Supplemental Materials: Examining the Causes and Consequences of Medicine Theft Using Remote Tracking Technologies

Authors anonymized

February 4, 2025

Contents

1 Malawi Public Health System Context

Malawi has a government-sponsored healthcare system which consists of four types of health facilities: Health Centers, Rural Hospitals, District Hospitals, and Central Hospitals. Health Centers are facilities located at villages and responsible for primary care. These clinics have limited staff, composed of three employees: a nurse/midwife, a medical assistant, and a pharmacist. Together with Rural Hospitals, they are responsible for most of the health care in Malawi. They provide outpatient, maternity, and antenatal services. Secondary care is provided by the 26 District Hospitals across the country. In addition to primary care, District Hospitals are more equipped and staffed to provide other medical services, such as x-ray, ambulance, and laboratories. Finally Central Hospitals are the top tier of healthcare in Malawi and have more specialized services. There are only four of them in the country.

Health facilities are connected by a patient referral system. For instance, if patients cannot receive adequate treatment at local health centers, they will be referred to District Hospitals. Health facilities also share the same health commodity supply chain. They buy medicines and other medical supplies from Central Medical Stores. The Central Medical Stores Trust, established by the Malawian government in November 2010, is responsible for processing orders for medicines and medical supplies. CMST has three regional medical stores. We worked with the CMS South, located in Blantyre in the South Region.

Each central medical store processes orders for medicines and medical supplies from District Health Offices responsible for district hospitals and health facilities under its charge and Central Hospitals within their regions of operations. Health facilities send requests to the District Health Offices, and then the offices send the procurement orders to Central Medical Stores. The orders are shipped directly to health facilities.

There is a flow of information and health commodities between central medical stores and health facilities, as illustrated in the picture below. The Malawi Health Commodities Logistics Management System, the Ministry of Health's medical supply system of inventory management and recording for all medical supplies, is responsible for organizing the movement of health commodities. Health commodities are moved from the CMS to the regional medical stores, where they are packed and shipped for each health facility. The movement of health commodities is not only vertical. For example, health centers can also buy directly from District Hospitals. It's a common practice given the shortage of drugs in the country. However, because of issues with administrative capacity or collusion between corrupt health officials, many drugs often end up in clinics where they were not supposed to go. In Section 6.3.3, we discuss in more details how redeliveries work.

Figure S1: Movement of Health Commodities to Patients

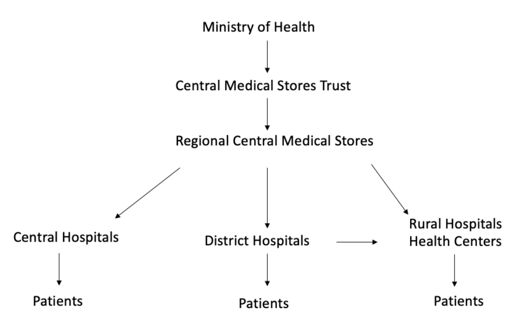

Malawi faces a critical issue: the lack of essential medicines in government health facilities. Stock-outs are common across health facilities and patients, who depend on the public health system, are often sent to buy drugs at private pharmacies and clinics. Many patients end up not receiving medication. According to a report by Oxfam, poor families spend up to 10% of their annual income on health care [1]. In the same study, Oxfam finds that only 9% of local health facilities in the country provided a full list of essential drugs for treating 11 common diseases. Our baseline survey indicates that 42% of respondents were told in the last months that their clinic could not provide needed drugs.

The reasons for drug shortages in Malawi are varied. Doctors and health authorities have blamed the centralized health delivery system, insufficient funding for drugs, the devaluation of the kwacha, lack of administrative capacity, and more recently the disruption of global system by the pandemic. In addition, drug theft has also been common and considered one of the main reasons for the acute drug shortage experienced by the country.

To address issues in the supply chain system and ensure access to medicines, in 2012, the Government of Malawi and donors developed a joint strategy for integrating the parallel supply chains into one supply chain managed by Central Medical Store Trust. The goal was to reform CMST to create the "necessary capacity and expertise to enable it to procure, store and distribute all essential medicines through one integrated supply chain system. [2]"

Donors and the government have established several initiatives to improve the delivery of medicines in Malawi. One of these is the campaign "I Speak Out Now", which encourages citizens to report drug theft via a hotline. Another is the Health Centre Advisory Committees (HCACs), a community-based program created by UKAid-funded Malawi Health Sector Programme Technical Assistance. The HCACs are groups of volunteers from the local community who monitor and oversee drug deliveries [3]. However, despite these efforts, drug and health commodities theft remains a problem in Malawi. The situation has become even more critical during the COVID-19 pandemic, with drug shortages becoming more common [4].

To address this issue, the Ministry of Health established the Drug Theft Investigations Unit (DTIU) in 2016. The DTIU is a government agency funded by the DFID, tasked with ensuring the security of public health procurement. Its primary strategy is to reduce drug theft by conducting audits, monitoring visits, and investigating those suspected of stealing drugs. The DTIU also collected and analyzed information provided by the hotline to target specific localities and support subsequent enforcement actions. Between August 2016 and April 2017, the DTIU investigated 62 individuals suspected of stealing and selling medicines from the public health system[2]. Of those, 16 were public health workers subsequently prosecuted for theft of medicines, according to a report by the Global Fund.

Despite the DTIU's efforts, it has had limited success due to insufficient resourcing and non-compliance with procurement regulations by local officials. Collaboration with the judicial system has also been challenging. To address these issues, we partnered with the DTIU to design interventions that do not require broad institutional reform or significant financial resources.

2 Timeline of Activities

This timeline outlines the sequence of activities conducted during the project in Malawi. Our study combined remote tracking data from an field experiment with survey and administrative data. As shown in Figure S2, baseline data collection started in March 2019. Our initial remote tracking data collection started a year later, in March 2020. However, due to the pandemic of COVID-19 activities had to been suspended. We redesigned the intervention to avoid in-person activities and rolled it out again in September 2021 alongside the remote tracking audit activities. The remote tracking audits lasted from October 2021 to January 2022. We completed the field activities with an endline phone survey with health officials.

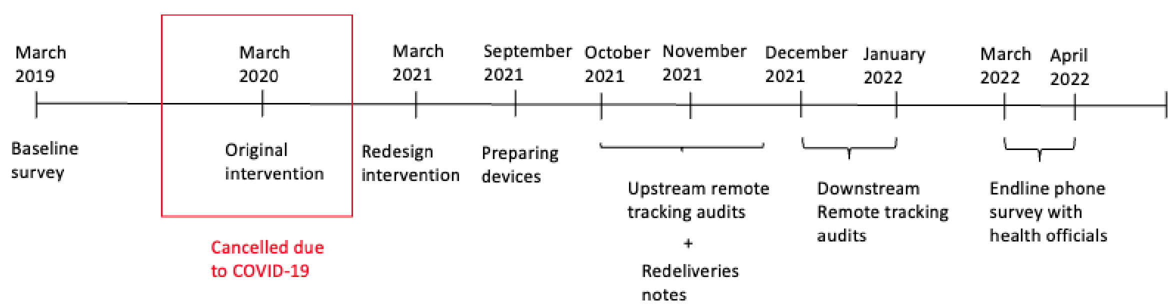

Figure S2: Timeline of Activities

3 Remote Tracking Audits

We conducted several tests in Malawi using different tracking devices to determine which would perform best for the project. We found that the OnAsset product was the best option for several reasons: (1) it was the most reliable product in terms of connectivity and communication; (2) it was easy to place and hide the OnAsset devices due to the design and size; and (3) the tracking protocol allowed for easy access to the API and data download.

3.1 Tracking Technologies

We use two types of tracking devices to track medications: (1) SENTRY 500, which is a cellular and GPS tracking device, and (2) SENTINEL, which is a bluetooth device that is automatically detected by any SENTRY 500 within its transmission radius of 300 meters. Figure S3 shows the devices. No installation is necessary for the devices, and they can simply be included in a shipped package. The SENTINELS were placed inside cartons or boxes of medications, while the SENTRY 500 units were placed inside trucks and used during follow-up mobile audits.

Figure S3: **OnAsset Technologies Devices.** The device in green is the SENTRY and in yellow SENTINELS.

The SENTRY 500 devices used Malawi's cell phone network to communicate and send location data. The SENTRY units capture and transmit data from itself and all surrounding SENTINELS to the server at regular intervals of 15 minutes. A technician consultant was responsible for running tests to improve the use of the devices during the experiment. These tests included calculating estimates such as how close the SENTRY needs to be to a SENTINEL in order to confirm the SENTINEL's existence with certainty and how long the devices' battery last. Additionally, we tested if different materials around the devices could affect the SENTRY's reading capacity. While the SENTINELS' batteries can last more than one year, SENTRY devices face a trade-off between battery life and reporting intervals. The battery can last up to three days after fully charged when using reporting intervals of 15 minutes. In our case, the SENTRY devices were only deployed for several hours when in delivery trucks, which did not exceed their battery life.

One challenge we faced was connectivity in Malawi. The SENTRY devices use both cellphone network and radio communication. Although many parts of Malawi do not have good connectivity, the SENTRY devices can store information and transmit the data to the server once they connect to the network. Thus, we were able to detect if the box/device arrived at the clinic at some point. For the audits of health facilities during the follow-up visits (Section 3.5), we also piloted tests with SENTRY devices in a backpack and SENTINELs within a box in an enclosed room.

3.2 Placing Tracking Devices in Government Warehouses and Delivery Vehicles

Prior to the placement of devices, all labels and identifying information were removed from SENTINEL devices. All SENTINELs were then relabeled with stickers identifying them as ”Temperature and Humidity Monitoring” devices and noting that they were property of the Malawi Ministry of Health and that the devices should not be removed (see example sticker in Figure S4). The stickers also provided a phone number which health officials could call with questions. The phone was manned by the Malawi Ministry of Health Directorate of Technical Services. We received only a couple calls, almost entirely from individuals who had found SENTINELs under suspicious circumstances.

Figure S4: SENTINEL Sticker during Auditing Deliveries

The devices were attached to medicine boxes in the Malawi Central Medical Stores (CMST) warehouse in Blantyre. The placement of devices was completed by a team of 14 people from a third-party logistics company (9), the warehouse (1), Malawi Ministry of Health (1), and the research team (3) during the course of normal packing operations.

Prior to the placement of the devices, all 14 people participated in a full-day training. The training explained how officials were to place the devices and introduced them to the cell phone application that they were to use in recording medicine and delivery details. With the exception of the research team and a manager from the Malawi Ministry of Health, none of the officials involved in the placement of devices were aware that the SENTINELs recorded location information or were intended to measure corruption. Officials were told (honestly) that the devices would record temperature and humidity and were part of a research program to improve medicine shipping and storage.

The deployment of the devices occurred during the normal medicine delivery cycles of October[1] and November of 2021. Using health facility orders provided to us by the Ministry of Health, and drawing on our pre-experiment survey and interview data, we sampled eight medicines that are both commonly ordered and commonly diverted (Table S1).

Table S1: Tracked Medicines

Medicine Description	**Anticipated Number of Tracked Boxes**
Amoxycillin,250mg,capsule,1000	89
Aspirin,300mg,tablet,1000	402
Ibuprofen,200mg,coated tablet,1000	786
Insulin zinc suspension (lente),100 IU/ml,10ml	3
Paracetamol,500mg,tablet,1000	519
Phenobarbitone,30mg,tablet,1000	166
Syringe Luer,10ml,disposable,hypoluer	235
Syringe,autodestruct,5ml,disposable	170

Officials were instructed to place SENTINELs on the lowest level of medicine packaging in an unobtrusive manner. In the case of pills, such as Amoxycillan the SENTINEL were taped to individual pill bottles. Where medicines were boxed (as in the case of syringes), SENTINELs were placed inside the box. An example of the device placement can be seen in Figure S5.

Figure S5: SENTINEL device placement

During the placement of the SENTINELs, officials were required to complete out a survey on their phone. This survey required the officials to read a QR code on the device and then provide the details on the medicine, including the batch number, destination and medicine type.

After the placement of SENTINELs, medicines were packaged into larger cardboard boxes and labeled for specific facilities. After all such boxes were sealed and labeled, the boxes were placed on trucks for delivery by a logistics company under contract with the Malawi Ministry of Health Central Medical Stores Trust. Pictures of a delivery truck and the CMST warehouse are shown in Figures S6 and S7.

For redundancy, officials were told to place two SENTRY devices at the rear of every delivery truck. These SENTRY devices were placed prior to the loading of any medicine boxes. Each delivery truck was restricted to a particular delivery route. Typically this encompassed a single district, though sometimes trucks visited multiple smaller districts.

3.3 Upstream Remote Tracking Audit Protocol

We accessed OnAsset API to download the data sent by the SENTRY devices. We have information when a SENTINEL was read by a SENTRY device, the SENTRY id, location, battery level, humidity, and luminosity. We conducted analysis at the SENTINEL level.

At the end of every day, we collected reports from the OnAsset portal with details on all the SENTINEL sightings identified by SENTRY devices for that day. There are two types of reports: (1) a sentry-sightings report, that shows the sightings for the SENTRY's last report, and (2) a sentry report, that contains a list of reports for each SENTRY device for every 15 minutes. For each report, we downloaded sightings that refer reports from SENTINELS. In other words, the reports contain information about every SENTINEL identified by the SENTRY devices with reporting time in the course of the trip. It also stores statistics by SENTINEL, such as battery, humidity, light, temperature, and rssi.

[1] These deliveries were delayed and were actually delivered in November.

Figure S6: Delivery Vehicle

3.4 Downstream Remote Tracking Audit Protocol

After the devices were sent and active in the ground, we ran follow-up visits to health facilities to find the exact location of the devices. Starting at the end of November 2021, our enumerators visited thirty nine health facilities in ten (10) districts namely Blantyre, Zomba, Thyolo, Phalombe, Mulanje, Machinga, Balaka, Mangochi, Neno, and Nsanje to check and verify if indeed the drugs that were ordered reached the facilities.

Upon arrival at the facility (but before entering the facility), enumerators informed time and identity of the facility. They placed charged and active SENTRY in a backpack and turned the SENTRY on. During the visits, they filled out a form with information about the district, the name of the facility, the location and the time. Then, they walked slowly around the health facility, completing a circle around the following areas, pausing for at least 10 minutes in each area:

1. All facility pharmacies
2. All areas where doctors or pharmacists are dispensing medicines
3. All warehousing facilities

Once the visit to the facility is complete, they immediately turned the sentry off, completed the remainder of the form, and sent a a message to the WhatsApp group with the time that enumeration at the clinic was completed. They were also requested to ask a receptionist or similarly employed official at the health facility about any locations where they might purchase medicines. Additionally they located all pharmacies and markets within 2 km of the facility and followed a similar protocol.

At the end, they were required to send a text to the whatsapp group with the time they left the area and submit the survey. Finally, we closely monitored the SENTRY reports. A research manager provided a report on the whatsapp group with details on all the SENTINEL sightings for that day. The report also mentioned if sighting information from ANY of the health facilities for that day are still missing.

3.5 Redelivery Remote Tracking Audit Protocol

To complement our analysis, we also obtained the full set of delivery notes from the District Health Offices (DHOs) for the months of October and November. These notes contain information on redeliveries, i.e. medications that were ship to a second facility. Health facilities can buy medicines directly from District Hospitals. In this case, the medicines go from Central Medical Stores to District Hospitals and, then, are sold to health facilities. The records detail the shipments from the DHOs to smaller health centres.

These 101 records were provided to us as scanned PDF documents. Figure S9 displays an example of a delivery note. To incorporate them in our analysis, we transcribed all the information from the reports to a spreadsheet and we merged with others datasets.

After collecting these delivery notes, we sampled 40 facilities that were identified as having received redeliveries. We then visited each of the 40 facilities and conducted a physical remote tracking audit using the protocol discussed in Section . This allows us to identify whether medicines identified has having been redelivered ended up in the correct facility.

Figure S7: CMST Warehouse

Figure S8: **Photo taken During Follow-up Visits**

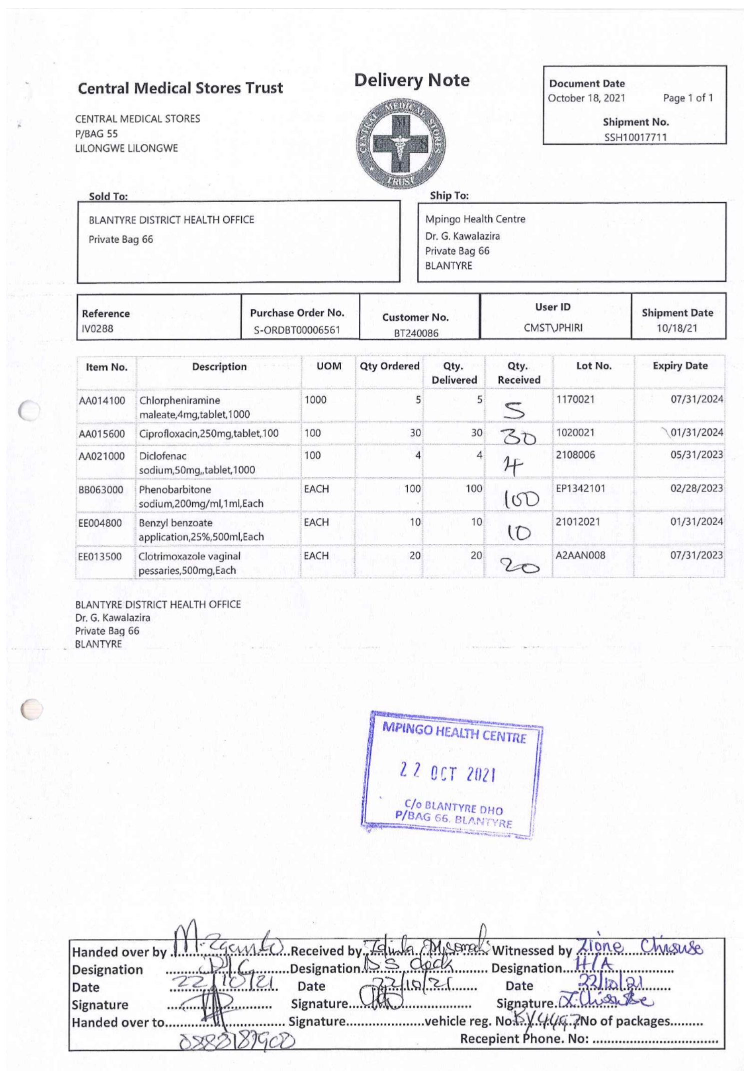

Central Medical Stores Trust

CENTRAL MEDICAL STORES
P/BAG 55
LILONGWE LILONGWE

Delivery Note

Document Date October 18, 2021	Page 1 of 1
Shipment No. SSH10017711	

Sold To:
BLANTYRE DISTRICT HEALTH OFFICE
Private Bag 66

Ship To:
Mpingo Health Centre
Dr. G. Kawalazira
Private Bag 66
BLANTYRE

Reference	Purchase Order No.	Customer No.	User ID	Shipment Date
IV0288	S-ORDBT00006561	BT240086	CMST\JPHIRI	10/18/21

Item No.	Description	UOM	Qty Ordered	Qty. Delivered	Qty. Received	Lot No.	Expiry Date
AA014100	Chlorpheniramine maleate,4mg,tablet,1000	1000	5	5	5	1170021	07/31/2024
AA015600	Ciprofloxacin,250mg,tablet,100	100	30	30	30	1020021	01/31/2024
AA021000	Diclofenac sodium,50mg,,tablet,1000	100	4	4	4	2108006	05/31/2023
BB063000	Phenobarbitone sodium,200mg/ml,1ml,Each	EACH	100	100	100	EP1342101	02/28/2023
EE004800	Benzyl benzoate application,25%,500ml,Each	EACH	10	10	10	21012021	01/31/2024
EE013500	Clotrimoxazole vaginal pessaries,500mg,Each	EACH	20	20	20	A2AAN008	07/31/2023

BLANTYRE DISTRICT HEALTH OFFICE
Dr. G. Kawalazira
Private Bag 66
BLANTYRE

MPINGO HEALTH CENTRE
22 OCT 2021
C/o BLANTYRE DHO
P/BAG 66. BLANTYRE

Handed over by .... Received by .... Witnessed by ....
Designation .... Designation DS Clerk .... Designation H/A ....
Date 22/10/21 .... Date 22/10/21 .... Date 22/10/21 ....
Signature .... Signature .... Signature ....
Handed over to .... Signature .... vehicle reg. No .... No of packages ....
Recepient Phone. No: ....

Figure S9: **Example of a Delivery Note**

4 Survey Activities

To better understand patterns of medicine theft, we conducted two complementary surveys: 1) an in-person survey of citizens and health centre advisory committees; 2) a phone-based survey of health facility officials.

4.1 Baseline Survey of Citizens and Health Centre Advisory Committees

In January-March 2019, we conducted an in-person survey with Malawian citizens in the catchment areas of 97 health facilities. The 97[2] health facilities were randomly sampled from a list of all of the health facilities in the Southern Region, excluding facilities that were affiliated with the Christian Health Association of Malawi (CHAM) and the Malawi Military, as these health facilities have additional, overlapping procurement processes.

At each sampled health facility, we recruited two distinct groups of participants for the survey. First, the enumeration team completed a random walk sampling protocol to recruit 35 individuals living within 10 km of the health facility. This resulted in a sample of 3,360 individuals. Second, the enumeration team used a purposive sampling procedure to recruit three individuals affiliated with the facility's Health Centre Advisory Committee. This resulted in a sample of 281 individuals.

The Health Centre Advisory Committee (HCAC) is a government-mandated oversight committee that is supposed to exist at each health facility in Malawi.[3] They are typically 10-member elected bodies that serve a social accountability function on behalf of the citizens in the facility's catchment area, though there is variation in their ability to hold health facility officials accountable, partially driven by variation in integration with other accountability structures [3]. They have four official duties: (1) bridging the communication gap between community and health staff, (2) inspection of facility conditions and drug stock, (3) formulating recommendations on facility equipment, and (4) complaint management [3].

The two survey instruments were similar, and both were conducted in person with trained Malawian enumerators and available in English and Chichewa. The citizens survey included 61 questions for the participant regarding their experiences at the facility, focusing on their observations and perceptions of theft. It also included a list experiment to measure the prevalence of public medication resale and 14 questions for the enumerator to collect information about the survey context. The HCAC survey included 20 questions for the participant regarding the operations and activities of the HCAC and 7 questions for the enumerator.

4.2 Endline Phone Survey of Health Facility Officials

In March and April 2022, we executed a phone-based survey of health facility officials. At each of the 97 facilities included in the baseline survey, the Malawi-based research managers used snowball convenience sampling to conduct a phone interview with two individuals: a pharmacist and one other official involved in medicine stocking and/or disbursement at the facility. Contact information could not be obtained for 12 facilities, so the final sample includes 172 officials from 85 facilities. The survey included 30 questions for the participant regarding their experiences at the facility, focusing on their observations and perceptions of theft, as well as four questions for the enumerator to collect information about the survey context. Surveys were conducted in English.

4.3 Descriptive Findings from Survey Data

One contribution of this project is to provide descriptive evidence regarding medications theft. In this sub-section, we present the findings of the baseline survey of citizens and Health Centre Advisory Committees and the endline phone survey of health facility officials.

First, we consider citizens' overall perceptions of theft at their local health facility. Figure S10 shows the results of a list experiment designed to anonymously elicit the probability that respondents have observed the sale of medicines in their community that should have been provided for free by government facilities. The results suggest that an estimated 50% of all citizens have observed the sale of medicines that should have been provided for free. Moreover, at approximately 80% of all health facilities, at least one citizen has observed the illegal sale of medicines.

One consequence of medicine theft is that citizens are unable to access medicines that they need. We asked citizens to estimate how often they are unable to access needed medicines from their health facility. On average, approximately 45% of needed medicines were unavailable at health facilities. We also asked this question of health facility officials and get a similar estimate (37%), though the correlation between the estimates from the citizens and officials is not statistically significant ($p = 0.21$). We show the distribution of perceptions of medication (un)availability in S11.

Of course, there are a number of reasons why medicines may not be available, and Malawi has suffered from chronically insufficient public health financing. Nonetheless, it is clear that many citizens attribute medicine unavailability (often called "stockouts") to theft. In Figure S12 we show the geographic distribution of perceived medicine unavailability (stockout) rates and perceived theft rates. It is clear that areas with higher stockout rates are also the areas where citizens perceive the highest rates of theft.

We also asked citizens a number of questions about the reasons that officials might engage in theft. Citizens perceive information and capacity gaps to be a main constraint. Fifty-three percent of citizens said that they lack effective ways to report the diversion of medications and cited this as a primary reason officials choose to steal medications in

[2] We originally sampled 100 facilities but three facilities were inoperable at the time of data collection.

[3] In reality, not all facilities mobilize HCACs, let alone utilize them.

Table S2: Summary Statistics of Respondents

Characteristic	Count	Percentage	
Gender			
Female	3188	95%	
Male	168	5%	
Marital Status			
Single	113	3.37	
Married	2501	74.52	
Married with multiple wives	145	4.32	
Separated	65	1.94	
Divorced	321	9.56	
Widowed	211	6.29	
Age Group			
0-20	189	5.63	
21-30	1196	35.64	
31-40	1103	32.87	
41-50	522	15.55	
51-60	204	6.08	
61-70	117	3.49	
71-80	24	0.72	
Older than 80	1	0.03	
Income			
Under 50,000 kwacha/month	2506	74.67	
50,000-100,000 kwacha/month	516	15.38	
100,000-200,000 kwacha/month	283	8.43	
200,000-400,000 kwacha/month	37	1.10	
400,000-1,000,000 kwacha/month	5	0.15	
Don't know	9	0.27	
	Mean	SD	Range
Family Size	4.91	1.85	1 - 24
Number of Kids	2.57	1.84	0 - 36

Malawi. Another reason cited was that patients simply have no way of knowing whether medications in their health facility were stolen.

In addition, it seems that formal accountability institutions are weak. For instance, Health Center Advisory Committees have been ineffective in overseeing deliveries of medicines. Figure S13 shows that nearly three-quarters (74%) of interviewed Health Center Advisory Committee (HAC) report witnessing a delivery at the health center in the past three months. However, very few witnessed three or more deliveries in three months, despite a mandate to witness each monthly delivery. Furthermore, approximately one-quarter (26%) of HAC respondents report witnessing no deliveries in the past three months.

Despite recognizing the weaknesses of accountability mechanisms, more than half of citizens believe that officials that steal medicines are highly likely to be caught and highly likely to face consequences (Figure S14a and Figure S14b). Among the officials, however, we observe a different reality. The median perceived probability of getting caught and facing consequences according to health facility officials is 51-75%. Only 20% of health facility officials perceive a 75-100% chance of getting caught, and 30% perceive a 75-100% chance of facing consequences (Figure S14c and Figure S14d).

Figure S10: Probability of Theft Being Observed

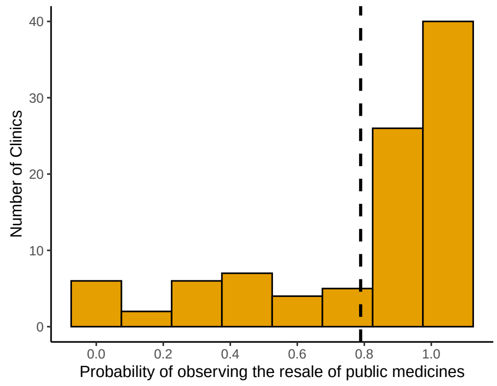

Note: This figure shows estimates of the facility-specific probability that the illegal sale of medicines has been observed. The figure is derived from a list experiment asking whether respondents have observed the sale of medicines that should have been provided for free.

Figure S11: Perceived Rates of Medication Unavailability by Citizens and Officials

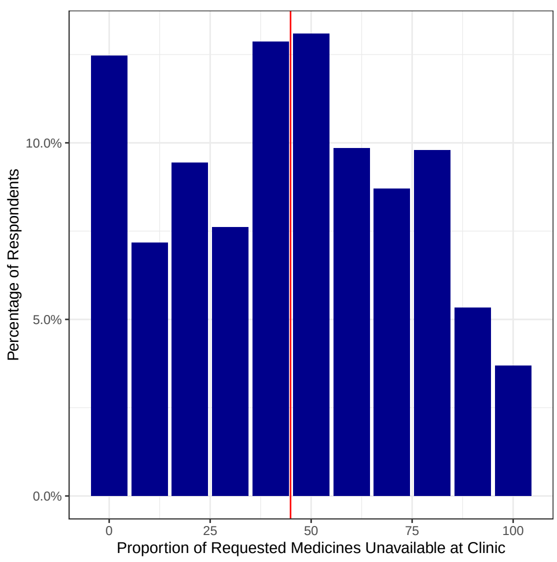

(a) Citizens' Perceptions of Medication Unavailability

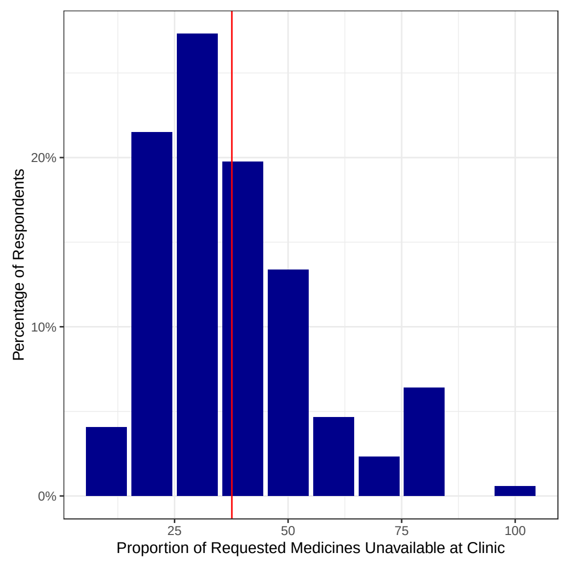

(b) Officials' Perceptions of Medication Unavailability

Figure S12: Perceived Rates of Theft and Stockouts

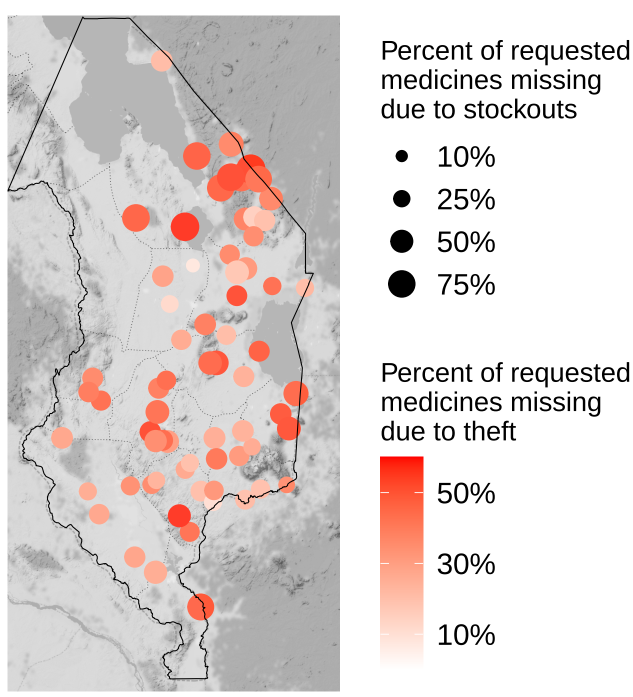

Note: The Perceived Stockout Rate indicates the percentage of times on average patients were unable to obtain needed medicines from their health facility in the last three months. The Perceived Theft Rate indicates the share of those unavailable medicines that the patient believed were stolen by public officials.

Figure S13: Health Centre Advisory Committees (HACs) Attendance at Deliveries

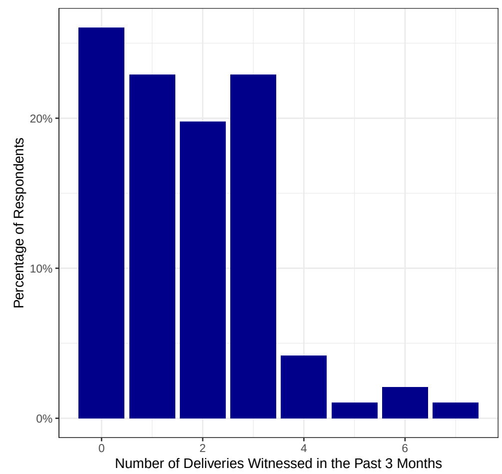

Figure S14: Perceived Rates of Theft Detection and Punishment by Citizens and Officials

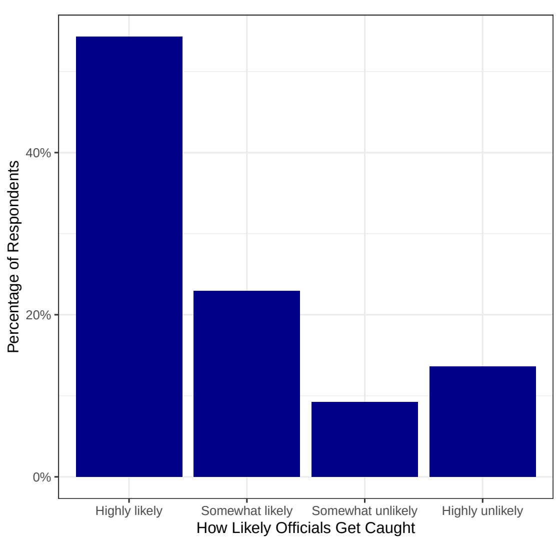

(a) Citizens' Perceptions of Theft Detection

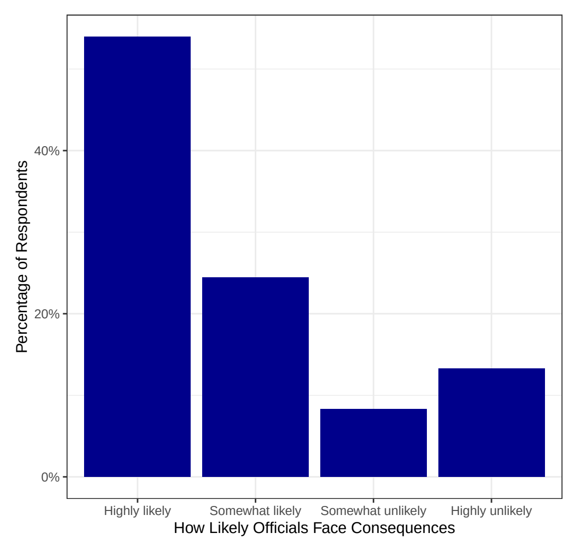

(b) Citizens' Perceptions of Theft Punishment

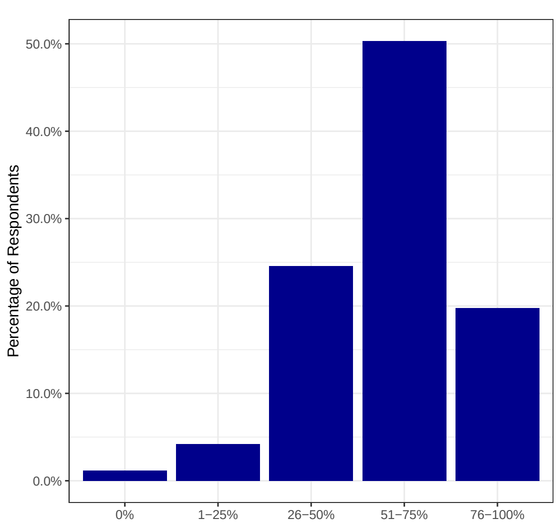

(c) Officials' Perceptions of Theft Detection

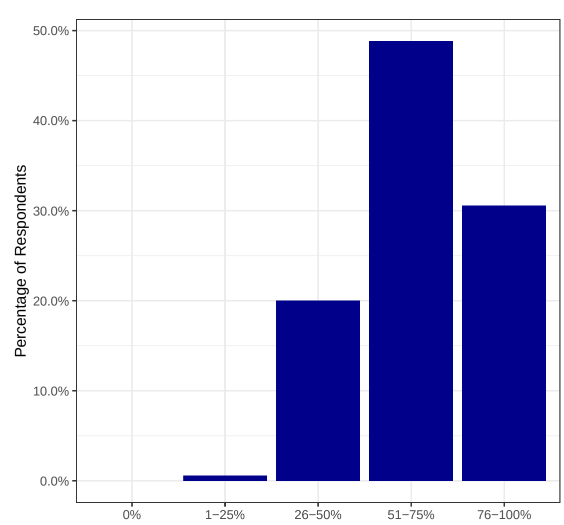

(d) Officials' Perceptions of Theft Detection

5 Administrative Data

Our partners in the Ministry of Health and the Drug Theft Investigations Unit (DTIU) also provided administrative data on medicine procurement and traditional audits:

1. Facility-Level Shipment Records: We obtained the full set of shipment manifests that were used to prepare the October and November deliveries, and used these records to construct variables regarding the volume and density of medicines during the shipment periods.

2. Delivery Notes: We also obtained the full set of delivery notes from the district health offices (DHOs) for the months of October and November. These records detail the shipments from the DHOs to smaller health centres (i.e., "redeliveries").

3. DTIU Audit Records: DTIU provided us with: a) traditional audit reports at the health facility level for 2016-2018; b) one-off reports of theft collected through the tipoff hotline for 2017-2019. See Section **??** for a discussion of how these audit findings correlate with the remote tracking audit findings.

5.1 Traditional vs. Remote Tracking Audits

As discussed above, we received tipline, investigation reports, and audit data from DTIU in 2018 and from MoH in 2023. We are uncertain whether we received all of the tipline, investigation reports, and audit data that exist, but we likely did not. (For example, the tips are numbered sequentially but we are missing some of the sequence of numbers.) This means we cannot really interpret a facility not being reported via a tip or not being audited as the absence of theft at that facility.

We subsequently coded these data to generate pre-study indicators of theft to compare to the remote tracking audit data at the health facility level. Tipline data were only included if they named an existing health facility. Many tips did not (and instead referred to a market or private shop). We have data on 27 tips for 24 named facilities in the Southern Region, all coming into the tip line between 2017 or 2018. Whether or not a tip ever came in about a health facility is not correlated with our facility-level measures of downstream theft ($r = 0.112$, $p = 0.2547$).

Similarly, audit data were only included if it was from an existing health facility. We have data on 45 audits for 45 facilities in the Southern Region, all likely occurring in 2016. In addition to concerns about missing audit data as described above, the scope of each audit is unknown. At each audited facility, at least one (and sometimes more) medicines are reported as missing. It is not clear if medicines that are *not* listed are not missing or were simply not audited. It is also not clear what the stock of the medicine was supposed to be (i.e., what percentage of the medicine stock is missing). The most commonly audited medicine was chloroquine phosphate (aka "LA"). There were no audit reports that did not include a line to report missing LA. Therefore, we base our analysis of the audit data on the quantity of LA missing and the value (in Malawi Kwacha) of this missing LA. (Note that this amount was not always provided.) These indicators positively correlate with our facility-level measures of downstream diversion, though the relationship is not statistically significant. Once one particularly high-diversion outlier is removed, however, the p-value on the correlation coefficient is under 0.2 for both variables (Figure S15a and Figure S15b).

Figure S15: Comparing Remote Tracking Audits Downstream Diversion Rates and Traditional Audits Missing Chloroquine Phosphate (LA) Rates

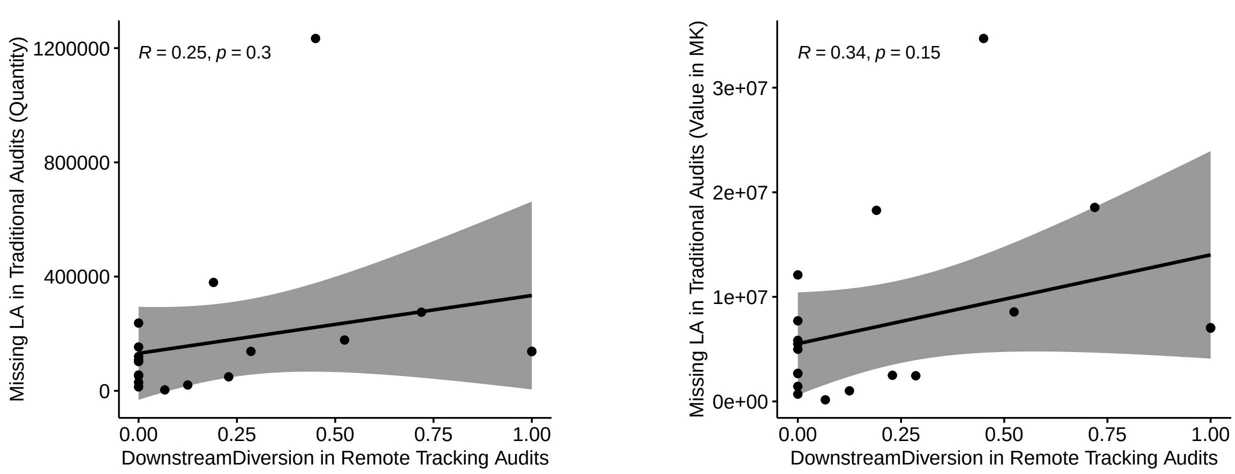

(a) DownstreamDiversion and Missing LA Quantity

(b) DownstreamDiversion and Missing LA Value

5.2 Remote Tracking Audit Costs

Table S3 includes all of the costs associated with running our remote tracking audits, estimated based on Malawi administrative data and project records. We note that the items "Equipment-Trackers" and "Printing Costs - Trackers Stickers" are startup costs: now that they have been paid, these costs would not need to be incurred again.[4] In total,

[4] This does not take into account some reasonable loss rate for the tracking devices. In our remote tracking activities, the SENTRY loss rate was zero, but the SENTINEL loss rate was much higher.

the total estimated costs of a remote tracking audit are US$119,700, or US$19,531 without the costs of the tracking devices and their stickers included.

Table S3: Remote Tracking Audit Costs

Cost Description	Unit	Unit Cost	Number Units	Total Amount (in MK)	Total Amount (in USD)	Notes
Communications	Total Cost	MK1800000	1	1800000	2222	Airtime for minutes and data used to coordinate the audits
Equipment - Trackers	Total Cost	US$98274	1	N/A	98273	80 parent tracking devices and 2400 child tracking devices
Printing Costs - Trackers Stickers	Total Cost	MK1536000	1	1536000	1896	Stickers for the tracking devices
DTIU Director salary	Monthly Cost	MK610000	0.75	457500	565	3 months of quarter-time work to coordinate the audits
DTIU/MOH officer salary	Monthly Cost	MK490000	8	3920000	4840	2 months of 4 people working full-time to execute the audits
CMST manager salary	Monthly Cost	MK812000	0.75	609000	752	3 months of quarter-time work to coordinate the upstream audits
Vehicle maintenance costs	Monthly Cost	MK300000	4	1200000	1482	Maintenance for 2 months of 2 vehicles
Fuel costs	Total Cost	MK6500000	1	6500000	8025	
Training costs	Total Cost	MK1070000	1	1070000	1321	Training DTIU/MOH officers on audit execution
DTIU office space	Monthly Cost	MK350000	0.75	262500	324	Overhead for DTIU office space for audit coordination and execution

6 Outcome Variables

Below we describe the coding of each of our outcome variables. Note that, primarily for reasons of clarity, we deviate slightly from out pre-analysis plan in the coding of these variables (see Section **??**).

6.1 Identifying Incorrect Deliveries from Upstream Audit (*Upstream Diversion*)

We define a successful delivery as one in which the medicine reached the facility for which it was intended during the delivery period. Specifically, we create a variable *UpstreamDiversion* which equals one if a SENTINEL is not delivered by the designated delivery truck to its intended destination. It equals zero if it was delivered elsewhere and is coded as missing if the delivery could not be confirmed.

More specifically, as discussed in Section 3.2, in the SENTINEL placement process, each SENTINEL was coded to a particular medicine and intended health facility destination while packing the medicines for delivery. We determine whether a medicine reaches this intended destination using the following procedure:

1. Using the OnAsset API, we determine the geographic coordinates of the final SENTRY sighting of the SENTINEL associated with the medicine. Since each SENTRY was associated with a delivery truck, this final

SENTRY sighting is the location at which a medicine left the delivery truck.

2. We then measure the distance between the final sighting determined in step 1 and the intended destination of the medicine. If this distance exceeds an uncertainty range (as defined below), we code this medicine as not delivered to its intended destination (*UpstreamDiversion* $= 1$). If this distance falls within the uncertainty range, we code this medicine as delivered to its intended destination (*UpstreamDiversion* $= 0$).

3. In cases where we lack sufficient data to identify the final location (e.g., due to a lack of cellular connectivity), we also rely on the downstream audit data. We code a medicine as having been delivered if we find during the downstream audits that one or more medicines intended for the same facility and delivered at the same time were found in the correct facility during the in-person audits. That is, if we can confirm that other medicines delivered at the same time and place were correctly delivered, we assume that all medicines delivered at the same time and place were correctly delivered (*UpstreamDiversion* $= 0$).[5]

One challenge with coding deliveries is that the coordinates of the final SENTINEL sighting cannot always be precisely determined. At any given sighting, each SENTRY provides coordinates based on GPS or cellular triangulation. Since satellite and cellular coverage in remote areas of Malawi is often incomplete, this means that we are sometimes unable to determine the location of the delivery; or we are only able to identify the delivery imprecisely (e.g., because of weak signal strength or triangulation off of few satellites or cellular towers). If we cannot obtain location coordinates, we first attempt to use the most recent valid coordinates as a way to identify the location of the delivery. In the case of 9.7% of tracked medicines, we cannot find a valid coordinates within 90 minutes of the delivery. We code such deliveries as indeterminate.

A related challenge is determining whether a medicine was delivered to the intended destination or another location. Because location coordinates are sometimes imprecise, we cannot always be confident that a delivery occurred in the location of the sighting. To address this challenge, we estimate the uncertainty associated with each SENTRY sighting. To derive context-specific estimates of uncertainty, we first created a training dataset of all SENTRY sightings where the true location was precisely known.

Using these data, we create a variable D_{ij} which measures the logged distance between the known true location (i) and the triangulated location (j) provided by a SENTRY (i.e., the precision of the location estimate). Our goal is to estimate the distribution $\hat{D_{ij}} = E(D_{ij}|X = x)$ where X are the characteristics of location i which we expect to matter for the location precision, D_{ij}. To estimate this equation we rely on a quantile regression forrest estimator [5], which provides a non-parametric way to estimate the expected value of different quantiles of $\hat{D_{ij}}$. Our vector of predictors, X include the density cellular towers in range of the true location, the signal strength of the SENTRY connection, the type of signal (GPS or cellular), the altitude of the true location, and the altitude of cellular towers.

Using the weights estimated from the quantile regression, we can estimate the likelihood that a triangulated location corresponds to a correct delivery (i.e., no upstream diversion). Formally, let a be the location of the intended delivery. Let b be the triangulated location provided by the last SENTRY sighting of a medicine's SENTINEL. We assume that the expected distribution of the distance between a and b can be well approximated by the prediction $\hat{D_{ab}} = E(D_{ab}|X = x)$ provided by our bagged sample weights. We code a correct delivery (*UpstreamDiversion* $= 0$) as one where $\hat{D_{ab}(95)} - D_{ab} >= 0$. Here $\hat{D_{ab}(95)}$ is the value within which our model predicts 95% of the values of D_{ab} to fall under the assumption that the delivery occurred at a. We code an incorrect delivery (*Upstream Diversion* $= 1$) as one where $\hat{D_{ab}(95)} - D_{ab} < 0$.

This estimated level of uncertainty appears to correctly characterize the true precision of triangulated locations. In out of sample tests, the mean of our prediction correctly explains 42% of the variance in D_{xy}. Further, D_{xy} is larger than $\hat{D_{xy}(95)}$ in less than 1.3% of cases. This suggests that our Type II error (where we code a delivery as not delivered to its intended location when it actually was) will be less than 2%.

In Figure S16 we plot the distance between intended and triangulated locations for deliveries that are coded as correct and incorrect. The number of cases where a Type II error is possible is exceedingly low: on average, cases where *Upstream Diversion* $= 1$ are delivered 37 kilometers away from their intended destination.

We are also concerned about the possibility of Type I error (where we code a delivery as delivered to its intended location when it was actually delivered elsewhere). This might be a particular concern if diversion is happening very close to the intended health facility and within the level of precision of the triangulated location. While this is possible, we don't see evidence that this represents a large proportion of the diversion in our study. In Table S4 we show how our estimates of *UpstreamDiversion* vary under different levels of estimated precision, $\hat{D_{ab}}$. If diversion was occurring nearby to intended delivery locations, we would expect that we would find higher rates of diversion where we can more precisely identify diversion. However, we see no evidence of such a pattern.

A final concern is that the rate of diversion might differ in cases of deliveries that are indeterminate (e.g., if officials steal at more remote locations at a higher rate). One way to evaluate whether indeterminate deliveries are more likely to be diverted is to compare the proportion of determined and indeterminate deliveries that were found in the downstream remote tracking audits. Of the 682 deliveries which could not be determined from coordinates alone, 531 of them were supposed to have been delivered to locations where we conducted a downstream remote tracking audit. Out of these 531 medicines, 361 (68%) were actually found at the correct location during the downstream audit. In comparison, of those medicines with determinate delivery locations, 65% were found in the downstream audits. This suggests that indeterminate deliveries are not substantially more likely to be diverted.

[5] Note this does not preclude some medicines from a delivery being misdelivered (this happens regularly). The assumption here is instead that if we can locate one or more correctly delivered medicines, we assume that other medicines in that same delivery, and *that share the exact same delivery time stamps*, were also correctly delivered.

Figure S16: Distances between intended and triangulated locations (D_{ab}) for medicines coded as delivered and not delivered

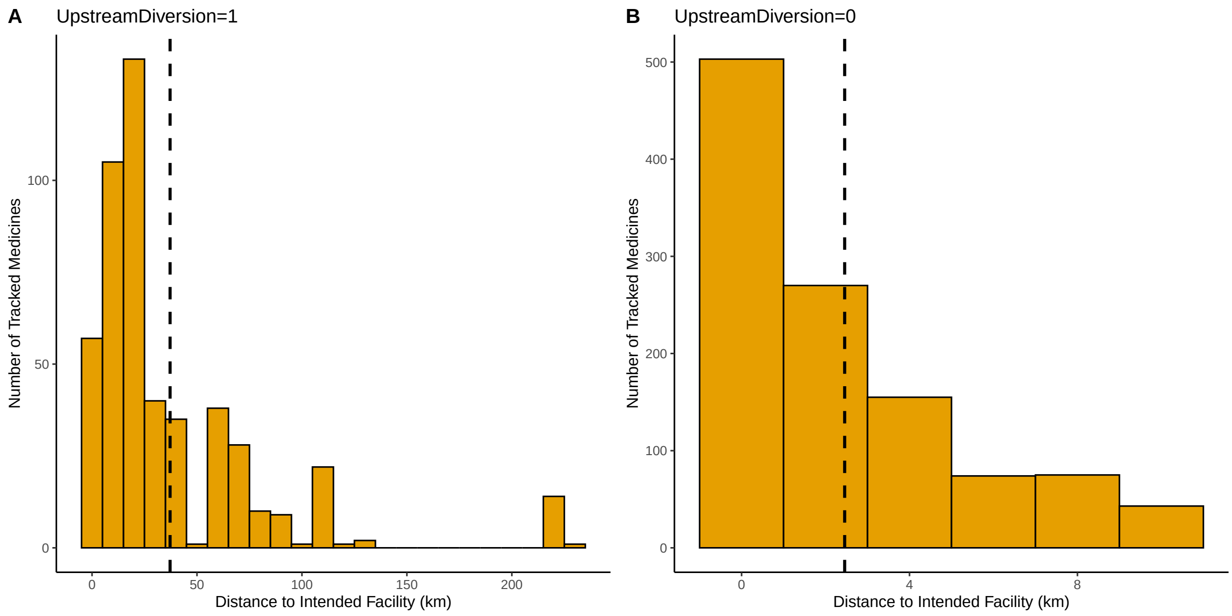

Note: This figure shows the difference between intended and triangulated locations (D_{ab}) of final medicine deliveries. Panel A shows the range of D_{ab} where deliveries are coded as incorrect (*Upstream Diversion*=1). Panel A shows the range of D_{ab} where deliveries are coded as correct (*Upstream Diversion*=0).

Table S4: *Upstream Diversion* for different levels of location precision, $\hat{D_{ab}}$.

Mean *UpstreamDiversion*	$\hat{D_{ab}}$	N Medicines
0.01	1 km	416
0.07	1.5 km	830
0.09	2 km	1311
0.10	2.5 km	1530
0.10	3 km	1760
0.11	3.5 km	1829
0.11	4 km	1858
0.12	4.5 km	1932
0.12	5 km	1935

Note: This table shows the mean delivery diversion rate (*UpstreamDiversion*) for samples of tracked medicines that are more or less precisely coded to a final delivery location.

6.2 Assessing whether a medicine is delivered to a valid health facility

For medicines that are not correctly delivered, we also attempt to identify whether the actual delivery location could have been a valid public health facility (see Figure 2 in the main text). To do this, we adopt the same procedure from Section 6.1 above used to validate deliveries for correct locations.

Specifically, let A be a vector of all health facility locations in Southern Malawi. Let b be the triangulated location provided by the last SENTRY sighting of a medicine's SENTINEL. For each $a \in A$ we calculate $\hat{D_{ab}(95)}$ using the quantile regression procedure described above. If $\hat{D_{ab}(95)} - D_{ab} >= 0$ is true for *any* $a \in A$, we assume that the delivery most likely went to a valid health facility. If $\hat{D_{ab}(95)} - D_{ab} < 0$ is false for *all* $a \in A$, we assume that the delivery most likely went somewhere other than a valid health facility. Consistent with our coding rules above, we also, code a medicine as going to a valid health facility if the medicine was found at a valid health facility during the downstream audit.

While our conclusions are those most consistent with the data, there is a risk of misclassification in this estimator. An important assumption that we make is that if it is consistent with our data that a delivery *could* have happened at a valid public health facility, then it actually happened at a public health facility. We cannot rule out the possibility, for instance, that deliveries occurred proximate to a public health facility, but not at the public health facility. Likewise it is possible that deliveries far from valid public health facilities, nonetheless happened at a public health facility. On average, deliveries classified as *not* at a public health facility were 5.2km from a facility. Deliveries classified as being at a public health facility were 2.6km away. Since we expect a small amount of error in our estimator, it is possible, but unlikely that some of these deliveries actually happened at a public health facility.

Since downstream audits rarely overlapped with missed deliveries, it is difficult to validate this estimator directly. One indication of the potential for misclassification is the extent to which our results vary under alternative confidence intervals for $\hat{D_{ab}}$. If we use a 90% confidence interval rather than a 95% confidence interval, our estimate of non-valid deliveries increases from 5% to 11%. If we use an 80% or 70% confidence interval, our estimate increase to 31% and 41%. Thus, even under these more relaxed assumptions, supply chain error — rather than theft — seem to explain the majority of missed deliveries.

In Section 6.3.3 below, we discuss some potential reasons for this high rate of incorrect deliveries.

6.3 Identifying Diversion from Downstream Audit

We are interested in whether medicines that arrive at the correct facility remain at that facility in the weeks following the delivery. As discussed in Section 3.5, audit teams conducted in-person tracking audits at a random sample 144 health facilities that were supposed to have received deliveries or redeliveries of tracked medicines.

6.3.1 *Downstream Diversion*

Using data from these audits, we create a variable called ***Downstream Diversion***. This variable equals one if a medicine was delivered to its intended destination (*Upstream Diversion = 1*) but was then not found in the in-person audits at the correct location, or at a redelivery location. It equals zero if the medicine was found at the correct location, or at a redelivery location. It is missing in cases where a medicine was not confirmed as correctly delivered, or when an intended facility was not audited.

6.3.2 *Private Theft*

In addition to identifying whether a medicine was diverted, we attempt to determine why a medicine went missing. In most cases, the reason why a medicine a medicine went missing cannot be determined since we fail to track their SENTINEL in any of our auditing activities. However there are specific cases where we can identify what happened.

One such case is when medicines are re-sold to private markets and pharmacies within the local community. As discussed in Section 3.5, we conducted remote tracking audits in the nearest private pharmacy and market to each of the 104 health facility included in the Downstream Remote Tracking Audits. Since not all health facilities have nearby markets or pharmacies, we audited a total of 143 private markets and pharmacies. It is important to note that we can only identify cases of theft that occur within these sampled facilities. It is possible (and likely) that private theft occurs through other channels.

Using these audits, we code *Private Theft* equal to one if medicine which is missing at the correct health facility was found at one of the audited private facilities. We code *Private Theft* as zero if the medicine is missing for another reason. Note this means that *Private Theft* is missing where *Upstream Diversion* or *Downstream Diversion* equals zero.

Private theft can occur as a consequence of either upstream or downstream diversion, depending on the likely perpetrator of such theft. We attempt to distinguish between these two channels in the main text. We conclude that most instances of private theft occur as a result of downstream diversion.

6.3.3 *Supply Chain Error*

A second reason why medicines might be diverted is that they went to the wrong facility due to error. There are multiple reasons why this might be the case. In some cases it seems likely that delivery teams made an error in the delivery process and provided medicines intended for one facility to another. This can happen, for instance, due to the mislabelling of medicine boxes during the loading process or errors identifying boxes during the unloading process. It is also possible that some of this error is due to intentional or intentional substitution of similar medicines by delivery teams. Due to the lack of cross-checking of paperwork, errors of this sort are easy to overlook.

There are other potential sources of error. As we discuss in Section , medicines are also sometimes redelivered from one facility to another. Such redeliveries are supposed to be documented in redelivery notes held by District Health Offices (DHOs). However compliance with these regulations are sometimes spotty and delivery notes are frequently incomplete. When this is the case, a medicine can end up a facility which made no official order was made.

It is also possible that medicines end up at the wrong facility as part of coordinated theft. For instance, corrupt officials may seek to divert medicines to those facilities that are particularly prone to theft. It is difficult two rule this out. However we are aware of no accounts of such coordination between delivery and clinic officials.

To code supply chain errors, we create a variable *Supply Chain Error* which equals one if a missing medicine is found at a facility which was not the original ordering facility and was not a facility for which we have a documented redelivery request. *Supply Chain Error* equals zero if a medicine is missing for another reason. This means that *Supply Chain Error is missing in cases where Downstream Diversion* equals zero.

6.4 Identifying Diversion during Redeliveries

As discussed in Section 6.3.3, medicines that are delivered to district health offices and hospitals (DHOs) are sometimes redelivered to other health facilities. To audit these deliveries, an audit team collected the delivery records (delivery notes) from all district health offices in the Southern Region. From these delivery notes, we identified 150 health facilities that had ordered tracked medicines from 11 DHOs and hospitals in October and November. We visited 97 of these facilities in the downstream audits. Our sample included all redelivery facilities in the districts of Balaka, Blantyre, Mangochi and Neno.

We are interested in what happens to these medicines eligible to be redelivered. To measure this, we first identify the sample of tracked medicines that were correctly delivered to the DHO, but were later found to be missing in the Downstream Remote Tracking Audit. This provides us with a sample of 81 medicines.

To identify which of these medicines were likely redelivered we attempt to match each of these 81 medicines with an entry on the delivery notes. Unfortunately there are no unique medicine serial numbers on medicines or delivery notes, so there is no way to identify exactly which medicine was redelivered. However by matching on medicine type, batch number and DHO name between our list of tracked medicines and transcribed delivery notes we are able to identify 66 tracked medicines which were eligible to be redelivered. While we cannot confirm that all these medicines were intended to be redelivered, this is the full population of tracked medicines which could plausibly have been redelivered according to this official documentation.

We summarize what happens to these 66 medicines eligible to be redelivered in Figure S17. Only 7 (11%) of these medicines actually reach an official ordering facility. 1 was found at facilities which were not on the delivery notes. The remainder were not found. As we show, in S18, if we focus on the four districts where we visited the population of redelivery-eligible facilities, we observe a slightly higher rate of redelivered medicines (17%).

In sum we think these data suggest a couple things about this delivery channel: First, the discrepancy between intended and actual locations of redelivery eligible medicines imply large errors in official records. One explanation that this discrepancy is due to falsification of entries in delivery notes. Such falsification could be an attempt to hide corruption. However this falsification may also be an attempt to hide a general lack of compliance with official record-keeping regulations. One plausible scenario is that delivery notes records were completed *after* rather than *before* the audit visit was announced. Such post-hoc compliance is consistent with the high proportion of these medicines that are in fact redelivered, but to locations that are not noted in official documentation. This suggests that official delivery notes may have been compiled after the audit request based on incomplete sources of data.

However there does seem to be evidence that the redelivery channel is a particular target for corruption. Compared to direct delivery channels, a higher proportion of medicines in this channel goes missing or ends up in private facilities. This is consistent with more anecdotal accounts of theft during redeliveries.

Some caution is warranted in the interpretation of theses data since the sample is small. Also, since we lack unique identifiers, we cannot confirm that all medicines in this analysis were intended for redelivery.

Figure S17: Location of medicines eligible for redelivery

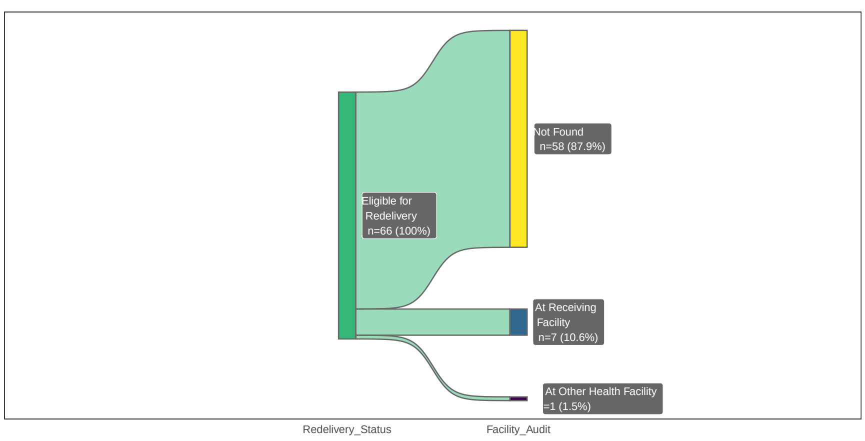

Figure S18: Location of medicines eligible for redelivery in Balaka, Blantyre, Mangochi and Neno

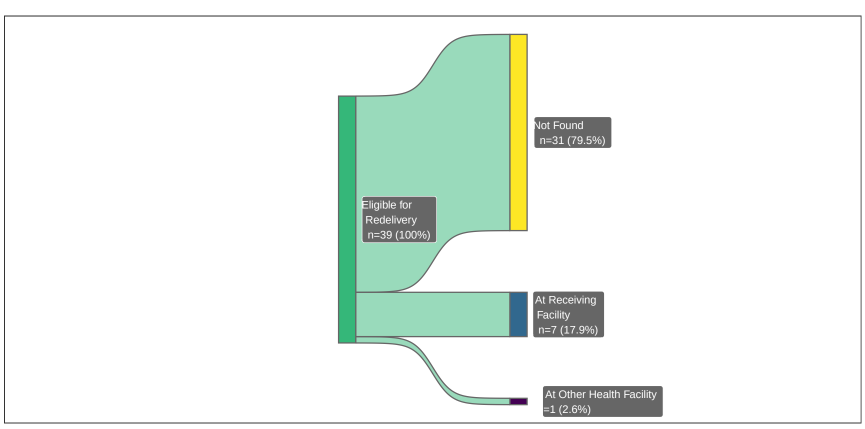

6.5 Index of Theft Perceptions

For the construction of the one-factor confirmatory factor analysis (CFA) index presented, a set of relevant variables was carefully selected from our Baseline Citizen Survey. These variables were chosen based on their theoretical relevance to measure perceptions of medicine theft. The variables are the following:

1. Did the clinic provide you or your family with these medicines?

2. In the past 3 months, has the clinic ever been unable to provide medicines that your family needs?

3. In general, how effective do you think clinic is at providing medicines that this community needs?
4. Agree or Disagree with the following sentence: The theft of medications keeps people from getting high-quality health care in my community.
5. Number of visits when you needed a medication and it was not available for you at the clinic pharmacy, and you believe it was because someone stole the medicine.
6. In the past 3 months, have you or someone in your household ever gone without medications because they were not available from clinic?
7. In the past 3 months, approximately how much in kwacha have you or someone in your household spent on purchasing medications that were not available at clinic?

Prior to analysis, variables were transformed into binary variables, with the exception of the last variable, which is an ordered variable that varies from 1-6. Each number refers intervals of money spent on medicines. These preprocessing steps were applied to maintain the compatibility of the variables within the CFA framework.

The resulting factor loadings from the one-factor CFA analysis were used to construct the index. Each variable's loading on the latent factor was treated as its weight in the index. The index score for each participant was calculated by summing the products of their individual variable scores and the corresponding factor loadings. This procedure yielded a single numerical value for each participant, representing their position on the latent construct. Then, we calculated clinic-level averages of the index, which were used in the analysis.

7 Summary Data

7.1 Summary of Outcome Variables

In Figures S19, S20, S21, S22 below we plot the distribution of our main outcome variables.

Figure S19: Distribution of *Downstream Diversion*

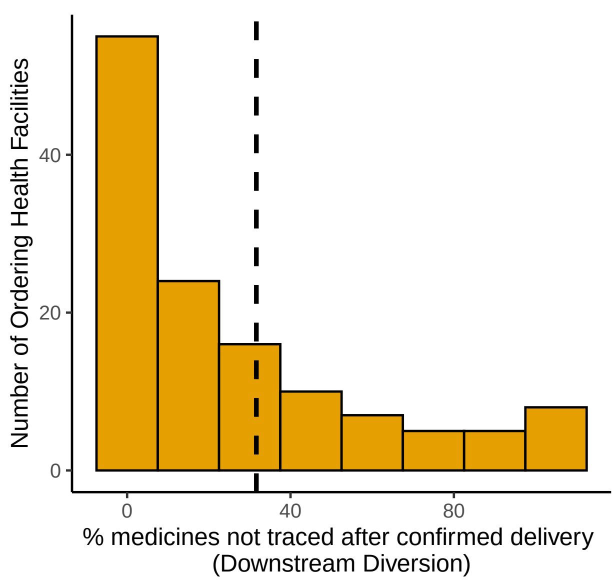

Figure S20: Distribution of *Upstream Diversion*

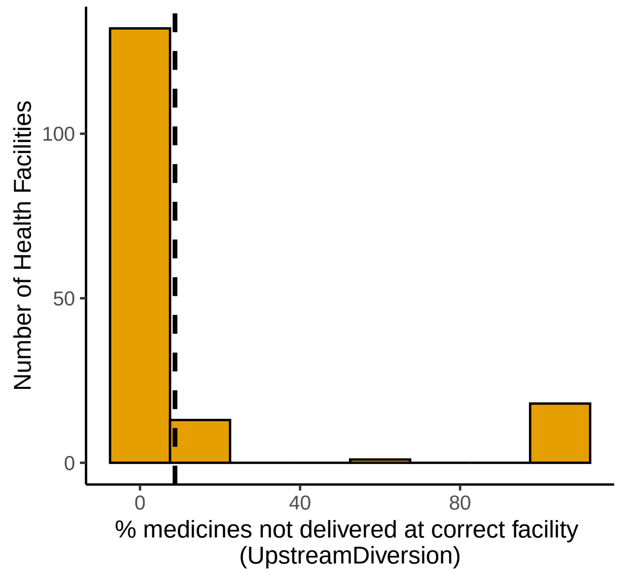

Figure S21: Distribution of *Private Theft*

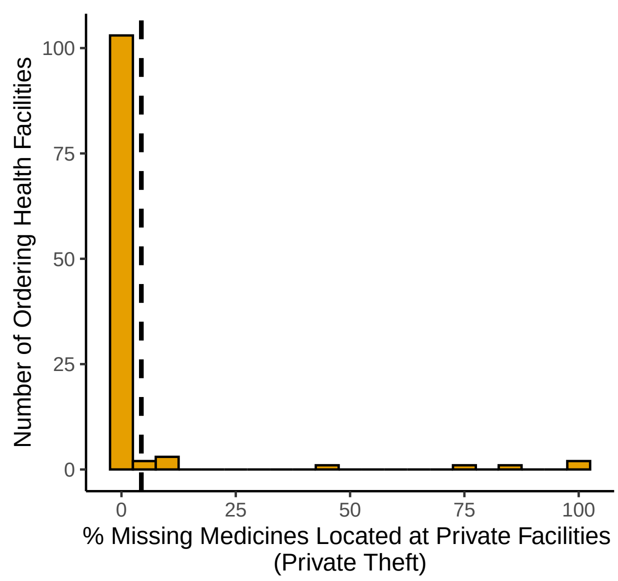

Figure S22: Distribution of *Supply Chain Error*

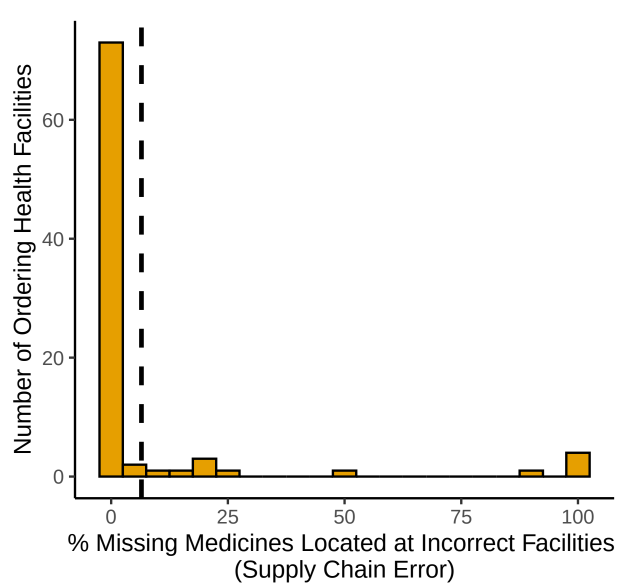

7.2 Diversion by Medicine Type

In Figure S23 we show the percentage of medicines diverted by medicine type. Amoxicillin, a common antibiotic, is the most commonly diverted medicine by far (over 88% diverted), followed by phenobarbitone and ibuprofen.

Figure S23: Rate of diversion by type of tracked medicine

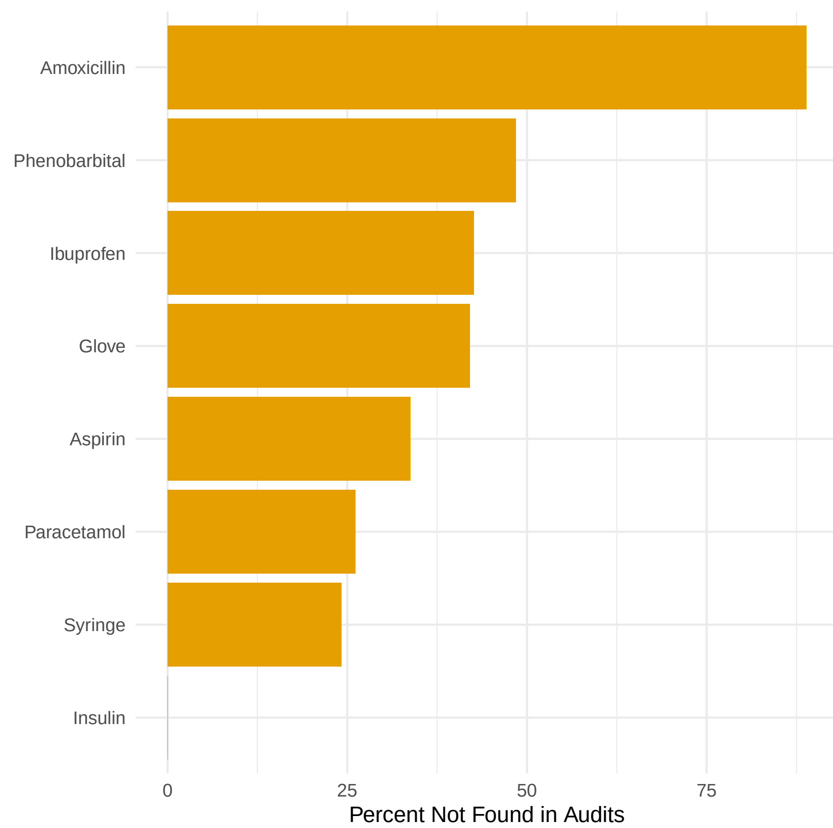

7.3 Diversion by District

In Figure S24 we show the percentage of medicines diverted by district. The most medicines were diverted in Zomba, followed by Balaka and Neno.

Figure S24: Rate of diversion by district

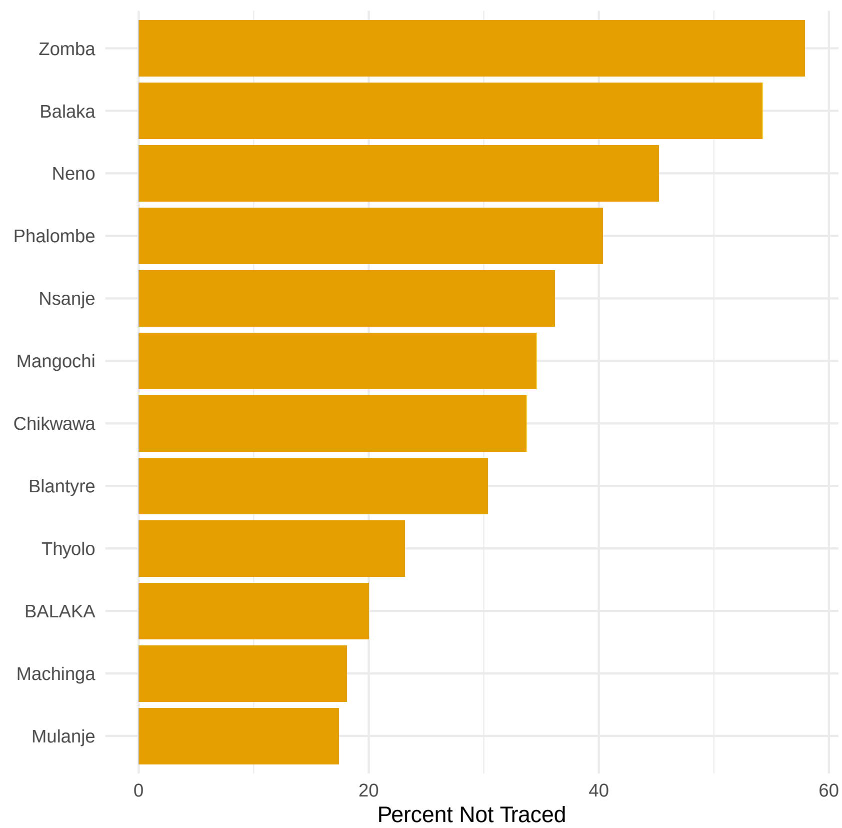

8 Additional Regression Analysis

8.1 Effects of Medicine Price on Diversion

We consider alternative ways to estimate the effects of medicine pricing on diversion. In Tables S5 and S6 we re-estimate the effect of medicine price on whether a sentinel was diverted using alternative specifications, including no controls (column 1), facility fixed effects (column 2) and district fixed effects (column 4). The results are consistent with those in the main text. Perhaps most interestingly, we see consistent results including facility fixed effects, suggesting that officials discriminate by price even within individual shipments.

Table S5: Effect of Medicine Value on Upstream Diversion

	Dependent variable:			
		Facility FE	with controls	controls + District FE
	(1)	(2)	(3)	(4)
Medicine Value (log)	0.007	−0.002	0.009	0.001
	(0.007)	(0.003)	(0.007)	(0.003)
Observations	1,911	1,911	1,608	1,313
R^2	0.001	0.828	0.034	0.118

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows the effect of medicine value (log wholesale price of the tracked medicine) on whether a sentinel went missing during upstream remote tracking. From a linear regression with facility clustered errors. Controls include facility type, log delivery size, and the log number of tracked units at a facility.

Table S6: Effect of Medicine Value on Downstream Diversion

	Dependent variable:			
		Facility FE	with controls	controls + District FE
	(1)	(2)	(3)	(4)
Medicine Value (log)	0.032*	0.032*	0.047***	0.054***
	(0.017)	(0.017)	(0.016)	(0.017)
Observations	1,634	1,634	1,409	1,312
R^2	0.010	0.493	0.162	0.211

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows the effect of medicine value (log wholesale price of the tracked medicine) on whether a sentinel went missing during downstream remote tracking. From a linear regression with facility clustered errors. Controls include facility type, log delivery size, and the log number of tracked units at a facility.

8.2 Effects of Community Monitoring on Diversion

In Table S7 and Table S8 we re-estimate the effects of community monitoring (whether HCACs monitor deliveries) on upstream and downstream diversion using alternative specifications. We exclude controls (column 1) and including district fixed effects (column 2 and 4). The results are similar in magnitude and significance to those in the main text.

Table S7: Effect of community monitoring on upstream diversion

	Upstream Diversion			
		District FE	with controls	controls + FE
	(1)	(2)	(3)	(4)
HCAC Monitoring	−0.098***	−0.068***	−0.138***	−0.113***
	(0.025)	(0.022)	(0.039)	(0.035)
Observations	2,057	2,018	1,735	1,735
R^2	0.019	0.081	0.074	0.134

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows the results of a linear regression estimating the effect of community monitoring (HCAC activity) on whether a sentinel was diverted. Controls include the number of tracked sentinels (log), the number of shipped medicines (log) and health facility type. Errors are clustered on health facility. P-values are from two-tailed t-tests.

Table S8: Effect of community monitoring on downstream diversion

	Downstream Diversion			
		District FE	with controls	controls + FE
	(1)	(2)	(3)	(4)
HCAC Monitoring	−0.158**	−0.110*	−0.076	−0.086
	(0.072)	(0.060)	(0.073)	(0.067)
Observations	1,791	1,755	1,550	1,550
R^2	0.020	0.076	0.178	0.217

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows the results of a linear regression estimating the effect of community monitoring (HCAC activity) on whether a sentinel was diverted. Controls include the number of tracked sentinels (log), the number of shipped medicines (log) and health facility type. Errors are clustered on health facility. P-values are from two-tailed t-tests.

8.3 Effects of Diversion on Respondent Medicine Access and Morbidity

In Table S9 and Table S10 we re-estimate the effects of upstream and downstream diversion on respondent stockout experiences using alternative specification. We exclude controls (column 1) and including district fixed effects (column 2 and 4). The results are similar in magnitude and significance to those in the main text.

In Table S11 and Table S12 we so similar estimates for the effect of diversion on whether the respondent experienced a worsening illness as a result of not receiving medicine from their local health facility. Again the results are consistent with those shown in the main text.

Table S9: Effect of upstream diversion on experience of stockouts

	Gone without medicines in the last 3 months			
		District FE	with controls	controls + FE
	(1)	(2)	(3)	(4)
Upstream Diversion	0.181***	0.182**	0.235***	0.298***
	(0.058)	(0.083)	(0.038)	(0.084)
Household Income	0.098***	0.093***	0.109***	0.090***
	(0.017)	(0.022)	(0.017)	(0.022)
Upstream Diversion*Household Income	−0.101***	−0.084**	−0.141*	−0.114
	(0.036)	(0.042)	(0.077)	(0.078)
Observations	2,451	2,451	2,348	2,348
R^2	0.037	0.116	0.061	0.129

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows a linear regression estimating the effect of respondent income (in standard deviations) and facility-level diversion on whether the respondent's household has been unable to get needed medicines from the clinic in the last three months. Errors are clustered on health facility. P-values are from two-tailed t-tests.

Table S10: Effect of downstream diversion on medicine expenses

	Spend on medicines			
		District FE	with controls	controls + FE
	(1)	(2)	(3)	(4)
Downstream Diversion	−0.052	−0.0001	0.021	0.145
	(0.225)	(0.228)	(0.227)	(0.233)
Times visited clinic	−0.031	−0.025	−0.021	−0.012
	(0.023)	(0.022)	(0.022)	(0.019)
Downstream Diversion*Times visited clinic	−0.015	−0.030	−0.040	−0.059
	(0.073)	(0.075)	(0.074)	(0.075)
Observations	2,082	2,082	2,018	2,018
R^2	0.003	0.013	0.007	0.023

Note: *p<0.1; **p<0.05; ***p<0.01

replace me

Table S11: Effect of upstream diversion on respondent morbitity

	Made illness worse			
		District FE	with controls	controls + FE
	(1)	(2)	(3)	(4)
Upstream Diversion	0.214***	0.165**	0.240***	0.247***
	(0.046)	(0.081)	(0.033)	(0.088)
Household Income	0.104***	0.087***	0.114***	0.085***
	(0.018)	(0.024)	(0.019)	(0.025)
Upstream Diversion*Household Income	−0.115***	−0.091*	−0.199***	−0.168***
	(0.041)	(0.048)	(0.046)	(0.050)
Observations	2,445	2,445	2,343	2,343
R^2	0.045	0.105	0.061	0.115

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows a linear regression estimating the effect of respondent income (in standard deviations) and facility-level diversion on whether the respondent's illness worsened as a result of lacking medicine. Errors are clustered on health facility. P-values are from two-tailed t-tests.

Table S12: Effect of downstream diversion on respondent morbitity

	Gone without medicines in the last 3 months			
		District FE	with controls	controls + FE
	(1)	(2)	(3)	(4)
Downstream Diversion	0.037	−0.004	0.071	−0.004
	(0.081)	(0.067)	(0.085)	(0.075)
Household Income	0.132***	0.110***	0.141***	0.109***
	(0.018)	(0.026)	(0.019)	(0.026)
Downstream Diversion*Household Income	−0.182***	−0.170***	−0.184***	−0.169***
	(0.054)	(0.054)	(0.057)	(0.056)
Observations	2,378	2,378	2,309	2,309
R^2	0.048	0.110	0.065	0.119

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows a linear regression estimating the effect of respondent income (in standard deviations) and facility-level diversion on whether the respondent's illness worsened as a result of lacking medicine. Errors are clustered on health facility. P-values are from two-tailed t-tests.

9 Experimental Intervention

The remote tracking audit activities described in the main manuscript were executed as part of broader impact evaluation with the goal to experimentally identify the deterrent effect of procurement oversight. As the COVID-19 pandemic prevented our originally designed intervention from proceeding, we redesigned the intervention to avoid in-person activities and rolled it out alongside the remote tracking audit activities described in the main manuscript. The large time lapse between baseline survey data collection and intervention rollout was not ideal, though unavoidable.

Specifically, across a random sample of approximately 2,400 boxes of medicines (the same sample as those tracked in the remote tracking audits), we randomly assigned half ($n = 1,258$) to a top-down monitoring message intervention. We placed a prominent and government-branded sticker containing information to be read by public officials involved in health services delivery (Figure S25). The remaining $n = 1,112$ boxes were randomly assigned to the control group, which did not receive a sticker. The goal of the intervention was to inform recipients that a particular box of medicines is being tracked by the Ministry of Health. Building on behavioral theories of crime, we expected that recipients of this intervention would update their beliefs regarding the probability and costs of interdiction and would be less likely to divert the medicines. The experiment results will be analyzed in a separate paper, but our pre-analysis plan with more details can be found at [LINK TEMPORARILY REDACTED FOR AUTHOR ANONYMITY].

Figure S25: Monitoring Sticker

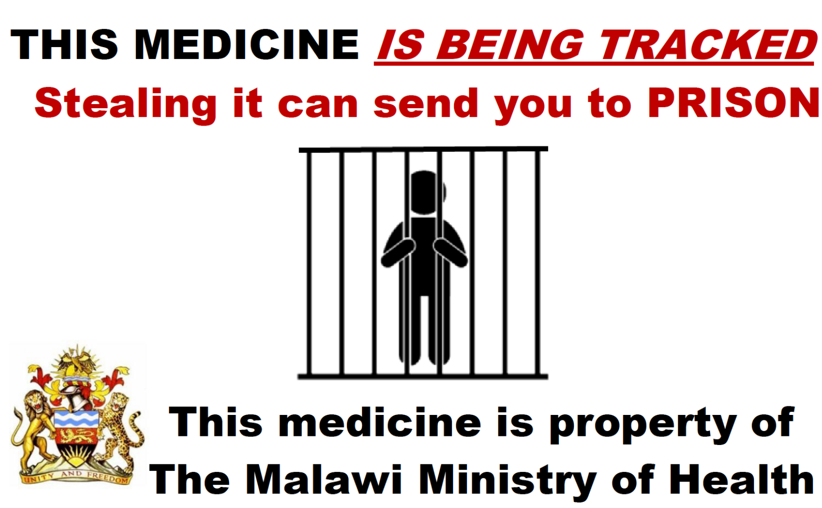

Ethical review for this project took place in 2018 at the London School of Economics and Political Science, and the project was approved in December 2018. In addition, the Malawi National Health Sciences Research Committee reviewed and approved the project twice: once for the original intervention design in March 2019 and once for the intervention redesign in July 2021.

10 Challenges

The results presented in the main manuscript require several caveats regarding challenges and acknowledgements of limitations. This section delineates these issues.

Measuring Diversion: Deliberate Malfeasance versus Administrative Weakness

There are several limitations of our measures of diversion. First, and most importantly, they do not perfectly operationalize theft. They partially capture theft and partially capture administrative weaknesses - for example, disorganization, insufficient record-keeping, or inadequate training. Of course, theft and administrative weakness may be correlated. Still, future research could consider how to use the tracking system we developed to more precisely identify theft distinctly from these administrative issues.

Device Manipulation

One important form of non-compliance is tampering. Though we took steps to obfuscate tracking devices, we have anecdotal evidence of at least a few instances where SENTRIES were damaged or simply switched off to stop their tracking capability. Similarly, some SENTINELs were unable to be found after leaving the delivery trucks, even once the logistics company managing the deliveries deployed a team to retrieve them. In the analysis above, we interpret these instances as diversion.

11 Pre-Analysis Plan

This study was pre-registered on the American Economics Review pre-registration platform prior to analysis on February 16, 2022. The hypotheses in this PAP focus on the experimental intervention, rather than on the non-experimental analysis in this paper. However we include the PAP here for transparency and to document our intentions with respect to measurement and our assumptions about statistical power.

Pre-Analysis Plan:
The Effects of Top-Down Monitoring on Medicine Theft

Ryan Jablonski[*], Brigitte Seim[†], Mariana Carvalho[‡], and Clark Gibson[§]

27 January 2022

Contents

[*]London School of Economics and Political Science
[†]University of North Carolina, Chapel Hill
[‡]University of California, San Diego
[§]University of California, San Diego

1 Preface

The study described in this pre-analysis plan was revised in response to the COVID-19 pandemic. The pre-analysis plan for the original study can be found here. Significant changes were made to the originally proposed study due to changes in the research and health environment in Malawi.

2 Motivation

We seek to uncover spatial patterns of medicine diversion[1] in local public health facilities, and identify how public health messaging can change patterns of diversion and impact citizens' access to medicines. While firm estimates are lacking, medicine theft is likely one of the leading causes of preventable disease in low income countries. According to Malawi Ministry of Health officials, 29% of spending on medicines and medical supplies disappears due to theft alone (Mphande, 2017). We estimated in 2019 that 75% of all communities in Malawi have experienced the theft of medical supplies (author surveys). Malawi is a particularly extreme outlier, but these estimates are not far off from estimates in other countries (Bate et al., 2010). Additionally, medicine theft costs governments and donors many millions each year between lost aid and spending on parallel procurement procedures.

Solving these procurement and corruption challenges is particularly vital in the current COVID-19 epidemic. Estimates from previous health crises suggest that ineffective health service delivery explains much regional variation in mortality rates and subsequent poverty in poorly institutionalized health settings (Christensen et al., 2020; Maffioli, 2018). Moreover, the rapid expansion of global health aid will likely create opportunities for higher rates of theft and corruption (Svensson et al., 2000; Thorp, 2020).

We will experimentally test a top-down monitoring messaging intervention (in the form of a sticker placed on medicine boxes) designed to reduce medicine diversion and increase the efficiency of public health procurement. Additionally, we introduce a novel measurement protocol that uses GPS and Bluetooth tracking devices to measure rates of diversion at different levels of the medicine supply chain. This allows us to move beyond existing studies of corruption that focus on changes in corruption

[1] In this pre-analysis plan, we use the word "diversion" to encompass medicines failing to arrive at their intended destination due to *either* theft *or* disorganization. Our scoping research for this project leads us to believe that theft is a common reason for diversion, but we cannot rule out the possibility that some instances of diversion are simply mismanagement due to disorganization, poor record-keeping, or other administrative gaps.

levels to present evidence demonstrating how anti-corruption interventions displace corruption to other parts of the system.

This low-cost, highly scalable intervention is immediately applicable to the government's efforts to improve health services delivery. Our implementation partner, the Drug Theft Investigation Unit (DTIU), is a DFID-funded government agency housed under the Malawi Ministry of Health responsible for ensuring the security of public health procurement. Like many similar organizations in the developing world, the DTIU has had only limited success. Officials complain of insufficient resourcing and the frequent non-compliance of local officials with procurement regulations. Collaboration between DTIU and the judicial system has also been fraught. We designed ouur intervention to circumvent these challenges. Unlike more traditional capacity building or auditing interventions, this intervention does not require broad institutional reform or a large input of financial resources. This treatment is also designed to speak to the effectiveness of similar public messaging campaigns sponsored by many donors, such as the ispeakoutnow.org campaign sponsored by Global Fund.

This intervention is also efficient in a crisis environment. Traditional interventions to prevent corruption involve audits, training, and local capacity building. During the COVID-19 pandemic, these measures would have put public officials at risk and undermined quarantine measures. This leads to weak procurement oversight during a time when corruption is particularly attractive. By offering a low-cost, low-interaction, and easily fielded intervention, we expect the lessons from this study to be useful for crisis procurement management generally.

3 Research Design

3.1 Intervention

In this study we will experimentally test the effects of a top-down monitoring intervention. In cooperation with local partners, we will randomly sample approximately 2,400 medicine boxes scheduled for delivery to local health facilities in Southern Malawi. Using data from previous surveys, we will weight this sample towards commonly ordered and commonly diverted medicines. On 1,200 of these boxes (one-half), we will place a prominent and government-branded sticker containing information to be read by public officials involved in health services delivery (Figure 1). Boxes in the control group will not receive a sticker. The goal of this message is to inform recipients that a particular box of medicines is being tracked by the Ministry of Health. Building on behavioral theories of crime, we hypothesize that recipients of this message will update their beliefs in the probability and costs of interdiction and will be less likely to divert the medicines. Also we expect the messaging to decrease "agency slack" among health services officials by increasing their knowledge of central government monitoring. We therefore expect more efficient delivery of medicines in this treatment group.

Figure 1: Monitoring Sticker

3.2 Sampling and Tracking

To capture treatment effects of the top-down monitoring intervention, we rely on long-range and short-range tracking devices manufactured by OnAsset, a US-based company specializing in supply chain visibility technology. Specifically, the tracking protocol involves two types of devices: (1) SENTRY 500; and (2) SENTINEL. The SENTRY 500 is a cellular and GPS tracking device. It is the "parent" device. The SENTINEL is a "child" Bluetooth device automatically read by any SENTRY 500 within their transmission radius. The SENTRY 500 units capture and transmit data from itself and all surrounding SENTINELS to the server at regular intervals. In addition to the location of the devices, the SENTRY also captures data on temperature, light, humidity, and battery. The data on the server can then be accessed via API.

The experiment and corresponding medicine tracking occurred during the medicine delivery cycles of October[2] and November of 2021. Using health facility orders pro-

[2]These deliveries were delayed and were actually delivered in November.

vided to us by the Ministry of Health, and drawing on our pre-experiment survey and interview data, we sampled eight medicines that are both commonly ordered and commonly diverted.

Next, in partnership with the Ministry of Health, we accessed the main warehouse that is responsible for delivery of medicines to public health facilitys in the Southern Region. A team of 14 people from the third-party logistics company (9), the warehouse (1), Malawi Ministry of Health (1), and the research team (3) then placed SENTINEL tracking devices in a random sample of the boxes for the medicines (Table 1). The SENTINEL devices were covered with a sticker stating each device was a "SHIPPING TEMPERATURE & HUMIDITY MONITORING DEVICE" and property of the Ministry of Health (Figure 2).

Table 1: Tracked Medicines

Medicine Description	**Anticipated Number of Tracked Boxes**
Amoxycillin,250mg,capsule,1000	89
Aspirin,300mg,tablet,1000	402
Ibuprofen,200mg,coated tablet,1000	786
Insulin zinc suspension (lente),100 IU/ml,10ml	3
Paracetamol,500mg,tablet,1000	519
Phenobarbitone,30mg,tablet,1000	166
Syringe Luer,10ml,disposable,hypoluer	235
Syringe,autodestruct,5ml,disposable	170

These counts are from preliminary data collected at the main warehouse and may change once the data are cleaned and duplicate records are eliminated.

Figure 2: Tracking Device Sticker

In addition to the unique device code and the medicine information, we also recorded the facility destination and treatment status for each box in which a SENTINEL was placed. The boxes in the October and November deliveries were intended for 171 facilities, ranging from large district hospitals to small local health centres. Treatment status was randomized across the whole sample of boxes, with approximately half (1258) of the boxes randomly assigned to treatment and approximately half (1112) randomly assigned to control. The same team that placed the SENTINELS inside the boxes placed the treatment stickers on the boxes. In a parallel process, the team placed two SENTRIES on each delivery truck.

Approximately two weeks after the deliveries were scheduled to occur, two teams of three people each[3] conducted a follow-up "trackers tracing" exercise in which they traveled to each intended facility destination with a SENTRY inside a backpack to scan for SENTINEL devices. While on the facility premises, they inconspicuously inquired about possible markets and private pharmacies where medicines were sold and brought the SENTRY to these locations to scan for SENTINELS once more.

In addition to the data we collected from these devices, we also conducted an in-person survey with citizens and health officials in the catchment areas of 97 health facilities (N=3,360) prior to this experimental intervention. These 97 facilities were re-sampled in the trackers tracing exercise. Data from this survey will be used to evaluate some of the conditional hypotheses below. The data may also be used for validation purposes and as control variables in some of our analyses.

[3]Each team involved one person from the research team and then two people not previously involved in the device and sticker placement at the warehouse.

4 Theory and Hypotheses

4.1 Theory

The intervention draws on the literature considering the effect of bureaucratic oversight on corruption. Typically, the oversight takes the form of an audit, and the corruption takes the form of embezzlement or theft of government funds and materials, particularly during the procurement and provision of public goods and services. Scholars in this literature theorize that audits can reduce corruption via two mechanisms. In one mechanism - a "behavioral effect" - the *threat* of an audit compels politicians to refrain from corruption. As long as politicians believe the probability of an audit is sufficiently high (and, importantly, that the probability of detection of corruption in said audit is sufficiently high), they will refrain from corruption. In the other mechanism, audits have a *selection effect*, whereby audits detect corruption, the corrupt are removed from their posts, and then corruption decreases. While there is empirical evidence supporting both mechanisms linking audits and corruption across diverse contexts, there are only a few studies assessing whether audits can reduce medicine diversion, and no experiments (that we know of) on the effects of audits on medicine diversion in developing countries.

The treatment message in our experiment is closely aligned with this literature. We anticipate that messages on medicine boxes that convey the Ministry of Health is tracking medicines will reduce rates of diversion of these boxes. More specifically, we anticipate the messages will heighten the perceived probability of theft detection and punishment, causing the behavioral effect discussed above and lowering the rate of diversion. While we will be unable to empirically identify any selection effect of the treatment message, because we will not be recollecting data after this (or any) audit has resulted in some officials' removal from office, we will be able to estimate the potential for such an effect by combining our pre-experiment survey and interview data with the data obtained in the intervention phase.

4.2 Main Hypotheses

Our main hypothesis relates to the overall effect of the intervention on medicine diversion.

HA1: The Monitoring Message will decrease the rate of medicine diversion.

4.3 Heterogeneous Effects

We also pre-specify two hypotheses regarding heterogeneous effects of the Monitoring Message. We expect that individuals decide to engage in medicine theft based on a cost-benefit calculation. The benefits of theft include the profits from the resale of the medicines. The costs of theft include the transaction costs associated with the theft and transfer/resale as well as the costs associated with being caught and (if caught) punished. Costs in the latter category are only incurred with some probability and the Monitoring Message is designed to increase posterior beliefs about these probabilities. We will consider heterogenous effects associated with this behavioural model, including the following:

HB1: The effect of the Monitoring Message on the rate of medicine diversion will be greater for medicines with lower financial value.

HB2: The effect of the Monitoring Message on the rate of medicine diversion will be greater at clinics where there is an active citizen oversight committee.

HB3: The effect of the Monitoring Message will be greater for *DownstreamDiversion* than *UpstreamDiversion*.

4.4 Non-Experimental Analysis

In addition to analysing the effects of treatment, we also plan to use these data to analyze overall patterns of diversion in the Malawi context. We will evaluate the ways in which patterns of diversion depend upon geographic features (especially road density, the proximity of national borders and distance from major towns), citizen oversight, theft perceptions from baseline surveys, and political alignment.

5 Measurement

5.1 Outcome Variables

Our primary goal is to measure the diversion of medicines. We anticipate measuring diversion in four mutually supportive ways:

1. *UpstreamDiversion* This variable will take on a value of one if a SENTINEL is not originally delivered to its intended destination. It will take on a value of

zero otherwise.

2. *MedicineNotTraced* This variable will take on a value of one if a SENTINEL cannot be found at any public health facility during the trackers tracing exercise (that is, after delivery has taken place). It will take on a value of zero otherwise. This variable can only be coded for facilities where trackers tracing occurred.

3. *DownstreamDiversion* This variable will take on a value of one if a SENTINEL was originally delivered to its intended destination (*UpstreamDiversion* equals 0) and then is subsequently not found at any public health facility during the trackers tracing exercise (*MedicineNotTraced* equals 1).

4. *TracedDiversion* This variable will take on a value of one if a SENTINEL was found at at a location other than that intended during the trackers tracing exercise. It will take on a value of zero otherwise. This variable can only be coded for facilities where trackers tracing occurred.

We see these measures as complementary, but varying in their power and precision. We note for instance that events #3 and #4 represent more direct measures of theft than events #1 and #2 and may be less subject to measurement error.

We will also consider a more conservative coding rule in which *UpstreamDiversion* and *TracedDiversion* are only coded as one when a medicine was delivered to or found at a location other than a government health facility.

In addition to the above coding rules, there will be exceptional cases where we code a diversion event or not for reasons listed above, and such events will be clearly listed in the supplemental material to our analysis. We iterate anticipated cases of this sort below:

1. A SENTINEL that was intended to be delivered to a district health office (DHOs) can sometimes be redelivered to other facilities in the district. As a result, such a medicine will only be considered diverted if: (1) there are no records of that medicine being re-delivered to another facility; or (2) there are such records, attempts were made to locate the medicine at the redelivered facility, and such efforts were unsuccessful. Because there is unmeasured uncertainty in the level of diversion from DHOs, we will estimate our main analysis with and without medicines targeted to DHOs.

2. A SENTINEL may be found at a location other than a clinic and then turned into public health authorities. In most cases, the medicine associated with this SENTINEL will be considered diverted, though we anticipate that the coding for such situations may be context specific.

5.2 Covariates

We plan to collect data on the covariates depicted in Table 2. These will be used to conduct balance, attrition and heterogeneous effect tests as discussed below. If we are unable to collect data for a given pre-specified covariate, that covariate will be excluded from analysis.

Table 2: Covariates

Variable	Level	Description
Medicine Type	SENTINEL (Medicine)	A set of dummy variables identifying the types of tracked medicines.
Health Facility District	Health Facility	A set of dummy variables identifying the districts of the intended health facility destinations.
Health Facility Distance	Health Facility	A continuous variable measuring the logged distance between the central medical store and the intended health facility destination.
Health Facility Type	Health Facility	A categorical measure indicating the classification of the facility.
Health Facility Medicine Count (log)	Health Facility	The total number of medicines that were delivered to the facility.
Health Facility Tracking Density	Health Facility	A continuous variable ranging from 0 to 1 that indicates the proportion of medicines sent to a facility that were tracked with a SENTINEL.
Health Facility Tracking Density	Health Facility	A continuous variable ranging from 0 to 1 that indicates the proportion of medicines sent to a facility that were tracked with a SENTINEL.
SENTRY Effects	SENTRY (Route)	A set of dummy variables identifying the SENTRY units reporting tracking data.
Departure Date	SENTRY (Route)	A set of dummy variables identifying the dates on which deliveries departed from central medical stores.

Additionally, in a restricted analysis limited to the 97 facilities with baseline data, we will include controls for respondents' perceptions of medicine theft, the consequences of medicine theft, the probability of being caught, and facility quality. We will also include a control variable for how active or not a citizen oversight committee is at the facility and the frequency with which citizens were told that medicines were unavailable at a facility. .

6 Analysis Details

6.1 Main Analysis

Our primary estimating approach is given by the linear probability model in Equation 1. Here y_i is a binary variable indicating whether medicine unit i was diverted (as described in Section 5). Let t_{1i} indicate whether unit i is assigned to the Monitoring Message treatment or control. Let u_j be a fixed effect for each route j. We will also show estimates with facility fixed effects.

Since the treatment is assigned at the unit level, we will not cluster the standard errors in the main analysis.

$$y_i = \beta_1 t_1 + u_j + e_i \tag{1}$$

We will also estimate Equation 1 with the inclusion of the covariates listed in Section 5.2.

6.2 Analysis of Heterogenous Effects

As discussed in Section 4.2 we anticipate that treatment effects might differ across theoretically relevant sub-groups. We will estimate these effects using the linear probability model in Equation 2 where g_i is the conditioning variable.

$$y_i = \beta_1 t_1 + \beta_2 t_1 g_i + \beta_3 g_i + u_j + e_i \tag{2}$$

6.3 Contingencies

6.3.1 Non-Compliance and Manipulation

One important form of non-compliance is tampering. We will take steps to obfuscate tracking devices (SENTINELS), however it is feasible that some devices will be identified and then damaged or removed. One form of tampering would be if medicines were repackaged in order to circumvent tracking. As discussed above, we will be able to identify such events from anomalous tracking data. In our main analysis we will consider tampering to be an instance of theft. We will also run our analysis separately on identifiable cases of tampering.

It is also possible that officials will remove all SENTINELS from a delivery, or disable the SENTRY units. Such units will be considered attritted.

Another form of non-compliance is if stickers are removed at some point during the procurement process, for instance due to damage or intentional tampering. It is not feasible to measure compliance at the individual unit level, but we hope to conduct an endline survey to determine if: (1) the messages are still affixed; and (2) were read and understood.

6.3.2 Missing Data and Attrition

It is possible that some tracked units will become faulty or non-responsive, or will be discovered and destroyed by health officials. It may also be infeasible to obtain data from some units due to coding errors in enumerator surveys or other technical faults. We will consider such devices to be attritted and their data will be excluded from analysis.

In some cases, it may be impossible to obtain data for one or more health facilities (for instance, due to a technical fault or due to the follow-up tracing team being unable to identify or access the location of a facility). We will consider medicines destined for such facilities to be attritted and they will be excluded from analysis.

We will conduct an omnibus test of balanced attrition using available pre-treatment covariates. If we do reject the null, we will conduct a bounds analysis to validate the robustness of our results to the missing at random assumption.

If we are unable to collect data for a given pre-specified covariate, that covariate will be excluded from analysis. If we are missing data for only a small portion of observations, and a reasonable imputation method is available for those, we will impute missing data rather than exclude the variable from analysis.

7 Power Analysis

For our main analysis, we assume that we will be able to successfully track 2,400 units. We consider several reasonable power scenarios in Figure 7. For instance, if we assume an overall population diversion rate of 10%, we obtain 80% power assuming an additive treatment effect of at least 0.04 (a 40% reduction relative to baseline diversion). While we lack firm data on baseline diversion rates, this number is conservative relative to audits conducted by Global Fund and others.

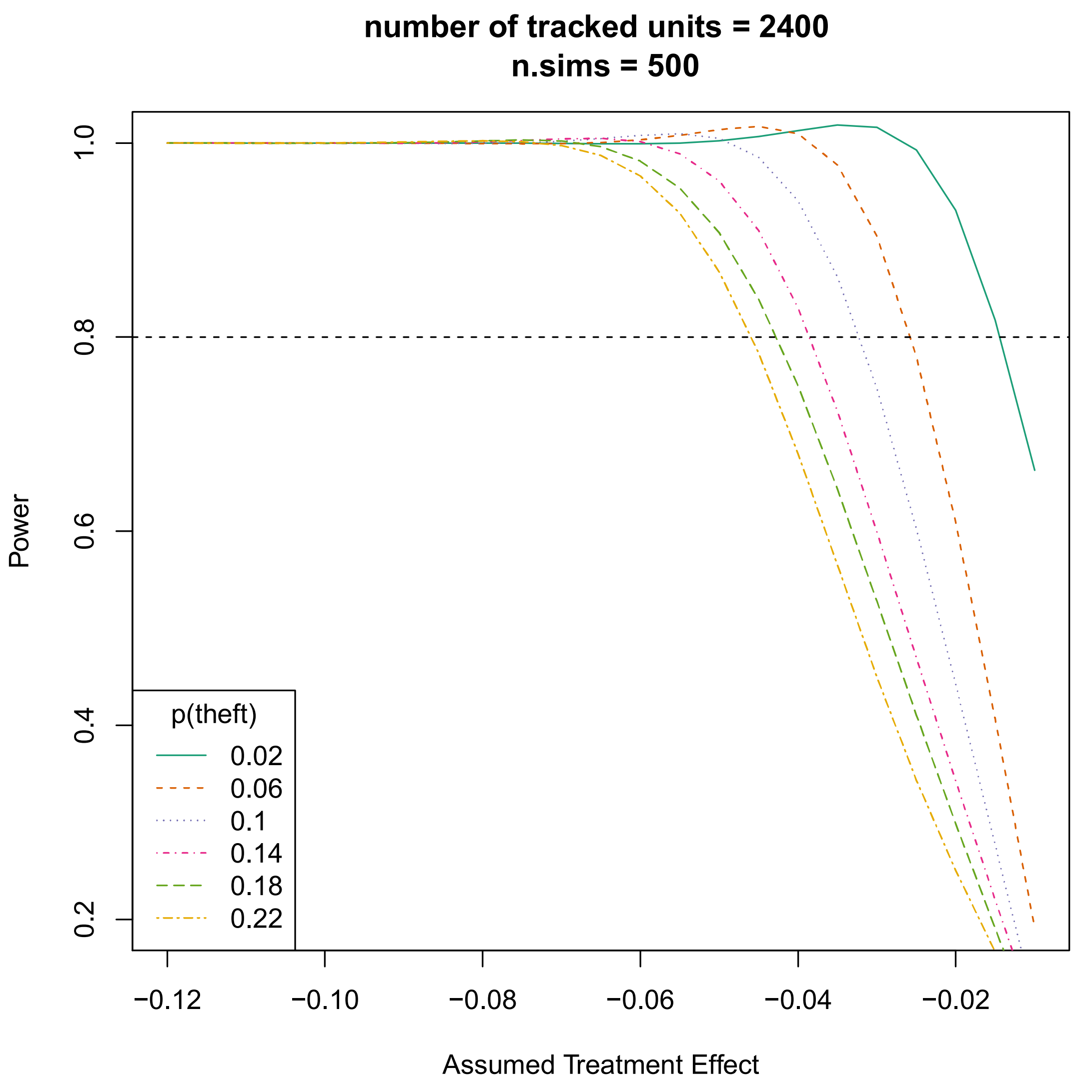
number of tracked units = 2400
n.sims = 500
Power
1.0
0.8
0.6
0.4
0.2
−0.12
−0.10
−0.08
−0.06
−0.04
−0.02
Assumed Treatment Effect
p(theft)
0.02
0.06
0.1
0.14
0.18
0.22

7.1 Replication Code

```
library(RColorBrewer)
set.seed(4567)

sims = 500                       # number of simulations
n.tracked = 2400           # number of sentinals/tracked units
pTreat = 0.5                     # prob of treatment
pStolens = seq(0.02,0.22,
               by=0.04)          # prob of theft
treatEffects = seq(-0.01,
            -0.12, by=-0.005)   #vector of assume treatment effects


power = matrix(nrow=length(treatEffects),ncol=length(pStolens))

k=0

#loop over p(stolen)
for (pStolen in pStolens) {
  k=k+1
  j=0
  #loop over effect sizes
  for (treatEffect in treatEffects) {
    j=j+1
    reject = rep(NA,sims)
    tStat = rep(NA,sims)
    y = rep(NA,n.tracked)
    x = rep(NA,n.tracked)

    # loop over simulations
    for (simul in 1:sims) {

      stolen    = rep(NA,n.tracked)
      treat    = rep(NA,n.tracked)

      # loop over drugs
      for (i in 1:n.tracked){
        # treatment assignment at unit level
        treat[i] = rbinom(n=1, size=1, prob=pTreat)
        #probability of theft
        pr=pStolen + (treat[i]*treatEffect)
        pr=ifelse(pr<0,0,pr)
        #indicator for whether the drug is stolen
        stolen[i]  = rbinom(n=1, size=1, prob = pr)
      }
```

```
      y = stolen
      x = treat

      fit.sim    = lm(y ~ x)
      tStat      = summary(fit.sim)$coefficients[2,3]
      reject[simul] = (tStat <= -1.96)

    }

    power[j,k] = mean(reject)

  }
}


cols=brewer.pal(n = length(pStolens), name = "Dark2")
par(mfrow=c(1,1))
plot(treatEffects, power[,1], ylim=c(0.2,1),xlab="Assumed␣Treatment␣Effect", cex=1,
  for (i in 1:length(pStolens)){
    lines(smooth.spline(treatEffects, power[,i], spar=0.5), col=cols[i], cex=2, yli
  }
abline(h=0.8, lty=2)
legend("bottomleft",lty = c(1:length(pStolens)),
         legend = c(pStolens),
         title = "p(theft)",
         col=cols)
title(paste("number␣of␣tracked␣units␣=", n.tracked, "\n␣n.sims␣=", sims))
```

References

Bate, R., Hess, K., and Mooney, L. (2010). Antimalarial medicine diversion: stock-outs and other public health problems. *Research and Reports in Tropical Medicine (1)*, pages 19–24.

Christensen, D., Dube, O., Haushofer, J., Siddiqi, B., and Voors, M. (2020). Building resilient health systems: Experimental evidence from sierra leone and the 2014 ebola outbreak. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-28).

Maffioli, E. M. (2018). The political economy of health epidemics: evidence from the ebola outbreak. *Available at SSRN 3383187*.

Mphande, H. (2017). Massive Drug Theft Hits Malawi Hospitals. *Nyasa Times*, Sept. 04, 2017.

Svensson, J. et al. (2000). Foreign aid and rent-seeking. *Journal of International Economics*, 51(2):437–461.

Thorp, A. (2020). To Defeat the Coronavirus, Stop Corruption. *Foreign Policy*, Apr. 06, 2020.

References

1. Oxfam. *Missing Medicines in Malawi Campaigning against 'stock-outs' of essential drugs* tech. rep. (Oxfam, 2012). https://oxfamilibrary.openrepository.com/bitstream/handle/10546/226732/pi-local-governance-malawi-25052012-en.pdf?sequence=1.
2. "The Global Fund". *Audit Report: Global Fund Grants in Malawi* tech. rep. (The Global Fund, 2016). https://www.theglobalfund.org/media/2665/oig_gf-oig-16-024_report_en.pdf.
3. Lodenstein, E. *et al.* "We come as friends": approaches to social accountability by health committees in Northern Malawi. *BMC health services research* **19,** 1–14 (2019).
4. Kateta, M. W. Malawi continues to face acute drug shortages due to COVID-19. *Devex* (2022).
5. Meinshausen, N. & Ridgeway, G. Quantile regression forests. *Journal of machine learning research* **7** (2006).