

Keep the Water Flowing: The Hidden Crisis of Rural Water Management

Chengcheng Zhai

Mendoza College of Business, University of Notre Dame, Notre Dame, IN 46556. czhai2@nd.edu

Rodney P. Parker

Kelley School of Business, Indiana University, Bloomington, IN 47405. rodp@iu.edu

Kurt Bretthauer

Kelley School of Business, Indiana University, Bloomington, IN 47405. kbretha@iu.edu

Jorge Mejia

Kelley School of Business, Indiana University, Bloomington, IN 47405. jmmejia@iu.edu

Alfonso Pedraza-Martinez

Mendoza College of Business, University of Notre Dame, Notre Dame, IN 46556. apedraz2@nd.edu

Problem definition: In rural sub-Saharan Africa, people rely on communal handpumps for clean drinking water. These handpumps frequently break down and require repairs. Operating under the assumption that preventive maintenance is too expensive, local NGOs often face difficult decisions about where to allocate their limited resources to reduce water point downtime. NGOs could invest in gathering more functionality information so they can use their scarce resources on targeted water point repairs. Alternatively, they could conduct more preventive maintenance, increase their repair capabilities, or source more reliable and cost-effective spare parts to reduce repair demand.

Methodology/results: To reduce water point downtime, we propose integrating preventive maintenance into existing programs. In collaboration with local NGOs, we conducted field research in Ethiopia and Malawi. We collected 47,240 observations of water point functionality from NGOs in Malawi, the Central African Republic, and Ethiopia. We then develop a Markov decision process (MDP) model based on real-world practices to optimize maintenance schedules for NGO mechanics. We apply the model to field data from the three countries and find that incorporating preventive maintenance can reduce water point downtime by an average of 41.4% (ranging from 7.1% to 61.9%), often with little to no increase in logistics cost. We further conduct numerical experiments to examine the role of functionality information. For example, we find that with high information availability, the reactive maintenance visitation approach is more effective only when repair demand is low. **Managerial implications:** Our findings challenge the belief that preventive maintenance is prohibitively expensive and advocate for its integration into NGO programs. We recommend that resource-constrained NGOs prioritize water point reliability, expand repair capacity, and reduce major repair costs over investing in extensive data collection.

History:

1. Introduction

“Much of Africa’s water supply infrastructure is failing for a simple and avoidable reason: lack of maintenance.” International Institute for Environment and Development (Skinner 2009)

In rural areas of sub-Saharan Africa (SSA), over 184 million people rely on communal water points with handpumps for clean drinking water (Foster et al. 2019). These handpumps, which

use a manual mechanism to extract groundwater, are well-suited for remote areas with limited electricity. However, over 50,000 of them are non-functional (Skinner 2009) due to insufficient maintenance (UPGro Consortium 2024), leading to what experts call the “hidden crisis” in rural water management.

To keep the water flowing, some NGOs in SSA implement maintenance programs. However, managing an effective program is challenging. Water points often use plastic and rubber parts that wear out quickly, requiring regular maintenance to avoid costly breakdowns. Remote locations make it difficult to monitor the water points, adding another layer of complexity. With limited budgets, NGOs must carefully consider every dollar spent and make difficult trade-offs when allocating resources. Cost-efficiency is even more important after the termination of the U.S. Agency for International Development (USAID), which provided almost half of the humanitarian funding until 2024 (Santos Porras 2025). This paper explores how to design an effective water point maintenance program.

A key challenge in designing these programs is the tension between preventive maintenance and reactive repairs. While preventive maintenance can reduce breakdowns and repair costs, many NGOs adopt a “fix-only-when-broken” approach. According to a senior officer of the *United Nations* Office for Project Services in Ethiopia when comparing to preventive maintenance, “most NGOs engage reactively on maintenance as an emergency response [because] they think it is too expensive and there is a lack of funding.” This approach often stems from systemic barriers such as short donor funding cycles (1–2 years), which are insufficient for the sustained commitments preventive maintenance requires (Corbett et al. 2022).

Within the operations management (OM) literature, the effectiveness of preventive maintenance depends on the context. For example, in humanitarian fleet management, preventive maintenance of motorcycles enables health workers to visit rural villages more frequently (Mehta et al. 2016). However, in healthcare operations, preventive care may underperform when resource trade-offs prioritize urgent needs over failure prevention (Örmeci et al. 2016, Grant et al. 2022). These divergent findings highlight the importance of context-aware decision-making (Corbett 2024). To respond to calls for contextualized operational analysis, we study the operations of three NGOs in three SSA countries. This leads to our first research question: *What is the value of incorporating preventive maintenance in maintenance programs to reduce water point downtime?*

The challenges of preventive maintenance are further compounded by the availability of functionality information. In resource-constrained settings like rural SSA, NGOs often lack complete

information about water point functionality. Approaches to collecting functionality information vary. Some NGOs invest in costly initiatives like call centers, while others rely on periodic visits without gathering functionality data.

While information is widely recognized in OM as a valuable asset that reduces uncertainty and improves decision-making (Lawrence 1999, Raju and Roy 2000), its benefits are less clear in resource-constrained environments (Corbett 2018). As Simon (1971) noted, “In an information-rich world, the wealth of information means a dearth of something else — a scarcity of whatever it is that information consumes.” Corbett (2018) similarly argues that information is only valuable when it leads to better decisions. Therefore, we aim to identify the conditions under which the availability of functionality information is beneficial and evaluate its impact on water point downtime and logistics cost. This leads to our second research question: *What is the value of functionality information in a water point maintenance program?*

To address these questions, we adopt the field-based approach of de Vries and Van Wassenhove (2020) and conduct comparative research with three NGOs in SSA, each with different practices. We conducted field research in Ethiopia (January 2020) and Malawi (June 2022). In Malawi, *Fisherman's Rest* runs an *outgoing* call center to contact communities and gather full functionality information, enabling reactive targeted repairs. In the Central African Republic (CAR), *Water for Good* follows a cyclic visitation schedule that does not rely on advance functionality information, ensuring that each water point is visited at least once per year as part of a blanket rotation. In Ethiopia, *Relief Society of Tigray (REST)* runs an *incoming* call center, relying on communities to report problems, thereby resulting in partial functionality information.

Through these partnerships, we collected a unique dataset on water point functionality in rural SSA, addressing a significant data scarcity in this region (Starr and Van Wassenhove 2014). The dataset contains 47,240 longitudinal records of water point functionality status, spanning nearly ten years in Malawi, five years in CAR, and three years in Ethiopia.

Connecting what we learned from the field and the OM literature, we develop an infinite-horizon discounted Markov decision process (MDP). In our proposed approach, NGOs conduct both preventive maintenance and reactive repairs on all water points they are scheduled to visit. Consistent with practices in CAR and Ethiopia, we group water points into clusters. In each period, the NGO observes the number of water points *reported* broken in each cluster (which is zero in the no-information setting) and how long it has been since the last visit. It then decides which cluster to

visit. Because many NGOs operate under tight budgets, the model minimizes both the expected water point downtime and the logistics cost of travel and part replacement.

To quantify the value of preventive maintenance (research question 1), we apply the MDP to data from each country. The results show substantial but varied benefits across nine regions and time frames. In Malawi, the reduction in downtime achieved by our optimization model ranges between 7.1% (with a 9.1% reduction in logistics cost) and 52.7% (with a 16.3% reduction in cost). In CAR, downtime reductions range between 41.7% (with a 18.8% increase in cost) and 54.7% (with a 15.7% increase in cost). In Ethiopia, downtime reductions range between 47.0% (with a 2.8% reduction in cost) and 61.9% (with a 6.3% reduction in cost). These results demonstrate that optimally scheduled preventive maintenance can substantially reduce downtime. Importantly, it can also lower logistics cost — a critical concern for NGOs operating under tight budget constraints.

To ensure the robustness of our findings, which rely on breakdown probabilities estimated using the Kaplan–Meier method with real-world data, we validate the survival model through temporal hold-outs and cross-validation. We also bound the MDP results using associated confidence intervals. Our cross-case research design supports context-specific managerial recommendations. It also allows us to examine how program effectiveness is shaped by current NGO practices, information availability, and local community capacities (de Vries and Van Wassenhove 2020, Corbett 2024), while also helping to bound the generalizability of our findings.

To understand the role of functionality information (research question 2), we conduct two numerical experiments on three clusters to examine the effects of parameters, such as the probability of breakdown and community reporting of breakdowns. Inspired by country-specific findings, these experiments use generated data, allowing the results to be extended to different contexts.

The first experiment compares the relative performance of reactive (e.g., Malawi) and cyclic (e.g., CAR) maintenance visitation approaches for reducing downtime. We find that the more effective visitation approach depends on both functionality information availability and repair demand. When information availability is high, the reactive approach is more effective under low demand, while the cyclic approach performs better under high demand.

The second experiment explores how NGO's prioritization of logistics cost and downtime affects outcomes. We find that when the focus is on reducing the logistics cost, increased functionality information (e.g., better call centers) has minimal impact on visitation schedules or on reducing downtime and cost. Thus, simply collecting more information may not ensure the most effective

program. Our results reconcile general OM principles with the practical constraints of capacity, attention, and cost in resource-limited settings.

Our work makes three contributions to the humanitarian operations literature. First, we provide one of the few analytical studies that quantify the value of preventive maintenance in humanitarian settings, drawing on extensive field data from rural SSA. This directly addresses Corbett's (2024) call to measure the benefits of harm prevention and to develop actionable operational benchmarks (Corbett et al. 2022). Second, we address Akkermans et al.'s (2018) call to explore sustainable infrastructure in resource-constrained settings by offering practical insights into the design of effective water point maintenance programs. Third, we challenge the common assumption that more data always improves outcomes, instead showing that its value is context-dependent.

Our findings further offer actionable managerial insights. First, NGOs can leverage our results to raise donor awareness about the need to fund maintenance programs that extend beyond the typical 1–2 year funding cycle, enabling the establishment of supply chains necessary for effective preventive maintenance. Without sustained investment, new water points risk becoming unreliable over time. Second, when faced with resource constraints, NGOs should prioritize reducing spare parts costs (e.g., efficient spare part sourcing), enhancing handpump reliability (e.g., using higher quality spare parts), and expanding repair capacity (e.g., expanding the mechanic team), rather than focusing on collecting additional functionality information.

2. Literature Review

Our research lies at the intersection of three streams of OM literature: humanitarian operations to improve water access, infrastructure asset management, and the value of information.

2.1. Humanitarian Operations to Improve Water Access

Humanitarian OM research on water access in developing countries primarily emphasizes equitable distribution and infrastructure development. For instance, Dawande et al. (2013) propose pricing policies that encourage farmers to optimize water usage under limited monitoring and resources. Zhai et al. (2023) examine the bottom-up decision-making system for building new water points, and optimize their locations in rural SSA to reduce travel distance to water points. Leveraging community involvement, the authors propose a more equitable centralized decision-making system for building new water points. Fink et al. (2022) identify factors influencing water point functionality, including mechanic response time and proximity to urban centers. However, these studies do not address operational strategies for maintaining existing water infrastructure or quantifying

the benefits of preventive maintenance, making our paper among the first to analyze water point maintenance strategies and their value for sustainable service delivery.

2.2. Managing Infrastructure Assets Under Incomplete Information

There is a tension over the effectiveness of preventive maintenance, as its benefits are context-dependent. Reliability theory, a key theory of operations research, emphasizes the importance of reducing system failures through preventive strategies (Barlow 1998, Aven and Jensen 1999, Nahmias and Olsen 2021). In humanitarian contexts, for example, Mehta et al. (2016) demonstrate how systematic preventive maintenance of motorcycles enables health workers to conduct more effective outreach in rural villages, improving access to healthcare. Similarly, Guajardo et al. (2012) show that performance-based contracts incentivize service providers to perform more frequent preventive maintenance, which significantly enhances the reliability of aircraft engines.

Preventive actions, however, are not always effective. For example, when hospitals face full near-term appointment capacity, they must choose between delaying preventive visits, risking costly patient outcomes, or using expensive surge capacity to schedule them sooner. Grant et al. (2022) show that when surge capacity costs greatly exceed readmission costs, or when patient demand is high, delaying preventive appointments may be optimal. Similarly, Örmeci et al. (2016) highlight the trade-off between addressing urgent diagnostic needs and offering preventive screenings, finding that hospitals should deprioritize preventive screenings once more than 50% of the population has been screened. These examples illustrate the potential of preventive maintenance to improve system performance and reduce costs when applied under the right conditions.

The challenges of such trade-offs become even more pronounced in resource-constrained settings like rural SSA, where NGOs managing water points often lack complete information about their functionality status. Unlike information-rich environments, such as fleet or hospital management, NGOs must allocate limited resources between repairing known broken water points and addressing those with unknown functionality status. This uncertainty, coupled with tight budget constraints, complicates decisions on whether and how to implement preventive maintenance effectively.

The OM literature further offers conflicting views regarding whether investing in functionality information justifies the effort. On one hand, OM research highlights the value of information in reducing uncertainty and improving decision-making (Lawrence 1999, Raju and Roy 2000). For example, advance demand information enhances supply chain performance, improving inventory utilization (Lee et al. 2000, Özer and Wei 2004). Conversely, some studies challenge the

assumption that more information leads to better outcomes. For example, Gallego and Özer (2001) demonstrate that advance demand information loses its operational value when the demand exceeds a certain threshold. Similarly, Ha and Tong (2008) find that the benefits of information sharing depend on the competitive environment, where it can either enhance or harm decision-making.

Our work builds on this tension in the literature to examine the conditions under which information becomes valuable and test the common NGO practice of collecting as much functionality information as possible. Despite the need for effective maintenance strategies in such challenging environments, OM research remains limited (Akkermans et al. 2018). Our study addresses this gap by comparing optimization model solutions with extensive field data collected from NGOs in three SSA countries. While all NGOs face resource constraints, they differ in their maintenance strategies, functionality information availability, and degree of community engagement.

3. Methods and Research Design

Our integrated approach (1) formulates an infinite-horizon discounted Markov decision process (MDP) to study NGO water point maintenance visitation decisions; (2) applies the MDP to field data collected in each country; and (3) runs numerical experiments with generated data to further examine decision trade-offs.

3.1. Optimization Model

We develop an infinite-horizon discounted MDP to model NGO maintenance programs. The model minimizes expected water point downtime and logistics cost by determining the optimal maintenance visitation schedules. The notation is summarized in Table 1. The NGO manages N water points, divided into N_c clusters. Each cluster $i \in \mathbb{N}_c = \{1, 2, \dots, N_c\}$ has W_i water points.

At the start of each period $t \in \mathbb{T} = \{1, 2, 3, \dots\}$, the NGO observes the state of each cluster i , $s_{it} = (\Omega_{it}, \Delta_{it})$, where Ω_{it} is the number of water points reported broken (i.e., in need of a major repair) in cluster i and Δ_{it} is the number of periods since cluster i was last visited given the current period t . The system state across all N_c clusters is a vector, $\mathbf{s}_t = (s_{it})$. We assume a single mechanics team visits one cluster per period.

The decision variable is $a_{it} = 1$ if cluster i is visited in t , and $a_{it} = 0$ otherwise. The action of the system is $\mathbf{a}_t = (a_{it})$. Under this action, the mechanic team conducts a tour to perform the necessary services for *all* the water points in the assigned cluster. A functioning water point receives *preventive maintenance*. A water point with reduced water flow requires a *minor repair*. A water point not yielding any water requires a *major repair*. We assume the team carries sufficient (small) spare

Table 1 Model Notation

N	Total number of water points
N_c	Number of clusters
W_i	Number of water points in cluster i
Ω_{it}	Number of reported broken water points in cluster i in period t
Δ_{it}	Number of periods since last visit to cluster i in period t
s_{it}	State variable: $s_{it} = (\Omega_{it}, \Delta_{it})$
a_{it}	Decision variable: $a_{it} = 1$ if cluster i is visited in period t , and $a_{it} = 0$ otherwise
$p(\Delta_{it})$	Probability that a water point in cluster i is broken in period t
a, b	Parameters that control the shape of $p(\Delta_{it})$
q	Probability that a community calls to report a major breakdown
d_i	Tour length of visiting cluster i
Γ	Maximum number of periods allowed without visiting a cluster
ρ_f	Per-kilometer fuel cost
ρ_r, ρ_m	Part replacements cost in major repairs and preventive maintenance, respectively
α	Weight on logistics cost in the per-period cost function
η	Discount factor for the infinite-horizon MDP

parts for the services required. From a modeling perspective, we treat preventive maintenance and minor repair services as equivalent: neither one results in water point downtime, as the pump continues to provide water, and the costs are identical, as both involve only the replacement of the same fast-wearing spare parts (Appendix A).

We assume that after servicing, all water points resume perfect functionality. After all visits have been completed, the communities not visited in t report any major breakdown with probability q . This aligns with feedback from our NGO partners, who suggest that communities typically take action only when there is no water flow at all. As *Fisherman's Rest* in Malawi explains: “[preventive maintenance and a minor repair] is not [local communities’] priority at that time because they are still getting the water. They would prioritize other options [such as buying food].”

The probability of calling, q , can also represent outgoing calls from the Malawian NGO to the community. The NGO does not know the functionality status of the water points not reported broken. We assume q is homogeneous across communities. When $q = 0$, $\Omega_{it} = 0$. We impose the constraint $\Delta_{it} \leq \Gamma$, as some NGOs require all water points to be visited within a particular time frame (e.g., a year in Ethiopia and CAR).

We define $p(\Delta_{it}) \in [0, 1]$ as the probability that a water point in cluster i is broken (i.e., requires a major repair) in period t . We assume that the longer a water point has not been visited, the more

likely it is to need a major repair. This aligns with Fink et al. (2022), who find that shorter intervals between visits improve functionality. We further assume water point breakdowns are independent of each other and independent of the reporting of breakdowns (q). However, we discuss the possibility of an inter-dependency of water point usage and likelihood of breakdown in Appendix B. At the end of t , the NGO updates the state, \mathbf{s}_{t+1} , according to the following transition functions:

$$\Omega_{i,t+1} = (\Omega_{it} + h) \cdot (1 - a_{it}) \quad (1)$$

$$\text{with } Pr(H = h) = \binom{W_i - \Omega_{it}}{h} (p(\Delta_{it}) \cdot q)^h (1 - p(\Delta_{it}) \cdot q)^{W_i - \Omega_{it} - h},$$

$$\Delta_{i,t+1} = \Delta_{it} \cdot (1 - a_{it}) + 1. \quad (2)$$

If cluster i is visited in period t ($a_{it} = 1$), then $\Omega_{i,t+1} = 0$ and $\Delta_{i,t+1} = 1$. If cluster i is not visited in period t , then Δ_i increases by one. There will be h newly reported broken water points, distributed as a binomial random variable following the assumption of independent water point breakdowns. As indicated by the transition functions, the current period's decision impacts the next period's state and, thus, the future periods' downtime and logistics cost.

3.2. Water Point Downtime and Logistics Cost

The NGO minimizes logistics cost and water point downtime over an infinite time horizon. Logistics cost, weighted by $\alpha \in [0, 1]$, includes travel and part replacement costs. Fixed expenses such as mechanic salaries, irrespective of repair activities, are excluded. Water point downtime, weighted by $1 - \alpha$, is the expected number of water points requiring major repairs in cluster i . We assume that all repairs are completed towards the end of t , thus, all broken water points experience downtime in period t regardless of the visitation decision.

The expected per-period cost is

$$\begin{aligned} g(\mathbf{s}_t, \mathbf{a}_t) = & \alpha \sum_{i=1}^{N_c} \left[a_{it} \rho_f d_i + a_{it} \rho_r (\Omega_{it} + (W_i - \Omega_{it})(1 - q) p(\Delta_{it})) \right. \\ & + a_{it} \rho_m (W_i - \Omega_{it} - (W_i - \Omega_{it})(1 - q) p(\Delta_{it})) \Big] \\ & + (1 - \alpha) \sum_{i=1}^{N_c} \left[\Omega_{it} + (W_i - \Omega_{it})(1 - q) p(\Delta_{it}) \right] \end{aligned} \quad (3)$$

where d_i is the tour length in kilometers (km) if cluster i is visited. We assume the mechanic team visits all water points in the assigned cluster in a fixed order, starting and ending at the NGO office. The tour length, d_i , is calculated using a traveling salesman algorithm. ρ_f is the fuel cost

per kilometer. ρ_r and ρ_m ($< \rho_r$) are the part replacement costs of a major repair and preventive maintenance (or minor repair), respectively. In the optimization model formulated below, each future period is discounted by $\eta \in (0, 1)$.

$$\begin{aligned} \nu^*(\mathbf{s}) &= \min_{(\mathbf{a}_t)} \mathbb{E} \left[\sum_{t=0}^{\infty} \eta^t g(\mathbf{s}_t, \mathbf{a}_t) \middle| \mathbf{s}_0 = \mathbf{s} \right] \\ \text{s.t. } & \sum_{i=1}^{N_c} a_{it} = 1, \forall t \in \mathbb{T}, \end{aligned}$$

Equations (1) and (2),

$$\Delta_{it} \leq \Gamma,$$

$$a_{it} \in \{0, 1\}, \forall i \in \mathbb{N}_c, t \in \mathbb{T}.$$

When $\alpha = 0$, we can analytically prove the structure of the optimal policy for three clusters. Appendix D.1 presents the proof. We consider steady state and thus remove the subscript t . Note that the theorem holds for any value of the probability of calling (q).

Theorem 1 *For three clusters ($N_c = 3$) where cluster 1 was visited in the previous period, and $\Delta_2 > \Delta_3$, $\alpha = 0$, and $W_i = W, \forall i$: (a) there exists a threshold, $\tau(\Omega_2)$ such that for $\Omega_3 < \tau(\Omega_2)$ it is optimal for the NGO to visit cluster 2 and for $\Omega_3 \geq \tau(\Omega_2)$ it is optimal to visit cluster 3; (b) for $\Omega'_2 > \Omega_2$, $\tau(\Omega'_2) \geq \tau(\Omega_2)$; and (c) $\tau(\Omega_2) \geq \Omega_2$.*

Theorem 1 shows that, in the three-cluster setting with $\alpha = 0$, the optimal policy collapses to a simple threshold rule on the reported broken water points for the two clusters not visited in the previous period. Specifically, for each fixed Ω_2 , there exists a nondecreasing threshold $\tau(\Omega_2) \geq \Omega_2$ such that the NGO should visit cluster 2 whenever $\Omega_3 < \tau(\Omega_2)$ and cluster 3 otherwise. This analytic insight gives NGOs a clear, implementable decision criterion by comparing the observed breakdown counts to a single threshold. It also motivates the two heuristic visitation approaches (one based on Ω , one on Δ) that we explore in Section 5.3.

3.3. Alternative Objective Function

Instead of tracking a count of broken water point days, we can consider a convex increasing downtime *deprivation cost* that increases with the duration of a breakdown. Specifically, we apply a cost function $f(j) = j^2$ for a water point that has been broken for j periods. This formulation reflects how economic, educational, and health losses compound: women and children spend more time collecting water, reducing income and study hours, and reliance on unsafe sources increases the risk of diseases such as diarrhea and cholera.

The full model and results of the convex deprivation cost formulation appear in Appendix D.4. Since the true deprivation cost depends on local conditions (e.g., income levels, water quality) and is difficult to measure consistently and reliably, we use downtime days, an operational metric, as our primary metric. However, Section 7.1 outlines how to translate downtime into expected health improvements based on existing literature.

3.4. Application of the Optimization Model

In Sections 4, 5, and 6, we first describe each NGO's program in detail and then apply our MDP model to the corresponding field data from Malawi, CAR, and Ethiopia. For each country, we use the Kaplan-Meier (K-M) estimator to estimate the probability of a water point breaking down after a given number of months since its last visit, denoted as $\hat{p}(\Delta)$. This nonparametric method is well-suited for time-to-event analysis with right-censored data, as some water points were never recorded as broken.

To ensure that our K-M estimates are generalizable beyond the observed data, we conduct both temporal hold-out validation and grouped k-fold cross-validation, with $k = 3, 5, 7, 10$, and 15. For each split, we estimate K-M curves separately on the training and test sets, and apply a log-rank test to assess whether the time-to-breakdown distributions differ significantly. Across splits, held-out survival curves are statistically indistinguishable at the 5% level (log-rank $p > 0.05$).

Appendix C.1 presents the estimated breakdown probability curves and the validation results. We ultimately fit the final model using the full dataset for two reasons: (1) K-M estimation benefits from a larger number of observed events, which increases statistical precision and robustness; and (2) our goal is not to predict individual water point failures, but to estimate the expected proportion of water points found broken during a cluster visit after a given number of months.

Note that our optimization model is built on a cluster-based framework. Depending on the country, each cluster contains a large number of water points, resulting in a wide range of possible values for Ω_{it} and creating a large state space. Therefore, to reduce state space, when applying the optimization model to each country, we redefine Ω_{it} as the *percentage* of water points reported broken in cluster i in period t , discretized into integer values.

Consistent with the NGOs' primary objectives, we set $\alpha = 0$ to focus only on minimizing water point downtime. The model thus assesses the effect of preventive maintenance on reducing downtime and its influence on the associated logistics cost. To enable cross-country comparisons, we standardize cost parameters: fuel at $\rho_f = \$0.074$ per kilometer, major repairs at $\rho_r = \$60.6$, and

preventive maintenance or minor repairs at $\rho_m = \$12.5$. Repair costs reflect average costs from Malawi's data in 2022, and fuel costs are based on a 2019 grant from REST, the Ethiopian NGO.

Tables 2 to 4 present the main results from Malawi, CAR, and Ethiopia, respectively in Sections 4 to 6. *Column (A)* reports the NGO's observed breakdowns; *Column (B)* replaces those observations with values calculated using the estimated breakdown probability $\tilde{p}(\Delta)$. *Column (C)* and *Column (D)* both use the same $\tilde{p}(\Delta)$, with (C) adding preventive maintenance to the NGO's existing schedule and (D) applying our optimized visitation schedule with preventive maintenance.

All performance changes are measured relative to Column (B) to keep the probability baseline consistent. Comparing (A) to (B) confirms that $\tilde{p}(\Delta)$ closely reproduces observed breakdowns. The difference between (B) and (C) then isolates the impact of incorporating preventive maintenance under the current schedule, while the difference between (B) and (D) quantifies the benefit of incorporating preventive maintenance with an optimized visitation schedule.

To ensure the validity, robustness, and generalizability of our numerical results with real-world data, we quantify uncertainty by using the lower and upper bounds of the 95% confidence intervals on $\tilde{p}(\Delta)$ derived from our K-M estimations. We further apply the optimization model across multiple regions and time frames, covering nine distinct scenarios. This broader evaluation helps mitigate potential biases that might arise from single-case analyses.

Having applied our model to field data from Malawi, CAR, and Ethiopia, we next turn to controlled numerical experiments to further explore the insights that emerged from those real-world data applications. Using generated data allows us to systematically vary key parameters that are difficult to isolate in real-world data, extending the interpretability and generalizability of our results across diverse contexts. Specifically, we examine: (1) the impact of communities' repair capabilities (Malawi; Section 4.3); (2) the impact of functionality information under two distinct maintenance visitation approaches (CAR; Section 5.3); and (3) the impact of functionality information under varying budget priorities (Ethiopia; Section 6.3).

4. Results: Malawi (Full Information)

In Malawi, the NGO *Fisherman's Rest* has established a robust infrastructure for systematically collecting water point functionality data. Malawi, which ranks 172/193 on the Human Development Index (United Nations Development Programme 2022), has enjoyed sustained peace and stable governance since gaining independence in 1964 (World Bank 2024). Established in 1997, *Fisherman's Rest* implements a broad spectrum of programs, including water management, school construction, girls' empowerment initiatives, and reforestation activities.

To keep the water flowing, *Fisherman's Rest* implements a water point maintenance program known as *Madzi Alipo* ("there is water"), managing a total of 1,758 water points. Launched in Traditional Authority (TA, equivalent to a county in the U.S.) Somba, Blantyre in 2013 and expanded to other TAs since 2018, the program focuses on repairing broken water points. Teams of at least two mechanics travel in a minivan stocked with spare parts. These teams visit multiple known broken water points, carry out repairs, and return to the base camp on the same day. The NGO does not perform preventive maintenance, but instead, encourages communities to maintain their own water points. To support local communities, the NGO offers community training programs where resources and capacity allow, equipping locals with essential maintenance skills.

A key aspect of the *Madzi Alipo* program is its *outgoing call center*, which contacts each community every three months to collect water point functionality information. This regular data collection enables the NGO to operate with **full information** ($q = 1$) when deciding which water points to visit. Under this approach, the NGO focuses exclusively on repairing broken water points, adopting a **reactive visitation approach**.

4.1. Field Research and Data

To better understand the maintenance program, the lead author conducted field research in collaboration with *Fisherman's Rest* in June 2022. This involved participating in three major repairs, two minor repairs, and one community training session across different areas within the Blantyre District. Additionally, the researcher observed call center operations and data recording processes. These firsthand experiences provided valuable insights into the practical implementation of the program, enabling us to model current practices more accurately and effectively. During this collaboration, *Fisherman's Rest* also shared a comprehensive dataset of 14,839 call center and mechanic visitation records collected between 2013 and 2022 (Appendix C.2).

4.2. Results of the Optimization Model

We apply our optimization model to TA Somba and the entire Blantyre District in 2019 (pre-COVID) and 2022 (post-COVID). Since our model is cluster-based and the NGO in Malawi conducts only day trips, we modify our approach as follows, using TA Somba in 2019 as an illustrative case: (1) we apply the K-means clustering algorithm to divide the 201 water points in TA Somba into four clusters (Figure A.2); (2) to maintain a consistent total number of services while limiting one cluster visit per period, we define each time period as four months; (3) to construct the visitation schedule used in Column (C), we examine the NGO's historical mechanic visitation data and,

for each cluster, identify the month in 2019 when the NGO visited that cluster most frequently. That month is then used to define the scheduled visit period for the cluster. Since there are four clusters but only three periods in a year, we further remove the cluster that was visited least frequently overall. This ensures that only one cluster is visited per period. Column (C) thus reflects a counterfactual scenario in which the NGO continues its historical visitation patterns, but adds preventive maintenance to the NGO's existing operations.

According to the data in Malawi, reported breakdowns never exceeded 15% of water points, so we restrict $\Omega_{it} \leq 15$ to ensure computational tractability. To estimate $\tilde{p}(\Delta)$, we use data recorded since 2019, when the NGO began systematic monitoring of water point functionality.

Table 2 first presents the results for TA Somba in 2019, as it is the area where the NGO has operated the longest. Comparing Column (D) to (B), incorporating preventive maintenance with an optimized schedule can potentially reduce downtime by 6,245 days (50.6%) and logistics cost by \$1,747 (34.0%). This reduction in downtime translates to 31 additional days of clean drinking water for each community in TA Somba in 2019.

Note that travel distances in Columns (C) and (D) are lower than those in Columns (A) and (B) because current practice involves daily round trips from the base camp, whereas our model consolidates visits into a single tour per period. Under current practice, the NGO visits an average of 1.64 water points per day, with a third quartile of two and a maximum of six. To more closely reflect this daily tour structure, we added Column (E), which assumes that the NGO visits two water points per day, the most common case, and starts and ends each day at the office. As a result, travel distance increases in Column (E). However, relative to the current practice presented in Column (B), incorporating preventive maintenance with an optimized schedule that minimizes downtime still decreases logistics cost by \$1,668 (32.3%).

Furthermore, under the lower-bound estimate of $\tilde{p}(\Delta)$, the optimization model reduces downtime by 5,961 days (51.6%) and associated logistics cost by \$1,423 (29.4%); under the upper-bound estimate, the model reduces downtime by 6,658 days (50.8%) and cost by \$2,177 (40.1%).

Appendix C.2.2 presents the results from TA Somba in 2022 (post-COVID) and the broader Blantyre district in both 2019 and 2022. In 2022, Fisherman's Rest managed 272 water points in TA Somba, an increase from 201 in 2019. Compared to current practices, incorporating preventive maintenance with an optimized cluster visitation schedule can reduce downtime by 709 days (8.6%) and logistics cost by \$1,648 (27.7%). From 2019 to 2022, the benefit of preventive maintenance in TA Somba declined. We posit that this decline could be attributed to the NGO's

Table 2 Results for 201 water points in TA Somba, Malawi, in 2019 (with K-M bounds)

	(A) Reported	(B) $\tilde{p}(\Delta)$	(C) With Preventive Maintenance	(D)		(E)		Lower		Upper	
				MDP One Tour	2WP/Day	$\tilde{p}(\Delta)$	MDP	$\tilde{p}(\Delta)$	MDP	$\tilde{p}(\Delta)$	MDP
Downtime (Days)	10,978	12,335	6,270	6,090		11,549	5,588	13,109	6,451		
# Major	60	63	26	24		57	25	69	20		
# Minor	104	101	151	154		107	152	95	162		
Distance (km)	761	761	132	137	1,208	761	131	761	145		
Logistics Cost (\$)	4,992	5,137	3,473	3,390	3,469	4,848	3,425	5,425	3,248		
Parts Cost (\$)	4,936	5,080	3,463	3,379	3,379	4,792	3,415	5,369	3,237		
(% of Total)	(98.88%)	(98.89%)	(99.71%)	(99.68%)	(97.41%)	(98.84%)	(99.71%)	(98.97%)	(99.66%)		
Travel Cost (\$)	56	57	10	11	90	56	10	56	11		
(% of total)	(1.12%)	(1.11%)	(0.29%)	(0.32%)	(2.59%)	(1.16%)	(0.29%)	(1.03%)	(0.34%)		

Note: From column (B) to (D), logistics cost decreases by \$1,747, of which \$1,701 (97.4%) is from lower parts costs and \$46 (2.6%) is from reduced travel distance. Column (E) assumes daily tours in which mechanics visit only two water points per day, resulting in higher travel distances.

community training program, which empowered local communities to conduct repairs independently. Community-led repairs increased substantially during this period, from 17 major and 35 minor repairs in 2019 to 43 major and 71 minor repairs in 2022. As a result, the need for preventive maintenance provided by the NGO may have diminished, reducing the effectiveness of preventive maintenance that relies on scheduled visits by NGO teams. A similar decline in benefit is observed when applying the model to the entire Blantyre district in 2019 vs. 2022.

4.3. Numerical Experiment: Incorporating Communities' Ability to Repair

The decline in the benefit of preventive maintenance observed in Section 4.2 suggests that community-led repair activity plays an increasingly important role in water point functionality in Malawi. Motivated by this empirical pattern, particularly the substantial rise in community-driven repairs between 2019 and 2022, we extend our model to explicitly incorporate communities' ability to address breakdowns independently.

Specifically, we introduce a parameter β representing the probability that a new breakdown is repaired by the community within the same period. If the community does not repair the breakdown, the issue remains unresolved until the NGO intervenes. After any community-led repairs have occurred, the NGO visits one cluster to carry out all necessary maintenance and repairs. This extension allows us to evaluate how increasing community repair capacity, as observed in TA Somba and the broader Blantyre district, may affect the relative value of preventive maintenance. The full model formulation is provided in Appendix D.2.

We conduct a numerical experiment using generated data with $N_c = 4$ clusters, each containing $W_i = 10$ water points. To isolate the effect on downtime, we set $\alpha = 0$ and assume full functionality information ($q = 1$), reflecting conditions observed in Malawi. Breakdown probabilities follow the logistic function $p(\Delta_{it}) = \frac{1}{1+e^{-a(\Delta_{it}-b)}}$, consistent with the handpump maintenance manual (Government of Malawi 2015). We fix $a = 1$ to represent gradual deterioration over time and vary b from 1.5 to 6.5. We consider values of $\beta \in \{0, 0.1, 0.2, \dots, 0.9\}$.

As shown in Figure A.4 (Appendix D.2), increasing β consistently lowers downtime across all b values. When b is small, indicating high water point repair demand relative to NGO capacity, even modest increases in β yield substantial gains. For example, with $b = 1.5$, increasing β from 0 to 0.1 reduces average downtime by 5.6%. At higher values of β , the percentage decrease becomes more pronounced: moving from $\beta = 0.8$ to $\beta = 0.9$ reduces downtime by 34.3%. In contrast, with $b = 6.5$ (i.e., lower repair demand relative to NGO capacity), each 0.1 increase in β leads to an average reduction of 2.8% in downtime, ranging from 2.5% to 3.2% per interval.

These findings suggest that improving community repair capacity is especially effective when water points are more prone to breakdowns. Even small gains in β can yield large reductions in downtime. When breakdowns are less frequent, the marginal benefit diminishes. However, such training remains valuable as it can serve as a sustainable exit strategy for NGOs, enabling them to transition out while preserving high functionality through localized ownership and repair capacity.

Although community training can play an important role in supporting self-repairs, the effectiveness and sustainability of such programs can vary across contexts due to differences in community capacity, turnover among trained individuals, and the local institutional environment. As a result, many NGOs prioritize complementary strategies, such as scheduled professional preventive maintenance, that offer more predictable cost, reliability, and logistical outcomes at scale.

4.4. Discussion of Results

In Malawi, our analysis indicates that part of the benefits achieved by our model in 2019 stems from preventive maintenance on water points at imminent breakdown that had not yet failed. Our results suggest a shift from the current practices of local day-trips focused solely on repairing broken water points to a more comprehensive approach that includes preventive maintenance for entire clusters. While *Fisherman's Rest* effectively addresses reported breakdowns, it often overlooks water points that have not been serviced for extended periods. This reactive approach aligns with the NGO's philosophy of encouraging community engagement and fostering independence, as evidenced by its training program. However, adopting a more proactive approach that incorporates preventive maintenance for at-risk water points could significantly reduce downtime and logistics cost.

5. Results: Central African Republic (No Information)

We now apply the model to Central African Republic (CAR), where the NGO *Water for Good* lacks the infrastructure or resources to systematically collect functionality information. CAR, which ranks 191/193 on the Human Development Index (United Nations Development Programme 2022), faces the constant threat of localized armed conflict throughout the country (U.S. Department of State 2024). Established in 2004, the NGO serves both urban and rural communities with limited access to clean water. Dedicated to preventing waterborne diseases, it also runs a popular radio program to raise awareness about hygiene, sanitation, and clean water practices. It has established two base camps in Berbérati and Bangui, each overseeing its own team of employees.

To keep the water flowing, *Water for Good* operates one of the longest-running maintenance programs, managing a total of 2,348 water points. The NGO has used electronic reporting systems to track spare parts across a supply chain spanning countries including France and Cameroon. The NGO manages its water points in groups, with the goal of visiting each group twice annually.

Every month, a team of two mechanics embarks on a 21-day pre-determined tour to visit multiple groups of water points. Mechanics stay overnight in the field and then take the rest of the month off. The team travels on trucks stocked with spare parts, tools, and at least 400 liters of fuel. It repairs broken water points but does not routinely carry out preventive maintenance, such as replacing fast-wearing parts on those that are still functional. While the NGO aims to visit all water points, logistical challenges, such as road inaccessibility or security concerns, can prevent this.

Water for Good does not have a call center, so it operates with **no information** (i.e., $q = 0$) about water point functionality status when deciding which water points to visit. Under this approach, the NGO adopts a **cyclic visitation approach**, relying exclusively on the time elapsed since the last visit to a group of water points.

5.1. Field Data

Because CAR is consistently rated as ‘Do not travel’ by the U.S. Department of State (2024), we did not conduct field research there. Nonetheless, *Water for Good* provided us with mechanic visitation data from 2018 to 2022, comprising a total of 14,817 records.

5.2. Results of the Optimization Model

Since the Berbérati and Bangui base camp operates independently, we apply the model to water points in each base camp separately. In practice, a monthly security report may prevent a tour from commencing due to localized armed conflict en route to or within the proposed visitation area. Thus, to look at a steady-state period, we focus on data from 2022 only.

Table 3 presents the results in the Berbérati area in CAR in 2022. There are 693 water points that were reached by the two mechanic teams operating out of the Berbérati base camp in 2022. We apply the K-means clustering algorithm to divide these water points into 12 clusters. With two mechanic teams operating from the Berbérati base camp, we assign six clusters to each team, evaluating all possible distributions to balance workload and select the assignment that minimizes this downtime. For $\tilde{p}(\Delta)$, we use all data from Berbérati since 2018.

In 2022 (presented in Column A), mechanics missed 73 visits to 45 water points due to logistical challenges discussed by the NGO. The NGO visited but did not perform preventive maintenance on 241 water points. These figures show that currently the NGO does not *routinely* conduct preventive maintenance in their visits. Comparing Column (D) to (B), incorporating preventive maintenance with an optimized schedule can potentially reduce downtime by 9,035 days (54.7%) at an additional logistics cost of \$3,418 (15.7%). This reduction in downtime translates to 13 additional days of clean drinking water for each community in 2022.

Furthermore, under the lower-bound estimate of $\tilde{p}(\Delta)$, the optimization model reduces downtime by 8,694 days (58.3%), but increases logistics cost by \$3,417 (16.3%); under the upper-bound, the model reduces downtime 9,337 days (51.7%) and increases logistics cost by \$3,430 (15.3%).

Table 3 Results for 693 water points in the Berbérati area, CAR, in 2022 (with K-M bounds)

	(A) Reported	(B) $\tilde{p}(\Delta)$	(C) With Preventive Maintenance	(D) MDP	Lower $\tilde{p}(\Delta)$	MDP	Upper $\tilde{p}(\Delta)$	MDP
Downtime (Days)	14,885	16,513	12,427	7,478	14,920	6,226	18,058	8,721
# Major	186	190	193	126	175	111	205	141
# Minor/ Preventive	771	767	1,078	1,342	782	1,357	752	1,328
Distance (km)	8,225	8,225	8,923	9,691	8,225	9,688	8,225	9,695
Logistics Cost (\$)	21,518	21,710	25,831	25,128	20,989	24,406	22,432	25,862
Parts Cost (\$)	20,909	21,102	25,171	24,411	20,380	23,689	21,823	25,145
(% of total)	(97.17%)	(97.20%)	(97.44%)	(97.15%)	(97.10%)	(97.06%)	(97.29%)	(97.23%)
Travel Cost (\$)	609	608	660	717	609	717	609	717
(% of total)	(2.83%)	(2.80%)	(2.56%)	(2.85%)	(2.90%)	(2.94%)	(2.71%)	(2.77%)

Note: From column (B) to (D), logistics cost increases by \$3,418, of which \$3,309 (96.8%) is from higher parts costs and \$109 (3.2%) is from increased travel distance.

Since we observe an increase in the logistics cost, in Appendix C.3.3, we formulate a MDP in which we minimize downtime subject to a budget constraint. We introduce a new state variable ϕ_t

that tracks the remaining budget at time t , then solve the problem via backward induction - from the terminal period (when no future logistics cost is incurred) back to the first period. We apply this model to the data from Berbérati. Under the budget constraint (with \$17 remaining at the end of month 12), the minimum downtime increases to 8,249 days, which is still a reduction of 50.0% from the current practice (compared to 54.7% without a budget constraint).

Appendix C.3.2 presents the detailed results from the Bangui area in 2022. We find that the optimization model, compared to current practice, could reduce downtime by 6,216 days (41.7%) at an increase of logistics cost by \$4,289 (18.8%). These results are consistent with those from the Berbérati area.

5.3. Numerical Experiment: The Role of Functionality Information Availability

Water for Good in CAR, which currently does not collect any functionality information in rural areas before making visitation decisions, has started a call center for urban areas. This effort is part of a broader trend in the water sector, where many NGOs are adopting call centers to improve functionality monitoring. However, as NGOs collect more functionality information, they often face challenges in effectively utilizing it, such as optimizing resource allocation and prioritizing tasks, which may require adjustments to their maintenance visitation approaches.

Therefore, to understand the value of information availability (research question 2), particularly its impact on different maintenance visitation approaches, we conduct a numerical experiment using generated data. This experiment compares, under varying levels of information, two maintenance visitation approaches that are easy to implement. The **reactive maintenance visitation approach**, based on current practices in Malawi, visits the cluster with the most reported broken water points (i.e., highest Ω). The **cyclic maintenance visitation approach**, based on current practices in CAR, visits the cluster that has not been visited for the longest time (i.e., highest Δ).

We consider $N_c = 3$, with $\alpha = 0$ and $q \in \{0, 0.1, \dots, 0.9, 1\}$. The probability of a breakdown is $p(\Delta_{it}) = \frac{1}{1+e^{-a \cdot (\Delta_{it}-b)}}$, with $a \in \{1, 2, 3, 4, 5\}$ and $b \in \{1.5, 2.5, \dots, 6.5\}$, consistent with the specification used in our other numerical studies. Without loss of generality, we consider the initial state $((0, 1), (\Omega_2, \Delta_2), (\Omega_3, \Delta_3))$ with $\Delta_2 > \Delta_3$.

We examine the percentage of all initial states in which the cyclic approach outperforms the reactive approach, and Figure 1 summarizes the key insights from this comparison. Detailed numerical results and analytical derivations describing each policy's structure are provided in Appendix D.3. We find that at low information availability (i.e., low q), the two maintenance visitation approaches perform similarly. However, at high information availability, when the repair demand is low (i.e.,

when $N_c < b$), the NGO has enough capacity to target reported breakdowns before additional breakdowns occur, making the reactive approach preferable. As the repair demand increases (i.e., $b < N_c$), all water points are likely to break down, whether they are reported or not. In such cases, a cyclic visitation approach becomes more effective. For robustness, we analyze the two maintenance visitation approaches using four clusters and find consistent results (Appendix D.3).

Figure 1 Summary of results in the comparison of the reactive versus cyclic maintenance visitation approach

		Availability of Functionality Information	
Repair Demand Relative to Capacity		Low	High
Low	Similar Performance	Reactive Approach Better	
	High	Cyclic Approach Better	

5.4. Discussion of Results

Our analysis reveals that part of the benefits of our optimization model comes from correcting the NGO's tendency to over-visit smaller clusters while neglecting larger ones. For example, the need for 197 additional visits in our model is from the NGO over-visiting three smaller clusters and under-visiting five larger ones. By making small adjustments to balance visits across clusters of different sizes, as suggested by our model, the NGO could reduce downtime.

Our numerical experiment further suggests that NGOs should implement a maintenance visitation approach based on both information availability and repair demand. We recommend first investing in greater maintenance capacity or reduced repair demand (e.g., pump reliability) *before* expanding functionality data collection. For *Water for Good* in particular, if the decision is made to collect functionality information in rural areas, it can initially maintain its cyclic visitation approach. Once the NGO can consistently reach a high percentage of communities via the call center, it should adapt its visitation approach based on repair demand relative to capacity. This requires the NGO to continuously monitor its capacity as the program expands.

6. Result: Ethiopia (Partial Information)

We now apply our model to Ethiopia, where the NGO *Relief Society of Tigray (REST)* relies on communities to reach out when a water point is broken. Ethiopia, which ranks 176/193 on the Human Development Index (United Nations Development Programme 2022), experienced a civil war in the Tigray Region from November 2020 to November 2022. Established in 1978, *REST* is one of the largest grassroots NGOs in Tigray. The NGO works on a wide range of programs, including providing access to clean water, promoting sanitation and hygiene, improving maternal and child health, and combating malnutrition.

To keep the water flowing, *REST* implements a maintenance program known as *Wahis Mai* (“water guardian”), managing a total of 8,015 water points. The NGO divides the region into eight zones, each operating independently. According to *REST*’s guidelines, mechanics in each zone embark on tours lasting up to 21 days, staying overnight in the field. During these tours, they perform all necessary repairs and preventive maintenance services. Mechanics travel in groups of at least two, equipped with SUVs that carry spare parts, food, and camping supplies. In addition, *REST* operates its own vehicle and spare parts warehouse.

A key aspect of the *Wahis Mai* program is its *incoming call center*, which allows communities to report pump breakdowns. This community-driven reporting system means the NGO operates with **partial information** ($0 < q < 1$) about water point functionality status when deciding which water points to visit, since not all communities report breakdowns due to factors like limited phone access. Under this approach, the NGO adopts a **hybrid visitation approach**, combining information from community reports with scheduled visits based on the time elapsed since the last visit. A tour may be altered if new phone calls are received during the tour.

6.1. Field Research and Data

The lead author conducted field research in Tigray, Ethiopia, in collaboration with *REST* in January 2020, and visited 11 water points. At one site, unbeknownst to the NGO, the water point was not functioning. The researcher observed an older woman collecting water from a muddy puddle, choosing not to travel further to a functioning pump. This observation shows that some community members value convenience over quality, which is also noted by our NGO partners in other countries. This incident amplifies the need for reliable maintenance of water points to prevent communities from resorting to dirty water sources. Through this collaboration, *REST* provided us with mechanic visitation data from March 2020 to August 2022, comprising a total of 17,584 records. Unfortunately, they could not share with us the record of the calls received.

6.2. Results of the Optimization Model

Since the civil war started in Tigray in November 2020, mechanics have only visited water points they could access safely, introducing potential bias into model comparisons. To account for this, we apply the optimization model separately to pre-conflict and during-conflict data. Table 4 presents the results for the Western Zone of Ethiopia from March to October 2020 (pre-conflict). During this period, 398 water points in rural areas of the Western Zone were actively in use.

Because the data provided did not specify cluster assignments, we first reconstructed the actual tours taken by mechanics based on visitation records and then grouped the water points into five clusters (Appendix C.4.2). Reported breakdowns never exceeded 20%, so to ensure computational tractability, we define each unit of Ω_{it} as two percentage points and restrict $\Omega_{it} \leq 10$ (e.g., when $\Omega_{it} = 1$, 2% of the water points are broken, and when $\Omega_{it} = 2$, 4% are broken). To estimate $\tilde{p}(\Delta)$, we use pre-conflict data from across the Tigray Region (i.e., before November 1, 2020).

Comparing Column (D) to (B), preventive maintenance with an optimized schedule can potentially reduce downtime by 10,201 days (47.0%) and logistics cost by \$350 (2.8%). This reduction in downtime translates to 26 additional days of clean drinking water for each community over eight months. Furthermore, under the lower-bound estimate of $\tilde{p}(\Delta)$, the optimization model reduces downtime by 9,772 days (51.8%) and the logistics cost by \$1,342 (11.2%); under the upper-bound, the model reduces downtime by 12,154 days (49.9%) and logistics cost by \$953 (7.5%).

Table 4 Results for 398 water points in the Western Zone, Ethiopia, Mar to Oct 2020 (with K-M bounds)

	(A) Reported	(B) $\tilde{p}(\Delta)$	(C) With Preventive Maintenance	(D) MDP	Lower		Upper	
Downtime (Days)	19,090	21,682	15,457	11,481	18,864	9,092	24,344	12,190
# Major	95	100	125	98	93	83	107	96
# Minor/Preventive	491	486	507	473	493	442	479	463
Distance (km)	2,634	2,634	1,637	1,739	2,634	1,299	2,634	1,475
Logistics Cost (\$)	12,089	12,330	14,034	11,980	11,993	10,651	12,667	11,714
Parts Cost (\$)	11,895	12,135	13,912	11,851	11,798	10,555	12,472	11,605
(% of total)	(98.40%)	(98.42%)	(99.13%)	(98.92%)	(98.37%)	(99.10%)	(98.46%)	(99.07%)
Travel Cost (\$)	194	194	122	129	195	96	195	109
(% of total)	(1.60%)	(1.58%)	(0.87%)	(1.08%)	(1.63%)	(0.90%)	(1.54%)	(0.93%)

Note: From column (B) to column (D), logistics cost decreases by \$350, of which \$284 (81%) is from lower parts costs and \$66 (19%) is from reduced travel distance.

Due to the civil war, no data is available beyond November 2020 in the Western Zone. We therefore apply our model to the Seneale Zone for 2020 and 2021, where data remain available during the conflict (Appendix C.4.4). The Seneale Zone includes Mekelle, the capital of the Tigray Region and the NGO's headquarters. In 2020, the NGO conducted 1,017 visits to 625 water points; in 2021, it made 1,038 visits to 602 water points. Of the 625 water points visited in 2020, 138 were not visited in 2021. Conversely, 115 water points visited in 2021 had not been visited previously. These shifts reflect the changing accessibility of areas due to the conflict, as well as the NGO's decision to serve new water points for humanitarian reasons.

In 2020 (pre-conflict), incorporating preventive maintenance with an optimized schedule can reduce downtime by 18,256 days (48.7%) and logistics cost by \$1,051 (4.9%) relative to current practice. In 2021, it could reduce water point downtime by 33,297 days (61.9%) and logistics cost by \$1,775 (6.3%). This enhanced effect likely stems from the NGO's expansion to previously unserved water points, where scheduled preventive maintenance yields larger gains. However, fully realizing these benefits in 2021 would require 235 additional visits, which could be challenging amidst a war.

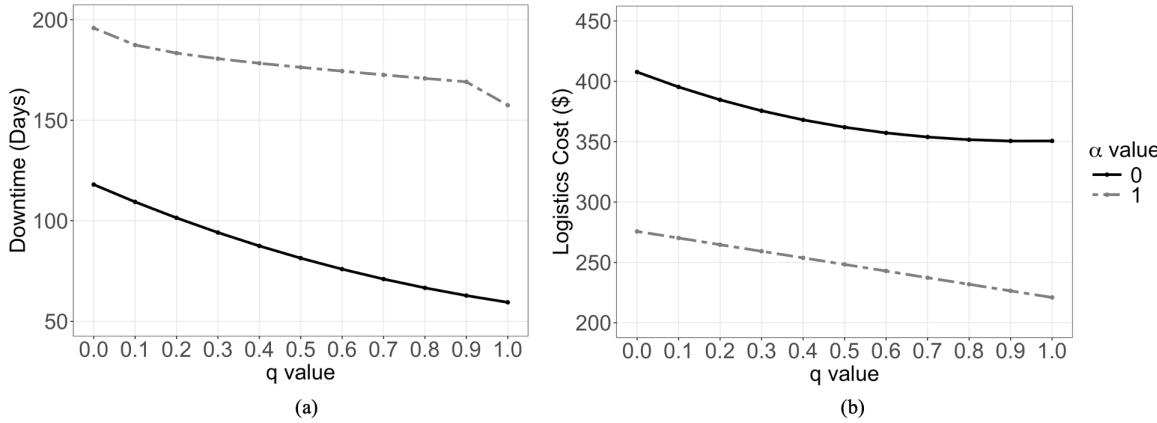
6.3. Numerical Experiment: The Impact of Functionality Information under Varying Budget Priorities

REST in Ethiopia continues to face limited access to external aid, even after the official end of the civil war (World Health Organization 2022). This financial concern is not unique to Ethiopia; many NGOs in the water sector operate under tight budgets. Achieving universal access to clean drinking water is estimated to require \$37.6 billion annually, nearly three times the current level of global investment (Hutton and Varughese 2016).

Motivated by this reality, we turn to our second research question: What is the value of functionality information, particularly when NGOs must balance minimizing water point downtime against managing logistics cost? We conduct a numerical experiment using generated data for three clusters to explore the effects of varying budgetary concerns and cost parameters. We are interested in how the availability of functionality information (i.e., q) affects the water point downtime and logistics cost as financial constraints tighten (i.e., α). We set $\alpha \in \{0, 0.1, \dots, 0.9, 1\}$ and the cost ratio ($\rho_r : \rho_m$) at 4 : 1, 5 : 1, 6 : 1, and 7 : 1 and assume equal travel distances for the three clusters.

Figure 2 shows changes in downtime (a) and logistics cost (b) as q increases, comparing $\alpha = 0$ (i.e., minimize only downtime) to $\alpha = 1$ (i.e., minimize only logistics cost). When $\alpha = 0$, increased information leads to a steeper decline in downtime, while the slope for logistics cost remains

Figure 2 Example of the changes in water point downtime (a) and logistics cost (b) as q increases for $\alpha = 0$ vs. $\alpha = 1$ when $a = 1$, $b = 2.5$, $(\rho_r : \rho_m) = (4 : 1)$, and $s_0 = ((0, 1), (0, 2), (0, 3))$



similar across α levels. These results suggest that functionality information is most valuable when financial constraints are moderate or low, enabling more informed and targeted decisions. However, as financial constraints tighten, NGOs tend to prioritize cost savings, reducing the operational advantage of detailed functionality data. Under severe constraints, the marginal benefit of such information diminishes, offering limited improvement in downtime or logistics cost reduction.

We further find that, with the same q value, downtime begins to increase significantly at lower α values when major repairs are more expensive. For example, when $q = 1$, downtime begins to increase significantly at $\alpha = 0.3, 0.2$, and 0.1 for cost ratios of $4:1$, $5:1$, $6:1$, and $7:1$, respectively. Collectively, these results suggest NGOs should prioritize reducing financial constraints, such as by lowering major repair costs, over collecting more functionality information.

6.4. Discussion of Results

In Ethiopia, our results indicate that preventive maintenance remains beneficial. In the Western Zone, part of the model's impact arises from reducing excessive visits to certain clusters. Although we lack direct data on community calls to the NGO, we use information collected during NGO mechanic visits to identify 144 cases in which communities reported having previously sought help from third parties (e.g., private mechanics). Notably, 41.7% of these visits occurred in a single cluster along a major road, which were visited four times by the NGO, while our model recommends only two. These findings underscore the need for NGOs with call centers to avoid neglecting less vocal or harder-to-reach communities. This aligns with Corbett (2018), who discusses how human biases can undermine decision-making even with ample data. Our model addresses this imbalance by distributing visits to both communities that report issues and those that do not.

Our numerical experiment further emphasizes the value of functionality information in optimizing maintenance decisions, particularly when minimizing downtime is the priority and financial constraints are moderate. However, in highly financially-constrained environments, the advantages of more functionality information become less pronounced. For NGOs operating with tight budgets, it may be more effective to prioritize reducing the costs of major repairs (e.g., improving parts sourcing), rather than allocating limited resources to collecting more functionality information.

7. Country Summaries and Potential Health Impacts

The NGOs operating in Malawi, CAR, and Ethiopia implement distinct water point maintenance programs, shaped by their local operational contexts and philosophies. To assess the value of preventive maintenance across these varied settings, we apply our optimization model to field data from each country. Table 5 summarizes the resulting changes in water point downtime and logistics cost across regions and time frames.

Table 5 Optimization results compared to current practices

Country	Region	Year	# Water Points	Downtime	Logistics Cost
Malawi <i>(Fisherman's Rest)</i>	TA Somba	2019	201	-6,245 days (-50.6%)	-\$1,747 (-34.0%)
	TA Somba	2022	272	-709 days (-8.6%)	-\$1,648 (-27.7%)
	Blantyre District	2019	728	-21,629 days (-52.7%)	-\$2,308 (-16.3%)
	Blantyre District	2022	1,502	-3,642 days (-7.1%)	-\$2,312 (-9.1%)
CAR <i>(Water for Good)</i>	Berbérati	2022	693	-9,035 days (-54.7%)	+\$3,3418 (+15.7%)
	Bangui	2022	991	-6,216 days (-41.7%)	+\$4,289 (+18.8%)
Ethiopia <i>(REST)</i>	Western Zone	2020	398	-10,201 days (-47.0%)	-\$350 (-2.8%)
	Seneale Zone	2020	625	-18,256 days (-48.7%)	-\$1,051 (-4.9%)
	Seneale Zone	2021	602	-33,297 days (-61.9%)	-\$1,775 (-6.3%)

Note: For the Malawi results presented in this table, we assume one large cluster tour per period. For the daily trip results, please refer to Table 2.

The column on downtime in Table 5 shows that, despite operational differences, incorporating preventive maintenance with an optimized cluster visitation schedule consistently reduces water point downtime across all nine cases analyzed. Seven cases exceed 40%. In Malawi, the improvement is driven by proactively targeting water points at imminent risk of failure, preventing future breakdowns. In Ethiopia, it results from a more balanced allocation of visits between communities that report breakdowns and those that do not. In CAR, the reduction is from a more even

rotation between smaller and larger clusters. In Malawi, improvements in 2022 are much smaller than in 2019 or in other countries, likely due to increased community self-repairs after training by *Fisherman's Rest*, which lowers the marginal value of NGO-led preventive maintenance.

The column on logistics cost shows the change in logistics cost associated with implementing preventive maintenance. In Malawi and Ethiopia, the logistics cost declines in every case, ranging from -2.7% to -33.7% and averaging -14.4% across the seven cases. In contrast, CAR sees an average logistics cost increase of 17.2%, reflecting the model's shift toward more adequate servicing of larger clusters. Nonetheless, such an increase is paired with an average downtime reduction of 48.2%, suggesting that the trade-off may be worthwhile.

While NGOs may be concerned about the potential cost of preventive maintenance, our findings show that logistics cost savings are achievable. For NGOs that operate with limited budgets, where every dollar matters, such reductions are especially valuable, freeing up resources that can be redirected to other essential activities. Even when logistics cost does increase, the substantial improvements in downtime may justify the investment. These results provide a compelling case for NGOs to pursue additional donor support to strengthen preventive maintenance programs.

Next, to explore the role of functionality information in NGO operations and provide actionable recommendations, we address our second research question in Sections 5.3 and 6.3 using two numerical experiments. Figures 1 and 2 show that functionality information is most valuable when NGOs have adequate repair capacity and prioritize reducing water point downtime over logistics cost. The effectiveness of a maintenance visitation approach depends not only on the availability of functionality information, which measures how often major breakdowns are reported by communities (note that communities do not always report such incidents), but also on repair demand, which reflects the true frequency of water point failures. The results suggest that strengthening repair capacity and pump reliability should take precedence over expanding functionality monitoring efforts.

For *Fisherman's Rest* in Malawi, which operates with high information availability, the results in Figure 1 show that a reactive maintenance visitation approach outperforms a cyclic approach only when repair demand is low. When water points become less reliable and require repairs more often, the advantage of the reactive approach in reducing downtime diminishes. Thus, Figure 1 indicates that maintaining high water point reliability and keeping repair demand low is important for the NGO to achieve reduced downtime under its current reactive approach. Encouragingly, in

response to findings from our study, the NGO began redesigning key parts in 2024, including a floating rod centralizer that prevents pipe damage, a common cause of major breakdowns.

For *Water for Good* in the Central African Republic, planned expansions of functionality information collection from urban to rural areas motivate a closer examination of the conditions under which a reactive maintenance visitation approach is effective. Figure 1 shows that increasing functionality information alone does not guarantee improved downtime: maintaining adequate pump reliability, and thus keeping repair demand within organizational capacity, is a prerequisite for successfully transitioning from a primarily cyclic approach toward a more reactive one. Furthermore, results from the numerical experiment on the impact of functionality information under varying budget priorities (Section 6.3) suggest that planned expansions in information collection should be accompanied by measures to reduce major repair costs. Lowering these costs helps contain repair demand and financial burdens, thereby enhancing the benefits of improved information availability.

For *REST* in Ethiopia, which suspended functionality data collection in many areas, including the Western Zone, due to the civil war, our findings suggest a clear operational path. As the NGO resumes operations, it should initially focus on strengthening visitation capacity, such as by reallocating staff from call centers to field teams, before reinstating large-scale functionality information collection (Figures 1 and 2).

7.1. Potential Health Impacts from Reductions in Water Point Downtime

When water points fail, communities often turn to unsafe sources, increasing the risk of water-borne diseases. In 2019, such diseases caused over 175,000 deaths among children under five in SSA (UNICEF 2023). Using data from Zambian water utilities and healthcare centers, Ashraf et al. (2021) estimate that each additional day without water access raises cases of diarrhea by one (13%), typhoid fever by 0.002 (23%), and respiratory infections by 2.4 (12%). Using these estimates, we project potential health gains implied by the downtime reductions in Table 6. These projections suggest that preventive maintenance programs have the potential to meaningfully reduce the incidence of these illnesses, highlighting the public health value of stronger maintenance program.

Given the challenges in obtaining community-level health data in SSA, the estimates from Ashraf et al. (2021) provide a reasonable proxy for the health impacts of our findings. Zambia serves as an appropriate reference point, as its rural infrastructure and healthcare challenges resemble those of our study countries. In addition, the neighboring Malawi shares similar geography and climate. These projections, however, are indicative rather than causal — they assume linear scaling, no seasonality, and comparable sanitation and hygiene contexts — and should be viewed as

order-of-magnitude guidance. When more precise health-impact data become available, the convex deprivation-cost model in Section 3.3 can be directly applied.

Table 6 Potential Reductions in the Number of Disease Cases from Improvements in Water Point Downtime

Optimization Model Results				Estimated Case Reductions, Ashraf et al. (2021)		
Country	Region	Year	Downtime	Diarrhea	Typhoid Fever	Respiratory Infections
Malawi <i>(Fisherman's Rest)</i>	TA Somba	2019	-6,245 days	-6,245	-12	-13,930
	TA Somba	2022	-709 days	-709	-1	-1,713
	Blantyre District	2019	-21,629 days	-21,629	-45	-52,270
	Blantyre District	2022	-3,642 days	-3,642	-8	-8,802
CAR <i>(Water for Good)</i>	Berbérati	2022	-9,035 days	-9,035	-19	-21,835
	Bangui	2022	-6,691 days	-6,691	-14	-16,170
Ethiopia (REST)	Western Zone	2020	-10,201 days	-10,201	-21	-24,652
	Seneale Zone	2020	-18,256 days	-18,256	-38	-44,119
	Seneale Zone	2021	-33,297 days	-33,297	-69	-80,468

8. Concluding Remarks

Nearly 200 million people in rural SSA rely on communal handpumps for clean drinking water (Foster et al. 2019). When these water points break down, communities face disruptions in their only source of safe water, increasing the risk of life-threatening diseases. Our research shows that preventive maintenance is a practical, high-leverage solution. Integrating preventive maintenance into existing NGO programs can reduce water point downtime by an average of 41.2% (ranging from 7.1% to 61.9%), and, in two of three countries, this comes with lower logistics cost; where costs rise (CAR), the downtime gains are large enough to merit strong consideration.

While recent studies on preventive maintenance in humanitarian operations have largely focused on fleet management (Pedraza-Martinez et al. 2011, Mehta et al. 2016, Chen et al. 2021), our work shifts the focus to rural water infrastructure. Using field data from three countries, we offer robust evidence of the potential value of preventive maintenance for rural water points. Our cross-case research design across three NGOs provides donors and NGOs with context-specific, decision-relevant evidence that preparedness- and mitigation-like efforts can be fiscally viable (de Vries and Van Wassenhove 2020, Corbett et al. 2022).

Furthermore, as Starr and Van Wassenhove (2014) remarked, “the three key problems [in humanitarian OM research] are data, data, and data.” However, our study reveals a critical nuance: the

value of functionality information is not universal; it depends on pump reliability and the NGO's visitation strategy. This echoes the caution from Corbett (2018) that "we should be cognizant of some of [big data's] concomitant downsides."

Our findings could extend to other contexts like off-grid solar technology, which powered over 490 million people in 2021 (World Bank 2022). However, poor maintenance has caused widespread failures: by 2023, only 5% of solar panels in India and 20% in Ugandan healthcare centers remained functional (The Washington Post 2023).

Although our study uses substantial field data, it has its limitations. First, we assume water point breakdowns are independent, which may overlook cascade failures where one breakdown influences nearby water points. While we explore this possibility in Appendix B, limited data prevent us from establishing a clear causal relationship between water point breakdowns. Second, despite our best efforts, obtaining health data with the necessary granularity proved extremely challenging in sub-Saharan African countries, limiting our ability to evaluate downstream impacts such as disease reduction or improved well-being.

Our research opens several avenues for future study. While the comparison of 2019 and 2022 results in Malawi indirectly shows the benefits of community training, future work could explicitly and jointly optimize NGO maintenance scheduling and community training programs to further improve system performance. Additionally, some NGOs are developing and introducing sensors that can measure daily water flow in handpumps remotely. Future studies should quantify the operational value of sensors under realistic capacity and budget constraints. Taken together, these steps could help donors and implementers fund and operate maintenance programs that reliably keep the water flowing at a sustainable cost.

Acknowledgments

References

Akkermans H, van der Valk W, Van Wassenhove LN, Wynstra F (2018) All along the asset life cycle: Research opportunities for operations and supply chain management. *Journal of Operations Management* .

Ashraf N, Glaeser E, Holland A, Steinberg BM (2021) Water, health and wealth: The impact of piped water outages on disease prevalence and financial transactions in zambia. *Economica* 88(351):755–781.

Aven T, Jensen U (1999) *Stochastic Models in Reliability* (New York: Springer).

Barlow R (1998) *Engineering Reliability* (Society for Industrial and Applied Mathematics).

Carter R (2021) *Rural Community Water Supply: Sustainable Services for All* (United Kingdom: Practical Action Publishing).

Chen L, Kim SH, Lee HL (2021) Vehicle maintenance contracting in developing economies: The role of social enterprise. *Manufacturing & Service Operations Management* 23(6):1651–1668.

Corbett CJ (2018) How sustainable is big data? *Production and Operations Management* 27(9):1685–1695.

Corbett CJ (2024) Om forum — the operations of well-being: An operational take on happiness, equity, and sustainability. *Manufacturing & Service Operations Management* 26(2):409–430.

Corbett CJ, Pedraza-Martinez AJ, Van Wassenhove LN (2022) Sustainable humanitarian operations: An integrated perspective. *Production and Operations Management* 31(11):4393–4406.

Dawande M, Gavirneni S, Mehrotra M, Mookerjee V (2013) Efficient distribution of water between head-reach and tail-end farms in developing countries. *Manufacturing & Service Operations Management* 15(2):221–238.

de Vries H, Van Wassenhove LN (2020) Do optimization models for humanitarian operations need a paradigm shift? *Production and Operations Management* 29(1):55–61.

Erpf K (2007) *Installation and Maintenance Manual for the Afridev Handpump (Revision 2–2007)*. Rural Water Supply Network (RWSN) / SKAT, St. Gallen, Switzerland, URL https://www.pseau.org/outils/ouvrages/skat_afridev_installation_and_maintenance_manual_2007.pdf.

Fink EL, Chintalapati P, Lane A, Wester A, Javernick-Will A, Linden K (2022) Determinants of rural hand-pump functionality through maintenance provision in the Central African Republic. *PLOS Water* 1(6):e0000024.

Foster T, Furey S, Banks B, Willetts J (2019) Functionality of handpump water supplies: a review of data from sub-Saharan Africa and the Asia-Pacific region. *International Journal of Water Resources Development* 36(5):1–15.

Gallego G, Özer O (2001) Integrating replenishment decisions with advance demand information. *Management Science* 47(10):1344–1360.

Government of Malawi (2015) *Community Based Management Training Manual* (Malawi).

Grant B, Gurvich I, Mutharasan RK, Van Mieghem JA (2022) Optimal dynamic appointment scheduling of base and surge capacity. *Manufacturing & Service Operations Management* 24(1):59–76.

Guajardo JA, Cohen MA, Kim SH, Netessine S (2012) Impact of performance-based contracting on product reliability: An empirical analysis. *Management Science* 58(5):961 – 979.

Ha YA, Tong S (2008) Contracting and information sharing under supply chain competition. *Management Science* 54(4):701–715.

Hutton G, Varughese M (2016) The costs of meeting the 2030 sustainable development goal targets on drinking water, sanitation, and hygiene. *Open Knowledge Repository* .

Lawrence DB (1999) *The economic value of information* (New York: Springer), ISBN 0387987061.

Lee LH, So CK, Tang SC (2000) The value of information sharing in a two-level supply chain. *Management Science* 46(5):626–643.

Mehta KM, Rerolle F, Rammohan SV, Albohm DC, Muwono G, Moseson H, Sept L, Lee HL, Bendavid E (2016) Systematic motorcycle management and health care delivery: A field trial. *American Journal of Public Health* 106(1):87 – 94.

Nahmias S, Olsen TL (2021) *Production and Operations Analytics* (Waveland Press), 8th edition.

Örmeci EL, Güneş ED, Kunduzcu D (2016) A modeling framework for control of preventive services. *Manufacturing & Service Operations Management* 18(2):227–244.

Özer O, Wei W (2004) Inventory control with limited capacity and advance demand information. *Operations Research* 52(6):988–1000.

Pedraza-Martinez AJ, Stapleton O, Van Wassenhove LN (2011) Field vehicle fleet management in humanitarian operations: A case-based approach. *Journal of Operations Management* 29:404–421.

Raju JS, Roy A (2000) Market information and firm performance. *Management Science* 46(8):1075–1084.

Santos Porras B (2025) Usaid and the uncertain future of international cooperation. URL <https://www.ie.edu/insights/articles/usaid-and-the-uncertain-future-of-international-cooperation/>.

Simon HA (1971) Designing organizations for an information-rich world. Greenberger M, ed., *Computers, Communication, and the Public Interest*, 37–72 (Baltimore, MD: The Johns Hopkins Press).

Skinner J (2009) Where every drop counts: Tackling rural africa's water crisis. Technical report, International Institute for Environment and Development (IIED).

Starr MK, Van Wassenhove LN (2014) Introduction to the special issue on humanitarian operations and crisis management. *Production and Operations Management* 23(6):925–937.

The Washington Post (2023) India joins rush to renewables, but its rural solar systems fall off grid. URL <https://www.washingtonpost.com/world/2023/07/31/india-solar-energy/>.

UNICEF (2023) Triple Threat: Water, Sanitation and Hygiene (WASH) in Humanitarian Contexts. URL <https://www.unicef.org/media/137206/file/triple-threat-wash-EN.pdf>, accessed: 2024-10-18.

United Nations Development Programme (2022) Human development insights. URL <https://hdr.undp.org/data-center/country-insights#/ranks>.

UPGro Consortium (2024) Hidden crisis: Unravelling current failures for future success in rural groundwater supply. URL <https://upgro.org/consortium/hidden-crisis2/>.

US Department of State (2024) Central African Republic. URL <https://travel.state.gov/content/travel.html>.

Wilson DJ (2019) The harmonic mean p-value for combining dependent tests. *Proceedings of the National Academy of Sciences USA* 116(4):1195–1200.

World Bank (2022) Off-grid solar market trends report 2022: State of the sector.

World Bank (2024) Malawi overview. URL <https://www.worldbank.org/en/country/malawi/overview>.

World Health Organization (2022) Crisis in northern Ethiopia.

Zhai C, Brethauer K, Mejia J, Pedraza-Martinez A (2023) Improving drinking water access and equity in rural sub-Saharan Africa. *Production and Operations Management* 32(9):2921–2939.