2010, Dupas 2014, Tarozzi et al. 2014), to de-worming drugs (Kremer and Miguel 2007) and water-disinfectants (Ashraf et al. 2010). Our work complements this literature by looking at demand for health-related *information* that can be exploited by households to devise risk mitigation strategies. In Bihar, India, another location with a severe As problem, Barnwal et al. (2016) estimated very high elasticity of demand, with uptake of tests falling from 69% to 22% of households when the price increased from INR10 to INR50, where the latter was about equivalent to daily per capita income. They also found that sales repeated two years later were substantially higher, suggesting that experience with the tests were an important factor. In our study area, and despite low prices and widespread awareness about the As problem coupled with little information about actual safety of drinking water, we found that only about one in five households purchased the test. Conditional on learning about the unsafe status of one's drinking water, we estimate that about 36% of households stated that they switched to a different source at the time of our return visit.

The second objective was to determine whether the existence of informal, within-village solidarity networks could be leveraged to increase demand for testing and especially well-sharing. A large literature documents the importance of village networks to cope with shocks, including health shocks, see Fafchamps (2011) for a review. In our case, and in an additional subset of 47 villages, we study demand for As tests —at the same price of BDT45— to self-formed *groups* of up to ten individuals, where group members were asked to sign an informal agreement according to which those with safe wells would share their well with others whose well water was found to be unsafe. The agreement was not binding legally, but our prior was that it would increase rates of switching from unsafe sources through two mechanisms. First, by making sharing more likely through a form of soft-commitment and, second, by facilitating the spread of information about the safety of wells, thereby increasing the salience of safe options within the village. Although we find that demand for tests remained about the same as with sales that did not involve informal water-sharing agreements, we estimate that switching rates were about 25 percentage points higher. Unfortunately, although the result is striking, data limitations do not allow us to probe conclusively whether the increase in switching was driven by sharing within groups of buyers.

In an additional subset of 15 villages, individuals who purchased a test at the usual price were also given a metal placard of a color depending on the As level: blue for As below 10ppb, green if between 10 and 50, and red if 'unsafe', that is, above the Government threshold of 50ppb. The placard was then attached to the well spout. Similar metal placards have been used routinely in testing campaigns, both as a reminder about the safety of the well water, and as a means to facilitate the spread of information about which wells are safe within a village. However, the cost of the placards (about BDT80) is high enough to increase significantly the total cost of testing campaigns. It was thus important to determine whether they made any difference relative to the alternative solution (adopted in the two experimental arms described earlier) of informing the household via a simple and inexpensive laminated card to be kept in the house, with the indication of the test result. While we find that demand for testing was barely affected by the concurrent offer of a metal placard, the switching rates were almost doubled (66 vs. 36 percent) relative to those recorded when test results were communicated

via a simple laminated card.

Additionally, we describe the findings from two other experimental arms, although the very small sample size (six villages per group) means that estimates are imprecise. Anecdotal evidence from the fieldwork related to earlier work suggested that willingness to pay for information may be reduced if households may receive 'bad news', especially when these are not accompanied by clear and effective mitigation strategies. For instance, in Malawi Thornton (2008) found that the fraction of individuals tested for HIV, at a time when anti-retroviral therapies were not widely available yet, testing prevalence remained very low if incentives were not provided. In six villages, we thus offered the As tests at the same price but with the agreement that the price would be paid only in case of 'good news'. This more than doubled demand relative to the simple offer (48 vs. 20 percent), although of course the comparison between the two sale mechanisms is confounded by the effective reduction in the 'perceived price' induced by the 'pay-for-good-news' contract, at least for everyone except the rare individuals who felt completely sure that their water was safe. The difference in demand is so large to be significant even at the 1% level, despite the small number of villages in this experimental arm. Even when we doubled the price at BDT90 in another group of six villages, demand remained substantially larger (34 percent).

Finally, we assess the cost-effectiveness our selling schemes by assuming a fixed budget and examining the income gains that would arise by switching from an unsafe well to safe one —relying on estimates of the effect of reduced exposure to arsenic on income by Pitt et al. (2015). We first show that charging a small fee for a test allows to more than double the amount of wells that can be tested, relative to blanket testing, and is over three times as cost-effective —due to higher switching rates. Relative to this simple offer, promoting informal water sharing agreements and fees-for-good-news (at BDT45) are 43% and 71% more cost-effective, respectively. Despite the fact that metal placards have been widely used in testing campaign and that we estimate that they led higher switching rates, we find that their high cost renders them less cost-effective than our alternative of providing cheap laminated cards with test results.

The paper proceeds as follows. In the next section we provide some additional background information on the extent of the arsenic problem in our study area and

describe our experimental design. In Section 3.3, we describe our data collection protocol, present selected summary statistics, and show that by chance the means of a number of covariates were not balanced at baseline, highlighting the importance of controlling for baseline characteristics in our estimates (the adjusted and unadjusted estimates remain overall substantively similar). In Section 3.4 we discuss the results from our interventions. Next, in Section 3.5 we use our estimates together with evidence from earlier literature to gauge the cost-effectiveness of our interventions, also when compared to blanket testing campaigns as carried out in several areas in the past. Finally, we conclude in Section 3.6.

3.2 Study design

This study was carried out in Sonargaon, a sub-administrative unit (or *upazila*) of Narayanganj district, located approximately 25 kilometers south-east of the capital Dhaka. According to the 2011 Census of Bangladesh, Sonargaon had a population of about 400,000, and administrative records at the time of the study listed a total of 365 villages, in a 171 squared kilometers territory. Sonargaon is located in a part of the country where arsenic contamination of shallow tubewell water is widespread. According to a blanket testing conducted between June 1999 and June 2000, under the supervision of the Bangladesh Arsenic Mitigation and Water Supply Program (BAMWSP), in about 90 percent of villages 40 percent or more of tubewell water had arsenic levels above the Bangladesh standard of 50ppb, while the median proportion of unsafe wells was a high 86% (Chowdhury et al., 2000).[2]

For this paper, we first selected all 128 villages in Sonargaon with more than 10 wells and with a share of unsafe wells between 40 and 90 percent, according to the BAMWSP blanket testing campaign. In each study village, surveyors would walk across the whole area identifying all wells, regardless of their ownership status or of evidence of having been tested before. For privately owned wells, the

[2]Blanket testing in Sonargaon was carried out by the Bangladesh Rural Advancement Committee (BRAC), a partner NGO of BAMWSP. A total of 25,048 tubewells were tested for arsenic, however, the test kits utilized did not allow for reliable readings of arsenic concentration above 0ppb and below 100ppb.

surveyors would identify which household owned the well, while for public wells, the surveyors would identify which household was the main well caretaker or user. Surveyors would then conduct a home visit, where they would explain the risk of consuming arsenic-contaminated tubewell water to an adult —typically the most senior woman— and then offer to test the well for a fee. Additionally, surveyors recorded the approximate GPS location of all wells, regardless of whether the owner bought a test or not. Finally, surveyors would administer a short household questionnaire and distribute color-coded laminated cards with identification numbers. All cards included the following informative messages: (i) that arsenicosis is not a communicable disease, (ii) that arsenic cannot be removed by boiling water, (iii) that testing tubewell water for arsenic is necessary, and (iv) that the Bangladeshi safety standard for arsenic concentration in water is 50ppb. Black cards were given to households who did not buy a test. These further included a message encouraging households to test their well water for arsenic to avoid health-risks associated with the consumption of arsenic-contaminated water.

When a test was purchased, tubewell water was tested using Arsenic Econo-Quick (EQ) test kits, which have been shown to be reliable when used in the field, and can deliver results within ten minutes, see George et al. (2012) for details. The tests cost USD0.35 for volume purchases, although the total cost per test was estimated to be about USD2.3 per test, for a testing campaign that also covered the costs of trained personnel and metal placards to be attached to the tubewell spouts (Barnwal et al., 2016). Owners of wells with an arsenic concentration below 10ppb were given a blue card, those whose well had a concentration between 10ppb and 50ppb were given a green card, and those whose well had a concentration above 50ppb were given a red card. These cards also displayed the well's test result (in ppb), rounded to the nearest integer in the following sequence $\{0, 10, 25, 50, 100, 200, 300, 500, 1000\}$. Owners of unsafe wells (arsenic concentration above 50ppb) were encouraged face-to-face to switch to a safe (blue or green) well, while owners of wells with concentrations below 10ppb were encouraged to share their well water with their neighbors. Green card holders were both encouraged to share their water and to switch to a safer (blue) well, if possible.

Our experimental variation comes from differences in selling schemes for arsenic tubewell water tests across villages. The experimental arms were as follows.

In a first group of villages, which we term group A, surveyors offered to test tubewell water for a fee of BDT45 (about USD0.60). This fee was expected to cover the salary of the testers and their supervisor. In particular, of the BDT45 charged per test, testers kept BDT30 to cover their transportation expenses and salary, and handed over the remaining BDT15 to their supervisor. The price was determined assuming that a field worker would test about 15 wells/day for 20 days/month, leading to a monthly salary of BDT9,000 (USD115/month), which is roughly what village-health workers were paid for blanket testing in the neighboring Araihazar in 2012-2013 (van Geen et al. 2014). According to the same scenario, the supervisor of 10-15 workers would earn BDT45,000-67,500 (USD578-867), a range that spans what he earned while supervising the testing in Araihazar in 2012-2013. Across all experimental arms, the cost of the field kits (USD0.35/test) was covered by the project.

In a second group of villages (B), groups of at most 10 neighbors were gathered together and asked (i) if they wanted to test their water for a fee of BDT45 each and (ii) if they wanted to sign an agreement to share any safe water among the signators —before conducting any testing. The agreement had no legal standing, and was meant to serve as a soft commitment device. Anecdotal information form the field indicated that people were afraid of formally committing to share their well, but that verbal agreements took place instead. All neighbors that committed to sharing their tubewell water were allowed to see the test results from all other neighbors in the group. Both well owners that committed to share their well (formally or verbally) and those who did not could purchase the test.

In a third group of villages (C), household were offered to test tubewell water for BDT45. But, for households that bought the test, a color-coded stainless steel placard was attached to the well's pump-head. Placards displayed both in text and color whether the arsenic concentration was below 10ppb (blue), between 10ppb and 50ppb (green), or above 50ppb (red). Further, as shown in Figure 3.1, they displayed two hands holding drinking cups, one hand holding a drinking cup, and a large cross over a hand holding a drinking cup, respectively. The split of the test fee between the tester and supervisor in groups B and C was the same as in group A. The project further gave a bonus of BDT12 to testers per every household that signed the well-sharing agreement or that verbally committed to sharing their well

in group B. It also covered the cost of metal placards (BDT80/placard) in group C.

In addition, our experimental design included two experimental arms (D and E), with few villages each, where payment for a test was only required if the water tested safe. In groups D and E, the fees were BDT45 and BDT90, respectively. Further, when a sale was made and the well tested safe, the tester would keep BDT30 and BDT60, respectively, and hand over the difference to the supervisor. The intervention's protocol stipulated that sales offers should be made in villages assigned to group E first, in order to avoid having to increase the price in some villages later in the project, which could cause discontent.

All 128 villages considered in this study were stratified by union (a larger administrative unit) and by high or low arsenic contamination —which was determined by whether the share of unsafe wells in the BAMWSP testing campaign was below or above the median.[3] Then, using a pseudo-random number generator, 50 villages were assigned to treatment arms A and B each, and 16 villages were assigned to treatment arm C. The remaining 12 villages were stratified by high or low arsenic status only, with half of them being randomly assigned to treatment arm D and the other half to E. There were two deviations from the data collection protocol. First, in the programming of the mobile application used for data collection, 27 villages were assigned by mistake a treatment different from the original one. The partial re-assignment of treatments was due to a data-entry error and is, in principle, unrelated to village characteristics. Second, in five cases, surveyors were unable to differentiate a village from the one adjacent to it. While we have data from households in these five villages and the ones adjacent to them, we can only distinguish pairs. Both villages in each pair received the same treatment. The spatial distribution of villages in our intervention is displayed in Figure 3.2. As expected, the randomization —with the caveats described above— led to large spatial variation in treatments.

[3]Unions are the third smallest administrative unit, after mouzas (groups of 2 to 3 villages), and villages. Our 128 villages come from 9 unions.

3.3 Data

Between December 2015 and June 2016, surveyors approached all households owning a well and offered them an arsenic test under various selling schemes. Concomitantly, surveyors administered household (baseline) questionnaires and recorded information on sales and test results —for households that bought a test. Additionally, GPS coordinates of all well owning households were recorded, and surveyors further gathered information on whether there were any visible labels attached to a well indicating whether it is safe or unsafe with respect to arsenic. A total of 14,052 households were approached for a test offer and 14,046 completed, at least partly, the baseline questionnaire. We did not collect any data from households who did not own a well.

Follow-up surveys were conducted between August 2016 and January 2017.[4] The average time elapsed between surveys was about eight months, and 86.4% of the households had their follow-up interview between seven and nine months after the baseline interview. Both the baseline and the follow-up surveys recorded whether a household used their own well for cooking and drinking. Throughout this paper, unless otherwise noted, we restrict out analysis to households that use their own well at baseline, which is the set of households for which we can ascertain whether well-switching happened.[5] In our sample, 12,734 well owners (90.7%) used their own well for cooking and drinking at baseline. Of these, 787 households (6.2%) had repeated household IDs due to errors in ID assignment at baseline. Additionally, there were 1,119 households interviewed at baseline (8.8%) for which we could not identify survey responses, either due to attrition or error in ID tracking. These household were excluded from the well-switching analysis.

We present selected summary statistics measured at baseline in Table 3.1. Row 1 shows that over 90% of well owners used their own well for cooking and drinking, and this is the sample for which we summarize all other variables. 85% of households have a male household head, 27% of the household heads are wage-

[4]Unlike the baseline survey, where the wages of surveyors and the supervisor were covered mainly from test fees, the cost of the follow-up survey was paid for by the project.

[5]For households that do not drink form their own well at baseline, we cannot systematically record well-switching.

workers, 42% are self-employed, with the remaining largely occupied in domestic activities. Household heads have low levels of educational attainment on average, with the majority having only primary schooling or less. They also tend to be of low socio-economic status, as only 17% of the houses surveyed have concrete roofs (an indicator of wealth), with the rest having tin or mud roofs (in rare cases). Further, households in Sonargaon are small in size, with an average of 3.6 members, and 1.5 children.

Despite the blanket test by BAMWSP between 1999 and 2000, a large majority of households (76%) still does not know whether their well is safe or unsafe with respect to arsenic. We find that 7% of households report having an unsafe well, and the remaining 17% report having a safe well.[6] However, only 0.3% of the households had a well visibly labeled as safe. Further, using geographic information systems (GIS), we estimate that the average well owner had about 0.02 wells labeled as safe within 50m, out of an average of nearly 12 wells. Further, the average well owner in our study area lives in a village where 74.7% of the wells tested by BAMWSP were unsafe with respect to arsenic.

About 98.6% of all wells are privately owned. The wells in our sample are on average relatively shallow (176 feet) and relatively new (9.11 years) —which suggests that a significant proportion of wells were installed after the BAMWSP blanket testing. The average installation cost of wells in our sample is BDT7,474. This implies that the cost of BDT45 for the test charged in treatments A, B, C, and D (in case of a "safe" result), represents about 0.6% of the installation cost.

Column 9 of Table 3.1 shows the p-value for the null hypothesis of equality of means across treatment arms. While we cannot reject the null hypothesis for 12 of the 19 household characteristics considered at the 95% confidence level, we can reject it for the remaining seven. These include likely important characteristics such as the education of the household head, roof quality, household size, number of children, knowing the safety status of one's well with respect to arsenic, and whether one's well is visibly labeled as safe with respect to arsenic. For instance,

[6]Of the households who reported to know the safety status of their well with respect to arsenic, about 2% had their well last tested at the time of the BAMWSP blanket testing campaign and 93% had their well last tested after the BAMWSP campaign. The remaining 5% plausibly did not have their well previously tested and were likely reporting subjective beliefs about their well safety status.

while the 21% of household heads in out study area have no schooling, this number drops to about 6% in treatment C. Because these characteristics may affect test purchase decisions and switching behavior upon receiving news of having an unsafe well, we will control for all relevant household characteristics available. The bottom panel of Table 3.1 shows that the percentage of households that could not be matched from the baseline to the follow-up survey is not different across treatments, neither if one only considers households that were unmatched because of repeated IDs nor if one only considers those with unique IDs, but that were lost at follow-up.

3.4 Results

In this section, we first estimate the effect of our intervention's test selling schemes on uptake rates. In particular, we assess whether selling modalities that promote well sharing through informal agreements (B), public posting of results through placards (C), or whose prices depend on the test outcome (D and E), had different uptake rates compared to a simple test-for-fee scheme (A).

We hypothesize that the expected benefit from buying a test is increasing in the following key elements: (i) the net present value of consuming safe water, (ii) the probability that one's well is unsafe, and (iii) the probability that one will be able to switch to a safe well, conditional on the test results indicating that one's well is contaminated. Because treatments B and C likely reduce the cost of searching for a safe well upon finding out that one's well is unsafe, we expected demand to be higher for these treatments relative to A. The expected price to be paid by any well owner in treatment D is equal to BDT45 times the (subjective) probability of the well being safe. Because this product is always smaller or equal than the price charged in treatment A (BDT45), we expected that this would lead to higher uptake relative to A. Since the price charged in treatment E, when a well turns out to be safe, is BDT90 whether there is a positive average effect on test uptake will depend on the fraction of households who believe their well to be safe. Further, treatments D and E allowed households to avoid the "psychic cost" of paying for bad news (i.e., learning that one's well is unsafe), which we expected to further incentivize demand.

We then examine whether well-switching rates, among households that found that their well is unsafe, were different across experimental arms. Differential well-switching rates could arise for two reasons. First, differential test uptake rates which lead to differential number of known safe wells in the vicinity of each household —and thus switching opportunities. Second, due to treatment specific characteristics with potential direct impacts on switching, like informal commitments to share water or public provision of information about the status of tested wells —which might increase the salience of the risks associated with arsenic contamination.

3.4.1 Demand

To estimate the average treatment effect of selling schemes B, C, D, and E, relative to A, we estimate the following equation

$$buy_{hv} = \beta^B B_v + \beta^C C_v + \beta^D D_v + \beta^E E_v + \gamma X_{hv} + \delta_{ul} + \epsilon_{hv}, \tag{3.1}$$

where buy is equal to one if a household bought a test at baseline (buy), and zero otherwise, B, C, D, and E are indicator variables for the respective treatments, X is a set of predetermined household and tubewell characteristics (X), and ϵ is an error term. h and v are indices for household and village, respectively. To account for our stratified design, we further include in equation 3.1 union-by-high or low arsenic status fixed effects (δ_{ul}).

Table 3.2, column 1, reports OLS estimates of the linear probability model described in equation 3.1, without the inclusion of controls. Throughout, standard errors are presented in parenthesis and clustered at the village level. At the bottom of column 1, we report the average test uptake rate in group A, which is 22.8%. Our estimates indicate that neither incentivizing well sharing by means of an informal agreement (B) nor the use of costly metal placards to post test results (C) led to significantly higher demand.[7] However, in treatment D, where the expected price to be paid by households was always smaller or equal to the price

[7] We note that only 44% of the households who purchased a test under selling scheme B agreed to share their water ex-ante.

charged in treatment A, we find a large and significant increase in test uptake rates of 27.9 percentage points. This result is consistent with previous findings in South Asia documenting a large sensitivity of demand for arsenic tests with respect to price (Barnwal et al., 2017). In treatment E, where the expected price need not to be smaller that the fee charged in treatment A, we find a positive and marginally insignificant effect (pvalue 0.106) on test uptake of 10.2 percentage points.

Table 3.2, column 2, further includes controls for household and tubewell characteristics which might be relevant in determining demand for arsenic tests. The estimated effect of treatment D on test uptake remains both qualitatively and quantitatively the same, while the estimated effect of treatment E raises to 14.5 percentage points and becomes significant at the 99% confidence level. While we do not attach a causal interpretation to the correlations between controls and test uptake, most of these correlations have the expected sign. For instance, well owners that think that their well is safe stand to gain little from buying a test, accordingly, we find that those who do not know the status of their well or that think it is unsafe are 15 and 12 percentage points more likely to buy a test, respectively. Demand for a test is also higher among wealthier households, i.e., those with a concrete roof (6 percentage points), households whose heads are more educated (19 percentage points), and bigger households (4.5 percentage points per every additional member).

Because selling schemes D and E required payments only if the test results turned out to be safe, households sufficiently uncertain about their well's safety had disproportionately higher incentives to buy a test, relative to households with similar levels of uncertainty in treatment A. We then examine the heterogeneous effect of our treatments by whether a household reported to have a safe well, not to know their well's status, or to have an unsafe well at baseline. To do this, we interact our treatment indicators with an indicator for not knowing one's well status and an indicator for reporting to have an unsafe well. Columns 3 and 4 of Table 3.2 show that households who did not know the status of their well were about 40 percentage points more likely to buy a test in treatment D, relative to A. For households who thought (or knew) that their well was unsafe, test uptake was about 60 percentage points higher in treatment D. Columns 4 further shows that the effect of treatment E on test uptake is also heterogeneous, and that households

that reported having unsafe wells saw a significant increase in test uptake rates of 19.3 percentage points, relative to A. Column 3, which does not include controls, shows a marginally insignificant effect (p value 0.113) of treatment E on test uptake of 17.5 percentage points, for households who report having an unsafe well. The fact that a significant effect on uptake is found for households that did not know the status of their well in treatment D but not E may be explained by the fact that the higher price to be paid in case the well turned out to be safe in E, relative to D, likely disincentivized a larger set of households with intermediate levels of uncertainty.

Finally, Columns 5 and 6 of Table 3.2 examine the heterogeneous effect of our treatments across richer (i.e., with a concrete roof) and poorer (i.e., with a mud or tin roof) households. In treatment D, where the expected price is smaller or equal to the required payment in A, we would expect higher uptake among the poor who face, by definition, tighter financial constraints relative to A. If households in treatment E attribute a sufficiently high probability to their well being unsafe, we would also expected higher uptake among the poor. Our estimates show that the poor have about a 31 percentage points higher probability of buying a test in treatment D, relative to A. Similarly, the poor have a higher probability of buying a test in treatment E, relative to A, of about 15 percentage points. We find no significant effect of tretaments D or E on test uptake rates among the non-poor.

3.4.2 Responses to test results

Of the 12,734 households who used their own well for cooking and drinking and who were offered an arsenic test, 3,263 (25.6%) bought a test under one of our selling schemes. Testing campaigns, like ours, improve households' knowledge about the potential health risks they face and allow them to take preventative measures to abate these risks. In the case of arsenic, testing campaigns further facilitate preventative actions, like switching to a safer well, by revealing the distribution of safe wells. Well-switching is a relatively easy avoidance measure, in the common occurrence that the large spatial variability of arsenic, even within small areas, leads to unsafe tubewells being often at short distances from a safe one (van Geen et al., 2002). Figure 3.3 plots the spatial distribution of safe, unsafe, and

untested wells in Sonargaon and confirms that the spatial distribution of arsenic in our study area was indeed large, even within small areas. Visual inspection of this general pattern becomes more apparent by looking at each village at a time. As an illustration, Figure 3.3 also provides an magnified view of the spatial distribution of wells in the village of Pirojpur (treatment A), which further confirms the stochastic occurrence of arsenic in ground water.

Based on the test results from our intervention, we estimate that the average well owner had about 1.8 and 3.8 neighbors that bought a test within 25m and 50m, respectively. We further estimate that the average well owner had about 1.6 and 3.1 wells that tested safe for arsenic within 25m and 50m, respectively. The latter results highlight the increased well-switching opportunities brought about by testing. Importantly, interventions that lead to higher test uptake have a higher potential for increasing well switching opportunities by uncovering more safe wells. Table 3.3 presents estimates of equation 3.1, where the dependent variable has been replaced by the number of well owners that bought a test within 25m (column 1), and 50m (column 2), and the number of wells that tested safe within 25m (column 3), and within 50m (column 4). Column 1 indicates that treatments D and E increased the number of wells tested within 25m by 2.2 and 0.8, respectively, relative to treatment A. Column 2 further indicates that these differences were about twice as large when considering a wider radius of 50m. We find no significant effects for treatments B and C. We further find that treatments D and E increased the number of wells that tested safe within 25m by 1.3 and 0.6, respectively. When considering a radius of 50m, the estimated effects go up to 1.4 and 1.3 safe wells, respectively. We find no such significant effects for treatments B and C.

For well switching to be feasible, at least three conditions must be met: (i) households must be within a moderate distance from a safe well, (ii) households must know the location of safe wells, and (iii) households must be able to access safe wells. If the availability of safe wells within a moderate distance were a limiting factor for well switching in testing campaigns, then we should expect increased well-switching rates in treatments D and E. However, if the limiting factors are knowledge of the location of safe well and access to these wells, then we should expect increased switching rates in treatments B (which publicized test

results within a group of neighbors and created soft well-sharing commitment) and C (where all test results were public).

We now examine the effect of selling schemes B, C, D, and E on well switching. To that end, we construct an indicator variable ($switched$) which is equal to one if a household reported to cook and drink using water from their own well at baseline but not at follow up, and zero otherwise. Because we focus on well owners that use their own well at baseline, this variable is always well defined. We further restrict our attention to households whose well tested unsafe and who thus had incentives to switch. For completeness, we also analyze the effect of the various selling schemes on switching for households who did not buy a test, and who could have had incentives to switch to a known safe well.[8] Our analysis relates our switching variable to indicator variables for treatments B, C, D, and E, household and tubewell characteristics (X), union-by-low or high arsenic status fixed effects, and an error term (η)

$$switched_{hv} = \beta^B B_v + \beta^C C_v + \beta^D D_v + \beta^E E_v + \gamma X_{hv} + \delta_{u,l} + \eta_{hv}. \quad (3.2)$$

Table 3.4, column 1, shows that of the households whose well tested unsafe in the test-for-fee selling scheme (A), 36.7% switched wells by the time of the follow-up survey —about 8 months later on average. This rate is largely aligned with previous findings on the effect of information provision about arsenic content in tubewell water on well-switching (Opar et al. 2007; Chen et al. 2007; Madajewicz et al. 2007; Bennear et al. 2013; Balasubramanya et al. 2014; George et al. 2012; Inauen et al. 2014). Column 1 further reports estimates of the average treatment effect of the various selling schemes on well switching —without controls— for households whose well tested unsafe. Our estimates indicate that households in treatments B and C were 24.1 and 32.8 percentage points more likely to switch wells, respectively, and that these effects are significant at the 95% confidence level. However, we do not find a significant effect of treatments D or E on well switching. Column 2 adds household and tubewell controls, and the estimates of

[8] Among the 1,881 households whose well tested safe, only 11 had switched wells at follow up. We exclude these households from our switching analysis.

the effect of treatments B and C remain largely unchanged. However, we find that treatment D increases switching, relative to A, by 25.7 percentage points, and that this difference is significant at the 90% confidence level (p-value 0.066). These results show that information about the location of safe wells and access to these wells are limiting factors of switching. The results also suggest that the availability of safe wells is also a limiting factor. The lack of an effect of treatment E on switching, vis-à-vis the significant effect of treatment D, can be explained by two factors. First, the smaller effect of treatment E on test uptake and thus on the possibility to reveal safe wells. Second, a selection effect. Treatment E increased test uptake only among those who believed their well to be unsafe, while treatment D further increased uptake among those who reported not to know the status of their well. To the extent beliefs by well owners were right, treatment E would have uncovered less safe wells than treatment D, providing reduced switching opportunities (see Table 3.3).

We next explore whether the effect of our selling schemes varies across households with differential number of wells that tested safe, within a short distance. We do this as distance from a safe well has been previously documented to significantly hinder switching (Opar et al. 2007; Pfaff et al. 2017). To this purpose, we augment equation 3.2 by interacting our treatment dummies with the number of safe wells within 25m (or 50m), and control for the interaction of our treatment dummies with the number of households that bought a test within 25m (or 50m). Column 3 of Table 3.4 shows that the effect of treatments B, C, and D do not vary depending on the number of wells that tested to be safe within 25m. Our estimates indicate that the effect of treatment E on switching varies negatively with the number of safe wells nearby. However, as is shown in Column 4, none of the interaction terms are significant when adding additional controls. When considering a larger radius around each household, results for treatments B, C, and D remain qualitatively similar (columns 5 and 6). The negative interaction term of treatment E and number of safe wells is now significant (even when adding controls). The latter results must be taken with a grain of salt. While switching rates in treatments A, B, C, and D range from 37% to 72%, providing a substantive amount of variation, the switching rate in treatment E is 9%.

Columns 7 and 8 of Table 3.4 show estimates of the average treatment effect of

selling schemes B and C on switching, relative to treatment A, among households who did not buy a test. We exclude treatments D and E from this analysis as only one household, out of all non-buyers, switched wells in each of these experimental arms. While we find no significant effect of treatment B on switching among non-buyers, we find that treatment C had a small and significant effect of about 1.4%.

3.5 Cost-effectiveness

One of the objectives of this paper was to understand the sustainability of a market for arsenic tests, by offering them at a price of BDT45, which covered the wages of testers and their supervisor. Such a market —which could involve public subsidies— has the potential to help people hedge the risks associated with the consumption of arsenic-contaminated water by providing them with information about the quality of their well and revealing the location of safe water sources. In this section, we assess the cost-effectiveness of blanket testing and of our selling schemes by assuming a fixed budget and examining the income gains that would arise by switching from an unsafe well to safe one.

This exercise faces at least three limitations. First, our experimental design did not consider free provision of tests, thus, we rely on the estimate by Pfaff et al. (2017) of the effect of blanket testing on well switching from the neighboring upazila of Araihazar. Pfaff et al. (2017) examine well switching following the BAMWASP blanket testing campaign that swept Araihazar in 2003, and focus on villages where arsenic-mitigating campaigns had not been previously in place. They estimate a switching rate of 27% for households that learned that their well was unsafe, which is about 10 percentage points lower than the one we estimate in our test-for-fee selling scheme (A). The difference is likely explained by the fact that our estimates are based on households who bought a test, i.e., that have a higher willingness to pay for environmental quality information, and that may therefore be more willing to take health protective actions —like switching— upon receiving adverse news. We cannot rule out, however, that underlying preferences for clean water are different in the two study areas, and that this is driving the difference in switching rates. Second, our surveys did not collect data on labor supply, wages, or home production. We thus rely on recent

estimates by Pitt et al. (2015), who quantify the effect of reducing arsenic exposure among Bangladeshi prime-age men on earnings, to assess the gains from switching to safe water sources (USD54/year). Assuming that the value of female production is half of that of men, as Pitt et al. (2015) do, the yearly gains from switching from an unsafe well to a safe well per households are estimated to be USD81/year.[9] This estimate ignores all additional gains that might arise through improved health later in life. Third, we do not have objective data on the safety status of the wells households switched to, and assume that all households that switched wells did so to a safe one. We discuss the implications of this assumption below.

The cost of a blanket campaign like that of BAWMSP is about USD0.91/test (USD0.35 for the field kit and USD0.56 for the tester's wage and logistics). With a budget of USD1,000, one could thus test about 1,100 wells. Throughout, we will assume that 26% of all tested wells are unsafe —consistent with our findings from Sonargaon. Further, assuming that switching among those who learn that their well is unsafe happens at a rate of 27% (Pfaff et al., 2017), blanket testing would lead a total of almost 80 switchers. Assuming a discount rate of 8%, a blanket testing campaign with a USD1,000 budget, would lead to income gains of about USD67,000 over the next 20 years, rendering blanket testing highly cost-effective. Our calculations also indicate that over 98% of all switchers would have to switch to an unsafe well for blanket testing not to be cost-efficient, which is very unlikely.

While free provision maximizes switching opportunities within a given locality, charging a low fee could allow for a larger coverage —at the cost of reducing uptake among those with low willingness to pay. Given the same budget constraint considered above, the lower cost of treatment A (USD0.35/test) would lead to the testing of over 2,800 wells. Further, given our estimated effect of treatment A on switching, and the underlying incidence of arsenic, we find that about 270 households would switch wells. The implied income gains over the next 20 years are in the order of USD 230,000, which is USD163,000 larger than those estimated

[9]This estimate is based a reduction of arsenic contamination to levels observed in uncontaminated countries. Switching to a safe well in our study area would likely lead to a similar reduction of arsenic levels, as 85.4% of the wells we tested in our sample had no detectable level of arsenic in them.

for a blanket testing campaign. The larger cost per test in the latter, relative to treatment A, can explain about 86% of the difference in income gains, with the remainder being due to differential switching rates —likely due to differences in preferences for clean water between buyers and non-buyers.

We now examine whether treatment B, which paid a bonus to testers for arranging the sharing of any safe water among local groups, increased cost-effectiveness. The bonus was of about USD0.15 per every household that agreed to share their water. As only about 44% of households in this experimental arm signed the agreement, the expected bonus per test sold was of about USD0.07. Given that this treatment was more expensive than A, we calculate that a budget of USD1,000 would lead to nearly 2,400 wells being tested —less than in A. However, given the higher estimated switching rate, we calculate that this selling scheme would result in about 380 switching households and income gains in the order of USD328,000.

A common feature of other testing campaigns has been the provision of costly metal placards to post test results (USD1.37 for test kit and placard vs. 0.35 for test kit). While these allow households to identify the location of safe wells, the provision of cheap laminated cards —which was in place in treatments A, B, D, and E— could suffice if households communicate with each other. While switching rates in treatment B and C were largely similar, the cost per test in treatment C is more than triple that of B, which means that a budget of USD1,000 would only allow for the testing of about 190 wells. Our estimated switching rate implies that about 124 households would switch wells, leading to increases in income of about USD107,000.

A budget of USD1,000 would lead to about 2,800 wells being tested either under treatments D or E. However, the estimated differences in switching rates imply that while treatment D would lead to income gains of about USD394,000, the grains from treatment E would be of about USD230,000. Our calculations further indicate that over 99% of households that switch wells would have to do so to an unsafe one for any of our experimental arms not to be cost-effective.

This analysis has so far been silent about the desirability of our different selling schemes from the testers' perspective, which is relevant to the sustainability of a private market for tests. In the design of our experimental arm A, we set the price of a test so that, in expectation, the testers and their supervisor would

have a wage about the same to what they made during a blanket testing campaign in neighboring Araihazar. The design assumed that testers would be able to test about 15 wells a day, which is about the number of wells they were expected to test in Araihazar. However, low test uptake rates in treatments A, B, and C, imply that testers might have had to exert more effort to earn the same wage. This, in turn, was less of an issue in treatment D and E, even though payments to testers were only made whenever the well tested safe. In treatment D, the higher uptake rate implies that the expected revenue from a test offer was BDT11, compared to BDT6 in A. While uptake in E was lower than D, the higher fee charged when a well tested fee implies that the expected revenue from a test offer was about BDT15.

3.6 Limitations and conclusions

In this paper we study the effect of different arsenic test selling schemes on test uptake and, conditional an learning that one's well is unsafe, their effect on well switching. We first show that when selling tests for BDT45 (USD0.60) —an amount sufficient to cover the wages of testers and their supervisor— about one in five households purchased a test. Our estimates indicate that neither placards used to post test results nor promoting informal water sharing agreements increased demand. Test uptake, however, was significantly increased by test offers that required the payment of a fee only when the water tested safe. Conditional on learning about the safety of one's tubewell water, informal agreements, metal placards, and fees-for-good news (at low prices) nearly doubled the rate at which households stopped using their contaminated well.

Our intervention has important implications for the scale-up of future testing campaigns. First, we show that for a fixed budget, charging a small fee for a test (treatment A) allows to more than double the amount of wells that can be tested, relative to a blanket testing campaign. Further, those that buy a test and that find out that their well is contaminated have higher switching rates than those who also got bad news but that were provided a test for free in other interventions —likely due to selection into buying. These results imply that charging a fee of BDT45 for a test can be over three times as cost-effective as a blanket testing campaign.

Further, promoting water sharing and charging a fee only for good news (at a low price) increase cost-effectiveness by 43% and 71%, respectively. Charging BDT90 for good news was as cost-effective as treatment A, while installing metal placards —due to their high cost— was in fact was half as cost-effective as A. Our calculations indicate that both promoting water-sharing and charging a low price for a test only if the well tests safe are the most cost-effective selling schemes among the ones we implemented. These results beg the question of whether an intervention that would both promote water-sharing and only charge for good news would be as cost effective as either one of them (or both combined). We leave this question for future research.

Our study, however, also faces a series of limitations. First, we do not know the objective arsenic concentration of the wells households switched to and are thus unable to quantify reductions in exposure to arsenic among switching households. Second, we lack data on biomarkers, including those for arsenicosis, and are thus unable to establish: (i) differential demand and switching behavior across households whose members have visible signs of arsenic poisoning and those whose members do not, and (ii) the effect of our intervention on health. Third, the lack of a free test provision experimental arm means that we are unable to assess the actual distribution of arsenic in Sonargaon, and are only able to gauge that distribution among buyers.

Figures and Tables

Figure 3.1: Placards

Notes: This figure displays the stainless steel placards attached to safe (bue), marginally safe (green), and unsafe (red) wells.

Figure 3.2: Sonargaon and treatment assignment

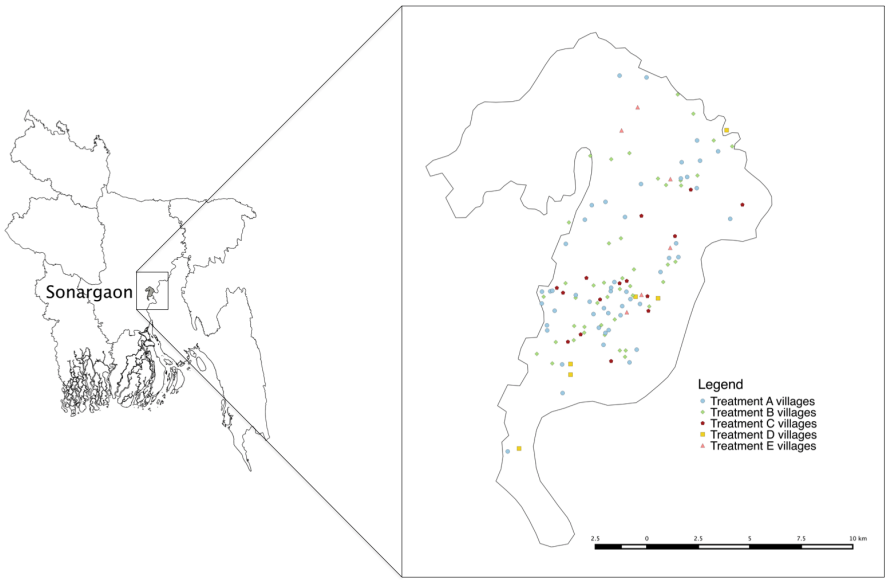

Notes: Author's illustrations. The symbol representing each village is placed at the mean latitude and longitude of all well-owners in that village, who were interviewed at baseline. The treatment associated with each village is the one implemented by surveyors (see Section 3.2).

Figure 3.3: Test uptake and results in Sonargaon

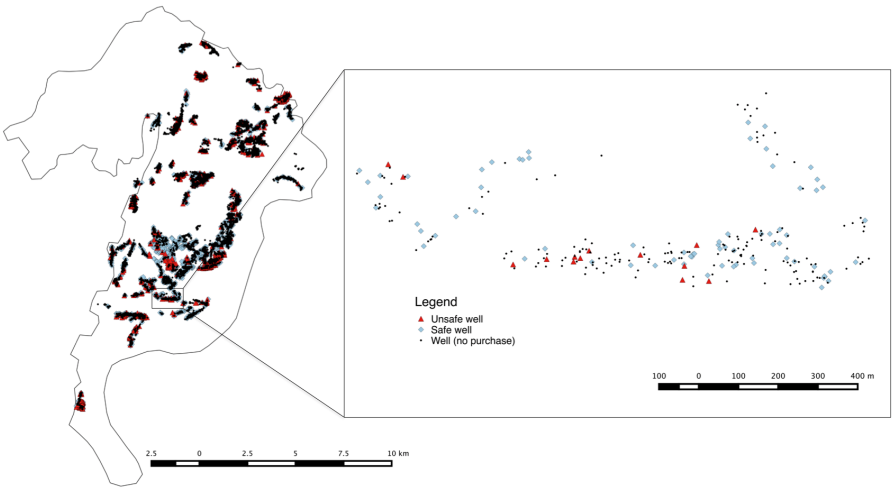

Notes: Author's illustrations. The figure in the box illustrates test upatke and results in Pirojpur, a village who was part of experimental arm A (BDT45 fee-for-test).

Table 3.1: Balance of household and tubewell characteristics at baseline

				Means by experimental arm					H_0: equal means†
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Obs	Mean	DS.	A	B	C	D	E	p-value
Drinks and cooks using own well water	14,046	0.907	0.291	0.931	0.879	0.891	0.943	0.904	0.170
H. head is male	12,734	0.845	0.361	0.843	0.845	0.857	0.842	0.845	0.927
H. head is wage-worker	12,099	0.272	0.445	0.247	0.300	0.375	0.162	0.172	0.220
H. head is self-employed	12,099	0.424	0.494	0.438	0.441	0.290	0.530	0.333	0.357
H. head has no schooling	12,099	0.211	0.408	0.260	0.168	0.055	0.303	0.410	0.000***
H. head has primary schooling only	12,099	0.324	0.468	0.319	0.334	0.319	0.364	0.225	0.022**
Concrete roof	12,555	0.172	0.377	0.175	0.181	0.134	0.156	0.209	0.002***
No. household members	12,555	3.599	1.410	3.597	3.484	3.405	4.013	4.483	0.000***
No. of children	12,734	1.450	1.071	1.438	1.368	1.383	1.730	1.984	0.000***
Status unknown (As)	11,751	0.756	0.430	0.682	0.804	0.904	0.758	0.596	0.000***
Status unsafe (As)	11,751	0.074	0.263	0.098	0.046	0.056	0.044	0.217	0.106
Well is visibly labeled safe	11,751	0.003	0.051	0.003	0.001	0.001	0.003	0.022	0.006
No. wells within 50m	11,939	11.708	14.538	10.046	14.302	11.111	12.202	8.043	0.105
No. visibly safe wells within 50m	11,939	0.023	0.165	0.028	0.004	0.007	0.041	0.153	0.014**
Share of unsafe wells (BAMWSP)	12,734	0.747	0.128	0.759	0.720	0.782	0.763	0.731	0.160
Private well	12,734	0.986	0.118	0.980	0.990	0.992	0.991	0.976	0.405
Well depth (100 feet)	12,734	1.761	1.076	1.829	1.752	1.729	1.590	1.588	0.925
Well age (years)	12,734	9.11	7.58	8.84	9.51	8.98	9.11	8.79	0.068*
Well cost (10,000 of BDT)	12,734	0.747	0.628	0.762	0.748	0.709	0.730	0.745	0.942
Attrition									
Umatched	12,734	0.150	0.357	0.151	0.137	0.114	0.267	0.135	0.594
Repeater household ID at baseline	12,734	0.062	0.241	0.068	0.064	0.040	0.048	0.073	0.954
Lost at follow-up	12,734	0.088	0.283	0.082	0.073	0.075	0.219	0.061	0.511

Notes: Except for the variable "Drinks and cooks using own well water", all other variables are summarized for household who use their own well for cooking an drinking at baseline. Differences in the number of observations across these variables is explained by missing entries during the data collection. † The F-test for equality of means across treatment arms controls for union-by-high or low arsenic status fixed effects. Standard errors are clustered at the village level. Significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 3.2: Test uptake analysis

	(1)	(2)	(3)	(4)	(5)	(6)
B	0.009	0.001	0.029	0.032	-0.001	-0.014
	(0.038)	(0.018)	(0.030)	(0.031)	(0.034)	(0.016)
C	0.044	0.032	0.005	0.028	0.036	0.020
	(0.034)	(0.030)	(0.037)	(0.040)	(0.032)	(0.029)
D	0.279***	0.282***	-0.018	-0.072	0.312***	0.310***
	(0.030)	(0.020)	(0.099)	(0.090)	(0.028)	(0.020)
E	0.102	0.145***	0.132**	0.094*	0.138***	0.165***
	(0.063)	(0.039)	(0.064)	(0.055)	(0.044)	(0.034)
Status unknown		0.150***	0.103***	0.121***		0.149***
		(0.033)	(0.021)	(0.019)		(0.031)
B·Status unknown			-0.036	-0.042		
			(0.033)	(0.035)		
C·Status unknown			0.017	0.004		
			(0.047)	(0.050)		
D·Status unknown			0.382***	0.419***		
			(0.117)	(0.105)		
E·Status unknown			0.033	0.024		
			(0.070)	(0.069)		
Status unsafe		0.120***	0.011	0.060***		0.120***
		(0.030)	(0.019)	(0.022)		(0.030)
B·Status unsafe			-0.004	-0.025		
			(0.036)	(0.036)		
C·Status unsafe			0.051	0.029		
			(0.035)	(0.039)		
D·Status unsafe			0.604***	0.610***		
			(0.107)	(0.112)		
E·Status unsafe			0.175	0.193*		
			(0.109)	(0.107)		
Concrete roof		0.061***		0.063***	0.092***	0.044**
		(0.017)		(0.015)	(0.023)	(0.021)
B·Concrete roof					0.049	0.082***
					(0.032)	(0.028)
C·Concrete roof					0.052	0.082
					(0.050)	(0.066)
D·Concrete roof					-0.199*	-0.191*
					(0.106)	(0.098)
E·Concrete roof					-0.180**	-0.093*
					(0.069)	(0.048)

H. head is male		-0.094***		-0.094***		-0.095***
		(0.019)		(0.018)		(0.019)
H. head is		0.003		0.003		0.004
wage-worker	(0.015)		(0.015)		(0.016)	
H. head is		0.054***		0.054***		0.055***
self-employed	(0.017)		(0.015)		(0.017)	
H. head has		-0.113***		-0.112***		-0.115***
no schooling	(0.013)		(0.013)		(0.012)	
H. head has primary		-0.076***		-0.073***		-0.076***
schooling only	(0.011)		(0.011)		(0.011)	
No. household		0.045***		0.045***		0.045***
members	(0.008)		(0.008)		(0.008)	
No. of children		0.000		0.001		0.000
		(0.008)		(0.008)		(0.008)
Well is visibly		-0.180**		-0.179**		-0.183**
labeled safe	(0.080)		(0.080)		(0.079)	
No. wells		-0.002**		-0.002**		-0.002**
within 50m	(0.001)		(0.001)		(0.001)	
No. visibly safe		-0.015		-0.001		-0.007
wells within 50m	(0.034)		(0.030)		(0.034)	
Share of unsafe		0.012		0.048		0.030
wells (BAMWSP)	(0.095)		(0.097)		(0.090)	
Well depth		0.024**		0.025***		0.024**
		(0.010)		(0.009)		(0.010)
Well age		0.000		0.000		0.000
		(0.001)		(0.001)		(0.001)
Well cost		0.011		0.013		0.013
		(0.018)		(0.017)		(0.016)
Observations	12,734	10,376	11,751	10,376	12,555	10,376
R-squared	0.149	0.146	0.110	0.162	0.159	0.151
Share buy (A)	0.228	0.172	0.170	0.172	0.229	0.172

Notes: The dependent variable is an indicator for test purchase at baseline. Estimation method is OLS. All regressions include union-by-high or low arsenic status fixed effects. Standard error are clustered at the village level. Significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 3.3: Well-switching opportunities

	(1)	(2)	(3)	(4)
	No. buy (25m)	No. buy (50m)	No. safe wells (25m)	No. safe wells (50m)
B	0.071	0.576	0.084	0.655
	(0.320)	(0.536)	(0.313)	(0.516)
C	-0.022	-0.069	-0.045	-0.149
	(0.302)	(0.692)	(0.283)	(0.643)
D	2.215***	4.363***	1.277*	1.401*
	(0.712)	(0.771)	(0.702)	(0.828)
E	0.834***	1.891***	0.612**	1.261*
	(0.253)	(0.639)	(0.255)	(0.640)
No. wells within 25mz	0.878***		0.857***	
	(0.073)		(0.075)	
No. wells within 50m		0.822***		0.801***
		(0.106)		(0.112)
Observations	11,939	11,939	11,939	11,939
R-squared	0.919	0.871	0.917	0.867
Mean Dep. Var.	1.837	3.823	1.574	3.083
Mean Dep. Var. (A)	0.670	1.967	0.521	1.525

Notes: The dependent variables are (i) Cols1-2: Number of well owners that bought a test within 25m and 50m (excluding the respective household) and (ii) Cols3-4: Number of wells tested safe (<50ppb) within 25m and 50m (excluding the respective household). Estimation method is OLS. All regressions include union-by-high or low arsenic status fixed effects. Standard error are clustered at the village level. Significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 3.4: Well-switching analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
B	0.241**	0.253**	0.332***	0.264**	0.313**	0.285**	0.009	0.008
	(0.096)	(0.118)	(0.112)	(0.127)	(0.121)	(0.141)	(0.008)	(0.007)
C	0.328***	0.293***	0.404***	0.288**	0.395***	0.338**	0.016	0.014*
	(0.116)	(0.104)	(0.138)	(0.129)	(0.138)	(0.134)	(0.011)	(0.008)
D	0.069	0.257*	0.154	0.252*	0.144	0.316*		
	(0.141)	(0.138)	(0.172)	(0.136)	(0.194)	(0.159)		
E	-0.114	0.012	-0.053	0.003	-0.049	0.049		
	(0.133)	(0.121)	(0.157)	(0.126)	(0.152)	(0.124)		
No. safe wells (25m)			0.051	0.119				
			(0.040)	(0.085)				
B·No. safe wells (25m)			-0.072	-0.059				
			(0.088)	(0.150)				
C·No. safe wells (25m)			-0.001	-0.028				
			(0.075)	(0.106)				
D·No. safe wells (25m)			-0.036	-0.107				
			(0.044)	(0.090)				
E·No. safe wells (25m)			-0.167**	-0.208				
			(0.082)	(0.139)				
No. safe wells (50m)					0.032	0.060		
					(0.030)	(0.036)		
B·No. safe wells (50m)					-0.085*	-0.060		
					(0.046)	(0.059)		
C·No. safe wells (50m)					0.008	-0.014		
					(0.052)	(0.059)		
D·No. safe wells (50m)					-0.021	-0.063		
					(0.036)	(0.043)		
E·No. safe wells (50m)					-0.126***	-0.132**		
					(0.045)	(0.053)		

No. buy (25m)			-0.033	-0.067				
			(0.037)	(0.077)				
B·No. buy (25m)			0.046	0.013				
			(0.086)	(0.134)				
C·No. buy (25m)			0.026	0.025				
			(0.081)	(0.098)				
D·No. buy (25m)			0.038	0.064				
			(0.040)	(0.081)				
E·No. buy (25m)			0.068	0.091				
			(0.045)	(0.095)				
No. buy (50m)					-0.015	-0.014		
					(0.026)	(0.045)		
B·No. buy (50m)					0.062	0.015		
					(0.044)	(0.059)		
C·No. buy (50m)					0.005	-0.017		
					(0.033)	(0.044)		
D·No. buy (50m)					0.019	0.013		
					(0.028)	(0.047)		
E·No. buy (50m)					0.046	0.035		
					(0.030)	(0.051)		
Observations	806	602	715	602	715	602	7,603	6,600
R-squared	0.391	0.454	0.468	0.459	0.474	0.460	0.019	0.032
Sample	Drink & unsafe	Drink & unsafe	Drink & unsafe	Drink & unsafe	Drink & unsafe	Drink & unsafe	Drink & no test	Drink & no test
Controls	N	Y	N	Y	N	Y	N	Y
Share switched (A)	0.367	0.356	0.262	0.356	0.262	0.356	0.030	0.026

Notes: The dependent variables is well-switching at follow-up. Estimation method is OLS. All regressions include union-by-high or low arsenic status fixed effects. Standard error are clustered at the village level. The controls are the same ones included in Table 3.2. Significance: *** p<0.01, ** p<0.05, * p<0.1.

Bibliography

Adler, R. F., G. J. Huffman, A. Chang, R. Ferraro, P.-P. Xie, J. Janowiak, B. Rudolf, U. Schneider, S. Curtis, D. Bolvin, et al. (2003). The version-2 global precipitation climatology project (GPCP) monthly precipitation analysis (1979-present). *Journal of hydrometeorology 4*(6), 1147–1167.

Agarwal, N., C. Banternghansa, and L. T. Bui (2010). Toxic exposure in America: Estimating fetal and infant health outcomes from 14 years of TRI reporting. *Journal of Health Economics 29*(4), 557–574.

Ahmed, M. F., S. Ahuja, M. Alauddin, S. J. Hug, J. R. Lloyd, A. Pfaff, T. Pichler, C. Saltikov, M. Stute, and A. van Geen (2006). Ensuring safe drinking water in Bangladesh. *Science 314*, 1687–1988.

Almond, D., K. Y. Chay, and D. S. Lee (2005). The costs of low birth weight. *The Quaterly Journal of Economics* (3), 1031–1083.

Arceo, E., R. Hanna, and P. Oliva (2016). Does the effect of pollution on infant mortality differ between developing and developed countries? Evidence from Mexico City. *The Economic Journal 126*(591), 257–280.

Argos, M., T. Kalra, P. J. Rathouz, Y. Chen, B. Pierce, F. Parvez, T. Islam, A. Ahmed, M. Rakibuz-Zaman, R. Hasan, G. Sarwar, V. Slavkovich, A. van Geen, J. Graziano, and H. Ahsan (2010). Arsenic exposure from drinking water, and all-cause and chronic-disease mortalities in Bangladesh (HEALS): a prospective cohort study. *Lancet 376*, 252–258.

Ashraf, N., J. Berry, and J. Shapiro (2010). Can higher prices stimulate prod-

uct use? Evidence from a field experiment in Zambia. *American Economic Review 100*(5), 2383–2413.

Ashraf, Q. and O. Galor (2013). The "Out of Africa" hypothesis, human genetic diversity, and comparative economic development. *The American Economic Review 103*(1), 1–46.

Ashraf, Q. and S. Michalopoulos (2015). Climatic fluctuations and the diffusion of agriculture. *Review of Economics and Statistics 97*(3), 589–609.

Babcock, B. A. and D. A. Hennessy (2006). Getting more corn acres from the Corn Belt. *Iowa Ag Review 12*(4), 4.

Bakke, B., A. J. De Roos, D. B. Barr, P. A. Stewart, A. Blair, L. B. Freeman, C. F. Lynch, R. H. Allen, M. C. Alavanja, and R. Vermeulen (2009). Exposure to atrazine and selected non-persistent pesticides among corn farmers during a growing season. *Journal of Exposure Science and Environmental Epidemiology 19*(6), 544–554.

Balasubramanya, S., A. Pfaff, L. Bennear, A. Tarozzi, K. M. Ahmed, A. Schoenfeld, and A. van Geen (2014). Evolution of households' responses to the groundwater arsenic crisis in Bangladesh: Information on environmental health risks can have increasing behavioral impact over time. *Environment and development economics 19*(5), 631–647.

Barnwal, P., A. van Geen, J. von der Goltz, and C. K. Singh (2016). Cost-sharing in environmental health products: Evidence from arsenic testing of drinking-water wells in Bihar, India. Working Paper.

BBC (2015). Ethiopia Profile - Overview. *BBC News*. Published December 4. `http://www.bbc.com/news/world-africa-13349399`.

Behrman, J. R. and M. R. Rosenzweig (2004). Returns to birthweight. *The Review of Economics and Statistics 86*(2), 586–601.

Bennear, L., A. Tarozzi, H. B. Soumya, A. K. Matin, A. Pfaff, and L. van Geen (2013). Impact of a randomized controlled trial in arsenic risk communication

on household water-source choices in Bangladesh. *Journal of Environmental Economics and Management 65*, 225–240.

Besley, T. and T. Persson (2011). The logic of political violence. *The Quarterly Journal of Economics 126*(3), 1411–1445.

Bharadwaj, P., J. Eberhard, and C. Neilson (2010). Do initial endowments matter only initially? Birth weight, parental investments and academic achievement in school. *University of California at San Diego, Department of Economics*.

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

Blattman, C. and E. Miguel (2010). Civil war. *Journal of Economic Literature 48*(1), 3–57.

BLS (2016). Quarterly census of employment and wages. *US Department of Labor-Bureau of Labor Statistics*. `http://www.bls.gov/cew/datatoc.htm`.

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

CDC (2007). Hospital stays, hospital charges, and in-hospital deaths among infants with selected birth defects–United States, 2003. *Centers for Disease Control and Prevention. MMWR. Morbidity and mortality weekly report 56*(2), 25.

Chabra, S. and C. A. Gleason (2005). Gastroschisis. *NeoReviews 6*(11), e493–e499.

Chassang, S. and G. Padró i Miquel (2009). Economic shocks and civil war. *Quarterly Journal of Political Science 4*, 211–228.

Chay, K. Y. and M. Greenstone (2003a). Air quality, infant mortality, and the Clean Air Act of 1970. Technical report, National Bureau of Economic Research.

Chay, K. Y. and M. Greenstone (2003b). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession. *The Quarterly Journal of Economics 118*(3), 1121–1167.

Chen, Y., J. H. Graziano, F. Parvez, M. Liu, V. Slavkovich, T. Kalra, M. Argos, T. Islam, A. Ahmed, M. Rakibuz-Zaman, R. Hasan, D. Levy, A. van Geen, and H. Ahsan (2011). Arsenic exposure from drinking water and cardiovascular disease mortality: A prospective cohort study in Bangladesh. *British Medical Journal Forthcoming*.

Chen, Y., A. van Geen, J. H. Graziano, A. Pfaff, M. Madajewicz, F. Parvez, A. I. Hussain, V. Slavkovich, T. Islam, and H. Ahsan (2007). Reduction in urinary arsenic levels in response to arsenic mitigation efforts in Araihazar, Bangladesh. *Environmental Health Perspectives 115*(6), 917–923.

Chernozhukov, V. and C. Hansen (2008). The reduced form: A simple approach to inference with weak instruments. *Economics Letters 100*(1), 68–71.

Chevrier, C., G. Limon, C. Monfort, F. Rouget, R. Garlantézec, C. Petit, G. Durand, and S. Cordier (2011). Urinary biomarkers of prenatal atrazine exposure and adverse birth outcomes in the PELAGIE birth cohort. *Environmental Health Perspectives 119*(7), 1034–41.

Chowdhury, A., M. Hossain, R. Nickson, M. Rahman, M. Jakariya, and M. Shamimuddin (2000). Combating a deadly menace: Early experiences with a community-based arsenic mitigation project in Bangladesh. *BRAC Research and Evaluation Division, Research Monograph Series* (16).

Chowdhury, A. M. R., A. Bhuiya, M. E. Chowdhury, S. Rasheed, Z. Hussain, and L. C. Chen (2013). The Bangladesh paradox: Exceptional health achievement despite economic poverty. *The Lancet 382*(9906), 1734–1745.

Ciccone, A. (2011). Economic shocks and civil conflict: A comment. *American Economic Journal: Applied Economics 3*(4), 215–227.

Cohen, J. and P. Dupas (2010). Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics 125*(1), 1–45.

Coneus, K. and C. K. Spiess (2012). Pollution exposure and child health: Evidence for infants and toddlers in Germany. *Journal of Health Economics 31*(1), 180–196.

Couttenier, M. and R. Soubeyran (2014). Drought and civil war in Sub-Saharan Africa. *The Economic Journal 124*(575), 201–244.

Crost, B., C. Duquennois, J. Felter, D. I. Rees, et al. (2015). Climate change, agricultural production and civil conflict: Evidence from the Philippines. *Manuscript*.

Cruze, N. B., A. L. Erciulescu, B. Nandram, et al. (2016). Developments in model-based estimation of county-level agricultural estimates. *Proceedings of the Fifth International Conference on Establishment Surveys*.

Currie, J. (2009). Healthy, wealthy, and wise: Socioeconomic status, poor health in childhood, and human capital development. *Journal of Economic Literature 47*(1), 87–122.

Currie, J. (2011). Inequality at birth: Some causes and consequences. *American Economic Review 101*(3), 1–22.

Currie, J., J. Graff Zivin, K. Meckel, M. Neidell, and W. Schlenker (2013). Something in the water: Contaminated drinking water and infant health. *Canadian Journal of Economics/Revue canadienne d'économique 46*(3), 791–810.

Currie, J., M. Neidell, et al. (2005). Air pollution and infant health: What can we learn from California's recent experience? *The Quarterly Journal of Economics 120*(3), 1003–1030.

Currie, J., M. Neidell, and J. F. Schmieder (2009). Air pollution and infant health: Lessons from New Jersey. *Journal of Health Economics 28*(3), 688–703.

Currie, J. and J. F. Schmieder (2009). Fetal exposures to toxic releases and infant health. *American Economic Review 99*(2), 177–83.

Currie, J. and R. Walker (2011). Traffic congestion and infant health: Evidence from E-ZPass. *American Economic Journal: Applied Economics 3*(1), 65–90.

Currie, J., J. G. Zivin, J. Mullins, and M. Neidell (2014). What do we know about short-and long-term effects of early-life exposure to pollution? *Annu. Rev. Resour. Econ. 6*(1), 217–247.

Curwin, B. D., M. J. Hein, W. T. Sanderson, M. G. Nishioka, S. J. Reynolds, E. M. Ward, and M. C. Alavanja (2005). Pesticide contamination inside farm and nonfarm homes. *Journal of Occupational and Environmental Hygiene 2*(7), 357–367.

Curwin, B. D., M. J. Hein, W. T. Sanderson, C. Striley, D. Heederik, H. Kromhout, S. J. Reynolds, and M. C. Alavanja (2007). Urinary pesticide concentrations among children, mothers and fathers living in farm and non-farm households in Iowa. *Annals of Occupational Hygiene 51*(1), 53–65.

Dal Bó, E. and P. Dal Bó (2011). Workers, warriors, and criminals: Social conflict in general equilibrium. *Journal of the European Economic Association 9*(4), 646–677.

DeSimone, L. A., P. A. Hamilton, and R. J. Gilliom (2009). Quality of water from domestic wells in principal aquifers of the United States, 1991-2004: Overview of major findings. Technical report, US Geological Survey.

Donahue, S. M., K. P. Kleinman, M. W. Gillman, and E. Oken (2010). Trends in birth weight and gestational length among singleton term births in the United States: 1990–2005. *Obstetrics and Gynecology 115*(2 Pt 1), 357.

Duffield, J. A., I. M. Xiarchos, and S. A. Halbrock (2008). Ethanol policy: Past, present, and future. *South Dakota Law Review 53*, 425.

Duhigg, C. (2009). Debating how much weed killer is safe in your water glass. *New York Times*, August 22, 2009, accessed May 13, 2016. `http://www.nytimes.com/2009/08/23/us/23water.html?_r=1`.

Dupas, P. (2012). Health behavior in developing countries. *Annual Review of Economics 3*, 425–449.

Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica 82*(1), 197–228.

Dupas, P. and E. Miguel (2016). Impacts and determinants of health levels in low-income countries. NBER Working Paper 22235.

Duranton, G., P. M. Morrow, and M. A. Turner (2014). Roads and trade: Evidence from the US. *The Review of Economic Studies 81*(2), 681–724.

EPA (2016). Atrazine monitoring program data and results. *US Environmental Protection Agency, Washington, DC*. `https://www.epa.gov/ingredients-used-pesticide-products/atrazine-monitoring-program-data-and-results`.

Ewing, A. C., S. R. Ellington, C. K. Shapiro-Mendoza, W. D. Barfield, and A. P. Kourtis (2016). Full-term small-for-gestational-age newborns in the US: Characteristics, trends, and morbidity. *Maternal and Child Health Journal*, 1–11.

Fafchamps, M. (2011). *Risk Sharing Between Households*, Volume 1A of *Handbook of Social Economics*, Chapter 24, pp. 1255–1280. San Diego and Amsterdam: North-Holland.

Fearon, J. D. (2007). *Institutions and Economic Performance*, Chapter Economic Development, insurgency, and civil war. Harvard University Press.

Fernandez-Cornejo, J., R. F. Nehring, C. Osteen, S. Wechsler, A. Martin, and A. Vialou (2014). Pesticide use in US agriculture: 21 selected crops, 1960-2008. *US Department of Agriculture-Economic Research Service, Economic Information Bulletin* (124).

Fischer, G., F. O. Nachtergaele, S. Prieler, E. Teixeira, G. Tóth, H. van Velthuizen, L. Verelst, and D. Wiberg (2012). Global agro-ecological zones (gaez v3. 0): Model documentation. *International Institute for Applied systems Analysis (IIASA), Laxenburg. Rome, Italy: Austria and the Food and Agriculture Organization of the United Nations (FAO)*.

Flanagan, S. V., R. B. Johnston, and Y. Zheng (2012). Arsenic in tube well water in Bangladesh: Health and economic impacts and implications for arsenic mitigation. *Bulletin of the World Health Organization 90*(11), 839–846.

George, C. M., A. van Geen, V. Slavkovich, A. Singha, D. Levy, T. Islam, K. M. Ahmed, J. Moon-Howard, A. Tarozzi, X. Liu, P. Factor-Litvak, and J. Graziano (2012). A cluster-based randomized controlled trial promoting community participation in arsenic mitigation efforts in Bangladesh. *Environmental Health 11*(1), 41.

George, C. M., Y. Zheng, J. H. Graziano, S. B. Rasul, Z. Hossain, J. L. Mey, and A. Van Geen (2012). Evaluation of an arsenic test kit for rapid well screening in Bangladesh. *Environmental science & technology 46*(20), 11213–11219.

Gilliom, R. J., J. E. Barbash, C. G. Crawford, P. A. Hamilton, J. D. Martin, N. Nakagaki, L. H. Nowell, J. C. Scott, P. E. Stackelberg, G. P. Thelin, et al. (2006). *Pesticides in the nation's streams and ground water, 1992-2001*. Number 1291. US Geological Survey.

Gilman, S. D., S. J. Gee, B. D. Hammock, J. S. Vogel, K. Haack, B. A. Buchholz, S. P. Freeman, R. C. Wester, X. Hui, and H. I. Maibach (1998). Analytical performance of accelerator mass spectrometry and liquid scintillation counting for detection of 14c-labeled atrazine metabolites in human urine. *Analytical Chemistry 70*(16), 3463–3469.

Gleditsch, N. P., P. Wallensteen, M. Eriksson, M. Sollenberg, and H. Strand (2002). Armed conflict 1946-2001: A new dataset. *Journal of peace research 39*(5), 615–637.

Grossman, H. I. (1991). A general equilibrium model of insurrections. *American Economic Review 81*(4), 912–21.

Guiteras, R. (2009). The impact of climate change on Indian agriculture. *Manuscript, Department of Economics, University of Maryland, College Park, Maryland*.

Harari, M. and E. La Ferrara (2013). Conflict, climate and cells: A disaggregated analysis. *CEPR Discussion Paper No. DP9277*.

Hart, C. E. (2006). Feeding the ethanol boom: Where will the corn come from? *Iowa Ag Review 12*(4), 2.

Hartzler, B. and M. Owen (2005). 2005 Herbicide Manual for Ag Professionals. *Iowa State University Extension and Department of Agronomy*. `http://www.weeds.iastate.edu/reference/wc92/`.

Howard, G., M. Ahmed, A. Shamsuddin, S. Mahmud, and D. Deere (2006). Risk assessment of arsenic mitigation options in Bangladesh. *Journal of Health, Population and Nutrition 24*(3), 346–355.

Hutson, S. S. (2004). *Estimated use of water in the United States in 2000*. Number 1268. US Geological Survey.

Inauen, J., R. Tobias, and H.-J. Mosler (2014). The role of commitment strength in enhancing safe water consumption: Mediation analysis of a cluster-randomized trial. *British journal of health psychology 19*(4), 701–719.

IPCC (2014). *Climate Change 2013: The Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.

Jayachandran, S. (2009). Air quality and early-life mortality evidence from Indonesia's wildfires. *Journal of Human Resources 44*(4), 916–954.

Johnston, R., S. J. Hug, J. Inauen, N. I. Khan, H.-J. Mosler, and H. Yang (2014). Enhancing arsenic mitigation in bangladesh: Findings from institutional, psychological, and technical investigations. *Science of the Total Environment 488*, 477–483.

Jones, A. M. (2016). Increasing prevalence of gastroschisis—14 states, 1995–2012. *MMWR. Morbidity and mortality weekly report 65*.

Kalnay, E., M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, et al. (1996). The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American meteorological Society 77*(3), 437–471.

Kanamitsu, M., W. Ebisuzaki, J. Woollen, S.-K. Yang, J. Hnilo, M. Fiorino, and G. Potter (2002). Ncep-doe amip-ii reanalysis (r-2). *Bulletin of the American Meteorological Society 83*(11), 1631–1643.

Kastenberg, Z. J. and S. Dutta (2013). Ventral abdominal wall defects. *NeoReviews 14*(8), e402–e411.

Ki-moon, B. (2007). A climate culprit in Darfur. *The Washington Post*. Published June 16. `http://www.washingtonpost.com/wp-dyn/content/article/2007/06/15/AR2007061501857.html`.

Kirby, R. S., J. Marshall, J. P. Tanner, J. L. Salemi, M. L. Feldkamp, L. Marengo, R. E. Meyer, C. M. Druschel, R. Rickard, and J. E. Kucik (2013). Prevalence and correlates of gastroschisis in 15 states, 1995 to 2005. *Obstetrics and Gynecology 122*(2 0 1), 275.

Knittel, C. R., D. L. Miller, and N. J. Sanders (2016). Caution, drivers! children present: Traffic, pollution, and infant health. *The Review of Economics and Statistics 98*(2), 350–366.

Kovarik, B. (1998). Henry Ford, Charles F. Kettering and the fuel of the future. *Automotive History Review 32*, 7–27.

Kraemer, S. (2000). The fragile male. *BMJ: British Medical Journal 321*(7276), 1609.

Kremer, M. and E. Miguel (2007). The illusion of sustainability. *Quarterly Journal of Economics 122*(3), 1007–1065.

Liebman, J., L. Katz, and J. King (2004). Beyond treatment effects: Estimating the relationship between neighborhood poverty and individual outcomes in the MTO experiment.

Lobell, D. B., W. Schlenker, and J. Costa-Roberts (2011). Climate trends and global crop production since 1980. *Science 333*(6042), 616–620.

Lubinsky, M. (2012). Hypothesis: Estrogen related thrombosis explains the pathogenesis and epidemiology of gastroschisis. *American Journal of Medical Genetics Part A 158*(4), 808–811.

Luechinger, S. (2014). Air pollution and infant mortality: A natural experiment from power plant desulfurization. *Journal of Health Economics 37*, 219–231.

MacDorman, M. F., S. E. Kirmeyer, and E. C. Wilson (2012). Fetal and perinatal mortality, United States, 2006. *National vital statistics reports: from the Centers for Disease Control and Prevention, National Center for Health Statistics, National Vital Statistics System 60*(8), 1–22.

Madajewicz, M., A. Pfaff, A. van Geen, J. Graziano, I. Hussein, H. Momotaj, R. Sylvi, and H. Ahsan (2007). Can information alone change behavior? Response to arsenic contamination of groundwater in Bangladesh. *Journal of Development Economics 84*(2), 731–754.

Mallows, C. L. (1986). Augmented partial residuals. *Technometrics 28*(4), 313–319.

Marshall, J., J. L. Salemi, J. P. Tanner, R. Ramakrishnan, M. L. Feldkamp, L. K. Marengo, R. E. Meyer, C. M. Druschel, R. Rickard, R. S. Kirby, et al. (2015). Prevalence, correlates, and outcomes of omphalocele in the United States, 1995–2005. *Obstetrics and Gynecology 126*(2), 284–293.

Mattix, K. D., P. D. Winchester, et al. (2007). Incidence of abdominal wall defects is related to surface water atrazine and nitrate levels. *Journal of Pediatric Surgery 42*(6), 947–949.

McArthur, J., D. Banerjee, S. Sengupta, P. Ravenscroft, S. Klump, A. Sarkar, B. Disch, and R. Kipfer (2010). Migration of as, and 3 H/3 he ages, in groundwater from West Bengal: Implications for monitoring. *water research 44*(14), 4171–4185.

Miguel, E. and S. Satyanath (2011). Re-examining economic shocks and civil conflict. *American Economic Journal: Applied Economics 3*(4), 228–32.

Miguel, E., S. Satyanath, and E. Sergenti (2004). Economic shocks and civil conflict: An instrumental variables approach. *Journal of Political Economy 112*(4).

Mittal, A. (2010). *Energy-water nexus: Many uncertainties remain about national and regional effects of increased biofuel production on water resources*. DIANE Publishing.

Munger, R., P. Isacson, S. Hu, T. Burns, J. Hanson, C. F. Lynch, K. Cherryholmes, P. Van Dorpe, and W. J. Hausler Jr (1997). Intrauterine growth retardation in Iowa communities with herbicide-contaminated drinking water supplies. *Environmental Health Perspectives 105*(3), 308.

Nachtergaele, F. and M. Petri (2013). Mapping land use systems at global and regional scales for land degradation assessment analysis-version 1.1.

Obama, B. (2015). Remarks by the President at the United States Coast Guard Academy Commencement. *The White House - Office of the Press Secretary*. Published May 20. https://www.whitehouse.gov/the-press-office/2015/05/20/remarks-president-united-states-coast-guard-academy-commencement.

Ochoa-Acuña, H. and C. Carbajo (2009). Risk of limb birth defects and mother's home proximity to cornfields. *Science of the Total Environment 407*(15), 4447–4451.

Ochoa-Acuña, H., J. Frankenberger, L. Hahn, and C. Carbajo (2009). Drinking-water herbicide exposure in Indiana and prevalence of small-for-gestational-age and preterm delivery. *Environmental Health Perspectives 117*(10), 1619.

Olsen, I., S. Groveman, M. Lawson, R. Clark, and B. Zemel (2010). New intrauterine growth curves based on United States data. *Pediatrics 125*(2), e214.

Opar, A., A. Pfaff, A. A. Seddique, K. Ahmedc, J. H. Graziano, and A. van Geen (2007). Responses of 6500 households to arsenic mitigation in Araihazar, Bangladesh. *Health and Place 13*(1), 164–172.

Oreopoulos, P., M. Stabile, R. Walld, and L. L. Roos (2008). Short-, medium-, and long-term consequences of poor infant health an analysis using siblings and twins. *Journal of Human Resources 43*(1), 88–138.

Padgitt, M. (1994). *Crop Rotations*, Chapter 4.2, pp. 143–150. US Department of Agriculture-Economic Research Service, Washington, DC.

Pfaff, A., A. Schoenfeld, K. Ahmed, and A. van Geen (2017). Reduction in exposure to arsenic from drinking well-water in Bangladesh limited by insufficient testing and awareness. *Journal of Water Sanitation and Hygiene for Development*, washdev2017136.

Pitt, M. M., M. R. Rosenzweig, and N. Hassan (2015). Identifying the cost of a public health success: Arsenic well water contamination and productivity in Bangladesh. Working Paper 21741, National Bureau of Economic Research.

Rahman, A., L.-Å. Persson, B. Nermell, S. El Arifeen, E.-C. Ekström, A. H. Smith, and M. Vahter (2010). Arsenic exposure and risk of spontaneous abortion, stillbirth, and infant mortality. *Epidemiology 21*(6), 797–804.

Roth, G. W. (1996). Crop Rotations and Conservations Tillage. *Penn State College of Agricultural Sciences - Cooperative Extension*. `http://extension.psu.edu/plants/crops/soil-management/conservation-tillage/crop-rotations-and-conservation-tillage`.

Royer, H. (2009). Separated at girth: US twin estimates of the effects of birth weight. *American Economic Journal: Applied Economics 1*(1), 49–85.

Sambanis, N. (2002). A review of recent advances and future directions in the quantitative literature on civil war. *Defence and Peace Economics 13*(3), 215–243.

Sanchez, T. R., D. Levy, M. H. Shahriar, M. N. Uddin, A. B. Siddique, J. H. Graziano, A. Lomax-Luu, A. van Geen, and M. V. Gamble (2016). Provision of well-water treatment units to 600 households in bangladesh: A longitudinal analysis of urinary arsenic indicates fading utility. *Science of the Total Environment 563*, 131–137.

Sanders, N. J. and C. Stoecker (2015). Where have all the young men gone? Using sex ratios to measure fetal death rates. *Journal of Health Economics 41*, 30–45.

Schlenker, W. and D. B. Lobell (2010). Robust negative impacts of climate change on African agriculture. *Environmental Research Letters 5*(1), 4010.

Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of Sciences of the United States of America 106*(37), 15594.

Schnepf, R. and B. D. Yacobucci (2013). Renewable Fuel Standard (RFS): Overview and issues. In *CRS Report for Congress*, Volume 40155.

Seneviratne, S. I., N. Nicholls, D. Easterling, C. M. Goodess, S. Kanae, J. Kossin, Y. Luo, J. Marengo, K. McInnes, M. Rahimi, et al. (2012). Changes in climate extremes and their impacts on the natural physical environment. *Managing the risks of extreme events and disasters to advance climate change adaptation*, 109–230.

Shah, S. K., K. R. Aroom, P. A. Walker, H. Xue, F. Jimenez, B. S. Gill, C. S. Cox Jr, and S. D. Moore-Olufemi (2012). Effects of nonocclusive mesenteric hypertension on intestinal function: Implications for gastroschisis-related intestinal dysfunction. *Pediatric Research 71*(6), 668–674.

Shaw, G. M., W. Yang, E. Roberts, S. E. Kegley, A. Padula, P. B. English, and S. L. Carmichael (2014). Early pregnancy agricultural pesticide exposures and risk of gastroschisis among offspring in the San Joaquin Valley of California. *Birth Defects Research Part A: Clinical and Molecular Teratology 100*(9), 686–694.

Sissine, F. (2010). Renewable Fuel Standard Program (RFS2) regulatory impact analysis. Technical report, Technical Report EPA-420-R-10-006. Assessment and Standards Division, Office of Transportation and Air Quality.

Smith, A., E. Lingas, and M. Rahman (2000). Contamination of drinking-water by arsenic in Bangladesh: A public health emergency. *Bulletin of the World Health Organization 78*(9), 1093–1103.

Sneeringer, S. (2009). Does animal feeding operation pollution hurt public health? A national longitudinal study of health externalities identified by geographic shifts in livestock production. *American Journal of Agricultural Economics 91*(1), 124–137.

Stock, J. (2015). The Renewable Fuel Standard: A path forward. *Columbia, SIPA Center on Global Energy Policy. April*.

Stock, J. H. and D. W. Andrews (2005). Inference with weak instruments. *NBER Working Paper* (t0313).

Swan, S. H., R. L. Kruse, F. Liu, D. B. Barr, E. Z. Drobnis, J. B. Redmon, C. Wang, C. Brazil, J. W. Overstreet, S. for Future Families Research Group, et al. (2003). Semen quality in relation to biomarkers of pesticide exposure. *Environmental Health Perspectives 111*(12), 1478.

Tarozzi, A. (2016). Health interventions in low income countries: A (not so) low hanging fruit? Els Opuscles del CREI, no. 42.

Tarozzi, A., A. Mahajan, B. Blackburn, D. Kopf, L. Krishnan, and J. Yoong (2014). Micro-loans, bednets and malaria: Evidence from a randomized controlled trial in Orissa (India). *American Economic Review 104*(7), 1909–1941.

Theisen, O. M., H. Holtermann, and H. Buhaug (2011). Climate wars?: Assessing the claim that drought breeds conflict. *International Security 36*(3), 79–106.

Thornton, R. L. (2008). The demand for, and impact of, learning HIV status. *American Economic Review 98*(5), 1829–1863.

USDA (2010). Field crops: Usual planting and harvesting dates. *US Department of Agriculture-National Agricultural Statistics Service, Agriculural Handbook* (628).

USDA (2016a). Cropland data layer metadata. *US Department of Agriculture-National Agricultural Statistics Service, Washington, DC*. `https://www.nass.usda.gov/Research_and_Science/Cropland/metadata/meta.php`.

USDA (2016b). Published crop-specific data layer. *US Department of Agriculture-National Agricultural Statistics Service, Washington, DC*. `https://nassgeodata.gmu.edu/CropScape/`.

USDA (2016c). Quick Stats. *US Department of Agriculture-National Agricultural Statistics Service, Washington, DC*. `https://quickstats.nass.usda.gov`.

van Geen, A., E. B. Ahmed, L. Pitcher, J. L. Mey, H. Ahsan, J. H. Graziano, and K. M. Ahmed (2014). Comparison of two blanket surveys of arsenic in tubewells conducted 12 years apart in a 25^2km area of Bangladesh. *Science of the Total Environment 488*, 484–492.

van Geen, A., K. Ahmed, A. Seddique, and M. Shamsudduha (2003). Community wells to mitigate the arsenic crisis in Bangladesh. *Bulletin of the World Health Organization 81*(9), 632–638.

van Geen, A., H. Ahsan, A. Horneman, R. Dhar, Y. Zheng, I. Hussain, K. Ahmed, A. Gelman, M. Stute, H. Simpson, S. Wallace, C. Small, F. Parvez, V. Slavkovich, N. Lolocano, M. Becker, Z. Cheng, H. Momotaj, M. Shahnewaz, A. Seddique, and J. Graziano (2002). Promotion of well-switching to mitigate the current arsenic crisis in Bangladesh. *Bulletin of the World Health Organization 80*(9), 732–737.

van Geen, A., Z. Cheng, Q. Jia, A. A. Seddique, M. W. Rahman, M. M. Rahman, and K. M. Ahmed (2007). Monitoring 51 community wells in Araihazar, Bangladesh, for up to 5 years: implications for arsenic mitigation. *Journal of Environmental Science and Health Part A 42*(12), 1729–1740.

Villanueva, C. M., G. Durand, M.-B. Coutté, C. Chevrier, and S. Cordier (2005). Atrazine in municipal drinking water and risk of low birth weight, preterm delivery, and small-for-gestational-age status. *Occupational and Environmental Medicine 62*(6), 400–405.

von Uexkull, N. (2014). Sustained drought, vulnerability and civil conflict in Sub-Saharan Africa. *Political Geography 43*, 16–26.

Vrieling, A., J. De Leeuw, and M. Y. Said (2013). Length of growing period over Africa: Variability and trends from 30 years of ndvi time series. *Remote Sensing 5*(2), 982–1000.

Wallander, S. (2013). While crop rotations are common, cover crops remain rare. *US Department of Agriculture-Economic Research Service, Washington, DC*. Published March 04.

Waller, S. A., K. Paul, S. E. Peterson, and J. E. Hitti (2010). Agricultural-related chemical exposures, season of conception, and risk of gastroschisis in Washington State. *American Journal of Obstetrics and Gynecology 202*(3), 241–e1.

Wasserman, G. A., X. Liu, F. Parvez, H. Ahsan, P. Factor-Litvak, J. Kline, A. van Geen, V. Slavkovich, N. J. LoIacono, D. Levy, Z. Cheng, , and J. H. Graziano (2007). Water arsenic exposure and intellectual function in 6-year-old children in Araihazar, Bangladesh. *Environmental Health Perspectives 115*(2), 285–289.

Weidmann, N. B., D. Kuse, and K. S. Gleditsch (2010). The geography of the international system: The cshapes dataset. *International Interactions 36*(1), 86–106.

Westcott, P. C. (2007). US ethanol expansion driving changes throughout the agricultural sector. *Amber Waves 5*(4), 10.

Winchester, P. D., J. Huskins, and J. Ying (2009). Agrichemicals in surface water and birth defects in the United States. *Acta Paediatrica 98*(4), 664–669.

Zhang, J., M. Merialdi, L. D. Platt, and M. S. Kramer (2010). Defining normal and abnormal fetal growth: Promises and challenges. *American Journal of Obstetrics and Gynecology 202*(6), 522–528.