Thirsting for Solutions: the Impact of Drinking Water Scarcity on Migration in Ethiopia
— Preliminary Draft —

January 2025

Abstract

Rural Ethiopia faces persistent challenges in accessing drinking water, exacerbated by the escalating impacts of climate change. Yet, the effects of drinking water scarcity on migration remain poorly understood. This paper explores how wells failure influence migration decisions. Establishing a causal relationship is challenging due to the nonrandom nature of water insecurity. To overcome this, I leverage a new data base, documenting well locations and a one-time assessment of their functional status across Sub-Saharan Africa, to train an algorithm predicting well failures. Based on climate, geological, and hydrogeological factors, I generate monthly predictions of well status, allowing me to track functionality over time. I then link these predictions with the *Ethiopian Socio-Economic Surveys* spanning 2012 to 2016, using them as an exogenous proxy for drinking water access. Preliminary results show that migration increases with prolonged periods without a functional well in places facing drought. To further validate these findings, I incorporate two alternative measures of drinking water access: first, time-invariant data on well functionality; and second, an instrumental variable approach at the household level, using the median time to fetch water within an enumeration area. Both methods yield results consistent with the primary analysis.

1 Introduction

From 2012 to 2020, the World Economic Forum consistently ranked water crises among the top 5 risks in terms of impacts (Global Risks Report, 2020). Today, 2.3 billion individuals live in water-stressed countries (UN Water, 2021). Furthermore, the adverse effects of climate change are disrupting established weather patterns and precipitation (IPCC, 2023), creating an environment of increasing unpredictable water availability, and exacerbating water scarcity. Water is essential for life, health, and human dignity. Without sufficient water to meet basic needs, migration might emerge as an adaptation strategy. The literature on climate migration has mainly focused on the agricultural channel: lack of rainfall (or higher temperatures) lowers yields, decreasing agricultural income. This loss of income either traps the household in its origin (binding liquidity constraint) or pushes it to migrate (increase in incentives); see Cattaneo et al. (2019) for a review. This paper shifts the focus from the explored agricultural channel to a crucial yet understudied aspect: the impact of drinking water scarcity on migration. Specifically, I answer the question: Does difficult access to water for domestic use lead to migration?

To find a setting where migration might be affected by drinking water scarcity, I focus on Ethiopia. Notably, Sub-Saharan Africa stands out as the sole region where the number of people lacking access to water is increasing (JMP, 2021). Ethiopia and eight other countries [1] are home to 80% of under-served people in the region (JMP, 2021). Despite being depicted as the water tower of Africa, access to clean drinking water remains a challenge in Ethiopia. In 2016, 70% of the Ethiopian rural population lacked basic access to drinking water (JMP, 2023), meaning that they were either drinking from an unimproved source or walking 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 presents a compelling context for investigating the relationship between drinking water scarcity and migration.

This paper aims to understand the role drinking water scarcity plays in influencing migration. However, a causal relationship is difficult to establish due to nonrandom access to drinking water.

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

To address this challenge, the main analysis relies on predicted well failures as an exogenous source of variation. I have compiled the most comprehensive dataset available, to my knowledge, of known water points across Sub-Saharan Africa, documenting the location and functional status of wells at the time of monitoring. I use this information on the functional status to train an algorithm predicting wells' functionality and drying. By focusing on groundwater access, I am able to disentangle the effects of water scarcity on migration through agricultural impacts versus drinking water shortages. In Ethiopia, domestic water supply heavily relies on groundwater, while its use in irrigation remains limited. I link these predictions with migration outcomes from the *Ethiopian Socio-Economic Surveys* (ESS) spanning 2012 to 2016. The empirical strategy leverages both temporal and spatial variations in drought and in access to drinking water. Specifically, I examine migration patterns in Enumeration Areas (EAs) that experienced more predicted months without a functional well, contingent on the EA being affected by drought.

The main results of the paper rely on a machine learning approach, specifically a random forest algorithm, to predict well functionality using exogenous climate, geological, and hydrogeological data. The idea is to use the predictions as an exogenous proxy for drinking water access. The model is trained and optimized with a stratified split and cross-validation to predict monthly functionality status. Its performance is assessed using the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), yielding a high AUC score of 0.95, which indicates strong predictive power. The model is then used to predict monthly well functionality for each EA, with the number of months a well is non-functional as the main treatment variable. To further validate the model's performance in reflecting wells failure, I estimate the effects of the predicted number of months without a functional well on groundwater being the main drinking water source and the time taken to fetch water. The results suggest that while the predicted failure of wells does not significantly affect the primary water source, it increases the time required to collect water for those using groundwater as their main drinking water source. These findings provide further evidence that the algorithm successfully captures the drying up of wells and the resulting challenges in accessing water.

Using the predictions as an exogenous proxy for drinking water access, initial findings suggest that EAs with more months without a functional well experience higher levels of migration in the context of drought. The results show a significant negative coefficient for drought, which aligns with the literature on climate change and migration, suggesting that the migration observed in these areas is partially hindered by binding liquidity constraints (Cattaneo and Peri, 2016). However, the interaction between drought and the predicted number of months without a functional well significantly increases migration. In a comparison of two EAs affected by drought, the one with a greater number of months without a functional well will experience higher migration. This suggests that in drought-affected areas even though the agricultural channel does trap people in place, where drinking water is scarce, households are compelled to migrate. These findings emphasize the importance of considering water access as a crucial driver of migration, alongside more commonly studied factors like agricultural productivity.

To further validate the main results, I propose two alternative methods for measuring access to drinking water, both of which yield consistent findings with the primary analysis. First, I use the functionality status at the time of the report to classify wells as either functional or not functional. Since most wells are observed for functionality at a single point in time, this method classifies a well as non-functional if it was reported as such at any given time. I then estimate the impact of access to drinking water on migration by comparing EAs without access to a functional water point within an 8 km radius of the EA center to those with access. This alternative measure of water access aligns with the results of the main analysis, providing additional confidence in the robustness of the findings.

Finally, I adopt a more granular approach by transitioning from treatment at the EA level to the household level. In this method, I gauge access to water based on the time it takes to fetch water on the day preceding the interview. The variable is endogenous as it is likely to be correlated with unobserved characteristics of the household, also correlated with migration. Therefore, I rely on an instrumental variable approach: I use the median time taken to fetch water by the rest of the EA to instrument for a a household's time to fetch water. While not ruling out EA-level cofounders, it still removes household-level biases and allows for a household-level treatment and heterogeneity analysis. Importantly, the results from this household-level analysis are consistent with the main findings, reinforcing the robustness of the relationship between drinking water access and migration.

The main contribution of my paper is to shed new light on the climate migration nexus by exploring a different channel and investigating the impact of access to water for domestic use. The literature on climate-induced migration shows that, in developing countries, slow-onset events, such as rising temperatures, tend to lead to either voluntary and predominantly economically motivated migration (Bohra-Mishra et al., 2014) or immobility (Cattaneo and Peri, 2016). While there is no consensus on the impact of weather variability and climate shocks on migration, a constant theme is emphasized by several papers: the role played by the agricultural sector as a mediating channel. Cattaneo and Peri (2016) assume that poor and middle-income countries are more affected by weather variability because they are largely dependent on the agricultural sector. Cai et al. (2016)

find that temperature has a positive and statistically significant effect on international migration outflows only for agriculture-dependent countries. Using data on 108 countries from 1960 to 2010 and a two-stage least square estimation, Falco et al. (2019) claim this relationship is causal. They estimate that, on average, a climate-driven reduction in agricultural productivity of 1% from its decennial trend induces an increase in the emigration rate from about 2.5% to about 4.5% in the overall and poor country sample. However, other forces could be at play. In this paper, I am shifting the focus from the agricultural channel to a crucial yet understudied aspect: the impact of drinking water scarcity on migration. In addition, the literature has focused on marginal effects by linking weather variability and migration (Cattaneo et al., 2019), whereas most concerns about climate change and migration stem from water scarcity, desertification, etc. For those, thresholds are of particular importance. To my knowledge, this study is the first to analyze the lack of drinking water as a specific threshold.

The paper also contributes to the literature on access to groundwater and socio-economic outcomes; this literature began by studying the externality arising from groundwater use for irrigation in the United States (Brill and Burness, 1994; Gisser and Sanchez, 1980). As it developed, the literature focused on the institutions surrounding groundwater allocation and inter-sectoral effects of groundwater use in India (Aggarwal and Narayan, 2004; Banerji et al., 2012; Foster and Rosenzweig, 2008; Jacoby et al., 2004; Sekhri, 2013; Ryan and Sudarshan, 2020). The positive impact of groundwater irrigation on yields and its potential for adaptation to climate stress was shown focusing on the United States (Keskin and Hornbeck, 2014). Moving away from impacts on agriculture, (Sekhri, 2014) shows that increases in the cost of access to groundwater correlate with higher poverty and conflict; she also shows that groundwater scarcity results in an increase in sexual violence against women (Sekhri and Hossain, 2023). The closest papers to my study are Blakeslee et al. (2020) and Fishman et al. (2017). Blakeslee et al. (2020) studies the impact of drying wells directly on migration while Fishman et al. (2017) looks at labor reallocation to industry. Both papers show that some households turn to migration to offset the income effects of losing groundwater. However, in both their settings, groundwater is used for agricultural purposes. Therefore, their question differs from the channel I want to investigate (i.e., effects on basic needs).

My article also adds to the debate on “climate refugees”. Migration is a continuum, with forced migration at one end and economic migration at the other. When looking at climate-induced forced migration, existing analyses have mainly focused on migration due to sea level rise. Indeed, with such movements, there is little doubt about the involuntary aspect of the migration as the initial land disappears, see Burzynski et al. (2019) or Hauer et al. (2020). Migration because of lack of drinking water also has a major involuntary aspect. Therefore, my article contributes to the literature by providing evidence of climate-induced forced migration (not linked to sea level rise) and the selection into it. While there is a continuum, and no strict definition of forced or economic migration, whether the migration falls closer to one or the other has important implications. The type of migration may influence both the selection into migration (who migrates?) as well as the migratory response (to where? how?). Focusing on Dust Bowl migrants, Hornbeck (2020) provides evidence that climate-induced migration does not always follow typical patterns of economically driven migration. He shows that migrants from more affected counties were more negatively selected: migrants from more eroded counties had significantly lower levels of education than those from less affected areas or other contemporary migrants. And this difference persisted after migration, as those from the worst affected regions earned less than other migrants from the same period. Kleemans (2015) further illustrates how migration responses vary depending on the underlying motivation for migration. She shows with a dynamic migration choice model how "ex-post-risk-coping" migration (migration to cope with negative income shocks) and "investment" migration (migration to increase and diversify future expected income) may give way to opposite migratory responses to shocks. Despite growing interest in climate-induced migration, there remains limited understanding of who migrates in response to permanent environmental shifts. My paper seeks to address this gap. Identifying the key mechanisms and selection processes driving climate-induced migration is essential for developing evidence-based policy recommendations (Mbaye, 2017).

The paper is structured as follows: Section 2 describes the data. Section 3 introduces the empirical approach and identification strategy. Section 4 shows the main results relating to the impact of drinking water scarcity on migration. Section 5 contains a conclusion and 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 several climate and geological data.

2.1 Household Survey Data

I use the three waves of The Ethiopia Socio-Economic Survey (ESS) to get information on migration patterns, water access, and individual and household characteristics. The ESS is a part of the Living Standards Measurement Study-Integrated Surveys on Agriculture program (LSMS-ISA) of the World Bank.

There are three waves of the ESS, with data collected at multiple levels. The first wave (2011-2012) includes 290 rural and 43 small town enumeration areas (EAs). The following waves (2013-2014 and 2015-2016) expand the sample to include all urban areas. Together, the three waves create a panel data set of households from rural and small-town areas. Households interviewed in the first wave were tracked and re-interviewed in the second and the third wave. Out of the original sample (3776 households), around 98% of households were successfully re-interviewed in ESS3, indicating low levels of attrition (National Bureau of Statistics, 2014). The ESS employed a stratified two-stage sampling strategy. Ethiopia's regions served as strata. The sample is designed to be representative of rural and small town areas of Ethiopia. I construct a balanced panel of households over time, consisting of all 2494 households observed in all three waves, and that have complete information.

The main outcome variable is a count of the number of individuals that left the household permanently. In between waves ESS tracks individuals, so it asks whether each member of a household is still considered as part of it. If not, people who stayed behind report when, why and where (urban or rural area) the individual left. Using the reason why a member has left the household allows me to rule out death and divorce. Hence, I consider individuals who are no longer a member of their original household as permanent migrants.

Socio-economic determinants included in the analysis are at the individual level, household level, and EA level. Individual determinants include indicators for the individual being a female, age, education level, and the time taken to fetch water yesterday. Determinants at the household level include indicators for access to irrigation, main source of water in the rainy and dry season, time taken to fetch water in the rainy and dry season. EA level determinants include population size, presence of and price of connection to the water service.

2.2 Water Point Data

The data on water points originates from two key sources: Water Point Data Exchange (WPdx) and mWater, both of which play a crucial role in aggregating and monitoring rural water infrastructure. To my knowledge, this study is the first to leverage their information.

WPDx is a global, open-source database that compiles water point data from multiple countries, primarily in Sub-Saharan Africa. By standardizing information from governments, NGOs, and other stakeholders in a standardized framework, it supports the monitoring of rural water infrastructure. A key feature of WPDx 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.

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 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. Additionally, the platform supports dynamic data collection through mobile applications, allowing for continuous updates and long-term monitoring of water point functionality and service reliability.

By merging these datasets, I construct a non-exhaustive yet uniquely comprehensive panel of water wells in Ethiopia. To my knowledge, this represents the largest and most detailed dataset on water points in the country. The final dataset comprises 30,142 water points, incorporating both spatial and functional status information at the time of reporting.

2.3 Other Data

Basin level data – I use water basins as the spatial unit to predict well functionality. A water basin is an area where all surface water converges toward a common outlet point. For this study, I employ the HydroBASINS geographic data from the HydroSHEDS database, which provides globally consistent delineation of water basins. HydroBASINS further subdivides these into hierarchical sub-basins, based on the Pfafstetter coding system. This system organizes basins by topological relationships, enabling analysis at multiple spatial scales.

To capture localized patterns, I utilize the finest Pfafstetter level (level 12 of 12), which partitions sub-basins into average areas of approximately 100 km^2. HydroSHEDS data is also used to enrich the analysis with geologic and climatic attributes at the basin level, including variables such as natural water discharge, groundwater table depth, and soil clay fraction.

Weather data – The weather data come 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. Therefore, it is considered more reliable than data that rely only on weather stations, specifically in regions where observations are sparse and of low quality (Auffhammer et al., 2013). Glexner et al. (2020) show ERA5 performs well in East Africa. For example, they show that the spatial distribution of precipitation during extreme years is better represented in ERA5. Recent articles such as Kalkuhl and Wenz (2020), Kotz et al. (2022) also use ERA-5. I use monthly precipitation and evapotranspiration data (with fine resolution: 0.25° x 0.25°) to construct the Standardized Precipitation-Evapotranspiration Index (SPEI) at the EA level. The reference period I use in 1970-2018. And the main time scale is 12 months: 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.

The SPEI is a water balance measure normalized according to a log-logistic distribution (Vicente-Serrano et al., 2010; Begueria et al., 2014)[2]. 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 SPEI can measure drought severity in terms of intensity and duration. It can also identify the onset and end of drought episodes.

Hydro-Geological Data – I use the **Africa Groundwater Atlas** for the hydro-geological data. The Africa Groundwater Atlas is an online resource providing information on the hydrogeology and groundwater resources of 51 African countries. It is part of the Unlocking the Potential of Groundwater for the Poor (UPGro) Research Programme. UPGro aims to improve the evidence base around groundwater availability and management in sub-Saharan Africa and to enable the sustainable use of groundwater to benefit populations. The Atlas provides information on the geology (both bedrock and superficial/unconsolidated geology), with geological categories that reflect significant hydro-geological units. It also provides information on hydrogeology as a combined classification of aquifer type and productivity. The definition of aquifer type is in terms of the geological characteristics. The key feature of the aquifer type classification is the dominant way groundwater flows

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

through and is stored in aquifers. Aquifer productivity is estimated using borehole yield data as a proxy. The aquifer productivity categories used in the Atlas consists of a qualitative assessment because of the scale of these maps, the heterogeneity of most aquifers, and the limited availability of aquifer characteristics. The map produced is at a nominal scale of 1:5,000,000, which means that 1 cm on the map is equivalent to 50 km on the ground.

Since the Africa Groundwater Atlas provides coarse-resolution data, I supplement it with information from **MacDonald et al. (2012)**, who produce higher-resolution (5 km grid) quantitative maps of groundwater productivity, storage, and depth to groundwater. The groundwater productivity map estimates expected borehole yields across different hydrogeological units. Groundwater storage is derived by combining the saturated aquifer thickness with literature-based estimates of effective porosity for various aquifer types. Finally, depth to groundwater is modeled using an empirical, rules-based approach that accounts for rainfall, aquifer type, and proximity to rivers.

Lastly, I incorporate data from **Pelletier et al. (2016)** to assess groundwater potential, leveraging a high-resolution (30-arcsecond) global gridded dataset that maps soil thickness, intact regolith, and sedimentary deposits. This dataset estimates the depth of surface layers above unweathered bedrock, offering valuable insights into subsurface characteristics that influence groundwater availability. Areas with shallow bedrock are less likely to sustain productive aquifers due to their limited water storage capacity, whereas regions with deeper regolith and sedimentary deposits tend to have higher groundwater potential. By integrating this dataset, I refine the characterization of groundwater resources, improving the understanding of spatial variations in water availability.

3 Empirical Approach

3.1 Measure of drinking water access

Predicting well functionality - Measuring access to drinking water exogenously presents a key challenge in this paper. In Ethiopia, groundwater is the principal water source for domestic use (drinking, cooking, and cleaning). Consequently, data on wells offers critical insights into drinking water access. The dataset includes information on the location of water points and their functionality status at the time of the enumerator's visit, allowing me to determine whether a given well was functional during the reporting period. However, several limitations exist. First, the data does not comprehensively cover all wells, creating potential gaps. For example, individuals in the ESS data report using a well, but no corresponding water point is be identified in the WPdx or mWater datasets. Second, the timing of the water points survey often does not align with the ESS data collection periods. This temporal mismatch means that while I may have information on a well's functionality, it often falls outside the relevant socio-economic timeframe. Finally, the current functionality status of wells may be influenced by endogenous factors, such as local governance or community engagement, complicating causal interpretations.

To address these challenges and establish a time-varying, exogenous measure of well failure, I employ a machine learning approach, specifically a random forest algorithm. This model integrates diverse climate, geological and hydrogeological input data sources to predict wells' functionality status over time. Since only exogenous inputs are used to train the model, one could consider the predictions as an exogenous proxy of drinking water access. The model is trained using a stratified training and testing split, where I train on 70% of the data and then use the model to predict functionality status for 30% of the remaining data. It was trained with a 10-fold cross-validation approach to ensure robust performance evaluation. Optimization was guided by the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), chosen as the primary metric to balance sensitivity and specificity. The number of variables randomly sampled at each split (mtry) was set to 16, based on preliminary optimization. To address class imbalance in the dataset (see 1, weights were assigned to observations, giving the minority class (functional wells) a weight of 8 and the majority class a weight of 1.

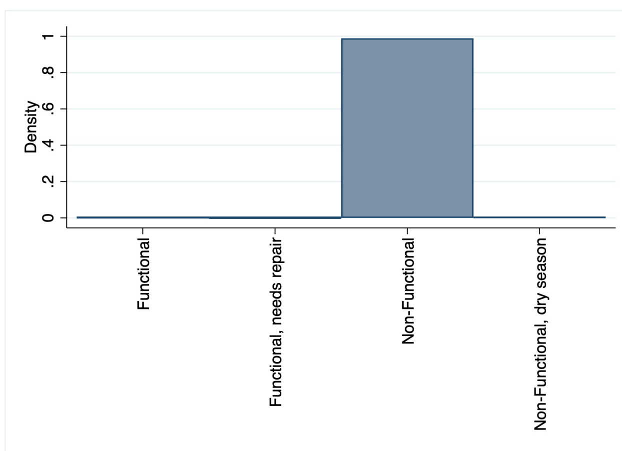

Figure 1: Functionality status of wells in Ethiopia

A common metric in machine learning for evaluating the performance of a predictive model across different thresholds is the ROC-AUC score. The Receiver Operating Characteristic (ROC) curve is a graphical representation that plots the true positive rate (sensitivity) on the y-axis against the false positive rate (1 - specificity) on the x-axis for various classification thresholds (see Figure 2). The Area Under the Curve (AUC) quantifies the overall ability of the classifier to distinguish between positive and negative classes. Intuitively, the AUC represents the likelihood that a randomly chosen positive instance will be ranked higher by the model than a randomly chosen negative instance. An AUC score of 1.0 indicates a perfect classifier, capable of completely separating the classes, while an AUC of 0.5 implies no discriminative power, equivalent to random guessing. The model performs well overall, with a high AUC (0.9486), indicating strong discriminatory power. Specificity (0.9771) is substantially higher than sensitivity (0.7071), indicating that the model is more effective at correctly identifying the control group (functional wells) while encountering greater difficulty with the treatment group, suggesting the presence of more noise or variability in the data associated with non-functional wells.

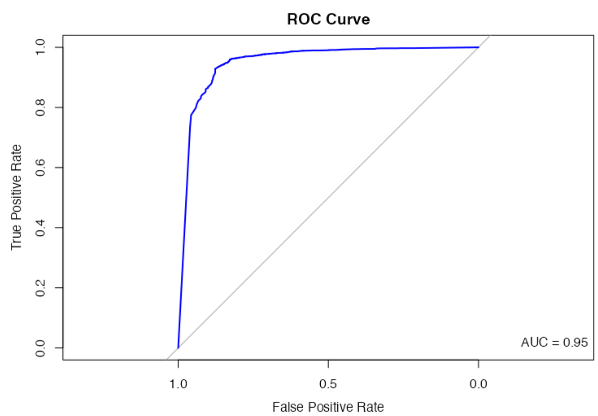

Figure 2

I then use this model to predict monthly status of wells for each EA. I, then create my main treatment which is the number of months a well in an EA (v) was not functional in the last two years ($NoWater_{v,t}$). To further validate that the model predicts correctly well failure, I check whether the predictions are associated with a decrease in the use of groundwater ($GW_{h,v,t}$) and an increased time to fetch water ($TimeWater_{h,v,t}$) [3]. Specifically, for the use of groundwater, I estimate:

$$GW_{h,v,t} = \alpha + \beta NoWater_{v,t} + \delta_{w\times t} + \delta_h + \epsilon_{h,v,t} \tag{1}$$

And for the time to fetch water:

[3] I include household fixed effects (δ_h) which account for any time-invariant properties of households. I also include woreda-year fixed effects ($\delta_{w\times t}$) which account for district time trends. Standard errors are clustered at the EA level.

$$TimeWater_{h,v,t} = \alpha + \beta_1 NoWater_{v,t} + \beta_2 GW_{h,v,t} + \beta_3 NoWater_{v,t} \times GW_{h,v,t} + \delta_{w\times t} + \delta_h + \epsilon_{h,v,t} \tag{2}$$

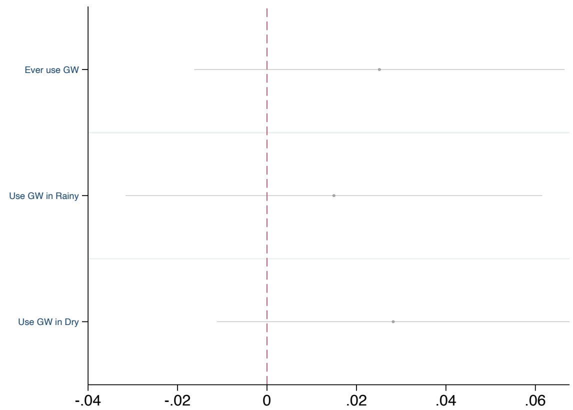

Figure 3: Marginal effect of an additional month without a functional well on groundwater use for drinking

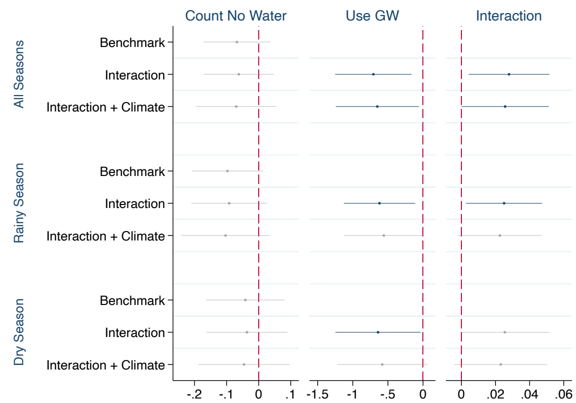

Figure 4: Marginal effect of an additional month without a functional well on time taken to fetch water

Figure 3 illustrates the relationship between the predicted functionality status of wells and the use of groundwater as the primary source of drinking water. Figure 4 depicts the effect of well functionality on the time spent fetching water. While an additional month without a functional well appears to have little impact on the declared primary water source, it significantly increases the time spent fetching water for households relying on groundwater as their main drinking source.

Water Point Treatment – My second measure of access to drinking water is also based on the data on water points. Both mWater and Water.Point.Data.Exchange provide information on the functionality status of water points at the time of the report. The ESS only provides scrambled coordinates of the EAs for anonymity reasons. I chose an 8km radius to allow for displacement of up to 5km of the true location of the cluster. I classify a village as having difficult access to drinking water ($NoWater_v = 1$) if there are no functional water points within 8km. Descriptive statistics can be found in the Appendix.

Time to water – My fourth measure of access to drinking water is based on the time to water reported by households. The variable is endogenous as it is likely to be correlated with unobserved characteristics of the household, also correlated with migration. I use the median time taken to fetch water by the rest of the EA to instrument for a household's time to fetch water. While not ruling out EA-level cofounders, it still removes household-level biases and allows for a household-level treatment and heterogeneity analysis.

3.2 Migration measure

I use different migration measures to assess migration resulting from drinking water scarcity.

My first measure of migration is based on reports of household members that stayed behind; it no longer means that the entire household has migrated. ESS asks individuals whether other members are still members of the household. Using another question on the reason why a member has left the household allows me to rule out death and divorce. Hence, I consider individuals who are no longer a member of their original household as permanent migrants. Despite not measuring entire household's migration, this measure is interesting because it speaks to the selection of the individual who left. Furthermore, it limits attrition concerns.

Then, I estimate migration as the number of cumulative months an individual has been away from her household during the last 12 months. It is a measure of seasonal migration. Again, it no longer means that the entire household has migrated. However, this measure allows me to explore if access to drinking water changes the patterns of migration.

For the purpose of the study, I also aggregate household measures of migration at the EA level. Descriptive statistics can be found in the Appendix.

3.3 Drought treatment

Based on the meta-analysis of different measures of weather variability in West Africa by Bertoli et al. (2022), I measure the occurrence and intensity of drought using the Standardized Evapotranspiration Index (SPEI). The baseline estimations use SPEI-12, which assesses hydrogeological drought. First, I use a cutoff for severe drought before the last interview. The $drought_{v,t}$ variable is a binary variable taking the value one if the EA has experienced a SPEI-12 below -2 at one point during the six months before the interview. As shown in Figure 5, there is only one occurrence when EAs experienced such a drought. It was just at the end of 2015. This drought lasted until the third wave of interviews: no EAs lived through one harsh month and then returned to better conditions. The only EAs that would have experienced only one month of drought would be those where the drought began just one month before the interview.

I also use continuous measures of drought. It allows me to have a measure of the longer-term impact of drought. Indeed, one shock might be less likely to drive people to migrate. The drought indicator is specified as the number of months with a drought in between waves or as the duration of the longest drought in between waves. In robustness tests, I also use the cumulative intensity of drought measured as the sum of the absolute value of the SPEI-12 in between two waves.

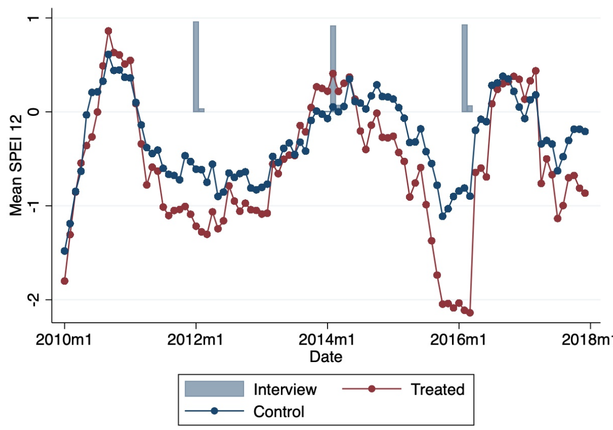

Figure 5: Mean SPEI12 and interview timing

3.4 Identification Strategy

My empirical analysis relies on a triple difference strategy at different levels depending on the specification. It can either be at the *Enumeration Area* (EA) level (v) or at the household level (h). Below, I will define every variable depending on j with $j = \{v, h\}$.

The independent variable is a drought variable, $Drought_{v,t}$, measured at the EA v level. The drought variable is interacted with a count of months where wells are predicted to not be functional, $NoWater_{v,t}$, of EA v or household h.

Functionally, the empirical approach consists of a quadruple difference with temporal and spatial variation in drought (i.e., is it a drought period? And is the EA affected by the drought?), and temporal and spatial variation in access to drinking water (i.e., for how many months were the EAs wells not functional?). Both the drought and access to drinking water measures collapse the spatial and temporal variations into one.

$$Mig_{j,t} = \alpha + \beta_1 drought_{v,t} + \beta_2 NoWater_{v,t} + \beta_3 drought_{v,t} \times NoWater_{j,t} + \delta_j + \delta_{w\times t} + \epsilon_{j,t} \quad (3)$$

The dependent variable $Mig_{j,t}$ indicates migration. I include household or EA fixed effects (δ_j), which account for any time-invariant properties of households or EA that could affect migration. I also include woreda-year fixed effects [4] ($\delta_{w\times t}$), which account for district time trends. Therefore, the coefficient β_2 measures the marginal increase in migration due to a supplementary month when the groundwater source was failing. The coefficient of interest is β_3, which estimates when experiencing drought, the change in migration for households/enumeration areas having difficult access to water compared to households/enumeration areas having an easier access to water. Standard errors are clustered at the EA level.

4 Results

4.1 Main results

The main results use the number of months a well in a community would have been non functional in the last 24 months based on the predictions of the algorithm to measure access to water. My preferred measure of migration is *Permanent migration* which counts the number of individuals who are reported to have left the EA. Looking at Table 1, we can see that the number of months a well is predicted to be non functional has a positive yet not significant impact on the number of people leaving the EA. Controlling for climate (sum of the absolute value of the SPEI in the last 24 months) does not change the direction of this effect. When interacting the experience of a drought with the number of months a well is predicted to be non functional, the coefficient becomes significant. The coefficient for drought is negative and statistically significant. It implies that EAs experiencing drought tend to have lower migration rates. Potentially, this could be explained by liquidity constraints: the literature on climate migration has showed that drought by reducing yields, reduces households'income preventing them to migrate. The interaction $drought_{v,t} \times countnowater_v$ is positive and significant, with a coefficient of 0.238. It suggests that EAs with more months without a functional well experience more migration when facing drought. The overall effect is still negative but the difference in migration between EAs with and without access to drinking water is significant.

[4] Districts in Ethiopia are called woredas. They correspond to the third level of the administrative divisions of Ethiopia

	Baseline		Interaction
	(1)	(2)	(3)
count_nowater	0.076	0.077	0.055
	(0.061)	(0.063)	(0.063)
Drought-intensity (lag 24)		−0.002	
		(0.048)	
drought=1			−5.372***
			(1.867)
drought=1 × count_nowater			0.238**
			(0.093)
Observations	279	279	279
R^2	0.605	0.605	0.608
EA Fixed effects	Yes	Yes	Yes
KebelexYear Fixed effects	Yes	Yes	Yes

Table 1: Results – Predicted functionality & Permanent Migration

Notes: Permanent migration is a count of the number of individuals who have left the *EA*. *Drought* is a binary measure of drought, reflecting if the *EA* has experienced a SPEI12 below −2. *Drought-intensity* is continuous measures of drought, reflecting the sum of the absolute value of the 12 months SPEI in the last 24 months. Robust standard errors are clustered on the enumeration area level and reported in parentheses. Significance at or below 1% (***), 5% (**) and 10% (*).

Turning to seasonal migration (see Table 2), the pattern is the same when interacting drought and the number of wells the EA spent without a functional well: EAs with more months without a functional well experience more migration when facing drought. The overall effect is still negative but the difference in migration between EAs with and without access to drinking water is significant. However when looking without the interaction, the coefficients are negative and non significant.

	Baseline		Interaction
	(1)	(2)	(3)
count_nowater	−0.070	−0.066	−0.095
	(0.105)	(0.106)	(0.109)
Drought-intensity (lag 24)		−0.017	
		(0.040)	
drought=1			−5.513**
			(2.629)
drought=1 × count_nowater			0.230**
			(0.111)
Observations	279	279	279
R^2	0.792	0.792	0.793
EA Fixed effects	Yes	Yes	Yes
KebelexYear Fixed effects	Yes	Yes	Yes

Table 2: Results – Predicted functionality & Seasonal Migration

Notes: Seasonal migration is the mean number of months individuals report being away from their households. *Drought* reflects a binary measure of drought, reflecting if the *EA* has experienced a SPEI12 below −2. *Drought-intensity* is continuous measures of drought, reflecting the sum of the absolute value of the 12 months SPEI in the last 24 months. Robust standard errors are clustered on the enumeration area level and reported in parentheses. Significance at or below 1% (***), 5% (**) and 10% (*).

4.2 Other measures of water access

Next, to validate the main results, I turn to other measures of water access.

Water Point Treatment – To check whether it is the algorithm that is driving the results, I provide results without relying on its predictions. Here $NoWater_v$ reflects whether the EA has access to at least one functional water point. Therefore, the measure does not vary in time. One potential threat to the identification stems from the variable $NoWater_j$ capturing not the effect of difficult access to drinking water but other characteristics of the community. I provide a balance table (Table 3) across communities without and with access to drinking water based on the measure of access to water by the functionality status of water points. It shows that there are few statistically significant differences between the two except for access to water related measures. Importantly, there is no significant difference in the likelihood of having an irrigation system or conflict in between the treatment and the control group.

Focusing on permanent migration measured at the household level (Table 4), the same patterns as the main results are observable: the coefficient for drought is negative and the interaction

$drought_{v,t} \times nowater_v$ is positive. However, the effect is only statistically significant when drought is measured as a count of drought months. Then, when the dependent variable is seasonal migration (Table 5), we observe again a negative coefficient for drought and a positive coefficient for the interaction with the measure of water access. Overall, when using access to a functional well within 8km rather that the algorithm's predictions, the results are similar. Moreover, those results are stable across different measures of drought : not only binary but also a count of the number of months a community experienced drought and the duration of the longest dry spell.

Time to water – Lastly, since I only have the geolocation of EAs, every treatment is at the EA level. To go down to household level, I rely on reports on the time taken to fetch water. Since this measure is likely to be endogeneous, I instrument it with the median time taken to fetch water in the community. The first stage regression (Column 1 in Table 6) confirms the relevance of median time to water as an instrument, with a highly significant coefficient, indicating its strong predictive power for time spent collecting water. However, the drought measure variable is not significant in the first stage, suggesting its effect on migration operates indirectly through water collection times. The results for permanent migration (Column 2 in Table 6) show no statistically significant effects when using the OLS model. In the second stage, only the interaction is significant, implying that there is a significant difference in migration patterns in between people that take more time to fetch water, in line with what was described earlier. For seasonal migration (Columns 5 – 7 in Table 6), the results are more nuanced. Indeed, they are reversed compared to main results but with the interaction being not significant. These results underscore the complexity of the relationship between water stress and migration, with nuanced gendered responses that merit further exploration.

[GENDER EFFECTS TO BE ADDED]

5 Conclusion

In conclusion, this paper delves into a critical yet understudied aspect of climate-induced migration: the impact of drinking water scarcity on migration patterns. Against the backdrop of increasing water crises globally and the particular challenges faced by Ethiopia, where access to clean drinking water remains a significant issue, this article investigates the nexus between water scarcity and migration. By shifting the focus from the commonly explored agricultural channel to the scarcity of drinking water, this research broadens our understanding of the drivers of migration in the face of climate change.

The paper contributes to the literature by shedding new light on the climate migration nexus and highlighting the importance of access to water for domestic use as a driver of migration decisions. By elucidating the role of drinking water scarcity in shaping migration patterns, my analysis underscores the complex interplay between environmental factors and human mobility. Moreover, it adds to the ongoing debate on climate refugees by providing evidence of climate-induced forced migration not directly linked to sea-level rise.

Understanding the mechanisms driving climate-induced migration is crucial for informing policy responses aimed at mitigating its adverse effects and supporting vulnerable populations. Ultimately, addressing the root causes of migration, including access to essential resources like clean drinking water, is essential for building resilience and promoting sustainable development in a changing climate landscape.

References

Aggarwal, R. and Narayan, T. (2004). Does inequality lead to greater efficiency in the use of local commons? the role of strategic investments in capacity. *Journal of Environmental Economics and Management*, 47:163–182.

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.

Banerji, A., Meenakshi, J., and Khanna, G. (2012). Social contracts, markets and efficiency: Groundwater irrigation in north india. *Journal of Development Economics*, 98:228–237.

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.

Bertoli, S., Docquier, F., Rapoport, H., and Ruyssen, I. (2022). Weather shocks and migration intentions in western africa: Insights from a multilevel analysis. *Journal of Economic Geography*, 22(2):289–323.

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.

Bohra-Mishra, P., Oppenheimer, M., and Hsiang, S. (2014). Nonlinear permanent migration response to climatic variations but minimal response to disasters. *Proceedings of the National Academy of Sciences*, 111(27):9780–9785.

Brill, T. and Burness, S. (1994). Planning versus competitive rates of groundwater pumping. *Water Ressources Research*, 30:1873–1880.

Burzynski, M., Deuster, C., Docquier, F., and de Melo, J. (2019). Climate change, inequality, and human migration. Technical report, LIDAM Discussion Papers.

Cai, R., Feng, S., Oppenheimer, M., and Pytlikova, M. (2016). Climate variability and international migration: The importance of the agricultural linkage. *Journal of Environmental Economics and Management*, 79:135–151.

Cattaneo, C., Beine, M., Frohlich, C. J., Kniveton, D., Martinez-Zarzoso, I., Mastrorillo, M., Millock, K., Piguet, E., and Schraven, B. (2019). Human migration in the era of climate change. *Review of Environmental Economics and Policy*, 13:189–206.

Cattaneo, C. and Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 122:127–146.

Falco, C., Galeotti, M., and Olper, A. (2019). Climate change and migration: Is agriculture the main channel? *Global Environmental Change*, 59.

Fishman, R., Jain, M., and Kishore, A. (2017). When water runs out: Adaptation to gradual environmental change in indian agriculture. *Retrieved from www.semanticscholar.org*.

Foster, A. and Rosenzweig, M. (2008). *Economic Development and the Decline of Agricultural Employment*, volume 47. Elsevier.

Gisser, M. and Sanchez, D. (1980). Competition versus optimal control in groundwater pumping. *Water Ressources Research*, 16:638–642.

Glexner, S., Demissie, T., and Tefera Diro, G. (2020). Did era5 improve temperature and precipitation reanalysis over east africa? *Atmosphere*, 11(9):996.

Hauer, M. E., Fussell, E., Mueller, V., Burkett, M., Call, M., Abel, K., McLeman, R., and Wrathall, D. (2020). Sea-level rise and human migration. *Nature Reviews Earth and Environment*, 1:28–39.

Hornbeck, R. (2020). Dust bowl migrants: Identifying an archetype. Working Paper 27656, National Bureau of Economic Research.

IPCC (2023). *Weather and Climate Extreme Events in a Changing Climate*, page 1513â€“1766. Cambridge University Press.

Jacoby, H., Murgai, R., and Rehman, S. U. (2004). Monopoly power and distribution in fragmented markets: The case of groundwater. *Review of Economic Studies*, 71:783–808.

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.

Keskin, P. and Hornbeck, R. (2014). The historically evolving impact of the ogallala aquifer: Agricultural adaptation to groundwater and drought. *American Economic Journal: Applied Economics*, 6:190–219.

Kleemans, M. (2015). Migration choice under risk and liquidity constraints. Working Paper 200702, Agricultural and Applied Economics Association.

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

Mbaye, L. (2017). Climate change, natural disasters, and migration. *IZA World of Labor*.

Ryan, N. and Sudarshan, A. (2020). Rationing the commons. Working Paper 27473, National Bureau of Economic Research.

Sekhri, S. (2013). Missing water: Agricultural stress and adaptation strategies in response to groundwater depletion in india. Technical report, Virginia Economics Online Papers 406.

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.

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.

A Tables

	Control			Treated			
	n	mean	sd	n	mean	sd	Diff
Time to water (hours)	45	0.30	0.53	438	0.20	0.29	-0.061
Source rainy season	45	5.86	1.67	438	4.64	1.49	-0.815**
Source dry season	45	5.57	1.52	438	4.57	1.44	-0.675*
Population	45	5547.93	2584.00	438	5918.20	3699.75	796.631
Age	45	25.51	2.64	438	25.23	3.35	-0.529
Age household head	45	49.19	5.79	438	46.41	5.94	-2.045
Religion	45	2.36	1.47	438	2.25	1.22	0.382***
Marital Status	45	1.91	0.21	438	1.89	0.25	-0.013
Distance to main road (km)	43	34.98	44.91	437	35.59	58.70	46.356*
Distance to nearest bus (km)	43	18.12	33.05	436	15.45	26.40	17.998*
Distance to Woreda town (km)	39	14.46	9.21	377	24.06	25.61	17.040***
Distance to main urban center (km)	44	61.07	37.42	432	78.34	94.72	67.337**
Distance to weekly market (km)	29	24.66	43.63	232	15.25	24.83	9.256
Distance to primary school (km)	44	0.67	2.03	437	1.42	15.16	-0.480
ACLED event	45	0.04	0.21	438	0.18	1.87	-0.021
ACLED event with fatalities	45	0.04	0.21	438	0.05	0.70	-0.021
Amount rain during growing season	44	2.45	0.66	424	2.21	0.75	-0.112
Start rain	44	2.55	0.55	424	2.38	0.72	-0.063
End rain	44	1.41	0.69	424	1.84	0.84	0.296**
Main crop	40	8.03	11.22	386	12.92	22.22	-2.330
Main planting month	42	9.45	1.15	420	8.72	2.23	-1.391***
Share of Agriculture within 1km	45	24.47	26.89	438	30.80	17.80	3.659
Presence irrigation scheme	44	1.34	0.48	424	1.33	0.47	0.015
Population using irrigation scheme	29	532.90	604.05	284	294.45	340.71	-368.231**
Evolution yields (last 2 years)	44	2.73	1.80	424	3.16	1.29	0.532**
Evolution crop revenue (last 2 years)	44	3.34	1.99	424	3.29	1.32	0.127
Evolution credit access (last 2 years)	44	4.02	2.25	423	4.00	2.00	1.032**
Evolution ability to repay loans (last 2 years)	44	4.41	2.64	421	3.82	2.10	0.115
Evolution livestock revenue (last 2 years)	44	3.34	1.46	422	3.50	1.27	0.122
Evolution pasture availability (last 2 years)	44	3.59	1.96	424	3.50	1.76	0.217
Evolution soil quality (last 2 years)	44	3.75	1.40	421	3.68	1.40	-0.085
Evolution non agricultural opportunities (last 2 years)	44	4.16	1.89	421	3.96	1.79	0.119

Table 3: Balance table

Note: Column 1 and 4 show the number of non-missing observations out of a total of 483 observations; 45 observations for the control group (with access to a functional well within 8km) and 438 observations for the treatment groups (without access to any functional well within 8km). Columns 2 and 3 (5 and 6) show the summary statistics respectively for the control (treated) group. Column 7 shows the coefficient from regressing the baseline variable on an indicator for treatment. Coefficients are estimated including region fixed effects. Standard errors are clustered at the enumeration area level.

	Binary		Count		Longest	
	(1)	(2)	(3)	(4)	(5)	(6)
Drought	−3.862	−14.750				
	(3.802)	(10.276)				
Drought=1 × Non_functional=1		14.321				
		(10.659)				
Drought-count (lag 24)			−0.089	−2.989*		
			(0.665)	(1.685)		
Non_functional=1 × Drought-count (lag 24)				3.213*		
				(1.921)		
Drought-maximal duration (lag 24)					−0.316	−2.930*
					(0.566)	(1.680)
Non_functional=1 × Drought-maximal duration (lag 24)						2.808
						(1.829)
Observations	483	483	483	483	483	483
R^2	0.544	0.545	0.544	0.545	0.544	0.545
Village Fixed effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
RegionxYear Fixed effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*

Table 4: Results – Water Point & Permanent Migration

Notes: Permanent migration counts the number of permanent migrants reported by household members that stayed behind. *Binary* is a binary measure of drought. *Count* and *longest* are continuous measures of drought. *Count* measures the number of months a community experienced a SPEI-12 below −2 in the last 24 months. *Longest* measures the longest dry spell a community experienced in the last 24 months. Robust standard errors are clustered on the enumeration area level and reported in parentheses. Significance at or below 1% (***), 5% (**) and 10% (*).

	Binary		Count		Longest	
	(1)	(2)	(3)	(4)	(5)	(6)
Drought	−1.517	−3.357				
	(5.760)	(5.250)				
Drought=1 × Non_functional=1		2.420				
		(9.065)				
Drought-count (lag 24)			−0.876	−1.966		
			(0.846)	(1.555)		
Non_functional=1 × Drought-count (lag 24)				1.207		
				(1.833)		
Drought-maximal duration (lag 24)					−0.752	−1.996
					(0.698)	(1.552)
Non_functional=1 × Drought-maximal duration (lag 24)						1.336
						(1.712)
Observations	483	483	483	483	483	483
R^2	0.483	0.483	0.484	0.484	0.484	0.484
Village Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
RegionxYear Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: Results – Water Point & Seasonal Migration

Notes: Seasonal migration counts the number of seasonal migrants reported by household members that stayed behind. Binary is a binary measure of drought. *Count* and *longest* are continuous measures of drought. *Count* measures the number of month a community experienced a SPEI-12 below −2 in the last 24 months. *Longest* measures the longest dry spell a community experienced in the last 24 months. Robust standard errors are clustered on the enumeration area level and reported in parentheses. Significance at or below 1% (***), 5% (**) and 10% (*).

	First Stage	Permanent Migration			Seasonal Migration		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		OLS	2SLS	2SLS (women)	OLS	2SLS	2SLS (women)
Drought-count (lag 24)	−0.014	0.002	−0.015	0.003	−0.005	0.001	−0.000
	(0.009)	(0.008)	(0.010)	(0.006)	(0.012)	(0.016)	(0.006)
median_time_yesterday_hours	0.905***						
	(0.163)						
time_yesterday_hours		0.007			−0.005		
		(0.010)			(0.009)		
Drought-count (lag 24) × time_yesterday_hours		0.003			0.006		
		(0.008)			(0.008)		
time_yesterday_hours (Linear prediction)			−0.232	−0.046		0.298*	0.099*
			(0.147)	(0.062)		(0.158)	(0.059)
Drought-count (lag 24) × time_yesterday_hours (Linear prediction)			0.291***	0.084**		−0.030	−0.030
			(0.083)	(0.041)		(0.162)	(0.060)
Observations	10421	10421	10581	10581	10421	10581	10581
R^2	0.410	0.476	0.469	0.420	0.418	0.417	0.385
Household Fixed effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
RegionxYear Fixed effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*

Table 6: Results – IV median time to water

Notes: Permanent migration counts the number of permanent migrants reported by household members that stayed behind. *Seasonal migration* counts the number of seasonal migrants reported by household members that stayed behind. *Count* is continuous measures of drought, it measures the number of month a community experienced a SPEI-12 below −2 in the last 24 months. Robust standard errors are clustered on the enumeration area level and reported in parentheses. Significance at or below 1% (***), 5% (**) and 10% (*).