Thirsting for Solutions: How Climate-Induced Well Failures drive School Absenteeism
— Preliminary Draft —

Lucile Dehouck*

June 2025

Abstract

Access to safe and reliable drinking water is essential for well-being, yet in many rural settings, water infrastructure is frequently unreliable. This paper examines the household-level consequences of well failure using a novel measure of predicted well functionality in Ethiopia. Leveraging a machine learning model trained on hydro-environmental variables, I construct an exogenous proxy for well failures and link it to detailed household survey data. I document a series of behavioral responses to those water access disruptions. Households adjust by switching away from groundwater toward more proximate, but potentially lower-quality sources such as surface water. Consistent with this shift, reported time spent collecting water declines, especially among adults. However, this reduction masks a reallocation of responsibility: more children are mobilized to fetch water, and this increased burden has downstream effects. Children in affected households are significantly more likely to miss school and are increasingly engaged in farming and casual work. These findings reveal the cascading effects of water access disruptions on household labor dynamics and children's time use, underscoring how environmental shocks can shape human capital accumulation and economic behavior.

*Paris School of Economics (PSE) and Université Paris 1 Panthéon Sorbonne – Email: lucile.dehouck@psemail.eu. I am very grateful to Katrin Millock for her continuous advice and support. This paper greatly benefited from discussions with and comments by Abebe D. Beyene, Tamma Carleton, Matthew Gordon, Solomon Hsiang, Kelsey Jack, Upmanu Lall, François Libois, David Margolis, Bluebery Planterose, and Wolfram Schlenker. I am also thankful for feedback from seminar participants at PSE, UC Berkeley, Columbia University, Policy Studies Institute, Dauphine, the Junior Seminar on The Economics of Migration, AERE, the 10th DENeB PhD Workshop, and the Environment and Climate Mobilities Network. Writing, code editing and bug diagnostics has been aided by the use of Open AI GPT models. This work benefitted from the financial support of Institut Convergences Migrations, funded by the CNRS, under reference number ANR-17-CONV-001, and from the Agence Nationale de la Recherche under the framework of the Investissements d'avenir programme reference ANR-17-EURE-001. All errors and omissions are my own.

1 Introduction

Today, 2.3 billion people live in water-stressed countries (UN Water, 2021). This crisis is exacerbated by climate change, which disrupts established weather patterns and makes water availability increasingly unpredictable (IPCC, 2023). Because drinking water is essential to survival, the loss of a household's primary safe water source may force families to reallocate time and labor toward water collection, with potentially far-reaching consequences for welfare and human capital. While much of the existing literature has focused on the agricultural impacts of climate shocks, far less is known about how disruptions in access to drinking water shape household decisions. To shed light on this overlooked channel, this paper asks: Do negative shocks to drinking water access – specifically, the failure of local wells – affect children's schooling and labor allocation within the household?

To study this question, I focus on Ethiopia. Along with eight other countries[1], Ethiopia accounts for 80% of the under-served population in Sub-Saharan Africa, the only region where the number of people lacking access to drinking water is still increasing (JMP, 2021). Despite its reputation as the "water tower of Africa," Ethiopia continues to face substantial challenges in delivering safe and reliable drinking water. In 2022, 60% of the Ethiopian rural population still lacked basic access to drinking water (JMP, 2023), meaning that they were either drinking from an unimproved source or walking for more than 30 minutes to fetch water. The country has a diverse climate and landscape, ranging from equatorial rainforest with high rainfall and humidity in the south and southwest, to desert-like conditions in the northeast, east, and southeast lowlands. However, natural variability in rainfall patterns and distribution, punctuated by extreme climatic events, has thrust many Ethiopian regions into extreme water scarcity conditions. In the past twenty years, droughts have led to ponds, wells, streams, and lakes drying up or becoming extremely shallow. Many people outside the cities collect water from these shallow water sources (JMP, 2023). In light of these challenges, Ethiopia offers a compelling context to examine how climate-induced disruptions to water infrastructure influence household labor allocation and child outcomes.

Previous research has shown that rural water systems in developing countries frequently fall into disrepair due to weak maintenance systems, limited financing, and unclear institutional responsibilities (Harvey and Reed, 2004; Giordano, 2009). Ethiopia is no exception, with similar challenges reported in the literature (Behailu et al., 2016; MacDonald et al., 2021; Pichon, 2019). The burden of coping with unreliable water supply disproportionately falls on women and girls, who are typically responsible for water collection in the household. This gendered division of labor

[1] Angola, Democratic Republic of Congo, Kenya, Madagascar, Mozambique, Sudan, Tanzania, and Uganda.

is well documented across contexts (Sorenson and Morssink, 2011; Crider and Ray, 2022), and it has wide-ranging implications for time use, education, and safety. However, despite a growing literature on climate shocks and human capital outcomes, the downstream effects of water point failures on time allocation, school participation, and household labor decisions remain poorly quantified. These gaps are particularly concerning given that climate change is projected to intensify the burden on women and girls (Carr et al., 2024). Causal evidence in this area is scarce, in part because infrastructure breakdowns are often endogenous to underlying social, geographic, or institutional conditions. This paper addresses that challenge by using predicted, climate-induced well failures as a source of plausibly exogenous variation in water access. I compile what is, to my knowledge, the most comprehensive dataset on rural water points in Sub-Saharan Africa, detailing the geographic location and functional status of wells at the time of monitoring. Using this dataset, I train a machine learning model to predict the likelihood of well failure based solely on exogenous climatic, hydrological, and environmental features. I then link these predictions to household-level outcomes from the *Ethiopian Socioeconomic Surveys* (ESS), a panel dataset covering the period 2012-2022. By combining predictive infrastructure data with climate variation and rich household-level information, this study estimates the causal impact of drinking water scarcity – operationalized as predicted well failure – on school absenteeism, intra-household time allocation, and labor outcomes. A key contribution of this approach is its ability to isolate the effects of domestic water access from broader agricultural impacts. This distinction is both contextually and empirically grounded. In Ethiopia, groundwater plays a critical role in providing domestic water, while its use for irrigation remains relatively limited [2]. Moreover, the wells in the dataset are explicitly designated for domestic use. These drinking water wells differ substantially from irrigation wells: they are typically shallower, yield lower volumes, and are designed to meet basic household needs rather than support crop production (Harter and Rollins, 2008). This design distinction reinforces the interpretation that observed behavioral responses reflect the effects of water insecurity for consumption and hygiene, rather than agricultural productivity shocks.

To have an exogenous shock on drinking water access, I rely on a machine learning approach, specifically, I use a random forest to predict well functionality. The model is trained and optimized with a stratified split and cross-validation. Its performance is assessed using the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), yielding an AUC score of 0.84, which indicates good predictive power. The model is then used to predict the functionality status of a representative well in each *Enumeration Area* (EA). These predictions provide a plausibly exogenous proxy for changes in water access. To further validate the model, I examine whether predicted well

[2] Ethiopia's UPGro Profile, accessed at `https://upgro.org/country-profiles/ethiopia/` on June 9th 2025

failures are associated with households reporting that the source they used failed to provide sufficient water in the last 2 weeks. Reassuringly, among households relying on groundwater, predicted failures are strongly associated with self-reported inability of the source to meet household water needs. In contrast, there is no such association among households relying on other types of sources, which are not subject to the same groundwater constraints. This supports the interpretation that the predicted well failure indicator captures real disruptions in water access, particularly for those dependent on wells. An additional concern is that the predictions may be capturing broader climate shocks rather than water access disruptions. To address this, I provide evidence that predicted well failures do not significantly affect agricultural outcomes such as quantity harvested, irrigation use, or farm size.

Having validated the predictive measure of well failure, I next examine how these disruptions in water access affect household water collection behavior. I document that households respond to well failure by switching drinking water sources. While reported changes in the main source are limited – possibly due to the timing and framing of survey questions – I find robust evidence of behavioral adjustment when using more immediate measures. In particular, well failure significantly reduces the likelihood of using groundwater as the source of water on the day of the interview and leads to a substantial increase in the use of surface water. This suggests that when wells fail, households turn to less reliable and potentially lower-quality alternatives to meet their water needs. Well failure also affects the time households spend collecting water. Using multiple recall periods, I find that predicted well failure is associated with a reduction in reported water collection time. These effects are most pronounced in reports from household heads, who indicate spending approximately 25-30 minutes less during the rainy season and 20 minutes less during the dry season. This pattern is consistent with a shift toward more proximate sources, such as surface water. However, more immediate, individual-level time-use measures appear noisier and mask opposing dynamics within the household: while adults reduce their time spent fetching water, more children are drawn into water collection. This redistribution of responsibility highlights important shifts in intra-household labor allocation.

This reallocation of water collection tasks has important consequences for children's education. I find that well failure increases the likelihood that children miss school. Specifically, households are 7-8 percentage points more likely to report that at least one child missed a week or more of school during the last semester – a relative increase of 70-80% over the baseline. The number of children missing school also rises significantly. These effects are not explained by migration and are robust to alternative specifications, underscoring that water access disruptions can meaningfully interfere with schooling.

Once children are withdrawn from school, they appear to be reallocated to household labor. I find that well failure increases the total number of hours spent farming, driven by a significant rise in children's participation. Adult labor remains unchanged, indicating a significant reallocation of labor responsibilities toward children. This pattern also extends to casual work, where the number of participating children increases. These results point to a cascading effect: well failure not only reshapes household water collection but also reallocates children's time toward work, with implications for both their education and labor outcomes.

The main contribution of this paper is to shed new light on the climate-education nexus by exploring a novel channel: the impact of access to water for domestic use on children's schooling and labor. A robust literature links climatic shocks to children's educational outcomes and labor supply in developing countries. Early studies in agrarian economies found that households often cope with adverse income shocks by withdrawing children from school. For instance, in rural India, Jacoby and Skoufias (1997) show that school attendance fluctuates with seasonal income shocks, implying that child education is used as a buffer when financial markets are incomplete. In Indonesia, Maccini and Yang (2009) find that women born in years with above-average rainfall attain significantly more schooling, are taller, and enjoy higher living standards in adulthood, illustrating that early-life climate conditions have persistent effects on human capital. These findings align with broader evidence that climatic volatility during childhood – such as droughts or floods – can impair human capital formation through multiple channels including lost schooling time, poorer nutrition, and illness. More recent work adds nuance by showing that the direction of climate's impact on schooling depends on children's age and the nature of the shock. Shah and Steinberg (2017) document that in India, increased rainfall raises agricultural wages, leading to better early-life outcomes but simultaneously decreasing school attendance among older children who are pulled into farm work. This underscores a recurring theme in the literature: when agricultural income improves, the opportunity cost of schooling rises for older children, reducing enrollment even as household resources increase. Conversely, negative shocks such as droughts often lead to reduced income and an increase in child labor, as shown by Beegle et al. (2006), who find that credit constraints exacerbate the educational impacts of climate shocks.

While this literature emphasizes the income and health channels linking climate shocks to education, it often treats water access as a background variable or control. Yet access to water for domestic use – particularly for drinking, cooking, and hygiene – plays a central role in shaping time use and labor allocation within households. In rural contexts, the burden of water collection frequently falls on women and school-aged girls, directly affecting educational attainment.

Some studies have explored the relationship between water infrastructure and education, but often lack plausibly exogenous variation to identify causal effects. Nauges and Strand (2017) documents that a significant negative relation between girls' school attendance and water collection activity in Ghana. Similarly, Koolwal and van de Walle (2013) show that easier access to water is associated to higher school enrollment and lower child illness, driven by time savings and reduced drudgery. These findings resonate with broader infrastructure studies: for instance, Devoto et al. (2012) find that Moroccan households are willing to pay for household water connections not primarily for health benefits, but for the convenience and time saved. This paper contributes to this literature by placing water fetching at the center of the climate-education nexus. Specifically, it distinguishes water access from the better-studied agricultural income and nutrition channels and shows how climate-induced variation in water availability – particularly well failures – affects children's schooling and labor. In doing so, the paper highlights how time burdens associated with water collection serve as a conduit through which environmental stress disrupts education. This is especially relevant in areas where water infrastructure is fragile or poorly maintained. While previous studies have focused on the benefits of improved water supply – such as reductions in child mortality (Galiani et al., 2005) – this study shifts attention to the consequences of losing access to water, an angle less explored in the literature. It also demonstrates that water reliability, not just access, matters for child outcomes under climate stress. Moreover, it leverages plausibly exogenous climate-induced well failures to provide a causal analysis of this mechanism.

This paper also builds on the literature documenting the welfare impacts of groundwater access. Sekhri (2014) finds that areas with accessible aquifers in India enjoy lower poverty rates and less disputes over irrigation water. Blakeslee et al. (2020) exploit quasi-random, small-scale variation in groundwater depletion within rural India – driven by geology and overextraction – to assess how long-term water scarcity affects households. They find that irrigation well drying leads to sharp declines in farm income and household wealth with limited evidence of on-farm adaptation but significant labor reallocation to off-farm work. Similarly, focusing on the Indian region Gujarat, Fishman et al. (2017) show that greater water scarcity leads to declines in irrigated agriculture and enhanced migration to cities, but only among dominant socio-economic groups. Sekhri and Hossain (2023) further highlight the social consequences of groundwater depletion, showing that falling water tables are associated with increased risk of sexual violence against women, as longer trips to fetch water increase exposure to risk. The paper adds to this literature by exploring how water scarcity affects children's human capital outcomes, focusing on schooling and labor. It introduces groundwater reliability as a determinant of intra-household time allocation and educational trajectories, an angle that has so far received limited attention. Moreover, most of the existing literature has focused on the Indian context (or developed countries). By relying on machine learning, I am able to study these outcomes in a data-scarce setting : Ethiopia – an important contribution, as many of the most vulnerable populations live in such environments.

Finally, this paper contributes to the expanding use of remote sensing and machine learning in empirical economics. The application of predictive methods has grown rapidly, providing researchers with powerful tools to overcome data limitations and enhance measurement in data-scarce settings. One well-known example is Hansen et al. (2013), who developed global deforestation maps using satellite imagery, enabling remote monitoring of environmental outcomes and generating over 8,000 citations since its release. More recent studies have leveraged satellite data and algorithms to measure variables of interest in development settings, from poverty estimation (Jean et al., 2016; Yeh et al., 2020) to monitoring economic activity (Potapov et al., 2022). To my knowledge, this is the first study to estimate climate-induced failures of drinking water wells using remote sensing. Senay et al. (2013) are close but they monitor waterhole drying, and use a different methodology. Moreover, I follow recent work such as Proctor et al. (2025) in using multiple imputation to account for potential measurement error in predicted data. By combining measures of climate, geological, and hydrogeological factors with ground-level water infrastructure data and household outcomes, this study offers a new empirical strategy to quantify the impacts of climate-induced water stress on children's education and labor. In doing so, it enriches both the methodological and the substantive literature at the intersection of climate, water, and human capital.

The remainder of the paper is structured as follows. Section 2 describes the data sources and construction. Section 3 presents key stylized facts on well failures. Section 4 outlines the empirical strategy and identification approach. Section 5 reports the main results, focusing on the effects of well failures on water collection behavior, school absenteeism, and the subsequent reallocation of children's labor. Section 6 concludes with a discussion of policy implications.

2 Data

The paper combines data from the Ethiopia Socio-Economic Survey and predictions on the functionality status of wells. The predictions are obtained by combining data on water points with different sources of hydro-environmental data.

2.1 Socio-Economic Data

This study draws on all five waves of the Ethiopian Socioeconomic Survey (ESS), conducted from 2011 to 2022[3]. The ESS is a nationally representative, multi-topic panel implemented by

[3] Wave 1 in 2011-2012, Wave 2 in 2013-2014, Wave 3 in 2015-2016, Wave 4 in 2018-2019 and Wave 5 in 2021-2022.

the Ethiopian Statistics Service in partnership with the World Bank's LSMS-ISA initiative. This longitudinal dataset provides comprehensive information on household demographics, consumption patterns, agricultural activities, income sources, labor market participation, asset ownership, and access to public services, hence enabling the analysis of rural livelihoods.

The ESS employs a two-stage stratified sampling design, selecting EAs within zones with probability proportional to size. An EA is typically made of approximately 150-200 households in rural areas. The original panel (Waves 1-3) maintained high response rates, with approximately 98% of households successfully re-interviewed in the third wave. Wave 4 marks the first wave of a new panel rather than a follow-up to previous waves. The refreshed panel (Waves 4-5) experienced higher attrition, largely due to security-related disruptions, including the exclusion of the Tigray region. Each wave includes a comprehensive household questionnaire covering demographics, education, health, labor, shocks and coping strategies, non-farm activities, housing, asset ownership, assistance, and access to credit. In rural areas, the survey incorporates detailed agricultural modules – post-planting, post-harvest, and livestock – capturing land use, input application, crop yields, and livestock holdings. Spatial data are anonymized through geographic displacement within buffers of up to 2 km in urban areas, 5 km in rural areas, and 10 km for a subset of households. I construct a panel of rural households over time which comprises 2,135 households observed in both Wave 4 and 5, and 3,222 households households observed across Wave 1 to 3.

Water access – The ESS records households' primary sources of drinking water across seasons, allowing for an assessment of both structural reliance and seasonal switching. Approximately 49% of households report using groundwater sources as their main source of drinking water throughout the year, underscoring the centrality of groundwater access and the potential implications of well failure for water security in this context (see Figure 1).

Importantly, while the ESS identifies a single primary source per season, a growing body of research suggests that households often draw on multiple sources simultaneously or sequentially to meet their needs. Studies from Sub-Saharan Africa show that households frequently engage in "source switching" to manage unreliability or to meet specific uses (Daly et al., 2021). This strategy reflects a form of adaptive behavior, but it also introduces vulnerabilities, especially when backup sources are distant, unsafe, or labor-intensive to access.

To further investigate this mechanism, I leverage the specialized water quality module included in Wave 3 of the ESS. Unlike the standard questions on households' primary water source, this module captures information on the actual source of water currently being used by the household at the time of the survey. This distinction is important, as it allows me to observe short-term adjustments in water use in response to well failures, rather than relying solely on reported "main" sources that may obscure switching behavior.

Time allocation for water collection is measured through two complementary sources: individuals time use questionnaires and household heads questionnaires. Individual time use questionnaires ask *"How many hours and minutes did [NAME] spend yesterday collecting water?"* while household heads responded to *"How much time does it take you to get to the source of water in the rainy(dry) season?"*. Table 1 reports the main descriptive statistics. Individual time use data aggregated at the household show that households spend an average of 1 hour and 20 minutes daily on water collection. Gender disparities in this task are pronounced: women devote approximately one hour daily to water collection, nearly triple the time spent by men (22 minutes). Children[4] also shoulder significant responsibility, averaging 36 minutes daily, with girls spending almost twice as much time as boys. Counting the number of members reporting positive time allocated to water collection, an average of 1.42 individuals per household participate. Household head reports reveal that most trips to water sources take between 16 and 30 minutes. This discrepancy with the aggregated household time may stem from several factors: underestimation by household heads (who are predominantly male), or the fact that reported travel times exclude waiting time at the water point and the return journey. Reported travel times increase slightly during the dry season.

School outcomes – For school outcomes, I focus on school absenteeism. It is measured in two ways-first, as a binary variable indicating whether any household member aged 7 to 18 missed school for more than a week in the month preceding the survey; and second, as a count variable capturing the number of such children within the household. Among households with enrolled children, 10% report at least one child missing school for more than a week in the survey month. Households report an average of 1.4 children missing school.

Time use – The ESS collects detailed time use information at the individual level, capturing how household members allocate their time across a range of daily activities. Respondents report the number of hours spent yesterday on fetching water and collecting firewood, as well as the time spent over the past seven days on activities like farming, casual and wage labor. This data enables a nuanced analysis of intra-household labor allocation and its response to external shocks, such as deteriorating access to clean water. Because the time use data is recorded at the individual level, it allows for the identification of shifts in responsibilities across age groups and between genders. The module offers a unique opportunity to examine how resource constraints – particularly water

[4]Throughout the paper, a child refers to an individual under the age of 15.

scarcity – reshape labor distribution within households, with a focus on the trade-offs between domestic duties and school attendance for children.

2.2 Wells Data

The data on water points used in this study are drawn from two major platforms: the Water Point Data Exchange (WPdx) and mWater. Both serve as centralized, open-access repositories for crowdsourced and institutional data on rural water infrastructure, widely used by governments, NGOs, and researchers to track water service delivery in low-income settings. By merging and harmonizing these datasets, I construct a non-exhaustive yet uniquely comprehensive inventory of drinking water wells in Ethiopia. To the best of my knowledge, this constitutes the largest and most spatially detailed dataset on rural water points currently available for the country. The final dataset includes 30,142 water points designated for drinking water use. Each entry contains geospatial coordinates and reported functional status at the time of monitoring. These wells span the national territory (Figure 2), with the majority recorded as functional at the time of survey (Figure 4). Most data points are from recent years, particularly 2021 (Figure 3), reflecting increased monitoring efforts during that period. It is important to emphasize that drinking water wells differ significantly from agricultural (irrigation) wells in both design and hydrological function. Drinking water wells are typically shallower, draw from upper aquifers, and are constructed to meet modest, consistent daily consumption needs-typically ranging from 20 to 50 liters per capita per day. In contrast, irrigation wells are deeper, require higher discharge volumes to meet crop water demands, and often involve mechanized pumping systems capable of extracting thousands of liters per hour (Harter and Rollins, 2008).

Water Point Data Exchange – WPdx is a global, open-access database that compiles water point data from multiple countries, primarily in Sub-Saharan Africa. By standardizing information from governments, NGOs, and other stakeholders, it supports the monitoring of rural water infrastructure. A key feature of WPdx[5] is its geospatial component, as it records the precise location (latitude and longitude) of individual water points, facilitating spatial analysis of water access. Additionally, it documents each water point's functional status at the time of assessment – whether functional, nonfunctional, or partially functional. The dataset also includes details on water source type (e.g., borehole, protected spring), extraction technology (e.g., hand pump, mechanized pump), management structures (e.g., community-based, privately managed), and construction date. In Ethiopia, WPdx has collaborated with various stakeholders, including the Millennium Water Al-

[5]The data was retrieved on `https://www.waterpointdata.org/access-data/` in November 2022.

liance, to enhance decision-making in water point implementation. The aim was that through using WPdx, decision makers can distribute resources efficiently and develop plans to improve rural water services.

mWater – mWater is a global, open-source platform designed to collect, manage, and analyze water point data. It was founded in 2012 in Tanzania with the mission to leverage mobile phone technology for monitoring water quality. Over time, it expanded to offer broader data collection, analysis, and visualization capabilities. With over 100,000 users in 180 countries, it is one of the leading platforms in the WASH (Water, Sanitation, and Hygiene) sector. mWater enables governments, NGOs, and researchers to monitor water infrastructure through a standardized and user-friendly system. The platform [6] provides geospatial data, recording the precise location (latitude and longitude) of individual water points, facilitating spatial analysis of water access and infrastructure coverage. mWater also tracks key attributes such as functional status (e.g., functional, nonfunctional, or partially functional), water source type, extraction technology, and management structures. Even though the platform supports dynamic data collection through mobile applications, allowing for continuous updates and long-term monitoring of water point functionality and service reliability, most of the wells were visited only once.

2.3 Hydro-environmental Data

Climate Data – The weather data comes from the European Center for Medium-Range Weather Forecasts (ECMWF) European Re-analysis fifth generation (ERA-5) dataset. Re-analysis data combine weather station data with forecast models. As a result, they are considered more reliable than station-only datasets, especially in regions with sparse or low-quality observational networks (Auffhammer et al., 2013). In particular, Glexner et al. (2020) find that ERA-5 performs well in East Africa. ERA-5 also offers a high spatial resolution: 0.25° x 0.25°, and has been used in recent studies such as Kalkuhl and Wenz (2020), Kotz et al. (2022).

Ethiopia's climate is marked by three main seasons (see Figure 5): Bega (dry season, from October-January), Belg (short rainy season, February-May) and Kiremt (the main rainy season, June-September). However, due to the country's diverse topography, rainfall patterns vary significantly across regions. For instance, the southwestern lowlands exhibit a unimodal rainfall regime, with rainfall typically occurring between February and November – differing from the standard seasonal cycle. There is a clear distinction in between highlands and lowlands in terms of precipitation regimes with lowlands experiencing much lower precipitations (see Figure 6).

[6] The data was retrieved on `https://portal.mwater.co/#/` in November 2022.

In order to measure drought, I combine monthly precipitation and evapotranspiration data to construct the Standardized Precipitation-Evapotranspiration Index (SPEI) at the Basin and EA level. The SPEI is a water balance measure normalized according to a log-logistic distribution (Vicente-Serrano et al., 2010; Begueria et al., 2014)[7]. The Water balance consists of the difference between precipitation and evapotranspiration. A one-unit deviation in the SPEI corresponds to one standard deviation from the long-run mean. The reference period I use is 1970-2022. The SPEI can measure drought severity in terms of intensity and duration. It can also identify the onset and end of drought episodes. This study primarily relies on a 12 months time scale: SPEI 12 represents a standardized measure of surface water balance over 12 months in relation to the expected surface water balance for the same period. There is variation both in time and in space in this measure (see Figure 7).

Groundwater Data – To complement the climate data, I draw on several groundwater and geophysical datasets. From Fan et al. (2013), I use high-resolution global maps of depth to groundwater, which provide insights into the accessibility and sustainability of groundwater resources. The MacDonald et al. (2012) dataset similarly offers estimates of groundwater depth but is specifically calibrated for African conditions, and may therefore align more closely with ground-truth observations on the continent (see Figure 8). To further characterize the subsurface environment, I use the African Groundwater Atlas, developed by the British Geological Survey, which compiles detailed country-level information on aquifer geology and groundwater productivity. Finally, I incorporate data from Pelletier et al. (2016), which provide globally consistent estimates of regolith depth – a key determinant of groundwater infiltration and storage potential.

Basin Level Data – HydroSHEDS (Hydrological data and maps based on SHuttle Elevation Derivatives at multiple Scales, Lehner et al. (2008)) is a globally consistent, high-resolution hydrological mapping product derived from NASA's Shuttle Radar Topography Mission (SRTM). One of its core components is the basins layer, which delineates water basins-areas where all surface water converges toward a common outlet point. For this study, I use the finest available Pfafstetter level (12 out of 12), which subdivides the landscape into small sub-basins with an average area of approximately 100 km^2. In addition to delineating hydrological boundaries, the basins layer serves as a repository for a wide range of environmental attributes, including elevation, terrain slope, karst extent, and more. These variables are spatially harmonized and linked to each basin, enabling comprehensive environmental analysis. For a full list of basin attributes, see Table 28.

[7] I use the R package SCI to calculate the SPEI.

3 Stylized facts on well failures

This section documents key empirical patterns in well failures across Sub-Saharan Africa, drawing on the harmonized dataset of water point observations. I focus on the prevalence of failure across regions and seasons, and correlations with environmental and technical characteristics.

Fact 1: Community-managed wells tend to be more functional – When comparing wells managed by communities, government entities, non-governmental organizations, and those with no identifiable management structure, community-managed wells stand out as the most reliably functional. Nearly 95% of community-managed wells are reported to be functional, compared to just 65% for wells without any identified management (see Figure 9).

Fact 2: Handpumps tend to be more functional – Wells equipped with handpumps report a functionality rate of approximately 90%, while those with motorized pumps are functional only about 80% of the time (see Figure 10). Although motorized pumps allow access to deeper groundwater sources and potentially more reliable supply, they are also more technologically complex and require regular maintenance. In the rural Ethiopian context, where technical capacity and spare parts may be limited, handpumps appear to be a more sustainable and appropriate technology.

Fact 3: Older wells tend to be less functional – Wells tend to become less functional as they age, with a general decline in reported functionality over time (see Figure 11). Interestingly, the data reveal a non-linear pattern: there is a noticeable dip in functionality beginning around five years after completion, followed by a partial recovery that lasts until around the 15-year mark. This suggests that wells may experience initial deterioration due to wear and tear, but some are subsequently rehabilitated or repaired, temporarily restoring their functionality [8]. After 15 years, however, the likelihood of a well being functional declines more sharply, indicating that long-term sustainability remains a challenge in the absence of consistent maintenance and infrastructure investment.

Fact 4: In drier places, wells tend to be less functional – Wells tend to be significantly less functional in drier areas, a pattern that holds whether dryness is measured using the SPEI (Figure 12) or total precipitation levels (Figure 13). In particular, locations with negative SPEI values – indicating sustained dryness or drought – exhibit notably higher rates of non-functionality: only around 70% of wells remain operational under such conditions. In contrast, in areas with positive SPEI values, nearly 95% of wells are reported to be functional. A similar pattern emerges when

[8] This interpretation aligns with accounts from stakeholders in Ethiopia

using cumulative precipitation as a metric: regions with lower rainfall levels consistently show higher rates of well failure. Notably, the association becomes more pronounced when considering precipitation lagged by two years (see Figure 14. This is particularly relevant, as lagged precipitation serves as a proxy for groundwater recharge: it is often estimated that it takes approximately two years for surface water to percolate through soil and rock layers and replenish aquifers.

Fact 5: In places where groundwater is deeper, wells tend to be less functional – The relationship between groundwater depth and well functionality varies depending on the data source. Using the groundwater depth estimates from MacDonald et al. (2012), a clear negative correlation emerges: wells in areas with deeper groundwater are significantly less likely to be functional (see Figure 16). This finding is consistent with the technical constraints of water extraction – deeper groundwater typically requires motorized pumps, which, as shown elsewhere, are less reliable and more prone to failure due to their complexity and maintenance needs. In contrast, when using the groundwater depth data from Fan et al. (2013), no clear relationship is observed between groundwater depth and well functionality (see Figure 15). This discrepancy may stem from differences in model design and calibration. While the MacDonald et al. (2012) dataset is specifically tailored to the African context, the Fan et al. (2013) model was trained primarily on data from high-income regions such as the U.S. and Europe. As a result, its predictions may be less accurate in low-resource settings like rural Ethiopia, where both environmental conditions and technological constraints differ substantially.

4 Empirical Approach

4.1 Predicting well failures

The well functionality data provide only a cross-sectional snapshot – each well is observed once with its status at the time of the survey. In this paper, I use machine learning techniques to extend this cross-section into a high-frequency monthly panel of predicted well functionality statuses. This approach offers three main advantages. First, it significantly expands the sample. In the raw data, only a limited number of wells were surveyed at the same time and in the same location as households, restricting the number of well-household matches. By predicting functionality status at a monthly frequency, I can generate plausible matches for more households over time. Moreover, even though the well dataset is the most comprehensive available for Ethiopia to my knowledge, it still does not capture the full universe of wells. In some EAs, households report using wells, yet no corresponding well is observed in the well data. The machine learning model helps bridge this gap by generating functionality predictions for a representative well in each EA and month, improving coverage and representativeness. Second, transforming the cross-section into a panel enables me to control for unobserved, time-invariant household characteristics that might confound cross-sectional estimates. This improves identification and allows for a cleaner estimation of the impact of water access disruptions. Third, the machine learning algorithm relies exclusively on hydro-environmental predictors – such as precipitation, groundwater depth, terrain slope, and other exogenous geophysical variables – meaning the predicted functionality outcomes are not correlated with household characteristics or behaviors. As a result, the predicted well status can be interpreted as an exogenous proxy for groundwater infrastructure reliability, strengthening causal inference.

Building the model – This predicted panel is generated by fitting a machine learning model to the observed cross-sectional data, using only exogenous hydro-environmental covariates as predictors. The modeling process occurs in three key steps (see Appendix C.2 for more details). In the first step, I apply a lasso regression to identify the most important variables that influence well functionality. The lasso model is particularly well-suited for this step due to its ability to handle high-dimensional data and perform variable selection. Next, I create interactions between the selected variables, hypothesizing that combinations of environmental factors may have compounded effects on well functionality. These interactions are then passed through another lasso regression, which helps identify the most significant interactions that further refine the model's predictive power. Finally, I train a random forest model using a randomly split sample of the data, employing 10-fold cross-validation to ensure robust and reliable predictions. The random forest algorithm allows for the capture of complex non-linear relationships between the predictors and well functionality, and cross-validation minimizes the risk of overfitting, providing an unbiased estimate of model performance.

Model performance – The model exhibits strong predictive performance, as illustrated by the ROC curve in Figure 17. The Area Under the Curve (AUC) is 0.84, indicating a high overall ability to discriminate between classes. The model achieves an accuracy of 0.74, meaning that it correctly classifies 74% of the observations. Moreover, its sensitivity (true positive rate) is 0.86, suggesting that the model is particularly effective at identifying when wells are functional. This performance is especially relevant in my setting because it means that the control group (those with functional wells) is relatively well captured. The treated group (those without a functional well) is a bit more noisy. The curve lies well above the diagonal line representing random guessing, confirming that the model performs substantially better than chance.

Moving away from statistical measures of model performance, Table 2 validates the use of predicted well failure as a proxy for actual water access disruptions by examining its correlation with self-reported well failure. Columns 1-3 show that predicted well failure is positively but not significantly associated with reported failure, suggesting limited predictive power when all water sources are pooled. However, Columns 4-6 introduce heterogeneity by groundwater use and reveal that the predictive model performs significantly better among groundwater users. The interaction term between predicted failure and groundwater use is positive and statistically significant at the 1% level, with an estimated effect of approximately 0.47-0.49. This implies that households relying on groundwater are substantially more likely to report a failure (by 42.3 percentage points) when their local well is predicted to fail. In contrast, for non-groundwater users, predicted failure does not meaningfully affect reported outcomes. These results provide strong support for the model's external validity and suggest that predicted well failure serves as a relevant proxy for changes in drinking water access, particularly among households that depend on groundwater.

4.2 Identification strategy

This paper estimates the causal effect of well failures on household outcomes by exploiting the timing of predicted climate-induced well failures in a panel fixed-effects framework. The comparison is between the same households before vs. after they experience a (predicted) well failure, while controlling for any time-invariant differences between households and common time trends. By using predictions based solely on climate and environmental shocks, the aim is to ensure that the variation in well failures is plausibly exogenous and not driven by household behavior, reporting patterns, or other endogenous factors.

The baseline estimating equation for a given outcome $Y_{h,t}$ (for household h in survey wave t) is:

$$Y_{h,t} = \beta \cdot WellFail_{c,t} + \gamma_1 Climate_{c,t} + \gamma_2 X_{h,t} + \delta_h + \delta_{z \times t} + \delta_m + \epsilon_{h,t} \tag{1}$$

where $WellFail_{c,t}$ is an indicator that household h within EA c experienced a predicted well failure in the month leading up to survey t [9]. The coefficient β is the parameter of interest, measuring the impact of a well failure on the outcome. Wells failure could be due to other endogeneous variable, to limit those concerns I add household and community controls as well as climate controls. I also include household fixed effects δ_h to control for all time-invariant characteristics of the household (such as baseline wealth, location, or any fixed propensity to use children's labor). I also include region-time fixed effects $\delta_{z \times t}$ to control for any time-varying zone level (second-level administrative units in Ethiopia) characteristics (such as conflict or economic trends). Standard errors are clustered at the EA level.

[9] Since the ESS provides geolocation only at the EA level, indicators for well failures and climate climate conditions are constructed at the EA level.

4.3 Threats to the identification

A potential concern is that the prediction model might inadvertently incorporate factors that are themselves outcomes or closely correlated with them – for example, if poorer communities are less able to maintain wells, community wealth could drive both well failure and downstream outcomes. This concern is mitigated by the fact that the prediction model is explicitly designed to capture climate-driven well failures and includes only exogenous hydro-environmental variables. Moreover, I include household fixed effects to absorb all time-invariant household characteristics, observed or unobserved, and zone-wave fixed effects to control for time-varying shocks or policies at the zone level.

Since the machine learning model relies solely on hydro-environmental variables, one potential concern is that it may capture not only well failures but also broader climate shocks – particularly those affecting agricultural outcomes. First, this is unlikely since climate variables used to predict well failure are aggregated at the basin level, whereas agricultural shocks typically operate at a more localized scale. Moreover, the main predictors are climate variables lagged interacted with hydrogeological features that are therefore not likely to correlate with factors affecting agricultural outcomes. I show in Figure 18; that the pair-wise correlation is low between climatic variables at the EA level and the predictions. Second, I control for climate at the EA level directly in the inference stage. Specifically, I include the mean 12-month Standardized Precipitation Evapotranspiration Index (SPEI) over the two years preceding each survey wave (waves occur roughly every two years). This ensures that the estimated coefficient on predicted well failure, β, captures the effect of losing access to a well above and beyond the general impact of drought. In other words, the identifying variation comes from comparing households that all faced similarly dry conditions – some of whom are predicted to lose access to a well, and others not – thus isolating the role of the well failure itself. To further probe the validity of this interpretation, I estimate Equation 2 using a range of agricultural outcomes as dependent variables. If the predicted well failure measure were merely proxying for broader environmental shocks, one would expect it to explain variation in agricultural outcomes as well. Reassuringly, I find no significant relationship between predicted well failure and harvest quantity, irrigation use, reported crop shocks, farm size, number of plots, or average plot size. These results, presented in Appendix C.1, support the interpretation that the measure captures disruptions in drinking water access rather than agricultural shocks.

Another concern relates to the recent literature on the appropriate use of remote sensing and machine learning-based variables in applied work (Proctor et al., 2025; Sanford et al., 2025). Until now, such variables have often been used directly as outcomes or covariates in causal inference models. However, machine learning models inevitably contain prediction errors, and if those errors are systematically correlated with variables in the inference model, they can bias the estimated effects. In a first step, since I have access to some ground-truth observations, I examine the structure of my machine learning variable ($WellFail_{c,t}$). I find that prediction errors are not random: they are correlated with distance to town and to the nearest road and to the consumption quintiles of households. To account for this, I include these variables as controls in my main specification. I also assess the robustness of my results using multiple imputation – a method designed to address concerns about error structure in predictive models (Carleton et al., 2025) – and find that the results remain stable. The results of these robustness checks are reported in Appendix C.3.

5 Results

The results are organized in two parts: household water collection behavior and child education outcomes. The evidence paints a consistent picture of how families adapt to a loss of their water source and the consequent burden on children.

5.1 Water collection

Switching away from groundwater: As a first-stage check, I examine whether households switch away from groundwater as their main drinking water source following a predicted well failure.

At first glance, I find no clear evidence that households change their main reported source in response to a recent failure. Table 3 investigates whether households switch away from groundwater as their main drinking water source following a predicted well failure for the rainy and dry seasons. Across all specifications, I find no statistically significant evidence that well failure leads households to report a different main water source. During the rainy season (Columns 1-3), the estimated effects are near zero and insignificant, likely reflecting the greater availability of alternative sources such as rainwater or surface water, which may reduce the salience of well failures. In the dry season (Columns 4-6), the coefficients on well failure are negative but remain statistically insignificant.

One possible explanation for the lack of a clear response in earlier specifications is that the timing of the predicted well failure and the outcome variable may not fully align. Households are asked to report their general main source of drinking water for the rainy and dry seasons, rather than their source at a specific point in time. As a result, a well failure occurring in the month prior to the survey may not influence how households report their usual source – especially if the failure is perceived as temporary or if households are accustomed to intermittent service disruptions. To explore this further, I restrict the sample to households interviewed during the dry season and examine whether a predicted well failure affects the main source they report using for that season. In this setting, I find negative and statistically significant effects across all model specifications (Columns 4-6 of Table 4). The estimates suggest that a predicted well failure reduces the likelihood of reporting groundwater as the main drinking water source by approximately 13 percentage points. Reassuringly, I do not find a significant effect among households interviewed during the dry season but reporting their main source for the rainy season (Columns 1-3). However, I do not find a significant effect among those interviewed during the rainy season (Table 5). This asymmetry likely reflects the fact that well failures are more disruptive and more salient in the dry season, when fewer substitute sources are available.

I also examine whether households exposed to longer durations without access to a functional well are less likely to report groundwater as their main drinking water source. As shown in Table 6, the estimated coefficients are negative but small and statistically insignificant in all specifications, both for the rainy and dry seasons (Columns 1-6). This lack of significance may reflect the routine nature of such disruptions: if well failures are frequent and temporary, households may not revise their stated main source, especially if they expect the well to become functional again. Alternatively, reported main source may reflect a long-term preference or default option rather than short-term availability.

To address this limitation, I leverage data from wave 3 of the survey, which includes a water quality module that asks about the source of water used on the day of the interview – a more immediate and behaviorally sensitive measure. Using this outcome, I find a consistent and statistically significant negative effect of predicted well failure on groundwater use. The results from Table 7 show that predicted well failure reduces the probability of using groundwater as the source of water on the day of the interview by approximately 0.25 percentage points, significant at the 10% level. This effect is sizable: it represents a 37% reduction relative to the mean use of groundwater (66%). This provides strong evidence that well failures lead to meaningful behavioral adjustments in water use, even when such changes are not captured in households' reports of their general main source. This finding reinforces the interpretation that well failures do affect household water behavior, even if such disruptions are not always reflected in how households report their general main source.

Focusing on this cross sectional report of the day-of source of water, I find that household switch toward surface water. The results show a statistically significant increase in the likelihood of using surface water: households experiencing well failure are approximately 31 percentage points more likely to report surface water as their source of the day (Columns 1-3) of Table 8). Given that the mean use of surface water is only 16%, this represents a nearly 200% increase relative to the mean, highlighting a substantial behavioral response.

In contrast, there is no significant increase in the use of trucked or piped water sources (Columns

4-9 of Table 8). These results suggest that, when wells fail, households are most likely to turn to less reliable and potentially less safe sources like surface water, rather than shifting to formal or purchased alternatives. This substitution pattern underscores the vulnerability of affected communities and reinforces the interpretation that predicted well failure meaningfully captures drinking water disruptions.

Time spent fetching water: Next, I examine how the time taken to fetch water evolves following a predicted well failure. Table 9 explores this relationship using three different measures of water collection time: time spent fetching water the previous day (Columns 1-3), time spent during the rainy season (Columns 4-6), and time spent during the dry season (Columns 7-9).

Across all specifications, predicted well failure is associated with a reduction in reported water collection time, although the magnitude and statistical significance of the effect vary by recall period and season. In Columns 1-3, where the outcome is time spent collecting water "yesterday," the coefficients are negative but not statistically significant. This may reflect that daily variation may obscure underlying behavioral shifts or inconsistencies between individual reports and household heads reports. In contrast, the results are stronger and statistically significant when using seasonal time-use data reported by household heads. In the rainy season (Columns 4-6), well failure is associated with a reduction of approximately 0.42 to 0.48 hours (25-30 minutes) in water collection time, significant at the 5% level. Similarly, in the dry season (columns 7-9), predicted well failure is associated with a decrease of about 0.33 hours (20 minutes), statistically significant in two of the three specifications (all except when additional controls are included). These findings align with earlier results indicating that households respond to well failure by switching from groundwater to more accessible sources such as surface water. Since surface water is typically closer, the observed reduction in time reflects a trade-off: households reduce their travel burden but may face increased health risks due to lower water quality.

To better understand how these shifts affect household dynamics, Table 10 investigates the intra-household allocation of water collection time, using the "time yesterday" data to examine the share of this burden carried by adults (Columns 1-3) versus children (Columns 4-6). The results reveal a clear reallocation of responsibility. Predicted well failure significantly reduces the adult share of water collection by approximately 10 percentage points. This decline is primarily driven by women reducing their water collection time (Graph 19). At the same time, the children's share increases by around 6 to 7 percentage points, significant at the 10% level. Given that children typically account for 25% of collection time, this reflects an increase of about 25% in their involvement. Taken together, these findings suggest that well failure not only prompts a change in water source and a reduction in collection time but also shifts the burden of water collection toward children.

This reallocation may help explain the earlier pattern: seasonal water collection time, reported by the household head (typically an adult), shows a significant reduction, while the "yesterday" measure, which reflects total household time, shows no clear effect – possibly because it masks offsetting changes between adults and children. These results underscore the broader social and intra-household consequences of water insecurity, highlighting its impact not only on water access but also on household labor dynamics and child time use.

Table 11 investigates the role of children in household water collection in response to predicted well failure, using three distinct measures: the number of child fetchers per household member (Columns 1-3), time spent collecting water per child fetcher (Columns 4-6), and total child water collection time per household member (Columns 7-9). The first panel (Columns 1-3) shows that predicted well failure is associated with a statistically significant increase in the number of children involved in water collection relative to household size. The estimated coefficient is approximately 0.019 and significant at the 10% level in two of the three specifications. Given that the mean of the dependent variable is 0.10, this corresponds to nearly a 20% increase in the intensity of children's involvement when a well fails. In contrast, Columns (4-6) show no significant change in the time spent per child fetcher. The coefficients are negative but imprecisely estimated, with large standard errors and no statistical significance, suggesting that while more children are participating in water collection, the time burden per child remains unchanged. The third panel (columns 7-9) examines total child water collection time per household member, combining changes in both the number of participating children and their time contributions. Here, the coefficients are small and statistically insignificant, indicating that the overall burden on children – measured in terms of average water collection time per household member – does not increase substantially, despite more children being involved. Taken together, these results suggest that households respond to well failure by mobilizing a greater number of children to participate in water collection. However, the per-child time burden remains unchanged, and the total child water collection time per household member does not significantly increase. This pattern is consistent with the overall reduction in water collection time observed in earlier results. It supports the interpretation that, when wells fail, households redistribute the water-fetching responsibility among more children without increasing the total burden.

5.2 Impacts on children's education an time use

School absenteeism: The increase in children's involvement in water collection suggests that they must be reallocating time away from other essential activities. Table 12 provides evidence that one such trade-off occurs in education: children are significantly more likely to miss school when wells fail. Columns 1-3 show that the probability of a household reporting that any child missed at least one week of school during the previous semester increases by approximately 7 to 8 percentage points, a statistically significant effect at the 5% level. Given a baseline absenteeism rate of 10%, this represents a substantial relative increase of 70-80%. Columns 4-6 go further by examining the number of children in the household who missed school, and again reveal strong and statistically significant effects. The estimated coefficients range from 0.14 to 0.15, significant at the 1% level, indicating that the impact of well failure extends beyond the extensive margin and leads to more children within a household being affected. These findings highlight that water access disruptions not only reconfigure intra-household labor dynamics – by shifting water collection duties to children – but also impose meaningful educational costs. The increase in absenteeism is consistent with children being pulled out of school, at least temporarily, to help meet basic household water needs.

Importantly, these results are robust to several alternative specifications. They remain statistically significant when controlling for the number of months a well was predicted to be non-functional over the past two years (Table 13), which captures cumulative exposure to water access disruptions. The results are also stable across different climate controls (Tables 14 and 15), suggesting that they are not merely driven by concurrent drought or rainfall anomalies. Moreover, I find no evidence that the observed increase in absenteeism is explained by migration patterns – predicted well failure does not significantly affect household migration indicators – reinforcing the interpretation that school absenteeism reflects a direct behavioral response to well failure, rather than broader household mobility or demographic change (Tables 16 and 17). Finally, I find no evidence of increased illness and diarrhea of children or adults following well failure, ruling out health deterioration as the primary mechanism (Table 18).

Reorganization of child labor: In response to well failure, children withdrawn from school appear to take on increased responsibilities not only in water collection but also in other household tasks. Table 19 shows that one such activity is agricultural work. In Columns 1-3, predicted well failure is associated with a statistically significant increase in the total number of hours spent farming. The estimated coefficients range from 2.87 to 2.88 hours, significant at the 1% level, representing a substantial increase relative to the baseline mean of 9.41 hours. This suggests that disruptions in water access have broader implications for household labor allocation, particularly in the agricultural sector. To examine who contributes to this additional farm labor, Columns 4-6 and 7-9 disaggregate farming time by adults and children. The results show that this increase is driven by greater reliance on children. While there is no significant change in adult participation, reliance on children rises significantly, with coefficients ranging from 0.066 to 0.067, statistically significant at the 5% level. Given a baseline mean of 0.23, this implies a nearly 30% increase in children's contribution to farm labor following a well failure. Further analysis indicates that this increase is concentrated on the extensive margin – that is, more children are participating in farming activities – rather than an increase in time per child (Table 20). This pattern aligns with the earlier finding that more children are missing school, suggesting a reallocation of time from education to labor. A similar pattern emerges for casual work (Table 21), where increased reliance on children is also driven by a rise in the number of participating children (Table 22). These consistent findings across sectors underscore how well failure triggers a chain of adjustments within households: from school absenteeism to increased child labor, both in farming and other informal economic activities. Taken together, the results provide evidence of a cascading set of effects: well failure increases school absenteeism and subsequently leads to higher child labor participation. Overall, these findings highlight how water insecurity can indirectly affect both children's time use and household livelihoods through shifts in intra-household labor dynamics

6 Discussion

This study provides novel evidence on the link between water infrastructure reliability and human capital outcomes in a low-income rural setting. The findings contribute to several strands of the literature. First, they add to the body of work on climate change and education by identifying a specific pathway: climate-induced resource shocks (drying of water sources) can lead to children missing school. Prior research on climate and schooling has noted that droughts and other shocks can reduce educational attainment (e.g., via income shocks or health impacts), but this paper pinpoints the mechanism of increased child labor in water collection as a driver of absenteeism. This complements studies like Nauges and Strand (2017) and Shah and Steinberg (2017), by providing direct causal evidence of what happens when water access collapses.

Second, the results speak to the literature on infrastructure and development. While many studies have looked at the benefits of new infrastructure (e.g., providing wells, piped water, electricity), far fewer have examined the effects of the failure of infrastructure. Rural water projects often suffer from sustainability problems, with broken pumps and dry wells being common (Sorenson and Morssink, 2011; Harvey and Reed, 2004; MacDonald et al., 2021). The findings show the significant welfare costs associated with such failures: not only do households lose a safe water source, but they also incur higher time costs and children's schooling is put at risk. This underscores that maintaining infrastructure may be just as important as building it, from a policy perspective. It also provides a quantitative sense of the value of a reliable water source – in educational outcomes – which can inform cost-benefit analyses of investments in maintenance or climate-proofing of water systems.

Third, methodologically, I demonstrate the value of using predictive analytics in development research. By training a machine learning model on water point data, I was able to create an instrument for a difficult-to-measure shock. This approach could be extended to other contexts: for example, predicting which areas will lose crops in a heatwave, or which power lines will fail in a storm, and using those predictions to study socio-economic impacts. This approach builds on recent work that uses novel data sources and algorithms to improve inference (Jean et al., 2016; Potapov et al., 2022). Even with messy real-world data, such a model can be a powerful tool to capture complex interactions (like well characteristics and drought) that create exogenous variation.

From a policy standpoint, the evidence highlights a need for climate-resilient rural water systems. In Ethiopia and similar settings, ensuring that wells do not easily run dry during droughts could have benefits that extend beyond just water provision. The results suggest there may be educational benefits: keeping children in school who would otherwise be diverted to fetching water. Moreover, the fact that community-managed wells had better functionality points to the importance of local engagement and maintenance capacity; supporting community management or providing maintenance services could reduce failure rates.

One limitation of this study is that I rely on predicted failures rather than observed failures in the household survey. This was a necessary innovation given data constraints, but it means the treatment is an approximation. There may be some households classified as having a "predicted well failure" who in reality found ways to cope – this would tend to attenuate the estimates. However, the strong and consistent results on outcomes and robustness tests validate that the predicted shock was meaningful to households.

In conclusion, as climate change accelerates, understanding these micro-level adaptation and impact pathways is crucial. The study shines light on how a seemingly local infrastructure issue – a broken well – can reverberate through a household's decisions and ultimately affect children's human capital formation. It calls for development policies that integrate water security with education and child welfare, ensuring that the gains from expanding school access and learning are not undermined by climate-induced resource shocks.

References

Auffhammer, M., Hsiang, S. M., Schlenker, W., and Sobel, A. (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, 7(2):181–198.

Beegle, K., Dehejia, R., and Gatti, R. (2006). Child labor and agricultural shocks. *Journal of Development Economics*, 81(1):80–96.

Begueria, S., Vicente-Serrano, S. M., Reig, F., and Latorre, B. (2014). Standardized Precipitation Evapotranspiration Index (SPEI) revisited: Parameter fitting, evapotranspiration, tools, datasets and drought monitoring. *International Journal of Climatology*, 34(10):3001–3023.

Behailu, B. M., Hukka, J. J., and Katko, T. S. (2016). Service failures of rural water supply systems in Ethiopia and their policy implications. *Public Works Management & Policy*, 22(2):179–196.

Blakeslee, D., Fishman, R., and Srinivasan, V. (2020). Way down in the hole: Adaptation to long-term water loss in rural India. *American Economic Review*, 110(1):200–224.

Carr, R. D., Kotz, M., Pichler, P.-P., Weisz, H., and Wenz, L. (2024). Climate change to exacerbate the burden of water collection on women's welfare globally. *Nature Climate Change*, 14:700–706.

Crider, Y. S. and Ray, I. (2022). Water and development: A gender perspective.

Daly, S. W., Lowe, J., Hornsby, G. M., and Harris, A. R. (2021). Multiple water source use in low- and middle-income countries: a systematic review. *Journal of Water and Health*, 19(3):370–392.

Devoto, F., Duflo, E., Dupas, P., Parienté, W., and Pons, V. (2012). Happiness on tap: Piped water adoption in urban Morocco. *American Economic Journal: Economic Policy*, 4(4):68–99.

Fan, Y., Li, H., and Miguez-Macho, G. (2013). Global patterns of groundwater table depth. *Science*, 339(6122):940–943.

Fishman, R., Jain, M., and Kishore, A. (2017). When water runs out: Adaptation to gradual environmental change in indian agriculture. *Accepted World Bank Economic Review*.

Galiani, S., Gertler, P., and Schargrodsky, E. (2005). Water for life: The impact of the privatization of water services on child mortality. *Journal of Political Economy*, 113(1):83–120.

Giordano, M. (2009). Global groundwater? Issues and solutions. *Annual Review of Environment and Resources*, 34:153–178.

Glexner, S., Demissie, T., and Tefera Diro, G. (2020). Did ERA5 improve temperature and precipitation reanalysis over East Africa? *Atmosphere*, 11(9):996.

Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., et al. (2013). High-resolution global maps of 21st-century forest cover change. *Science*, 342(6160):850–853.

Harter, T. and Rollins, L. (2008). *Water Well Design and Construction*. University of California, Division of Agriculture and Natural Resources. Publication 8086.

Harvey, P. A. and Reed, B. (2004). Rural water supply in Africa: Building blocks for handpump sustainability. *WEDC, Loughborough University, UK*.

IPCC (2023). *Weather and Climate Extreme Events in a Changing Climate*, pages 1513–1766. Cambridge University Press.

Jacoby, H. G. and Skoufias, E. (1997). Risk, financial markets, and human capital in a developing country. *Review of Economic Studies*, 64(3):311–335.

Jean, N., Burke, M., Xie, M., Davis, D., Lobell, D., and Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794.

JMP (2021). Report progress on household drinking water, sanitation, and hygiene. Technical report, UNICEF, WHO.

JMP (2023). Progress update on household drinking water, sanitation, and hygiene. Technical report, UNICEF, WHO.

Kalkuhl, M. and Wenz, L. (2020). The impact of climate conditions on economic production. evidence from a global panel of regions. *Journal of Environmental Economics and Management*, 103:102360.

Koolwal, G. B. and van de Walle, D. (2013). Access to water, women's work and child outcomes. *Economic Development and Cultural Change*, 61(1):27–56.

Kotz, M., Levermann, A., and Wenz, L. (2022). The effect of rainfall changes on economic production. *Nature*, 601:223–227.

Lehner, B., Verdin, K., and Jarvis, A. (2008). New global hydrography derived from spaceborne elevation data. *Eos, Transactions American Geophysical Union*, 89(10):93–94.

Maccini, S. and Yang, D. (2009). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review*, 99(3):1006–26.

MacDonald, A. M., Bonsor, H. C., Ó Dochartaigh, B. E., and Taylor, R. G. (2012). Quantitative maps of groundwater resources in Africa. *Environmental Research Letters*, 7(2):024009.

MacDonald, A. M., MacAllister, D. J., Kebede, S., Azagegn, T., Banks, E., and Bell, R. (2021). Evidence in the Horn of Africa of the resilience of rural water supply to drought. Technical report, UNICEF. Available via Natural Environment Research Council (NERC) Open Research Archive (E-print id 530771).

Nauges, C. and Strand, J. (2017). Water hauling and girls' school attendance: Some new evidence from Ghana. *Environmental and Resource Economics*, 66(1):65–88.

Pelletier, J. D., Broxton, P. D., Hazenberg, P., Zeng, X., Troch, P. A., Niu, G.-Y., Williams, Z. C., Brunke, M. A., and Gochis, D. (2016). A gridded global data set of soil, intact regolith, and sedimentary deposit thicknesses for regional and global land surface modeling. *Journal of Advances in Modeling Earth Systems*, 8.

Pichon, F. (2019). Rural water supply in Ethiopia: A political economy analysis. Research Report ODI Report, Overseas Development Institute (ODI), London. Available via EconStor, Id 12822.

Potapov, P., Turubanova, S., Hansen, M. C., Tyukavina, A., Zalles, V., Khan, A., Song, X.-P., Pickens, A., Shen, Q., and Cortez, J. (2022). Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century. *Nature Food*, 3(1):19–28.

Proctor, J., Carleton, T., and Sum, S. (2025). Measurement error in climate-economy research: New data and evidence from remote sensing. NBER Working Paper.

Sanford, L. C., Ayers, M., Gordon, M., and Stone, E. (2025). Adversarial debiasing for unbiased parameter recovery.

Sekhri, S. (2014). Wells, water, and welfare: The impact of access to groundwater on rural poverty and conflict. *American Economic Journal: Applied Economics*, 6:76–102.

Sekhri, S. and Hossain, M. A. (2023). Water in scarcity, women in peril. *Journal of the Association of Environmental and Resource Economists*, 10:1475–1513.

Senay, G. B., Velpuri, N. M., Alemu, H., Pervez, S. M., Asante, K. O., Kariuki, G., Taa, A., and Angerer, J. (2013). Establishing an operational waterhole monitoring system using satellite data and hydrologic modelling: Application in the pastoral regions of East Africa. *Pastoralism: Research, Policy and Practice*, 3(1):20.

Shah, M. and Steinberg, B. (2017). Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital. *Journal of Political Economy*, 125(2):527–561.

Sorenson, S. B. and Morssink, C., a. A. C. P. (2011). Safe access to safe water in low income countries: Water fetching in current times. *Social Science & Medicine*, 72(9):1522–1526.

Vicente-Serrano, S. M., Beguería, S., and López-Moreno, J. (2010). A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index. *Journal of Climate*, 23:1696–1718.

Yeh, C., Perez, A., Driscoll, A., Azzari, G., Burke, M., Lobell, D., and Ermon, S. (2020). Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. *Nature Communications*, 11(1):2583.

A Tables

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max	Observations
Wells					
Prediction	0.38	0.49	0.00	1.00	13939
well failure	0.02	0.15	0.00	1.00	14078
max streak (last 24 months)	124.23	48.34	5.00	156.00	13939
School					
any missing school	0.11	0.31	0.00	1.00	14078
nb missing school	0.15	0.51	0.00	7.00	14061
Water					
groundwater rainy	0.50	0.50	0.00	1.00	14035
groundwater dry	0.50	0.50	0.00	1.00	14031
surface rainy	0.21	0.40	0.00	1.00	14035
surface dry	0.19	0.39	0.00	1.00	14031
truck rainy	0.04	0.20	0.00	1.00	14035
truck dry	0.05	0.21	0.00	1.00	14031
piped rainy	0.25	0.43	0.00	1.00	14035
piped dry	0.26	0.44	0.00	1.00	14031
groundwater (day of)	0.66	0.47	0.00	1.00	3084
surface (day of)	0.16	0.37	0.00	1.00	3084
truck (day of)	0.02	0.13	0.00	1.00	3084
piped (day of)	0.14	0.34	0.00	1.00	3084
water source failure	0.17	0.38	0.00	1.00	7479
time water rainy (Household Head)	2.30	1.51	1.00	7.00	10300
time water dry (Household Head)	2.38	1.57	1.00	7.00	10278
water hours (Agregated Individual Reports)	1.30	2.81	0.00	75.00	14061
water hours adult (Agregated Individual Reports)	0.87	1.93	0.00	45.00	14035
water hours children (Agregated Individual Reports)	0.50	1.59	0.00	50.00	12199
nb water fetchers	1.42	1.06	0.00	11.00	14061
Other					
farm hours	47.15	52.17	0.00	432.00	14061
farm hours adult	32.19	37.39	0.00	353.00	14035
farm hours children	17.30	28.59	0.00	270.00	12199
nb farmers	2.40	1.89	0.00	14.00	14061
casual hours	1.86	9.24	0.00	193.00	14061
casual hours adult	1.67	8.48	0.00	193.00	14035
casual hours children	0.21	3.03	0.00	120.00	12199
nb casual	0.10	0.42	0.00	8.00	14061

Table 2: Impact of Predicted Well Failures on Household Reports of Water Source Failure

	All Sources			Heterogeneity by Groundwater Use		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.164	0.157	0.155			
	(0.132)	(0.134)	(0.140)			
well fail=1				−0.055	−0.059	−0.070
				(0.101)	(0.102)	(0.101)
Use groundwater=1				−0.147***	−0.144***	−0.137***
				(0.025)	(0.024)	(0.025)
well fail=1 × Use groundwater=1				0.466***	0.468***	0.493***
				(0.145)	(0.146)	(0.151)
spei12 24mo		−0.056	−0.048		−0.037	−0.031
		(0.037)	(0.036)		(0.034)	(0.034)
Observations	3077	3077	2891	3077	3077	2891
R^2	0.161	0.164	0.183	0.200	0.202	0.218
Mean of dep. var.	0.11	0.11	0.11	0.11	0.11	0.11
Mean gw users	0.66	0.66	0.66	0.66	0.66	0.66
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on reported water source failure. Columns (1)-(3) present average effects across all households, while Columns (4)-(6) explore heterogeneity by groundwater use. The dependent variable is a binary indicator equal to 1 if the household reports that its day-of water source was non-functional for at least one full day in the two weeks preceding the survey. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional. *Use groundwater* is equal to 1 for households that report using a groundwater source as their day-of water source. The interaction term captures the differential impact of well failure for groundwater users.

The regression includes zone and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, share of agricultural land within a 1km buffer, agricultural asset index, and the age, gender, and education level of the household head. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Impact of Predicted Well Failures on Reported Use of Groundwater as the Main Water Source

	Main Source – Rainy Season			Main Source – Dry Season		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.007	0.007	0.007	−0.048	−0.048	−0.048
	(0.071)	(0.071)	(0.071)	(0.078)	(0.077)	(0.077)
spei12 24mo		−0.038	−0.038		−0.051*	−0.051*
		(0.032)	(0.032)		(0.031)	(0.031)
Observations	13878	13878	13878	13874	13874	13874
R^2	0.634	0.634	0.634	0.627	0.628	0.628
Mean of dep. var.	0.50	0.50	0.50	0.50	0.50	0.50
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on households' reported use of groundwater as their main water source. Columns (1)-(3) present estimates for the rainy season, and Columns (4)-(6) for the dry season. The dependent variable is a binary indicator equal to 1 if the household reports using groundwater as its primary source during the corresponding season. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Impact of Predicted Well Failures on Reported Use of Groundwater as the Main Water Source – Sample Restricted to Households Interviewed in the Dry Season

	Rainy Season – Interviewed Dry Season			Dry Season – Interviewed Dry Season		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	−0.060	−0.056	−0.054	−0.131*	−0.126*	−0.133*
	(0.065)	(0.066)	(0.068)	(0.074)	(0.075)	(0.076)
spei12 24mo		−0.036	−0.036		−0.045	−0.050
		(0.037)	(0.037)		(0.035)	(0.035)
Observations	9325	9325	8944	9322	9322	8942
R^2	0.625	0.625	0.632	0.617	0.618	0.626
Mean of dep. var.	0.53	0.53	0.53	0.53	0.53	0.53
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on households' reported use of groundwater as their main water source, using a sample restricted to households interviewed during the dry season. Columns (1)-(3) present estimates for the rainy season, and Columns (4)-(6) for the dry season. The dependent variable is a binary indicator equal to 1 if the household reports using groundwater as its primary source during the corresponding season. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.
All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Impact of Predicted Well Failures on Reported Use of Groundwater as the Main Water Source – Sample Restricted to Households Interviewed in the Rainy Season

	Rainy Season – Interviewed Rainy Season			Dry Season – Interviewed Rainy Season		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.229	0.219	0.208	0.221	0.202	0.192
	(0.235)	(0.242)	(0.241)	(0.238)	(0.243)	(0.242)
spei12 24mo		−0.025	−0.032		−0.049	−0.057
		(0.072)	(0.072)		(0.072)	(0.073)
Observations	4240	4240	4228	4240	4240	4228
R^2	0.668	0.668	0.670	0.663	0.663	0.666
Mean of dep. var.	0.44	0.44	0.44	0.45	0.45	0.45
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on households' reported use of groundwater as their main water source, using a sample restricted to households interviewed during the rainy season. Columns (1)-(3) present estimates for the rainy season, and Columns (4)-(6) for the dry season. The dependent variable is a binary indicator equal to 1 if the household reports using groundwater as its primary source during the corresponding season. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.
All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Impact of Repeated Predicted Well Failures on Reported Use of Groundwater as the Main Water Source

	Main Source – Rainy Season			Main Source – Dry Season		
	(1)	(2)	(3)	(4)	(5)	(6)
Longest wt functional well	−0.001	−0.001	−0.001	−0.002	−0.002	−0.002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
spei12 24mo		−0.044	−0.048		−0.055*	−0.063**
		(0.032)	(0.032)		(0.030)	(0.031)
Observations	13701	13701	13305	13697	13697	13302
R^2	0.636	0.637	0.642	0.631	0.631	0.638
Mean of dep. var.	0.50	0.50	0.50	0.50	0.50	0.50
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of repeated predicted well failures on households' reported use of groundwater as their main water source, using a sample restricted to households interviewed during the rainy season. Columns (1)-(3) present estimates for the rainy season, and Columns (4)-(6) for the dry season. The dependent variable is a binary indicator equal to 1 if the household reports using groundwater as its primary source during the corresponding season. The main explanatory variable, *longest without functional well*, captures the longest consecutive streak (in months) in the last 2 years during which the representative drinking water well in the household's enumeration area was predicted to be non-functional.

All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Impact of Predicted Well Failures on Reported Use of Groundwater as the Day-of Water Source

	Source of the day		
	(1)	(2)	(3)
well fail	-0.259^{*}	-0.243^{*}	-0.246^{*}
	(0.146)	(0.142)	(0.139)
spei12 24mo		0.144^{***}	0.132^{***}
		(0.051)	(0.050)
Observations	3081	3081	2894
R^2	0.263	0.272	0.285
Mean of dep. var.	0.66	0.66	0.66
Fixed effects	✓	✓	✓
Climate		✓	✓
Controls			✓

Note: This table estimates the effect of predicted well failures on households' reported use of groundwater as their day-of water source. The dependent variable is a binary indicator equal to 1 if the household reports using groundwater for its water needs on the day of the interview. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

The regression includes zone and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, share of agricultural land within a 1km buffer, agricultural asset index, and the age, gender, and education level of the household head. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Impact of Predicted Well Failures on Reported Day-of Water Source Type

	Source of the day – Surface			Source of the day – Truck			Source of the day – Piped		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	0.312**	0.310**	0.307**	−0.017	−0.019	−0.020	−0.032	−0.038	−0.027
	(0.152)	(0.151)	(0.146)	(0.018)	(0.018)	(0.018)	(0.048)	(0.047)	(0.056)
spei12 24mo		−0.018	0.002		−0.016	−0.015		−0.056	−0.058
		(0.034)	(0.034)		(0.011)	(0.011)		(0.036)	(0.036)
Observations	3081	3081	2894	3081	3081	2894	3081	3081	2894
R^2	0.220	0.221	0.234	0.093	0.095	0.103	0.271	0.274	0.295
Mean of dep. var.	0.16	0.16	0.16	0.02	0.02	0.02	0.14	0.14	0.14
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table presents the estimated effects of predicted well failures on the type of water source households report using on the day of the interview. The dependent variables are binary indicators equal to 1 if the household reports using (1) surface water (columns 1-3), (2) trucked water (columns 4-6), or (3) piped water (columns 7-9) as their day-of water source. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

The regression includes zone and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, an agricultural asset index, and the age, gender, and education level of the household head. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Impact of Predicted Well Failures on Reported Time Spent Collecting Water

	Time Water – Yesterday			Time Water – Rainy Season			Time Water – Dry Season		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	−0.217	−0.217	−0.200	−0.462**	−0.481**	−0.416**	−0.334*	−0.338*	−0.293
	(0.172)	(0.171)	(0.174)	(0.188)	(0.193)	(0.200)	(0.195)	(0.197)	(0.204)
spei12 24mo		−0.026	−0.020		−0.164	−0.125		−0.034	−0.044
		(0.058)	(0.062)		(0.108)	(0.112)		(0.111)	(0.115)
Observations	13904	13904	13480	9766	9766	9212	9720	9720	9176
R^2	0.481	0.481	0.485	0.755	0.755	0.761	0.748	0.748	0.750
Mean of dep. var.	0.81	0.81	0.81	2.32	2.32	2.32	2.40	2.40	2.40
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the effect of predicted well failures on the time households report spending on water collection. The dependent variable in Columns (1)-(3) is the total time spent collecting water on the day prior to the interview, aggregated across all individuals in the household (in hours). Columns (4)-(6) and (7)-(9) use household heads' reports of average water collection time during the rainy and dry seasons, respectively. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.
All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Impact of Predicted Well Failures on the Allocation of Water Collection Time Between Adults and Children

	Time Water – Share Adults			Time Water – Share Children		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	-0.099^{***}	-0.099^{***}	-0.096^{**}	0.063^{*}	0.063^{*}	0.065^{*}
	(0.037)	(0.037)	(0.039)	(0.034)	(0.034)	(0.035)
spei12 24mo		0.005	−0.003		−0.008	−0.003
		(0.015)	(0.015)		(0.015)	(0.015)
Observations	13902	13902	13478	12195	12195	11796
R^2	0.529	0.529	0.538	0.563	0.563	0.571
Mean of dep. var.	0.62	0.62	0.62	0.25	0.25	0.25
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on the intra-household allocation of water collection responsibilities between adults and children. The dependent variable in Columns (1)-(3) is the share of total water collection time performed by adults on the day prior to the interview. Columns (4)-(6) use the corresponding share performed by children. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Impact of Predicted Well Failures on Child Participation in Water Collection and Time Allocation

	Nb Child Fetchers per Household Member			Water Time per Child Fetcher			Child Water Time per Household Member		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	0.019*	0.019*	0.018	−0.129	−0.130	−0.128	0.029	0.029	0.032
	(0.011)	(0.011)	(0.011)	(0.250)	(0.251)	(0.273)	(0.024)	(0.024)	(0.025)
spei12 24mo		−0.002	−0.000		0.051	0.040		−0.002	0.000
		(0.006)	(0.006)		(0.065)	(0.073)		(0.010)	(0.011)
Observations	11657	11657	11306	3449	3449	3263	11659	11659	11308
R^2	0.590	0.590	0.600	0.602	0.602	0.608	0.470	0.470	0.477
Mean of dep. var.	0.10	0.10	0.10	0.80	0.80	0.80	0.09	0.09	0.09
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the effect of predicted well failures on child involvement in water collection activities. Columns (1)-(3) estimate the impact on the number of child water fetchers per household member. Columns (4)-(6) report effects on the average water collection time per child fetcher (in hours). Columns (7)-(9) examine total child water collection time per household member (in hours). The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Impact of Predicted Well Failures on School Absenteeism

	Absenteeism: Any			Absenteeism: Nb		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.070**	0.070**	0.081**	0.140***	0.140***	0.151***
	(0.033)	(0.033)	(0.034)	(0.052)	(0.052)	(0.054)
spei12 24mo		0.018	0.020		0.016	0.019
		(0.016)	(0.017)		(0.028)	(0.028)
Observations	13930	13930	13480	13904	13904	13480
R^2	0.478	0.478	0.481	0.488	0.488	0.498
Mean of dep. var.	0.10	0.10	0.10	0.15	0.15	0.15
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on school absenteeism among children in the household. The dependent variable in Columns (1)-(3) is a binary indicator equal to 1 if any child in the household missed school for at least one week during the past semester. Columns (4)-(6) use as the outcome the number of children in the household who missed school during the same period. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: Impact of Predicted Well Failures on School Absenteeism – Controlling for Repeated Predicted Well Failures

	Count Months Non Functional			Longest wt Functional		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.160***	0.160***	0.175***	0.141***	0.142***	0.154***
	(0.057)	(0.056)	(0.058)	(0.053)	(0.053)	(0.055)
spei12 24mo		0.018	0.023		0.018	0.022
		(0.029)	(0.029)		(0.028)	(0.028)
Count 24 months	−0.005	−0.005	−0.006			
	(0.004)	(0.004)	(0.004)			
Longest wt functional well				0.001	0.001	0.001
				(0.001)	(0.001)	(0.001)
Observations	13680	13680	13262	13727	13727	13307
R^2	0.490	0.490	0.500	0.490	0.490	0.500
Mean of dep. var.	0.15	0.15	0.15	0.15	0.15	0.15
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the effect of predicted well failures on school absenteeism among children in the household. The dependent variable is a binary indicator equal to 1 if any child in the household missed school for at least one week during the past semester. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey. In Columns (1)-(3), the number of months in the past two years during which the representative drinking water well was predicted to be non-functional (*Count 24 months*) is included as a control. In Columns (4)-(6), the control variable is the longest consecutive streak of predicted non-functionality (*Longest wt functional well*).
All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Impact of Predicted Well Failures on School Absenteeism – Different Climate Controls

	Temp & Precip		Different SPEI					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
well fail	0.155***	0.146***	0.145***	0.158***	0.139***	0.151***	0.142***	0.154***
	(0.056)	(0.054)	(0.052)	(0.054)	(0.052)	(0.054)	(0.052)	(0.053)
tp	−0.000	−0.000						
	(0.000)	(0.000)						
temp	0.026	0.027						
	(0.022)	(0.024)						
SPEI12			0.023	0.028				
			(0.019)	(0.019)				
SPEI3					−0.007	−0.001		
					(0.018)	(0.018)		
spei3_24mo							0.060	0.068
							(0.059)	(0.059)
Observations	13680	13480	13904	13480	13904	13480	13904	13480
R^2	0.490	0.499	0.488	0.498	0.488	0.498	0.488	0.498
Mean of dep. var.	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Climate	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓

Note: This table reports the effect of predicted well failures on school absenteeism among children in the household. The dependent variable is the number of children who missed school for at least one week during the past semester. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey.

Columns (1)-(2) use monthly precipitation (*tp*) and temperature (*temp*) as climate controls. Columns (3)-(8) instead include various standardized precipitation-evapotranspiration indices (SPEI): 12-month SPEI in Columns (3)-(4), 3-month SPEI in Columns (5)-(6), and the average 3-month SPEI over the previous 24 months in Columns (7)-(8).

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 15: Impact of Predicted Well Failures on School Absenteeism – Different Climate Controls

	Temp & Precip		Different SPEI					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
well fail	0.081**	0.077**	0.074**	0.086**	0.071**	0.082**	0.072**	0.083**
	(0.034)	(0.034)	(0.033)	(0.034)	(0.033)	(0.034)	(0.032)	(0.033)
tp	−0.000	−0.000						
	(0.000)	(0.000)						
temp	0.014	0.013						
	(0.013)	(0.013)						
SPEI12			0.016	0.020*				
			(0.011)	(0.011)				
SPEI3					0.004	0.008		
					(0.010)	(0.010)		
spei3_24mo							0.047	0.051
							(0.033)	(0.033)
Observations	13706	13480	13930	13480	13930	13480	13930	13480
R^2	0.481	0.481	0.478	0.481	0.478	0.481	0.479	0.481
Mean of dep. var.	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Climate	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓

Note: This table reports the effect of predicted well failures on school absenteeism among children in the household. The dependent variable is a binary indicator equal to 1 if any child in the household missed school for at least one week during the past semester. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey.

Columns (1)-(2) use monthly precipitation (*tp*) and temperature (*temp*) as climate controls. Columns (3)-(8) instead include various standardized precipitation-evapotranspiration indices (SPEI): 12-month SPEI in Columns (3)-(4), 3-month SPEI in Columns (5)-(6), and the average 3-month SPEI over the previous 24 months in Columns (7)-(8).

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 16: Impact of Predicted Well Failures on Permanent Migration

	Any Permanent Migrant			Nb Permanent Migrant		
	(1)	(2)	(3)	(4)	(5)	(6)
well_fail	0.055	0.052	0.064*	0.121	0.116	0.122
	(0.036)	(0.036)	(0.036)	(0.080)	(0.080)	(0.079)
spei12_24mo		−0.023	−0.017		−0.043	−0.038
		(0.015)	(0.015)		(0.036)	(0.037)
Observations	10682	10682	10106	10682	10682	10106
R^2	0.745	0.745	0.750	0.762	0.762	0.767
Mean of dep. var.	0.33	0.33	0.33	0.58	0.58	0.58
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the estimated effect of predicted well failures on permanent migration outcomes. The dependent variable is either a binary indicator equal to 1 if any household member has permanently migrated (i.e., left the household), or a continuous variable measuring the number of such migrants. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey.

Columns (1)-(2) use monthly precipitation (*tp*) and temperature (*temp*) as climate controls. Columns (3)-(8) instead include various standardized precipitation-evapotranspiration indices (SPEI): 12-month SPEI in Columns (3)-(4), 3-month SPEI in Columns (5)-(6), and the average 3-month SPEI over the previous 24 months in Columns (7)-(8).

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 17: Impact of Predicted Well Failures on Seasonal Migration

	Any Seasonal Migrant			Nb Seasonal Migrant			Months Away		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well_fail	0.025	0.024	0.029	−0.029	−0.029	−0.020	0.243	0.061	0.286
	(0.049)	(0.049)	(0.047)	(0.086)	(0.085)	(0.084)	(0.926)	(0.940)	(0.911)
spei12_24mo		0.017	0.015		0.054	0.054		0.383	0.246
		(0.017)	(0.017)		(0.050)	(0.050)		(0.434)	(0.456)
Observations	13930	13930	13480	13904	13904	13480	719	719	679
R^2	0.494	0.494	0.505	0.528	0.529	0.537	0.691	0.691	0.702
Mean of dep. var.	0.17	0.17	0.17	0.31	0.31	0.31	2.90	2.90	
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the estimated effect of predicted well failures on health outcomes. The dependent variable varies across specifications: a binary indicator equal to 1 if any household member was reported to suffer from an illness or injury in the last 4 weeks in Columns (1)-(3), or if any child under 15 was reported sick (below 15) in Columns (4)-(6); and a binary indicator equal to 1 if any household member reported diarrhea in the last 4 weeks in Columns (6)-(9), or if any child under 15 reported diarrhea (below 15) in Columns (9)-(12). The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey.

Columns (1)-(2) use monthly precipitation (*tp*) and temperature (*temp*) as climate controls. Columns (3)-(8) instead include various standardized precipitation-evapotranspiration indices (SPEI): 12-month SPEI in Columns (3)-(4), 3-month SPEI in Columns (5)-(6), and the average 3-month SPEI over the previous 24 months in Columns (7)-(8).

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 18: Impact of Predicted Well Failures on Health Outcomes

	Impact on Health			Impact on Health - Children			Impact on Diarrhea			Impact on Diarrhea - Children		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
well fail	0.017	0.017	0.002	0.002	0.002	0.005	0.008	0.008	0.011	−0.005	−0.005	0.000
	(0.072)	(0.073)	(0.077)	(0.053)	(0.053)	(0.057)	(0.030)	(0.030)	(0.031)	(0.032)	(0.032)	(0.033)
spei12 24mo		−0.074*	−0.067		−0.033	−0.027		−0.019	−0.023*		−0.017	−0.021*
		(0.044)	(0.044)		(0.028)	(0.028)		(0.012)	(0.013)		(0.011)	(0.012)
Observations	13680	13680	13262	11480	11480	11135	13727	13727	13307	11518	11518	11171
R^2	0.544	0.544	0.552	0.490	0.490	0.498	0.432	0.432	0.440	0.440	0.440	0.446
Mean of dep. var.	0.62	0.62	0.62	0.29	0.29	0.29	0.08	0.08	0.08	0.07	0.07	0.07
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓			✓

Note: This table reports the estimated effect of predicted well failures on seasonal migration outcomes. The dependent variable is either: a binary indicator equal to 1 if any household member has spent more than three months away from the household (i.e., seasonal migration) in Columns (1)-(3); a continuous variable measuring the number of such migrants in Columns (4)-(6); or the number of months spent away per migration episode in Columns (7)-(9). The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey.

Columns (1)-(2) use monthly precipitation (*tp*) and temperature (*temp*) as climate controls. Columns (3)-(8) instead include various standardized precipitation-evapotranspiration indices (SPEI): 12-month SPEI in Columns (3)-(4), 3-month SPEI in Columns (5)-(6), and the average 3-month SPEI over the previous 24 months in Columns (7)-(8).

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 19: Impact of Predicted Well Failures on Farming Hours and their Allocation Between Adults and Children

	Time Farm			Time Farm – Reliance Adult			Time Farm – Reliance Children		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	2.874***	2.879***	2.881***	0.036	0.036	0.034	0.066**	0.066**	0.067**
	(0.990)	(0.983)	(1.059)	(0.037)	(0.037)	(0.038)	(0.031)	(0.032)	(0.033)
spei12 24mo		−0.616	−0.580		0.011	0.009		−0.011	−0.008
		(0.552)	(0.561)		(0.016)	(0.017)		(0.011)	(0.012)
Observations	13901	13901	13477	13903	13903	13479	12556	12556	12135
R^2	0.607	0.607	0.617	0.584	0.584	0.591	0.589	0.589	0.597
Mean of dep. var.	9.41	9.41	9.41	0.59	0.59	0.59	0.23	0.23	0.23
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the effect of predicted well failures on time spent on farming and the intra-household allocation of farm labor. The dependent variable in Columns (1)-(3) is the total time (in hours) the household spent on farming in the seven days before the interview. Columns (4)-(6) report the share of total farm time contributed by adults, and Columns (7)-(9) report the share contributed by children. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the householdâ€™s enumeration area is predicted to be non-functional at the time of the survey.
All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 20: Impact of Predicted Well Failures on Child Participation in Farming and Time Allocation

	Nb Child Farmers per Household Member			Farm Time per Child Farmer			Child Farm Time per Household Member		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	0.053***	0.053***	0.054***	4.301	4.301	4.497	1.901***	1.902***	1.941***
	(0.016)	(0.016)	(0.017)	(3.377)	(3.389)	(3.422)	(0.552)	(0.553)	(0.553)
spei12 24mo		−0.011	−0.011		−0.017	0.317		−0.054	0.004
		(0.007)	(0.007)		(1.334)	(1.366)		(0.223)	(0.231)
Observations	11657	11657	11306	4825	4825	4622	11659	11659	11308
R^2	0.638	0.638	0.649	0.656	0.656	0.666	0.589	0.589	0.596
Mean of dep. var.	0.16	0.16	0.16	22.56	22.56	22.56	3.06	3.06	3.06
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the effect of predicted well failures on child involvement in farming. Columns (1)-(3) estimate the impact on the number of child farmers per household member. Columns (4)-(6) report effects on the average farming time per child farmer (in hours). Columns (7)-(9) examine total child farming time per household member (in hours). The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.
All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 21: Impact of Predicted Well Failures on Casual Work Hours and their Allocation Between Adults and Children

	Time Casual Work			Time Casual – Reliance Adult			Time Casual – Reliance Children		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	0.046	0.046	0.139	−0.027	−0.027	−0.024	0.015**	0.015**	0.016**
	(0.246)	(0.248)	(0.246)	(0.023)	(0.023)	(0.024)	(0.006)	(0.006)	(0.006)
spei12 24mo		0.068	0.116		0.020	0.018		0.002	0.002
		(0.115)	(0.108)		(0.013)	(0.013)		(0.003)	(0.003)
Observations	13901	13901	13477	13902	13902	13479	13754	13754	13333
R^2	0.445	0.445	0.439	0.455	0.455	0.463	0.469	0.469	0.477
Mean of dep. var.	0.46	0.46	0.46	0.07	0.07	0.07	0.01	0.01	0.01
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the effect of predicted well failures on time spent in casual work and the intra-household allocation of casual labor. The dependent variable in Columns (1)-(3) is the total number of hours the household spent on casual work during the seven days preceding the interview. Columns (4)-(6) report the share of total casual labor hours contributed by adults, while Columns (7)-(9) report the share contributed by children. The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey.

All regressions include household, zone-wave, and month fixed effects. Specifications vary in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 22: Impact of Predicted Well Failures on Child Participation in Casual Work and Time Allocation

	Nb Child in Casual Work per Household Member			Casual Time per Child Casual Worker			Child Casual Time per Household Member		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
well fail	0.006**	0.006**	0.007**	0.000	0.000	0.000	0.108	0.108	0.120
	(0.003)	(0.003)	(0.003)	(.)	(.)	(.)	(0.087)	(0.087)	(0.091)
spei12 24mo		0.001	0.001		0.000	0.000		0.015	−0.006
		(0.001)	(0.001)		(.)	(.)		(0.030)	(0.026)
Observations	11657	11657	11306	10	10	10	11660	11660	11309
R^2	0.512	0.512	0.523	0.716	0.716	1.000	0.475	0.475	0.472
Mean of dep. var.	0.00	0.00	0.00	10.20	10.20	10.20	0.05	0.05	0.05
Fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓		✓	✓
Controls			✓			✓			✓

Note: This table reports the effect of predicted well failures on child involvement in casual work. Columns (1)-(3) estimate the impact on the number of child engaged in casual work per household member. Columns (4)-(6) report effects on the average casual work time per child casual worker (in hours). Columns (7)-(9) examine total child casual work time per household member (in hours). The main explanatory variable, *well fail*, equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional.

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Graphs

Figure 1: Main Reported Water Source

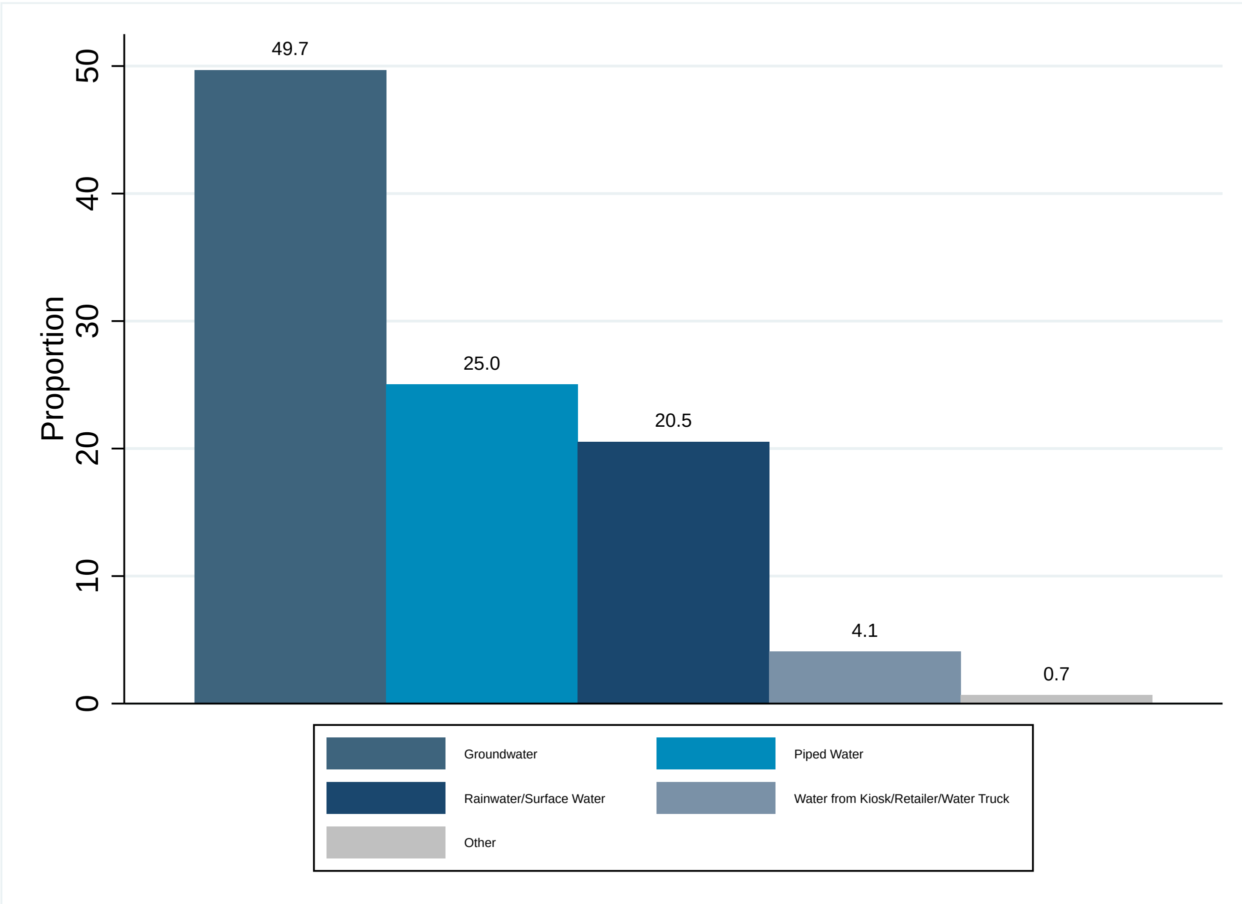

Note: The figure displays the distribution of households by their main reported source of drinking water, based on responses given for either the rainy or the dry season.
Source: Author's elaboration on ESS data.

Figure 2: Location of Water Points across Ethiopia

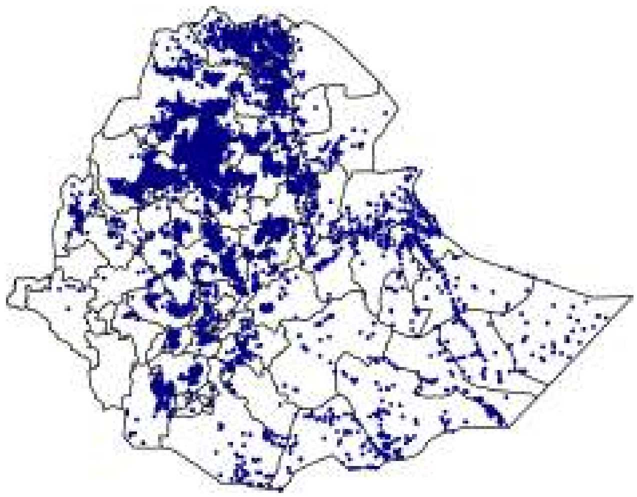

Note: The figure shows the geographic distribution of water points across Ethiopia, based on reported geocoordinates.
Source: Author's elaboration on mWater and WPDx data.

Figure 3: Report Year of Water Points

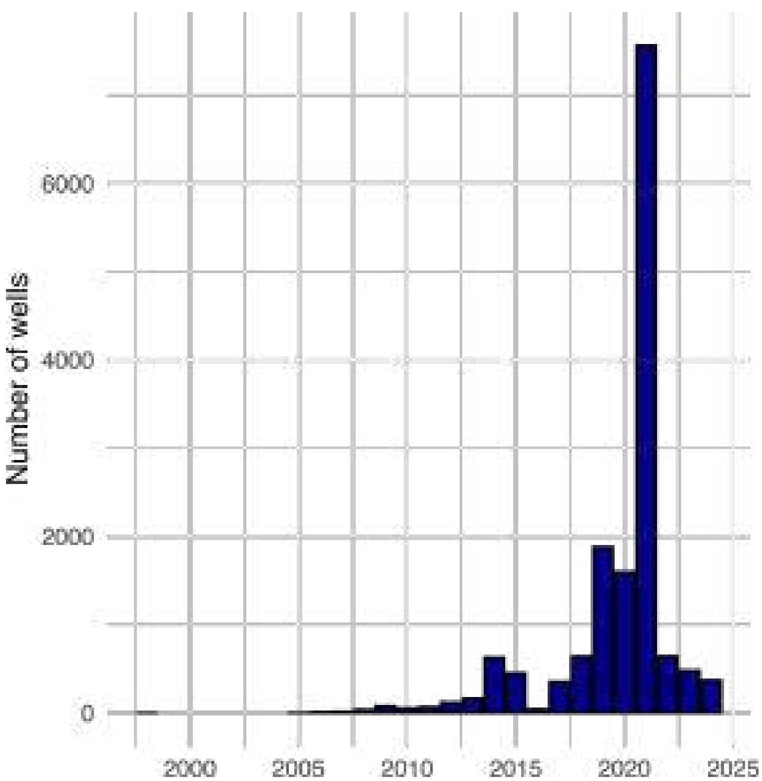

Note: The figure summarizes the year of monitoring of water points across Ethiopia.
Source: Author's elaboration on mWater and WPDx data.

Figure 4: Reported Functionality Status of Water Points

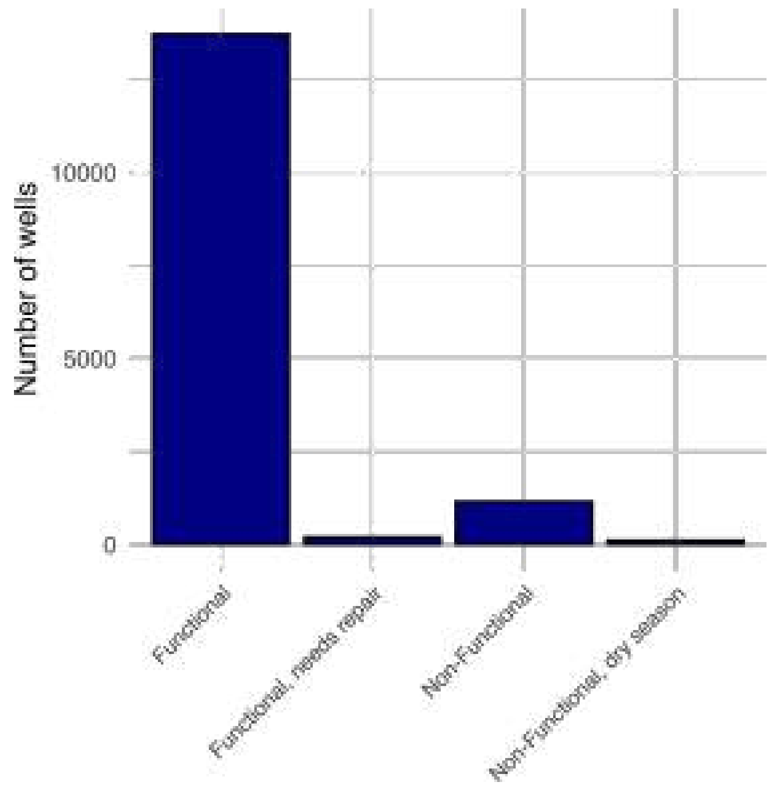

Note: The figure summarizes the reported functionality status of water points at the time of data collection.
Source: Author's elaboration on mWater and WPDx data.

Figure 5: Long Term Average of Monthly Precipitation

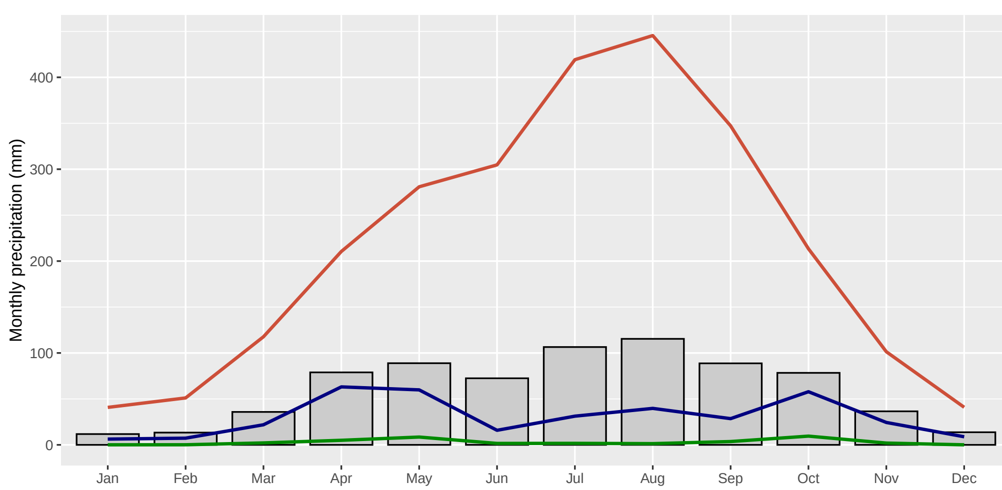

Note: The Figure represent the long-term average of the monthly precipitation (1970-2022) (mm) over Ethiopia. Red lines plot the 95th percentile of rainfall distribution, blue lines the 50th percentile, and green lines the 5th percentile.
Source: Author's elaboration on ERA-5 data.

Figure 6: Precipitation patterns in 2010 and 2015.

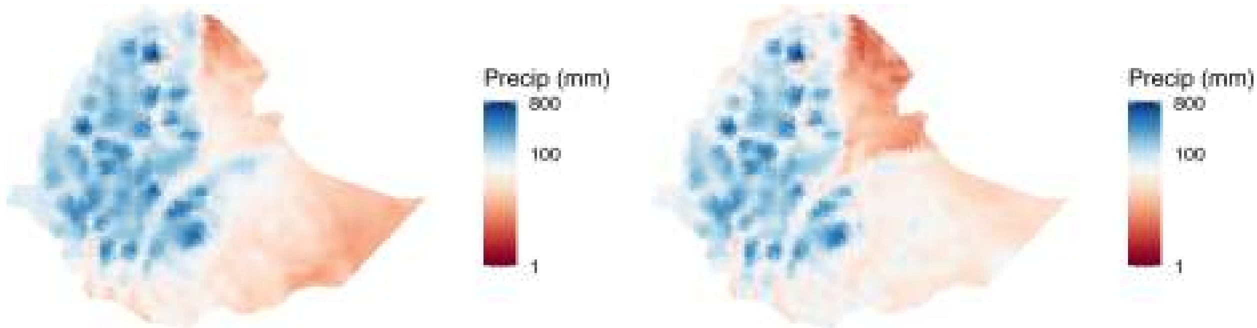

(a) 2010 **(b)** 2015

Note: The Figure represent the spatial distribution of yearly precipitation in 2010 and 2015.
Source: Author's elaboration on ERA-5 data.

Figure 7: 12 months SPEI in 2010 and 2015.

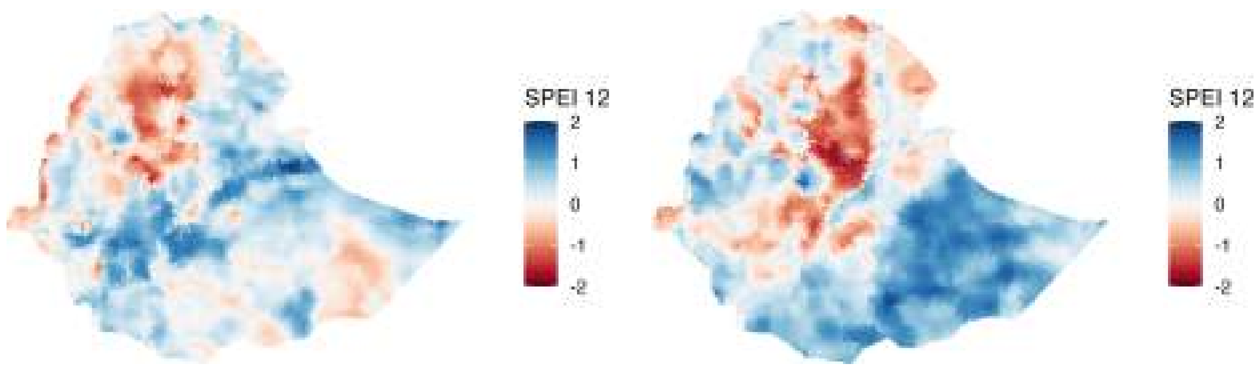

(a) 2010 **(b)** 2015

Note: The Figure represent the spatial distribution of yearly drought (measured with the 12 months SPEI) in 2010 and 2015.
Source: Author's elaboration on ERA-5 data.

Figure 8: Variation in access to groundwater

(a) Mean depth to groundwater (Fan et al., 2013)

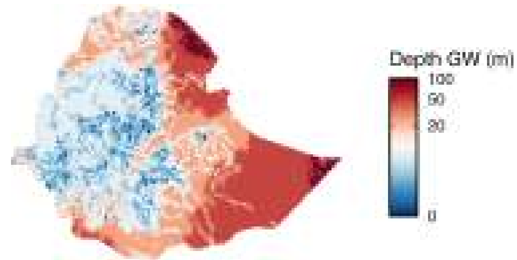

(b) Mean depth to groundwater (MacDonald et al., 2012)

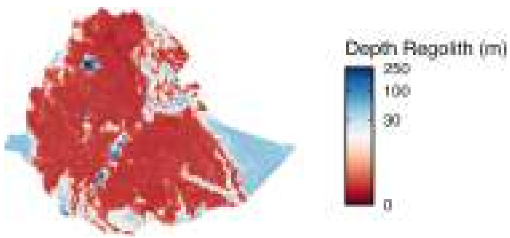

(c) Mean depth to regolith (Pelletier et al., 2016)

Note: The Figure represent the spatial distribution of groundwater.
Source: Author's elaboration on Fan et al. (2013), MacDonald et al. (2012), and Pelletier et al. (2016) data.

Figure 9: Functionality by Management

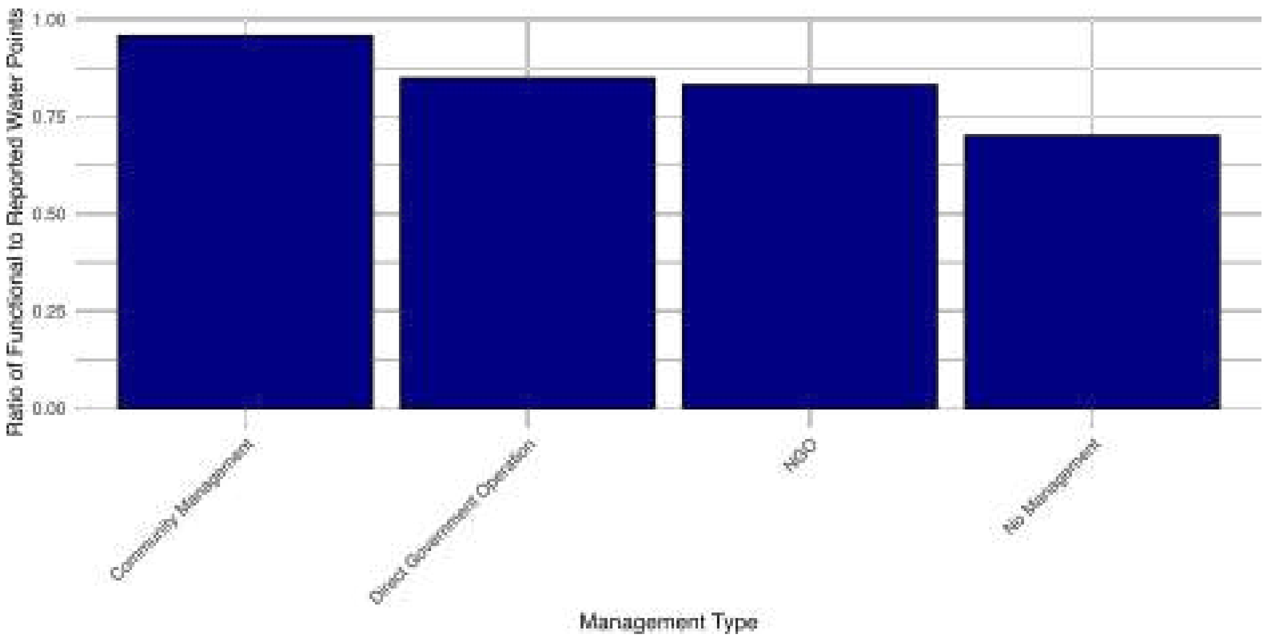

Note: The figure shows the proportion of functional water points by reported management type.
Source: Author's elaboration on mWater and WPDx data.

Figure 10: Functionality by Technology

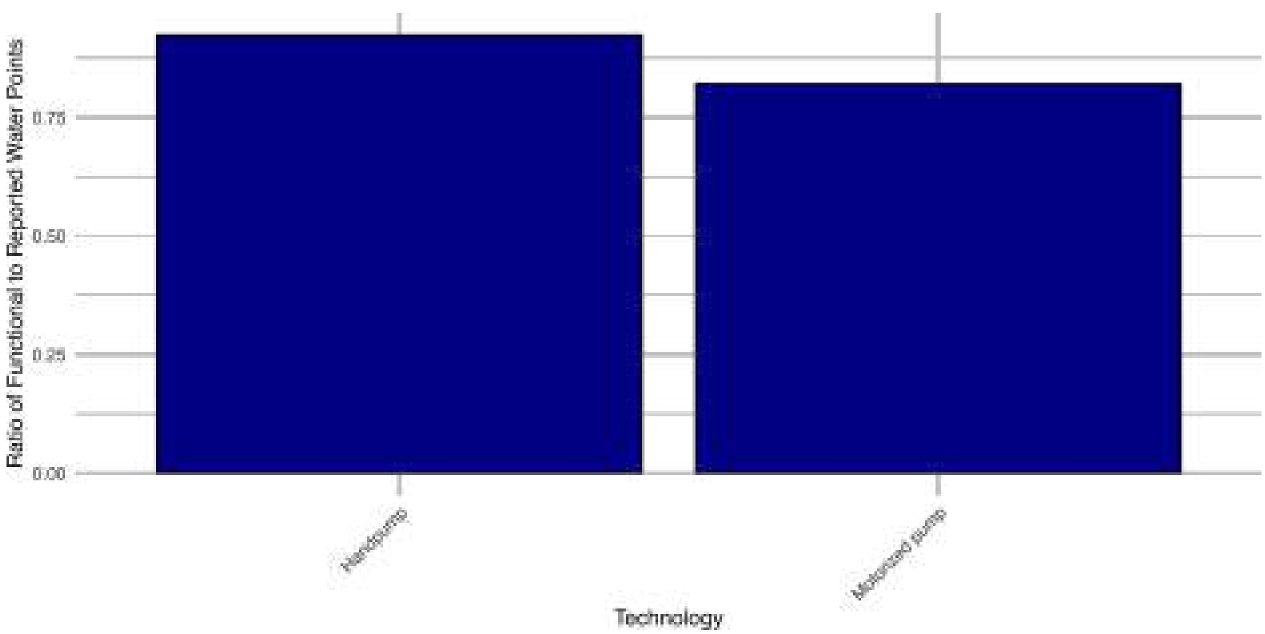

Note: The figure shows the proportion of functional water points by reported technology.
Source: Author's elaboration on mWater and WPDx data.

Figure 11: Functionality by Age

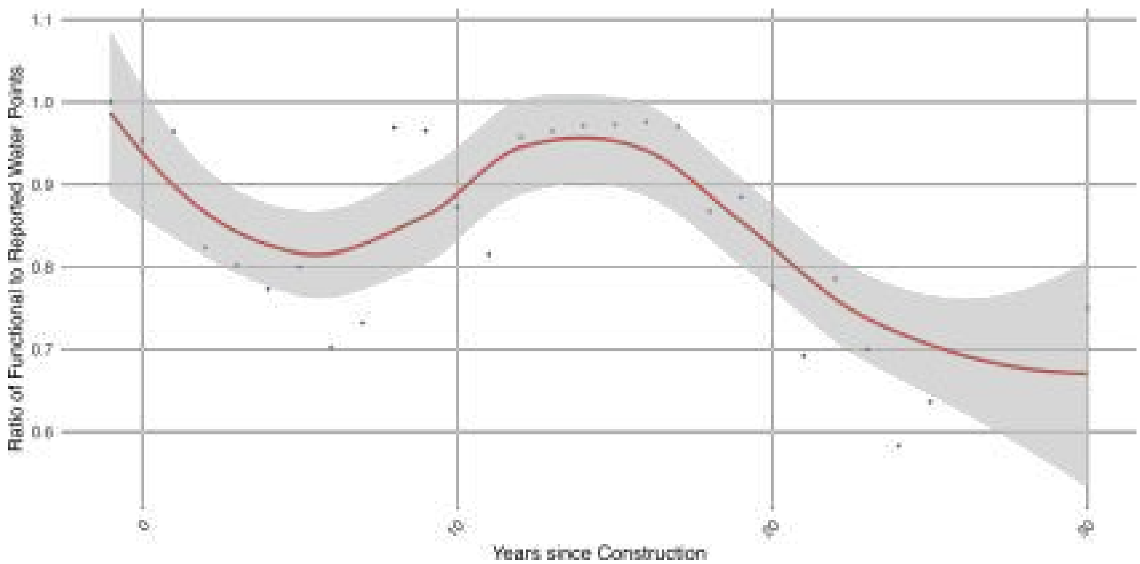

Note: The figure plots the ratio of functional to reported water points against years since construction. The red line represents a locally smoothed regression (LOESS), and the shaded grey area shows the 95% confidence interval around the fit.

Source: Author's elaboration on mWater and WPDx data.

Figure 12: Functionality by 6 SPEI in the last 12 months

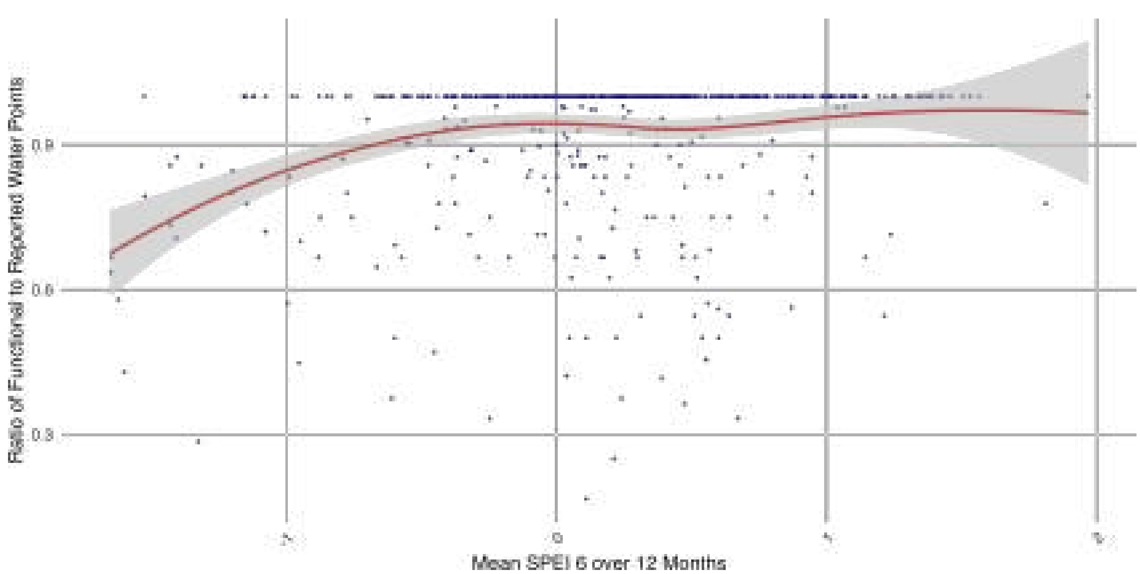

Note: The figure plots the ratio of functional to reported water points against the mean 6 months SPEI over the year. The red line represents a locally smoothed regression (LOESS), and the shaded grey area shows the 95% confidence interval around the fit.

Source: Author's elaboration on mWater, WPDx data and ERA-5 data.

Figure 13: Functionality by average monthly precipitation

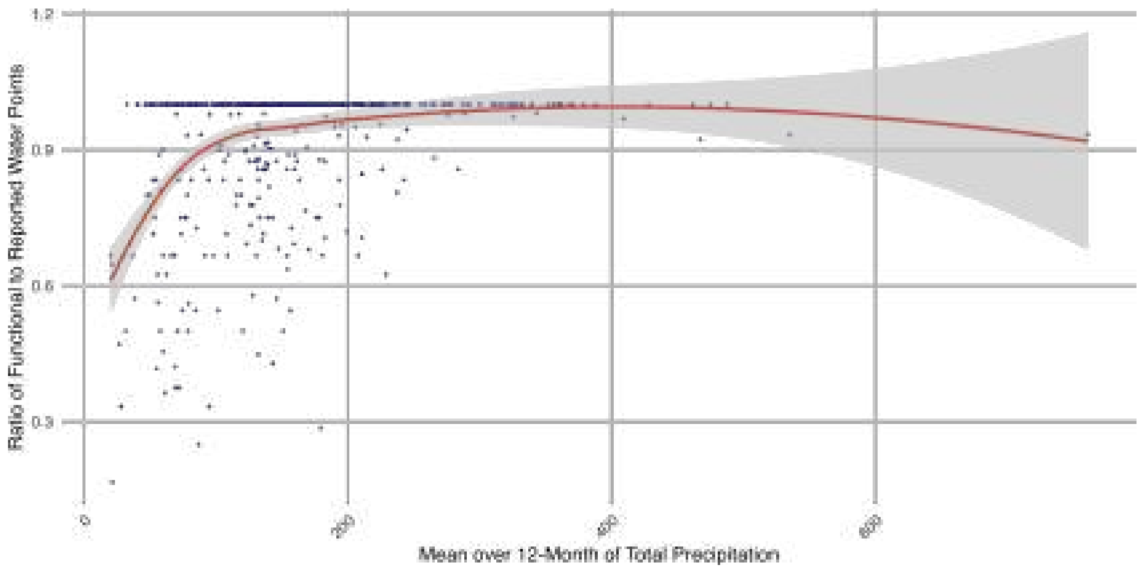

Note: The figure plots the ratio of functional to reported water points against the monthly precipitation. The red line represents a locally smoothed regression (LOESS), and the shaded grey area shows the 95% confidence interval around the fit.
Source: Author's elaboration on mWater, WPDx data and ERA-5 data.

Figure 14: Functionality by average monthly precipitations lagged two years

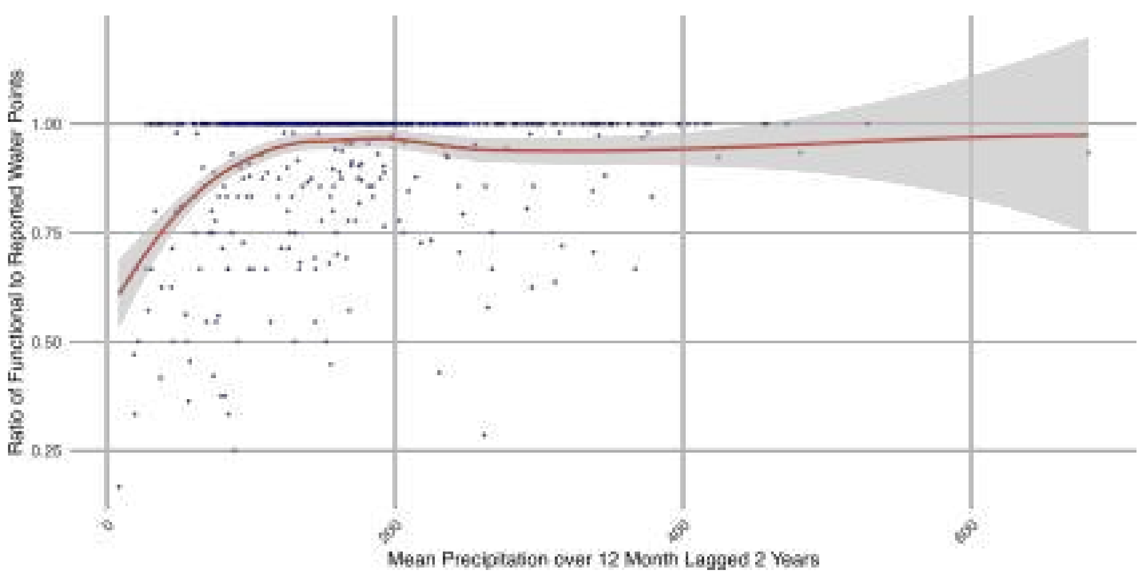

Note: The figure plots the ratio of functional to reported water points against the monthly precipitation from 2 years ago. The red line represents a locally smoothed regression (LOESS), and the shaded grey area shows the 95% confidence interval around the fit.
Source: Author's elaboration on mWater, WPDx data and ERA-5 data.

Figure 15: Functionality by groundwater depth (measured by Fan et al. (2013))

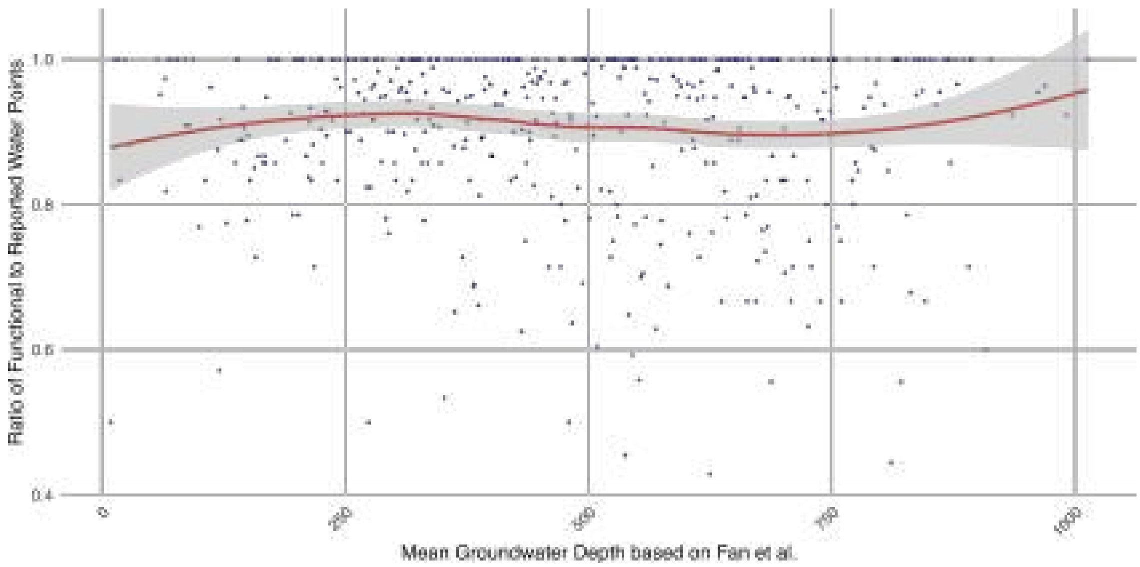

Note: The figure plots the ratio of functional to groundwater depth measured by Fan et al. (2013). The red line represents a locally smoothed regression (LOESS), and the shaded grey area shows the 95% confidence interval around the fit.
Source: Author's elaboration on mWater, WPDx, and Fan et al. (2013) data.

Figure 16: Functionality by groundwater depth (measured by MacDonald et al. (2012))

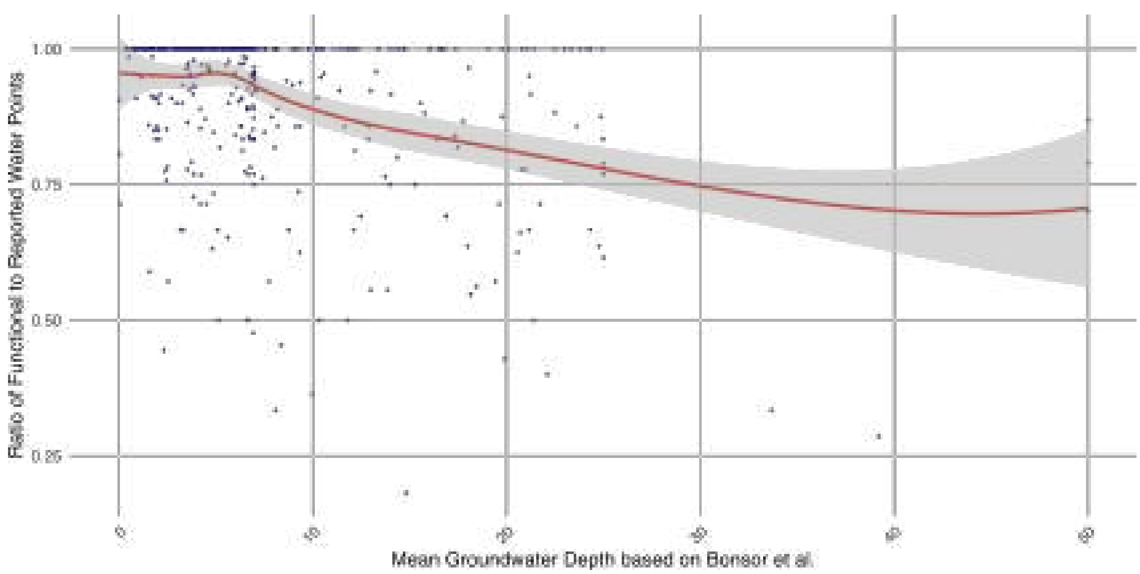

Note: The figure plots the ratio of functional to groundwater depth measured by MacDonald et al. (2012). The red line represents a locally smoothed regression (LOESS), and the shaded grey area shows the 95% confidence interval around the fit.
Source: Author's elaboration on mWater, WPDx, and MacDonald et al. (2012) data.

Figure 17: ROC Curve

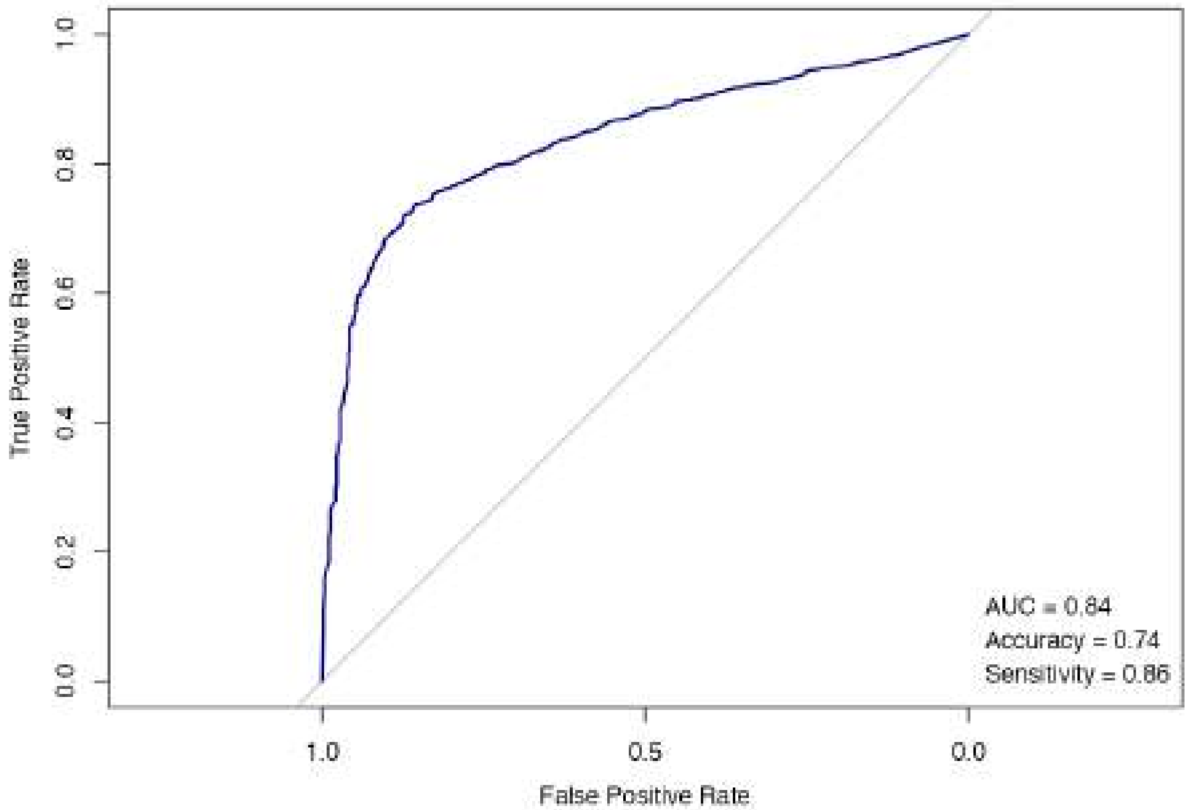

Note: The figure displays the Receiver Operating Characteristic (ROC) curve for the random forest model predicting well functionality. The blue line plots the true positive rate (sensitivity) against the false positive rate. The grey diagonal represents the performance of a no-skill classifier. The area under the curve (AUC) refers to the area under the Receiver Operating Characteristic (ROC) curve, which is commonly used to evaluate the performance of binary classification models. The AUC of 0.84 indicates strong overall discriminative ability. Accuracy, defined as the proportion of correctly classified observations, is 0.74. Sensitivity, or the true positive rate, measures the model's ability to correctly identify functional wells and is equal to 0.86.
Source: Author's elaboration.

Figure 18: Pair-wise Correlations

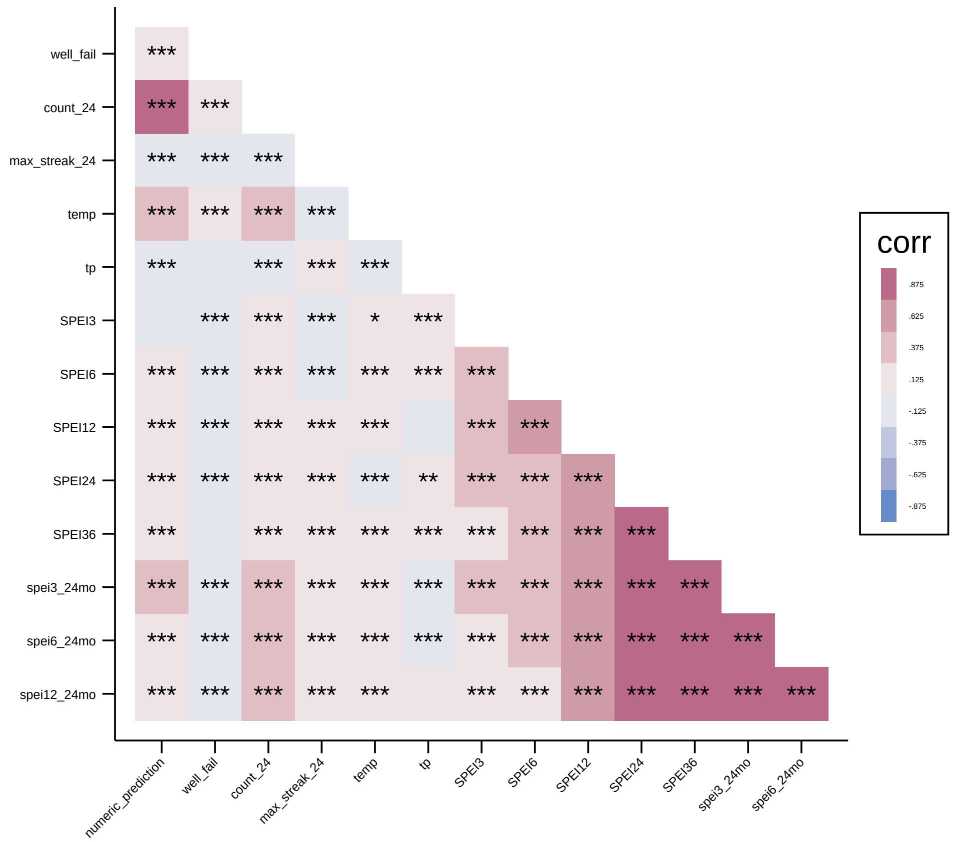

Note: The figure shows a pairwise correlation matrix of the variables included in the empirical specification. The colour intensity reflects the strength and direction of the correlations, ranging from -1 to 1. Each cell shows the correlation coefficient alongside its significance levels, with * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Source: Author's elaboration.

Figure 19: Impact of Predicted Well Failures on Reported Time to Water Gender Heterogeneity

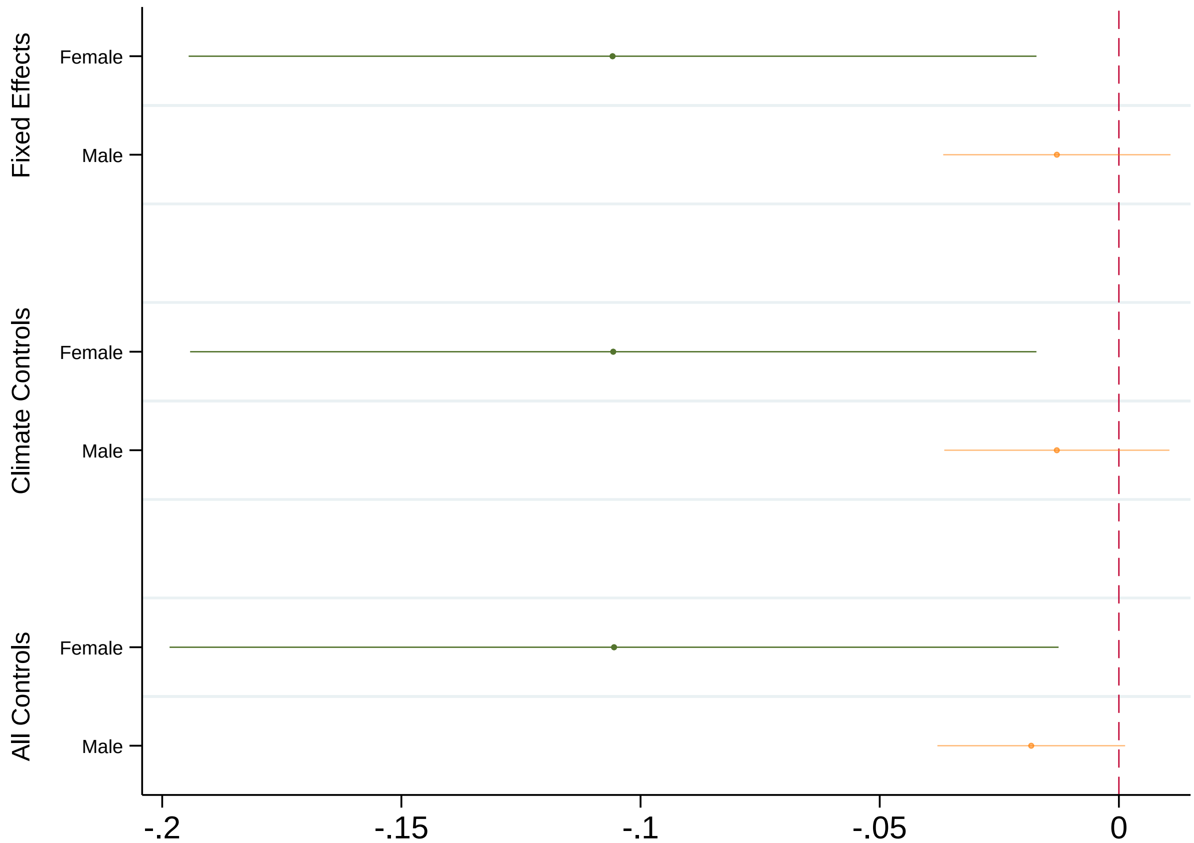

Note: The figure shows the estimated effects of predicted well failures on reported time spent collecting water, separately for men and women. Dots represent point estimates, and horizontal lines indicate 95% confidence intervals. Each row corresponds to a different regression specification: with only fixed effects, with climate controls, and with the full set of controls. Statistically significant effects at the 5% level are displayed in darker colors. *Source*: Author's elaboration based on household survey and predicted well failure data.

C Appendix

C.1 Agricultural outcomes

A potential concern is that the machine learning-based well failure indicator may inadvertently capture broader climatic stress – particularly shocks that affect agricultural production – rather than isolating disruptions in drinking water access. To assess this possibility, I examine whether the predicted well failure measure is associated with a range of agricultural outcomes. As a reminder, the baseline specification is given by:

$$Y_{h,t} = \beta \cdot WellFail_{c,t} + \gamma_1 Climate_{c,t} + \gamma_2 X_{h,t} + \delta_h + \delta_{z\times t} + \delta_m + \epsilon_{h,t} \tag{2}$$

where $WellFail_{c,t}$ is an indicator that household h within EA c experienced a predicted well failure in the month leading up to survey t [10]. The coefficient β is the parameter of interest, measuring the impact of a well failure on the outcome. Wells failure could be due to other endogeneous variable, to limit those concerns I add household and community controls as well as climate controls. I also include household fixed effects δ_h to control for all time-invariant characteristics of the household (such as baseline wealth, location, or any fixed propensity to use children's labor). I also include region-time fixed effects $\delta_{z\times t}$ to control for any time-varying zone level (second-level administrative units in Ethiopia) characteristics (such as conflict or economic trends). Standard errors are clustered at the EA level. In this appendix, all outcomes will be agricultural.

If the well failure indicator were simply capturing general environmental hardship, one would expect it to be significantly associated with declines in agricultural performance. However, I find no such relationship with harvest quantity, irrigation behavior, self-reported crop shocks, or measures of land use such as farm size and number of plots. These results reinforce the interpretation that the well failure prediction captures changes in drinking water access rather than agricultural productivity.

[10] Since the ESS provides geolocation only at the EA level, indicators for well failures and climate climate conditions are constructed at the EA level.

Table 23: Impact of Predicted Well Failures on Reported Harvest Quantity

	Well fail			Longest streak		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.042	0.038	0.046			
	(0.211)	(0.205)	(0.208)			
Longest wt functional well				−0.001	−0.001	−0.001
				(0.003)	(0.003)	(0.003)
spei12 24mo		−0.144*	−0.140*		−0.126*	−0.117
		(0.074)	(0.074)		(0.074)	(0.075)
Observations	9556	9556	9143	9448	9448	9043
R^2	0.727	0.727	0.733	0.728	0.728	0.733
Mean of dep. var.	708.67	708.67	708.67	699.08	699.08	699.08
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the relationship between predicted well failures and the quantity harvested, measured in kilograms. The dependent variable in all columns is the logarithm of the total quantity harvested. Columns (1)-(3) use a binary indicator, *well fail*, which equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey. Columns (4)-(6) instead use the length of the longest failure spell (*Longest spell without a functional well*) in the 24 months prior to the survey. All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 24: Impact of Predicted Well Failures on Irrigation

	Well fail			Longest streak		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	−0.048	−0.048	−0.047			
	(0.044)	(0.044)	(0.043)			
Longest wt functional well				0.000	0.000	0.000
				(0.000)	(0.000)	(0.000)
spei12 24mo		0.001	0.003		0.000	0.003
		(0.013)	(0.012)		(0.013)	(0.012)
Observations	8326	8326	7964	8255	8255	7897
R^2	0.694	0.694	0.703	0.695	0.695	0.704
Mean of dep. var.	0.11	0.11	0.11	0.12	0.12	0.12
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the relationship between predicted well failures and irrigation use. The dependent variable in all columns is a binary indicator, which equals 1 if the household reports irrigating at least one of its fields. Columns (1)-(3) use a binary indicator, *well fail*, which equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey. Columns (4)-(6) instead use the length of the longest failure spell (*Longest spell without a functional well*) in the 24 months prior to the survey.

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 25: Impact of Predicted Well Failures on Farm Size

	Well fail			Longest streak		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.941	0.942	1.103			
	(0.857)	(0.855)	(0.917)			
Longest wt functional well				−0.019	−0.019	−0.024
				(0.016)	(0.017)	(0.018)
spei12 24mo		−0.128	−0.064		0.016	0.063
		(0.195)	(0.204)		(0.187)	(0.206)
Observations	12201	12201	11789	12043	12043	11635
R^2	0.428	0.428	0.419	0.402	0.402	0.404
Mean of dep. var.	1.33	1.33	1.33	1.28	1.28	1.28
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the relationship between predicted well failures and farm size. The dependent variable in all columns is the size of the household's farmland, measured in square kilometers using GPS devices. Columns (1)-(3) use a binary indicator, *well fail*, which equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey. Columns (4)-(6) instead use the length of the longest failure spell (*Longest spell without a functional well*) in the 24 months prior to the survey.

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 26: Impact of Predicted Well Failures on the Number of Plots

	Well fail			Longest streak		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	−0.356	−0.356	−0.283			
	(0.384)	(0.384)	(0.382)			
Longest wt functional well				−0.002	−0.002	−0.003
				(0.007)	(0.007)	(0.007)
spei12 24mo		−0.038	0.020		−0.031	0.031
		(0.268)	(0.267)		(0.269)	(0.268)
Observations	13928	13928	13480	13751	13751	13307
R^2	0.856	0.856	0.861	0.855	0.855	0.860
Mean of dep. var.	8.65	8.65	8.65	8.71	8.71	8.71
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the relationship between predicted well failures and the number of plots. The dependent variable in all columns is the number of plots the household's farmland. Columns (1)-(3) use a binary indicator, *well fail*, which equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey. Columns (4)-(6) instead use the length of the longest failure spell (*Longest spell without a functional well*) in the 24 months prior to the survey.

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 27: Impact of Predicted Well Failures on Plots Size

	Well fail			Longest streak		
	(1)	(2)	(3)	(4)	(5)	(6)
well fail	0.186	0.186	0.205			
	(0.151)	(0.150)	(0.159)			
Longest wt functional well				−0.002	−0.002	−0.003
				(0.002)	(0.002)	(0.003)
spei12 24mo		−0.022	−0.014		−0.000	0.004
		(0.029)	(0.030)		(0.027)	(0.030)
Observations	12201	12201	11789	12043	12043	11635
R^2	0.420	0.420	0.413	0.396	0.396	0.398
Mean of dep. var.	0.16	0.16	0.16	0.15	0.15	0.15
Fixed effects	✓	✓	✓	✓	✓	✓
Climate		✓	✓		✓	✓
Controls			✓			✓

Note: This table reports the relationship between predicted well failures and plots size. The dependent variable in all columns is the mean size of plots in the household's farmland, measured in square kilometers using GPS devices. Columns (1)-(3) use a binary indicator, *well fail*, which equals 1 if the representative drinking water well in the household's enumeration area is predicted to be non-functional at the time of the survey. Columns (4)-(6) instead use the length of the longest failure spell (*Longest spell without a functional well*) in the 24 months prior to the survey.

All regressions include household, zone-wave, and month fixed effects. Specifications differ in the inclusion of climate controls (SPEI-12 over the past 24 months) and household-level covariates: distance to the nearest town and road, consumption quintile, percentage of agricultural land within a 1km buffer, and an agricultural asset index. Standard errors, clustered at the enumeration area level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.2 Details on the algorithm

The objective of the algorithm is to predict whether a well is functional or not. To remain agnostic about the potential drivers of well failure, I begin with a comprehensive set of 373 hydro-environmental variables. These variables are listed in Table 28. Each well is geolocated and linked to its corresponding basin, and all variables are constructed at the basin level. This spatial scale is chosen because it is difficult to determine which specific precipitation events directly influence individual wells. Variables marked with an asterisk ($*$) are sourced from the *BasinATLAS* database.

In the first step, I reduce dimensionality and identify the most relevant predictors using the Least Absolute Shrinkage and Selection Operator (LASSO) with cross-validation. This approach helps select the most predictive variables while minimizing the risk of overfitting. Specifically, I use the `cv.glmnet` function in `R`, which performs k-fold cross-validation (with $k = 10$ by default) across a grid of candidate values for the regularization parameter λ. The optimal λ is chosen as the value that minimizes the cross-validated prediction error (measured by deviance in the case of logistic regression), thus striking a balance between model complexity and predictive performance. Once this value is selected, I refit the LASSO model using the optimal λ and retain the variables with non-zero coefficients – these are considered the most informative features for distinguishing between functional and non-functional wells. The 235 variables selected through this procedure are listed in Table 29.

In the second step, I introduce interaction terms between the selected predictors and a set of 24 time-invariant variables (listed in Table 30). These variables are chosen based on their theoretical importance in shaping groundwater flows and stocks, as highlighted in the hydrology and geoscience literature. Because these characteristics do not vary over time, they are well-suited to capture spatial heterogeneity in how different locations respond to temporal changes in environmental conditions. Including these interactions allows the model to account for potential effect modification by underlying hydrogeological contexts. After generating these interaction terms, I once again apply a LASSO procedure to reduce dimensionality and isolate the most relevant predictors. This results in a final set of 383 variables, which are then used to train a random forest model.

To train and evaluate the predictive performance of the random forest classifier, I begin by splitting the dataset into a training set (70%) and a testing set (30%). This split is stratified by the outcome variable to ensure that both sets maintain the original distribution of functional and non-functional wells. The outcome variable is then explicitly re-leveled to set "functional" wells as the reference category for classification. Next, I address the class imbalance in the training data.

Since the number of functional wells substantially exceeds the number of non-functional ones, the model might otherwise perform poorly in detecting failures. To mitigate this, I apply a random downsampling procedure to the majority class within the training set: I randomly sample from the functional wells to match the number of non-functional wells, resulting in a balanced training sample. This ensures that the classifier gives equal importance to both classes during training. I then train a random forest classifier using the downsampled training set. I implement the model using the `caret` package in `R`, specifying 5-fold cross-validation to guard against overfitting and assess model performance more robustly. The classifier is trained to maximize the area under the ROC curve (AUC), which is particularly useful for imbalanced classification problems. I include a class weight of 10 for non-functional wells and 1 for functional wells to further correct for the original imbalance in prediction. The number of variables considered at each split (`mtry`) is pre-tuned and set to 18. This procedure ensures that the model is trained on a balanced dataset, uses rigorous cross-validation, and is evaluated on an independent test set that reflects the true distribution of well statuses.

Figure 20 shows the top-ranked variables in terms of their importance in the random forest model used to predict whether a well is functional or not. The variable importance scores are derived from the model's internal assessment of how much each feature contributes to improving prediction accuracy – higher bars indicate greater importance in distinguishing functional from non-functional wells. The top predictor is current precipitation interacted with the maximum groundwater depth in the basin measured by MacDonald et al. (2012). This suggests that rainfall's ability to influence well functionality is highly contingent on the depth of groundwater – likely because shallow aquifers recharge more easily. Similar patterns emerge with groundwater measured in Fan et al. (2013) and several SPEI interactions (e.g., `spei1_12mo:max_gw_bonsor`, `tp:mean_gw_fan`), confirming that climate anomalies (precipitation and drought) are more predictive when contextualized by the underlying groundwater regime. A notable feature of the variable ranking is the prominence of lagged climate variables. These lagged variables capture accumulated hydrological stress and suggest that well failure often reflects long-term climatic pressures rather than short-term shocks. Interactions between these lagged indices and soil or erosion indicators, such as `spei6_36mo_lag5:ero_kh_sav`, further imply that environmental degradation may exacerbate vulnerability to drought. In contrast, raw climate variables without spatial interactions appear less informative.

One potential concern is the possibility of information leakage, since the model is trained and tested across both space and time. Specifically, the testing set may contain observations from the same basin or the same month as those in the training set, which could lead to an overestimation of the model's predictive performance. While I am less concerned about this issue – given the robustness checks using reported water source failures (Table 2) – I acknowledge its relevance. For full transparency, I therefore report performance metrics under alternative data-splitting strategies: one where the train/test split is done at the basin level (Figure 21), and another where it is done by year (Figure 22). Results show that the model performs better when predicting across space than across time, suggesting stronger generalizability in the spatial dimension than in the temporal one.

Table 28: Underlying Data Sources for Model Inputs

Category	Attribute	Source Data	Citation	Columns
Climate	Current SPEI (1, 3, 6, 12, 24, 36)	ERA-5		SPEI.
	Mean SPEI over 12, 24, 60 months	ERA-5		spei.
	Lagged Mean SPEI (1-5 years)	ERA-5		spei.
	Current precipitation	ERA-5		tp
	Mean precipitation over 12, 24, 60 months	ERA-5		tp_.
	Lagged Mean precipitation (1-5 years)	ERA-5		tp_.
	Current temperature	ERA-5		temp
	Mean temperature over 12, 24, 60 months	ERA-5		temp_.
	Lagged Mean temperature (1-5 years)	ERA-5		temp_.
	Snow Cover Extent ∗	MODIS/Aqua	Hall & Riggs 2016	snw_pc_.
	Climate Zones ∗	GEnS	Metzger et al. 2013	clz_cl_.
	Climate Strata ∗	GEnS	Metzger et al. 2013	cls_cl_.
	Air Temperature ∗	WorldClim v1.4	Hijmans et al. 2005	tmp_dc_.
	Precipitation ∗	WorldClim v1.4	Hijmans et al. 2005	pre_mm_
	Potential Evapotranspiration ∗	Global-PET	Zomer et al. 2008	pet_mm_.
	Climate Moisture Index ∗	WorldClim & Global-PET	Hijmans et al. 2005	cmi_ix_.
	Global Aridity Index ∗	Global Aridity Index	Zomer et al. 2008	ari_ix_.
	Actual Evapotranspiration ∗	Global Soil-Water Balance	Trabucco & Zomer 2010	aet_mm_.
	Soil Water Content ∗	Global Soil-Water Balance	Trabucco & Zomer 2010	swc_pc_.
Hydrology	Groundwater Table Depth ∗	Global Groundwater Map	Fan et al. 2013	_fan
	Groundwater Table Depth		MacDonald et al. 2012	_bonsor
	Aquifer productivity	African Groundwater Atlas		EthHGComb
	Natural Discharge ∗	WaterGAP v2.2	D¨oll et al. 2003	dis_m3_.
	Land Surface Runoff ∗	WaterGAP v2.2	D¨oll et al. 2003	run_mn_.
	Inundation Extent ∗	GIEMS-D15	Fluet-Chouinard et al. 2015	inu_pc_.
	Limnicity (% Lake Area) ∗	HydroLAKES	Messager et al. 2016	lka_pc_.
	Lake Volume ∗	HydroLAKES	Messager et al. 2016	lkv_mc_.
	Reservoir Volume ∗	GRanD v1.1	Lehner et al. 2011	rev_mc_.
	Degree of Regulation ∗	HydroSHEDS & GRanD	Lehner et al. 2011	dor_pc_.
	River Area ∗	HydroSHEDS & WaterGAP	Lehner & Grill 2013	ria_ha_.
	River Volume ∗	HydroSHEDS & WaterGAP	Lehner & Grill 2013	riv_tc_.
Landcover	Potential Natural Vegetation Classes ∗	EarthStat	Ramankutty & Foley 1999	pnv_cl_.
	Potential Natural Vegetation Extent ∗	EarthStat	Ramankutty & Foley 1999	pnv_pc_.
Physiography	Elevation ∗	EarthEnv-DEM90	Robinson et al. 2014	ele_mt_.
	Terrain Slope ∗	EarthEnv-DEM90	Robinson et al. 2014	slp_dg_.
	Stream Gradient ∗	EarthEnv-DEM90	Robinson et al. 2014	sgr_dk_.
Soils & Geology	Depth to regolith		Pelletier et al. 2016	_regolith
	Rock type	African Groundwater Atlas		EthGLG
	Soil Erosion ∗	GloSEM v1.2	Borrelli et al. 2017	ero_kh_.
	Karst Area Extent ∗	Rock Outcrops v3.0	Williams & Ford 2006	kar_pc_.
	Lithological Classes ∗	GLiM	Hartmann & Moosdorf 2012	lit_cl_.
	Soil Water Content ∗	Global High-Resolution Soil-Water Balance	Trabucco & Zomer 2010	swc_pc_.
	Organic Carbon Content in Soil ∗	SoilGrids1km	Hengl et al. 2014	soc_th_.
	Sand Fraction in Soil ∗	SoilGrids1km	Hengl et al. 2014	snd_pc_.
	Silt Fraction in Soil ∗	SoilGrids1km	Hengl et al. 2014	slt_pc_.
	Clay Fraction in Soil ∗	SoilGrids1km	Hengl et al. 2014	cly_pc_.

Table 29: Selected Variables after Lasso

Category	Attribute	Definition
Climate	SPEI 1	Current 1 month SPEI; Binary variables of whether the current SPEI 1 was below -1, -1.5 and -2; Mean SPEI 1 over 24 and 60 months; Mean SPEI 1 over 24 months lagged 2 years; Mean SPEI 1 over 12, 24, and 36 months lagged 3 years; Mean SPEI 1 over 24, 36, and 60 months lagged 5 years.
	SPEI 3	Current 3 months SPEI; Binary variables of whether the current SPEI 3 was below -1, -1.5 and -2; Mean SPEI 3 over 12, 24, and 60 months; Mean SPEI 3 over 24 and 60 months lagged 2 years; Mean SPEI 3 over 12 and 36 months lagged 3 years; Mean SPEI 3 over 24, 36, and 60 months lagged 5 years.
	SPEI 6	Current 6 months SPEI; Binary variables of whether the current SPEI 6 was below -1, -1.5 and -2; Mean SPEI 6 over 12, 24, and 60 months; Mean SPEI 6 over 24 and 60 months lagged 2 years; Mean SPEI 6 over 24 months lagged 3 years; Mean SPEI 6 over 12, 24, 36, and 60 months lagged 5 years.
	SPEI 12	Current 12 months SPEI; Binary variables of whether the current SPEI 12 was below -1, -1.5 and -2; Mean SPEI 12 over 12 months; Mean SPEI 12 over 12, 24, and 60 months lagged 1 year; Mean SPEI 12 over 60 months lagged 3 years; Mean SPEI 12 over 12, 24, and 36 months lagged 5 years.
	SPEI 36	Current 36 months SPEI; Binary variables of whether the current SPEI 36 was below -1, -1.5 and -2; Mean SPEI 36 over 12 months; Mean SPEI 36 over 12 and 36 months lagged 1 year; Mean SPEI 36 over 12, 24, and 36 months lagged 3 years; Mean SPEI 36 over 12, 36, and 60 months lagged 5 years.
	Temperature (ERA-5)	Current mean temperature; Mean temperature over 12 months lagged 1 year; Mean temperature over 12 and 24 months lagged 2 years; Mean temperature over 12, 24, 36 and 60 months lagged 3 years; Mean temperature over 12, 24, 36, and 60 months lagged 5 years.
	Precipitation (ERA-5)	Current mean total precipitation; Mean precipitation over the year; Mean precipitation over 12 and 36 months lagged 1 year; Mean precipitation over 24 and 36 months lagged 2 years; Mean precipitation over 12 and 24 months lagged 3 years; Mean precipitation over 12, 36, and 60 months lagged 5 years.
	Climate Zones	Spatial majority in basin.
	Climate Strata	Spatial majority in basin.
	Temperature (WorldClim)	Annual average in total watershed upstream of basin; Annual minimum and maximum in basin; Long term monthly average in February, March, April, May, June, July, August, September, October and December.
	Precipitation (WorldClim)	Annual average in total watershed upstream of basin; Long term monthly average in March, April, May, June, July, August, September, October, and December.
	Potential Evapotranspiration	Annual average in the basin and in total watershed upstream of basin; Long term monthly average in January, February, April, June, July, August, September, October, and December.
	Actual Evapotranspiration	Long term monthly average in January, March, May, June, July, August, September, October, November, and December.
Hydrology	Groundwater Table Depth	Mean groundwater depth for the basin in Fan et al. 2013 and MacDonald et al. 2012; Minimum groundwater depth for the basin in Fan et al. 2013 and MacDonald et al. 2012; Maximum groundwater depth for the basin in Fan et al. 2013 and MacDonald et al. 2012.
	EthHGComb	Aquifer productivity.
	Natural Discharge	Annual basin average, minimum and maximum.
	Land Surface Runoff	Annual basin average.
	Inundation Extent	Annual basin minimum, maximum and long-term maximum; Annual minimum and long-term maximum in total watershed upstream of basin.
	Limnicity	Spatial extent in basin and in total watershed upstream of basin.
	Lake volume	Sum in total watershed upstream of basin.
	Reservoir volume	Sum in total watershed upstream of basin.
	Degree of Regulation	Index in basin.
	River Area	Sum in basin and in total watershed upstream of basin.
	River Volume	Sum in basin and in total watershed upstream of basin.
Landcover	Potential Natural Vegetation Classes	Spatial majority in basin.
	Potential Natural Vegetation Extent	Extent by class in basin and in total watershed upstream of basin.
Physiography	Elevation	Average, minimum, and maximum in the basin.
	Slope	Average in the basin and in total watershed upstream of basin.
	Stream Gradient	Average in the basin.
	Global Aridity Index	Average in basin and in in total watershed upstream of basin.
	Climate Moisture Index	Long term monthly average, in basin, in January, February, March, April, March, June, July, September, November, December.
Soils & Geology	Depth to regolith	Mean, minimum and maximum in basin.
	EthGLG	Rock type.
	Soil Erosion	Average in basin and in total watershed upstream of basin.
	Soil Water Content	Annual average in basin and in total watershed upstream of basin; Long term monthly average, in basin, in January, February, March, July, August, September, November, and December.
	Karst Area Extent	Spatial extent in basin and in total watershed upstream of basin.
	Lithological Classes	Spatial majority in basin.
	Organic Carbon Content in Soil	Average in basin and in total watershed upstream of basin.
	Sand Fraction in Soil	Average in basin and in total watershed upstream of basin.
	Silt Fraction in Soil	Average in basin and in total watershed upstream of basin.
	Clay Fraction in Soil	Average in basin and in total watershed upstream of basin.

Table 30: Selected Variables for Interaction

Category	Attribute	Source Data	Citation	Columns
Climate	Global Aridity Index ∗	Global Aridity Index	Zomer et al. 2008	ari_ix_
Hydrology	Groundwater Table Depth ∗	Global Groundwater Map	Fan et al. 2013	_fan
	Groundwater Table Depth		MacDonald et al. 2012	_bonsor
	Aquifer productivity	African Groundwater Atlas		EthHGComb
Landcover	Potential Natural Vegetation Classes ∗	EarthStat	Ramankutty & Foley 1999	pnv_cl_
Physiography	Elevation ∗	EarthEnv-DEM90	Robinson et al. 2014	ele_mt_
	Terrain Slope ∗	EarthEnv-DEM90	Robinson et al. 2014	slp_dg_
Soils & Geology	Depth to regolith		Pelletier et al. 2016	_regolith
	Rock type	African Groundwater Atlas		EthGLG
	Soil Erosion ∗	GloSEM v1.2	Borrelli et al. 2017	ero_kh_
	Karst Area Extent ∗	Rock Outcrops v3.0	Williams & Ford 2006	kar_pc_
	Lithological Classes ∗	GLiM	Hartmann & Moosdorf 2012	lit_cl_
	Sand Fraction in Soil ∗	SoilGrids1km	Hengl et al. 2014	snd_pc_
	Silt Fraction in Soil ∗	SoilGrids1km	Hengl et al. 2014	slt_pc_
	Clay Fraction in Soil ∗	SoilGrids1km	Hengl et al. 2014	cly_pc_

Figure 20: Main Model Variable Importance

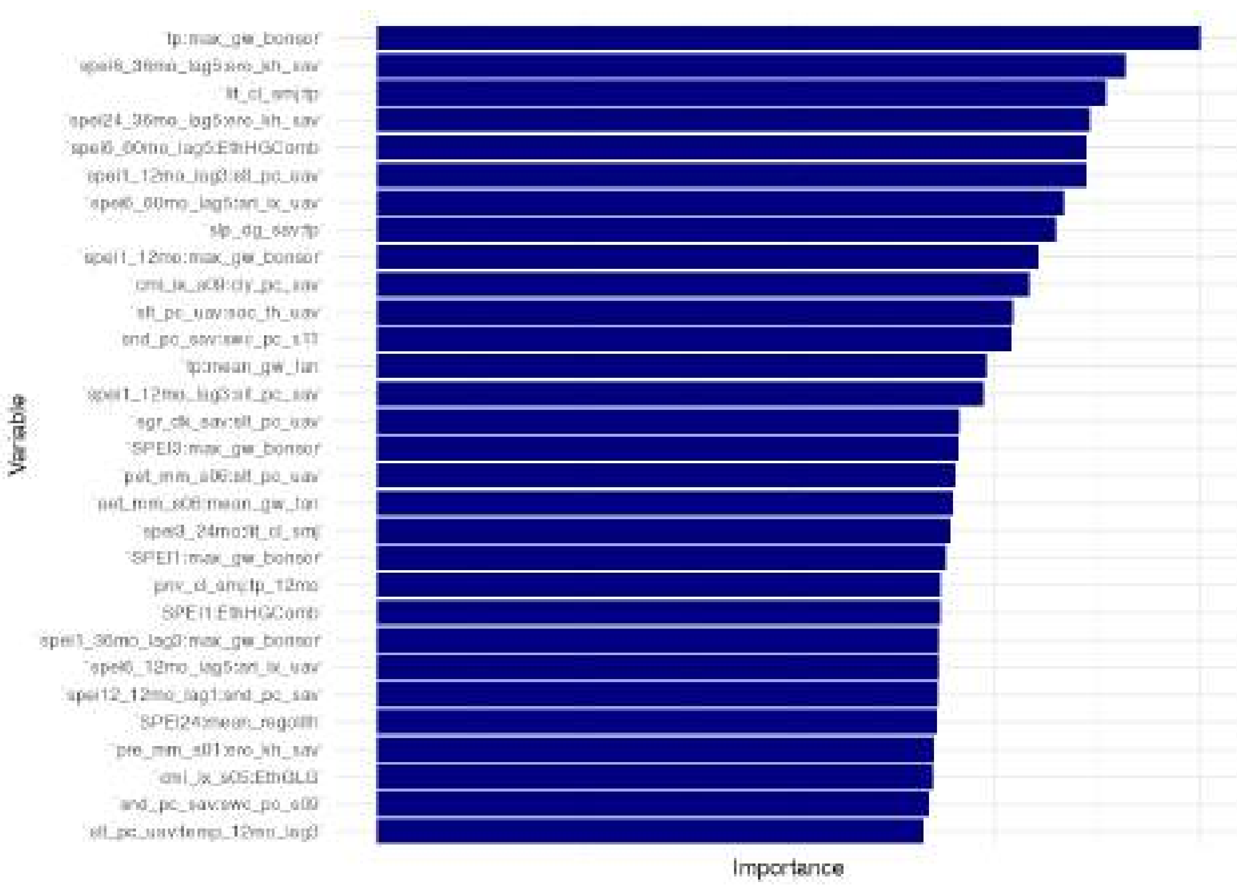

Figure 21: Basin Random Split Model Performance - ROC Curve

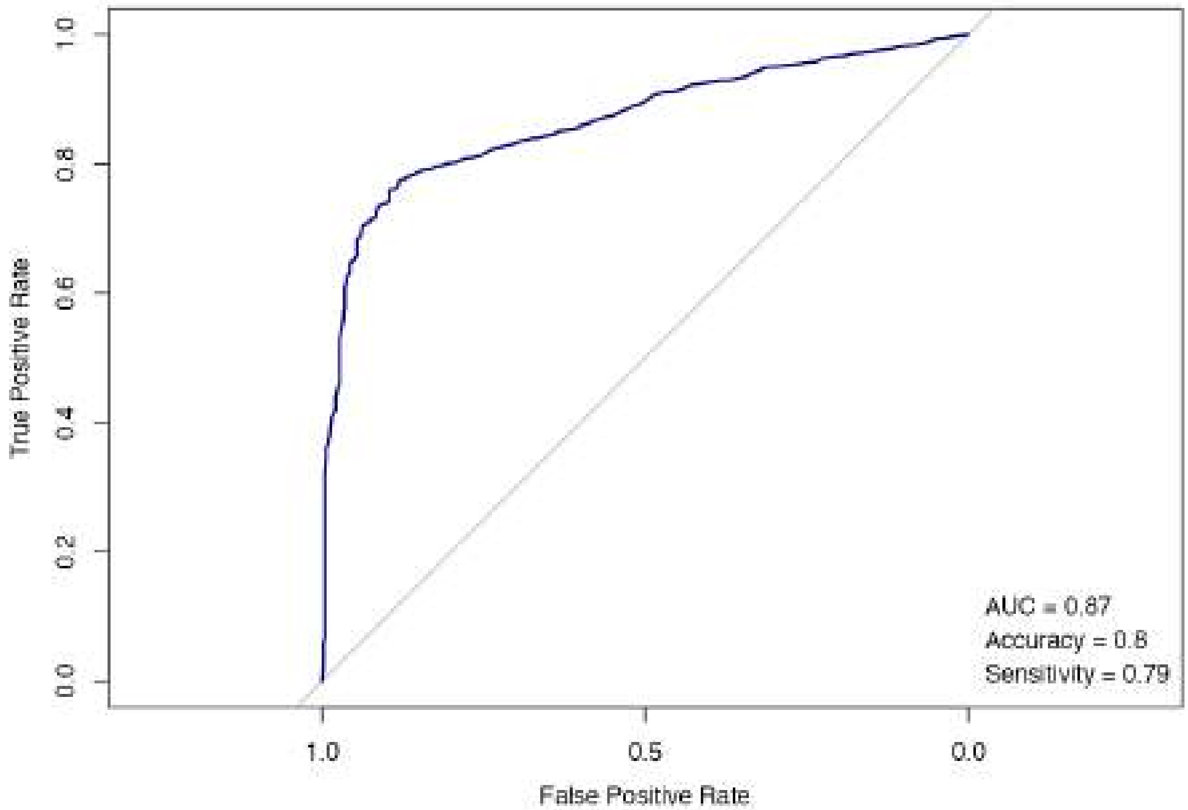

Figure 22: Year Random Split Model Performance - ROC Curve

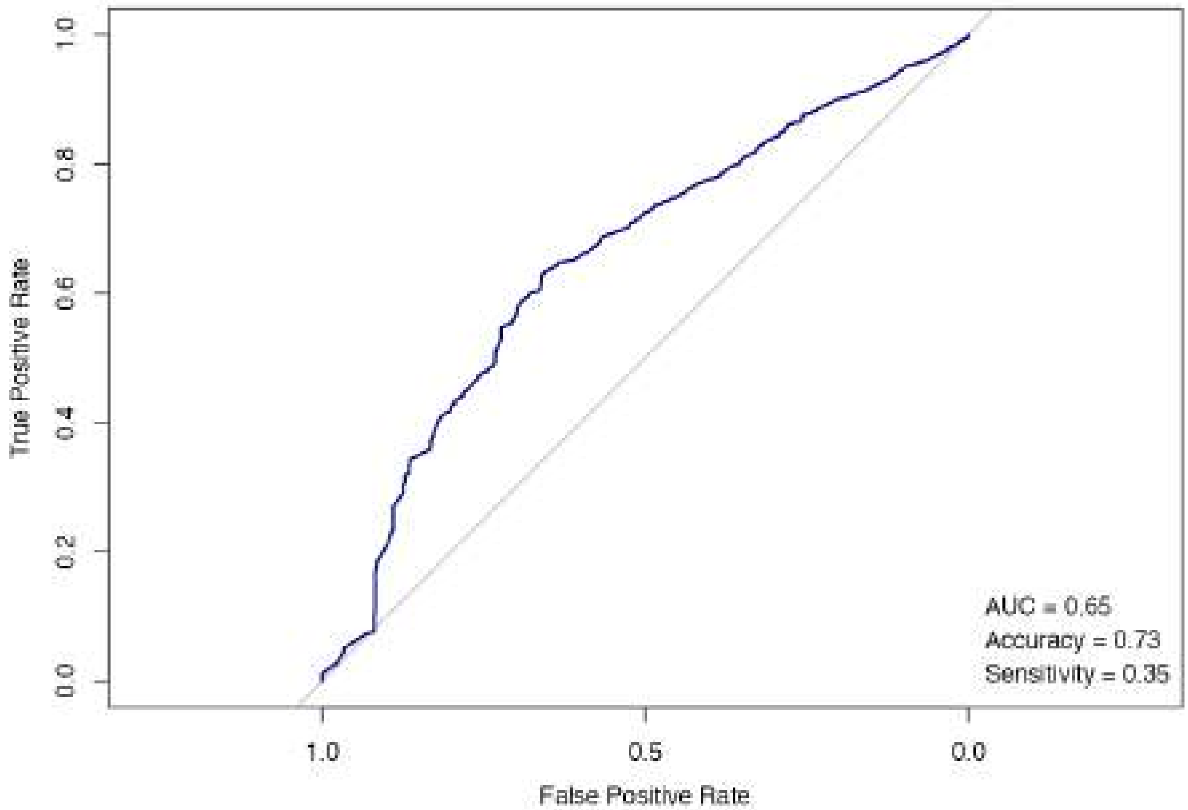

C.3 Details on measurement error of machine learning

[IN PROGRESS]