Unconditional Cash Transfers and Child Mortality: Experimental Evidence from Kenya*

Michael Walker ⓡ Nick Shankar ⓡ Edward Miguel

ⓡ Dennis Egger ⓡ Grady Killeen

March 20, 2025
Preliminary conference version; please do not circulate

Abstract

We estimate the intergenerational health impacts of large-scale unconditional cash transfers. One-time transfers of USD 1000 were provided to over 10,500 poor households across 653 randomized villages in Kenya. We collected regional census data on over 100,000 births, including on mortality and cause of death, and detailed data on household health behaviors. In the study's main finding, the cash transfer treatment leads to 55% fewer infant deaths before age one and 49% fewer child deaths before age five. Data on the location of health facilities, as well as the cause of death and transfer timing relative to birth, indicate that unconditional cash transfers and access to delivery services are complements in generating mortality reductions: the largest gains are estimated among households living close to health facilities who receive the transfer around the time of the birth, and treatment leads to a large overall increase in hospital deliveries. Infant and child mortality then largely revert to pre-program levels after cash transfers end. Despite not being the main aim of the original program, we show that unconditional cash transfers in this setting may be a cost-effective way to reduce infant and child deaths.

*AEA Trial Registry: AEARCTR-0000505, https://www.socialscienceregistry.org/trials/505. Walker: University of California, Berkeley, mwwalker@berkeley.edu; Shankar: University of California, Berkeley, nick.shankar@berkeley.edu; Miguel: University of California, Berkeley, emiguel@berkeley.edu; Egger: University of Oxford, dennis.egger@economics.ox.ac.uk; Killeen: University of California, Berkeley, gkilleen@berkeley.edu. Thanks to Jasmin Baier, Daniel Han, Layna Lowe, Prince Muraguri, and Maya Shen for excellent research assistance, and Carol Nekesa, Andrew Wabwire and REMIT Kenya for data collection and support, and audiences at the Pacific Conference for Development Economics, the Population Association of America Annual Meetings, Sociedad de Economistas del Uruguay (SEU) Annual Meeting, the Southern California Conference in Applied Microeconomics (SoCCAM), Maseno University, the University of California Berkeley, EAGxBerkeley, Meta, the University of Chicago, IMF, and the ASSA Meetings, for helpful comments and suggestions. Funding from Open Philanthropy, recommended by GiveWell, is supporting this work. The author order was certified randomized.

1 Introduction

Wealth is closely associated with health, across societies and across individuals within societies. The gradient between health and socioeconomic status is one of the most widely-documented correlations in the social sciences (Preston, 1975; Cutler et al., 2012; Lleras-Muney et al., 2024). Studies often show a concave relationship (Deaton and Paxson, 2004; Cutler et al., 2012), suggesting that poverty reduction in low-income settings may have particularly important implications for core health outcomes such as mortality. For instance, low and middle income countries still bear a disproportionate burden of child mortality (Burstein et al., 2019). However, experimental evidence on the poverty-mortality relationship has been limited due to the need for both large-scale data collection (to achieve adequate statistical power) and randomized variation in income shocks.

The rise of unconditional cash transfer (UCT) programs provides an opportunity to study the causal effect of exogenous income gains on mortality. Over 100 low- and middle-income countries have introduced UCT programs in the past two decades (Stedman, 2023), and a growing experimental literature has studied the effects of these programs on a wide range of development outcomes (Egger et al., 2022; Crosta et al., 2024). As UCTs become more established as an anti-poverty policy tool, interest has grown in understanding whether the benefits of UCTs spill over intergenerationally onto the children of recipients. Intergenerational effects may amplify the direct positive effects on recipients that have been documented in the short and medium-run, improving UCTs cost effectiveness. Yet despite the proliferation of studies on UCT programs (including many RCTs), it has remained challenging to experimentally examine impacts on one of the most central markers of child well-being — whether they survive infancy and their first five years of life. Randomized evaluations of UCTs to date have typically lacked the large sample size and longitudinal data collection necessary to precisely estimate impacts on relatively rare but important outcomes such as child mortality.

The recent spread of UCT programs worldwide has proceeded alongside — and at times in lieu of — investments expanding access to primary healthcare. Previous non-experimental studies suggest an association between UCTs and reduced child mortality (Richterman et al., 2023), and combined with the costs of maintaining health infrastructure, this raises the question of whether distributing cash is sufficient on its own to reduce child mortality (Blattman and Niehaus, 2014; Evans and Kosec, 2016; Stedman, 2023). Or are complementary investments in health infrastructure necessary to improve child survival in low-income settings? To illustrate, higher levels of income may not lead to major changes in birth outcomes if advanced delivery services and trained medical staff are not available to assist should com-

plications arise. Progress on this question is important for policymakers as they allocate scarce resources to improve health outcomes.

This article exploits large-scale experimental variation from an unconditional cash transfer program in rural Kenya, combined with a new census of over 100,000 births in the study region over more than a decade, to causally identify the effect of UCTs on infant and child mortality. The one-time cash transfer was roughly 75% of average annual income among eligible households, which represented approximately the poorest third of households in the study region, a rural county chosen by the implementing NGO for its high poverty levels. As such, it represented a large income shock for both recipients and a stimulus to the broader local economy, resulting in substantial living standard improvements for recipient households and others living nearby (Egger et al., 2022). The regional census of births undertaken by the research team was designed to have an adequate sample size to more precisely estimate the impact of this large economic shock on child survival.

Specifically, the census data includes detailed birth histories from all 106,871 women in the study area covering 2011 to 2023, as well as information on child mortality and survival. To better understand the underlying mechanisms behind changes in mortality, we collected verbal autopsies (VAs) for child deaths using the World Health Organization's 2022 methodology and assigned causes of death using the Institute for Health Metrics and Evaluation's machine-learning classification algorithm (World Health Organization, 2022a; Institute for Health Metrics and Evaluation, 2025). In settings such as Kenya in which physical autopsies are rare and vital records incomplete, VAs are considered the state-of-the-art method for determining causes of death at scale (Gacheri et al., 2014; Serina et al., 2015; Amek et al., 2014, 2018), and was validated in our particular study area by the Kenya Medical Research Institute.[1] We pre-specified the primary outcomes and experimental design (on the AEA RCT Registry, AEARCTR-0000505), which largely follows our previous work in leveraging both randomly-assigned treatment status and spatial variation in treatment intensity to study local spillover effects (Egger et al., 2022). We further collected extensive data on health facility access, pairing administrative data on the locations of Kenyan hospitals and clinics with new assessments of travel times collected by sending enumerators equipped with speedometers to travel local trails and roads, and for a representative sample of pregnancies obtained detailed information on antenatal, delivery, and postnatal healthcare utilization.

The central empirical finding of this study is that the cash transfer treatment led to large and significant reductions in infant and child mortality. Results from the study's primary specification, which accounts for spillover effects within- and across-villages, indicate that

[1]The Kenya Medical Research Institute (KEMRI), the country's flagship health research institution, also trained our field staff in performing VAs.

infant mortality fell by 23 deaths per thousand births among recipient households with a pregnancy during the transfer disbursal period. This represents a 55% decline in infant mortality relative to the mean among recipient household births in control villages, and is highly statistically significant. We find similar results for under-five mortality, or when estimating a reduced-form OLS specification that relies directly on the two-stage research design, which randomizes treatment at both the village level and subregion level (the econometric models are described in detail below).

The large-scale longitudinal nature of the data we collect enables us to assess the validity and robustness of the estimates in several ways. We document that these sizable experimental estimates are similar in magnitude to non-experimental differences in child mortality between eligible households and their wealthier ineligible neighbors (in this study's control group), as well as to the variation in mortality between the peak post-harvest season versus the pre-harvest "lean" season months in the decade we observe. Child survival thus appears very sensitive to current economic conditions in this setting, across multiple sources of variation in household economic conditions. The data also confirms expected patterns with respect to maternal age, and with the onset of the COVID-19 pandemic and concurrent lockdowns, which was associated with a rapid increase in child mortality. Furthermore, in a placebo check, we do not estimate any significant differences in infant or child mortality between treatment and control households in the pre-treatment period (2011-14).

An established literature in public health research confirms that child births and deaths are sufficiently consequential to be recalled by mothers and close relatives long after they occur (Rao et al., 2003; Lyons-Amos and Stones, 2017; Nareeba et al., 2021; Romero Prieto et al., 2021). As such, our birth census, which relies on family members to recall these vital events, utilizes the same methodology as used by prominent data sources on child health such as the Demographic and Health Surveys (Romero Prieto et al., 2021). Simple memory issues are unlikely to bias the results as they would introduce classical measurement error attenuating estimates, and the census data appears internally consistent across both cross-sectional and inter-temporal dimensions, as previously noted. Patterns of treatment effects also appear inconsistent with simple experimenter demand stories. For one, we do not find differences in child mortality between treatment and control households after the period of transfer disbursal (nor before, as noted above). We also find that some of the reported reductions in mortality during the cash disbursal period are driven by local cash transfer spillovers (where misreporting due to experimenter demand effects seems unlikely) rather than the direct effects of receiving cash alone.

What mechanisms could drive these large reductions? Several analyses indicate that access to antenatal and delivery services may have played a central role. First, when estimating a

dynamic specification that takes into account transfer timing relative to the timing of pregnancy, we find that gains in child survival are concentrated among women who received cash in the month they gave birth or shortly beforehand. Second, analyzing mortality reductions by cause of death, we find declines across most major cause categories but show that the largest share of the overall effect is concentrated in birth complications and neonatal deaths. Third, we document far larger reductions in child mortality in villages located closer to a health facility staffed by a physician, as assessed using our travel time estimates. This heterogeneity in treatment effects survives the inclusion of a rich set of covariates (accounting for a place's population, wealth, disease environment, geography, and market access) and remains significant if these covariates are interacted with treatment. Fourth, in a representative survey of censused households, we find that transfers are associated with a 70% increase in the rate of hospital deliveries. Taken together these patterns suggest that cash transfers can complement rather than substitute for investments in rural health infrastructure in this low-income setting.

The analysis suggests that the selection of women into pregnancy cannot explain the estimated reductions in child mortality. While we observe a modest transient increase in overall birth rates during the period of transfer disbursal of roughly 10%, several approaches designed to test for changes in the composition of births — made possible by the numerous waves of detailed household surveys we have earlier conducted in the study region — indicate that the characteristics of mothers did not substantively change.

We find stronger evidence of changes in behavioral mechanisms that directly increase the likelihood a child survives, including both access to healthcare (as noted above) and several other possible factors. One factor we study is the role of maternal labor supply, as female labor force participation during pregnancy is high in Kenya and previous non-experimental studies suggest that receipt of social assistance is associated with mothers working less (Amarante et al., 2016; Guldi et al., 2024). We find large reductions in female labor supply in households that receive treatment when a women is in the third trimester, indicating that cash transfers may have enabled women at a critical time in pregnancy to reallocate work hours to rest or other activities conducive to better child outcomes. Another consideration is the role of nutrition. In household surveys administered during the cash disbursal period, we observe that children in recipient households are less likely to have skipped meals or to have gone to bed hungry, and this general improvement in household nutrition could have affected both mothers and newborns.

This study aims to make several contributions. Most directly, it documents the effect of UCTs on one of the most basic indicators of human development — whether children survive their first one to five years of life. Experimental analysis of the impact of UCTs on

child mortality has been complicated by the need to collect detailed, long-term data on child survival from a large group of program beneficiaries. This study's expansive data collection efforts, which entailed visiting 653 villages multiple times to record information on over 100,000 births — combined with an underlying UCT randomized experiment — allow the present study to begin to fill this gap. While previous studies with different research designs have shown modest long-run health effects of cash transfers on adult recipients (Haushofer and Shapiro, 2018; Baird et al., 2019; Blattman et al., 2020), we show that UCTs can also have large inter-generational effects by lowering the mortality rates of children. Prior work has also shown that conditional cash transfers (CCTs) can be effective at lowering child mortality (Barham, 2011). But because these programs condition cash on behavior changes such as health checks that may directly affect mortality, the role of wealth effects versus behavior change to comply with the conditions for aid is unclear. This paper shows that the income effects of cash transfers are important and suggests that cash transfers can be effective without costly monitoring. The substantial child mortality reductions found in this study, coupled with the concentration of effects among mothers receiving cash near the month of birth, suggest that targeting assistance to such mothers may be highly cost-effective (echoing anthropometric gains for children documented in Baird et al. (2019)), and in fact comparably cost-effective to a number of WHO-recommended maternal and child health interventions.

Second, this paper aims to add to a literature on the relationship between economic status and health (Cutler et al., 2012). The program studied, which provided recipients with USD 1000 (USD 1871 PPP), was multiples larger in magnitude than the average level of assets at baseline held by beneficiaries. As such, it represented an large wealth shock in a high-poverty setting where the marginal return to wealth for health outcomes may be expected to be very large, in contrast to well-resourced settings where even sizable economic shocks, such as winning the lottery, may have smaller impacts on health (Apouey and Clark, 2015; Cesarini et al., 2016; Kim and Koh, 2021; Wyndham-Douds and Cowan, 2024). We provide novel experimental evidence that an income shock in low income settings such as ours can sharply reduce mortality for some of the most vulnerable members of the household, infants and children. We document that the socioeconomic gradient in child mortality is causal in nature and not just correlational in this setting. Yet we also document that these large effects dissipate when the cash transfers end, suggesting that while large temporary income shocks can result in considerable health improvements, other interventions may be needed for sustained gains.

Third, these results speak to the complementary roles cash and access to medical services may play in producing population health improvements. There is a striking pattern: the child mortality reductions associated with the randomized cash transfers are largely driven

by households who live near a health facility staffed by a doctor, even after accounting for an array of place and household covariates. This pattern is consistent with prior work indicating that quality health care plays an important causal role in mortality reductions (Okeke, 2023). Cash transfers may not be enough by themselves to fully realize child health targets in the absence of other health investments.

2 Intervention and Experimental Design

2.1 Context and Intervention

The NGO GiveDirectly (GD) provides unconditional cash transfers (UCTs) to poor households in rural Kenya, targeting households living in homes with thatched roofs as a basic means-test for poverty. In treatment villages, GD enrolls all households meeting its thatched-roof eligibility criteria ("eligible" households); approximately one-third of all households are eligible, and existing data indicate that they are indeed poorer on average than other local households (called "ineligible households"). For instance, measured consumption per capita among eligible households in control villages, for example, is 20% lower than among ineligible households in those villages, and gaps in household wealth are even larger. No households in control villages received transfers.

Eligible households enrolled in GiveDirectly's program received a one-time series of three transfers totaling USD 1,000 (USD 1,871 PPP) via the mobile money system M-PESA, where the three tranches were disbursed over the course of roughly 11 months. This is a one-time program and it was explained that no additional financial assistance would be provided to these households after their final transfer (and in fact none was provided). In total, the transfers constituted a shock of about 75% of household income for eligible households, and of 25% of local annual GDP at the time that they were distributed.[2] Villages were phased into treatment starting in late 2014 and throughout 2015, and the bulk of the payments were sent out during 2015 and 2016. The years 2015, 2016, and 2017 are thus the years most relevant for understanding the impacts of UCTs on child mortality as many women who received cash transfers while pregnant during 2016 gave birth in 2017.

2.2 Experimental Design

Treatment assignment was randomized at two levels, the village level and the sublocation levels. Within treatment villages, all households meeting GD's eligibility requirement received the UCT. The second sublocation level of randomization provided variation in local treatment intensity. Sublocations, an administrative unit directly above the village includ-

[2]More details on program design and implementation are described in Egger et al. (2022).

ing about ten villages on average, were randomly assigned to high or low saturation status. In high-saturation sublocations, two-thirds of villages were treated, while in low-saturation sublocations, only one-third of villages were treated. This generated substantial spatial variation in treatment intensity, which we use (as in Egger et al., 2022) to estimate spillover effects. In the analysis, we both directly use the research design (in terms of village and sublocation assignment) and an approach that utilizes all variation in local cash transfer exposure, taking advantage of the idiosyncratic variation in local village assignment and the fact that villages could be located near other sublocations with different saturation assignment.

Additionally, both treatment and control villages were randomly ordered for the program and data collection visits, allowing us to assign an "experimental start date" to each village in order to explore effects related to transfer timing. These approaches are described below in Section 3.2.

3 Data and Estimation

3.1 Data

The primary data source for the analysis is a household census carried out in all treatment and control study villages in Siaya, Kenya, as part of a third endline round of data collection. This endline 3 (EL3) household census took place from April to November 2023, and builds on the project's baseline household censuses and surveys (2014-15), endline 1 household surveys (2016-17), and endline 2 household censuses and surveys (2019-22).

The analyses focus on households that were present at baseline (i.e., the start of the program) and therefore have a clearly defined treatment and eligiblity status. The household census captures all such households that remained present in the study region in 2023, which encompasses over 90% of households present at baseline. Birth information for the approximately 6% of households that had migrated out of the study area is captured through a survey of a representative subsample of the baseline households, in which the study team endeavored to track all sampled households wherever they had moved across Kenya.[3]

The pre-analysis plan (PAP) pre-specified two primary outcomes: infant (under-one) mortality and child (under-five) mortality. For each measure, we focus on children born at least one (five) years before the time of data collection in order to have a consistent population for both the numerator (children who are deceased) and the denominator (e.g., children who are deceased plus children who have survived until the age of five). We thus examine effects

[3]Note to readers: data collection of the final wave of migrant household surveys is still ongoing as of the time of writing but this represents only a tiny share (less than 2%) of all child births in the sample and thus appears unlikely to change any of the results meaningfully.

on under-five mortality among children born through the end of 2017 (i.e., in the pre-period and in the treatment period of 2015-17), while for infant mortality we can estimate impacts among children born up to 2021, which also includes the post cash transfer period of 2018-2021.

In addition to collecting information on fertility and mortality, we also seek to assign causes of death for child mortality instances via verbal autopsies (VA) using the standard World Health Organization 2022 VA Questionnaire (World Health Organization, 2022a). Verbal autopsy is considered the state-of-the-art survey-based method for determining causes of death based on self-reported information (as physical autopsies are rarely performed in the study region or many other low- and middle-income regions and vital records are incomplete), with previous literature having validated the accuracy of VA methodology and the associated machine-learning classification algorithm, including in the Kenyan study region (Gacheri et al., 2014; Serina et al., 2015; Amek et al., 2014, 2018). Though the literature also notes the limitations of VA relative to the administrative data available in wealthy country settings, it is the most accepted approach to determining cause of death in settings such as the present one.[4]

Overall survey follow-up rates were very high in the EL3 household census: over 92% of households in the 653 study villages completed the birth history and child mortality modules. The response rate in the census activity was also nearly identical and not statistically different across the treatment and control villages. In total, across the birth census and representative surveys designed to collect information on migrant households, we collected information on approximately 101,000 births. When restricting attention to births in eligible households contemporaneous with the disbursal of cash transfers (2015-17), the population under study is nearly 6,000 births.

3.2 Estimation Framework

3.2.1 Primary Econometric Specifications

We first estimate the following reduced-form specification, focusing attention on births in households meeting the eligibility criteria for transfer receipt that were present at the time of the cash transfers:

$$\begin{aligned} y_{imhvs} &= \alpha_1 Treat_v + \alpha_2 HighSat_s \\ &\quad + \lambda_{t(i)} + \rho_{g(i)} + \lambda_{t(i)} \times \rho_{g(i)} + A_m + \delta M + \epsilon_{imhvs}, \end{aligned} \tag{1}$$

[4]The VA module includes a categorization for stillbirths in approximately 15% of cases. As noted below, the results are also robust to including only live births, which is sometimes the sample considered in analyses of infant and child mortality.

where y_{imhvs} is an outcome of interest (i.e., infant mortality) for a birth i in household h, located at baseline in village v and sublocation s. The variable $Treat_v$ is an indicator for residing in a treatment village at baseline and $HighSat_s$ is an indicator for being in a high-saturation sublocation. The specification includes child year of birth fixed effects, denoted by $t(i)$, and gender fixed effects $g(i)$. We control for maternal age through the inclusion of indicator variables represented by A_m, where m denotes one of five age groups (under 20, 20-25, 25-30, 30-35, or above 35).[5] Standard errors are clustered at the sublocation level. While much of the data is derived from the birth census of the study region, we also sample households that migrated from the region and therefore weight observations by inverse sampling probabilities to ensure estimates are representative of eligible households present at the time of the transfers.

Here α_1 captures the effect of the transfers on eligible households in treatment villages (relative to control villages) from two sources: the direct effect of treatment and the effects of any within-village spillovers. The coefficient α_2 estimates cross-village spillover effects based on the research design of high versus low saturation sublocations. This estimation of cross-village spillovers is relatively coarse as it does not utilize all experimental variation, implicitly assuming that all spillovers are contained within villages and sublocations. The sum of α_1 and α_2 denotes the total effect of the transfers from all three sources (direct effects, within-village spillovers, and cross-village spillovers). This linear combination of coefficients captures the effect of the transfers on eligible households in treatment villages in high-intensity sublocations relative to eligible households in control villages in low-intensity sublocations.[6]

Equation 1 is a straightforward and intuitive benchmark, yet it does not capture the full spatial dimension of spillovers. One-third of villages in low-saturation sublocations still received transfers, and there is additional variation caused by the idiosyncratic placement of treatment and control villages as well as sublocation boundaries. In the following analysis, we thus make use of the full spatial variation induced by the experimental design. These regressions allow us to compare areas that, due to the randomization of the intervention, received more cash relative to areas that received less. While we focus on recipients, using these models we can estimate impacts on non-recipients as well, as in Egger et al. (2022).

[5] M represents a vector of indicators denoting a missing value for a given covariate, which allows us to retain observations in order to maximize statistical power. Where a covariate is missing, we set the value equal to the covariate's mean, following Egger et al. (2022).

[6] We pre-specified a version of Equation 1 that clustered standard errors at the village level and only focused on α_1 while considering α_2 as non-primary. We then adjusted standard errors to cluster at the sublocation level as our analytic focus expanded to include α_2.

We utilize the following regression specification:

$$\begin{aligned} y_{imhvs} \;=\; & \beta_1 Amt_v + \sum_{r=2}^{\bar{R}} \beta_r Amt_{v,r}^{\neg v} + \rho_1 ShareElig_v + \sum_{r=2}^{\bar{R}} \rho_r ShareElig_{v,r}^{\neg v} \\ & + \lambda_{t(i)} + \rho_{g(i)} + \lambda_{t(i)} \times \rho_{g(i)} + A_m + \delta M + \epsilon_{imhvs} \end{aligned} \tag{2}$$

The key terms are β_1, which captures the effect within treatment villages from both direct receipt of the transfers and within-village spillovers (where cash transferred to the village is captured in Amt_v), and the β_r terms, which capture the effects of cash transfers in other villages (not v) at different bands of radius r ($Amt_{v,r}^{\neg v}$) away from village v. $ShareElig$ denotes the share of baseline households eligible for a transfer in an area, for which we control to facilitate more powerful instrumental variable sets as detailed below. The other terms are as in Equation (1). As in Egger et al. (2022), the maximum radius ($\bar{R}$) is found by estimating models with varying radii, then selecting the model that minimizes the Schwarz Bayesian Information Criterion (BIC).

The analysis focuses on the "average total effect" of the cash transfers on births in recipient households in high-saturation sublocations, which is defined as:

$$\begin{aligned} \widehat{\delta y^r} = & \hat{\beta}_1 \cdot \left(\overline{Amt}_v | i \text{ born in recipient household in a high-saturation sublocation}\right) \\ & + \sum_{r=2}^{\bar{R}} \hat{\beta}_r \cdot \left(\overline{Amt}_{v,r}^{\neg v} | i \text{ born in recipient household in a high-saturation sublocation}\right) \end{aligned} \tag{3}$$

If there are no effects of cash outside of the radius $\bar{R}$ (e.g., due to ambient effects in the study area such as regional inflation), this equation estimates the average effect of the cash transfers on recipient households in high-saturation sublocations compared to a counterfactual in which no cash was distributed (to the household or their neighbors).

The amount of cash an area received is a function of the share of households that were eligible at baseline ($ShareElig$), which is endogenous, and the share of eligible households that were treated. We therefore instrument for transfer amounts using the share of eligible households treated in an area and the share of eligible households treated interacted with the share of eligible households, which is a valid instrument since the estimates control for the proportion of households that were eligible in an area at baseline.[7] To account for spatial

[7]This instrument set is similar to Egger et al. (2022), which only uses the share of eligible households treated as an instrument. By logic similar to Abadie et al. (2023), this instrument vector is efficient under homogeneous treatment effects because it captures the true first stage. The original pre-analysis plan specified the instruments from Egger et al. (2022), but we noted that this was an active research area in applied econometrics and as such we might consider estimates based on recent econometric improvements. An amendment filed before the household survey data was analyzed documented our intention to switch to the

correlation in the estimates, we calculate spatial heteroskedasticity- and autocorrelation-consistent (HAC) standard errors using a positive definite kernel up to 10 km (Conley, 2008).[8]

As noted in Section 3.1, we possess complete child mortality data for all children born through the end of 2017, as all children born in 2017 or earlier were at least five years of age at the time the birth census commenced. We thus estimate effects on child mortality for children born during the cash transfer disbursal period (henceforth the "UCT period") of 2015-17, as well as in the pre-UCT period of 2011-14 as a form of placebo check. For infant mortality, the other pre-specified primary outcome, we are additionally able to estimate impacts during the years 2018-21, which comprise the period after transfers were sent (henceforth the "post-UCT period"). In addition to examining treatment effects across these three periods, we also examine impacts by birth year in some analyses.

3.2.2 Secondary Econometric Specification

Precise date-of-birth data from the household census enables us to examine whether the effect of cash varies depending on the timing of receipt.[9] We first restrict the sample to births where the transfer period — the experimental start date plus 11 months since the three transfers were distributed over that time frame to treated households — intersected with an age group g. The relevant age groups g are defined as the amount transferred between (i) 9 months and three years before birth, (ii) within 9 months before birth to birth (treatment in-utero), (iii) birth and up to 28 days after birth (treatment as neonate), (iv) 28 days after birth to one year after birth (treatment as infant), and (v) one and three years after birth (treatment as young child). For instance, if g is in-utero, then observation i is included in estimates of the effects of receiving cash in-utero if and only if the intersection of the transfer period and individual i's birth date minus 9 months is not null.

This allow for the estimation of treatment effects among eligible households in treatment and control villages in group g as follows:

$$\begin{aligned} y_{imhvs} \;=\; & \beta_g Amt_v + \sum_{r=2}^{\bar{R}} \beta_{g,r} Amt_{v,r}^{\neg v} + \rho_g ShareElig_v + \sum_{r=2}^{\bar{R}} \rho_{g,r} ShareElig_{v,r}^{\neg v} \\ & + \lambda_{t(i)} + \rho_{g(i)} + \lambda_{t(i)} \times \rho_{g(i)} + A_m + \delta M + \epsilon_{imhvs}, \end{aligned} \tag{4}$$

revised vector of instruments.

[8] We use the kernel $K_{ij} = 1(d_{ij} < 10) \cdot \left(1 - \frac{d_{ij}}{10}\right)^2$ where d_{ij} is the Euclidean distance in kilometers between i and j. We use this kernel, rather than a uniform kernel as in Egger et al. (2022), to ensure that standard errors are not complex-valued.

[9] In addition to the analysis presented below, we also pre-specified examining whether the effect of cash on child mortality changes over time; this work is still ongoing. These analyses was pre-specified as secondary in part because we have exact date-of-birth information for most, but not all, births in the household census.

We report group-specific estimates of the average total effect as defined in Equation 3 and conduct estimation and inference using the group-specific analogues to the spatial instrumental variables equation defined above.

Intuitively, this specification examines whether there are different effects on children that were exposed to cash in different times of their life, in utero, and after birth. Any effects of pre-birth exposure are more likely to work through effect on mothers (i.e., in terms of their nutrition, health and on the medical care they received). This perspective helps to motivate the mechanisms examined below through which cash could affect child mortality, and the non-uniform age groups are motivated by key stages of pregnancy and early life.

Because the transfers were distributed over 11 months, many births fall into multiple groups and so observations are included in multiple estimates. Therefore estimates of the effects of cash when the household received the transfer when a child was in-utero may also partially reflect the effect of receiving cash as a neonate, and similarly for other groups. Estimates should therefore be interpreted as the joint effect of cash exposure in age group g and adjacent pregnancy stages since the design of the experiment does not identify "unbundled" coefficients without additional assumptions.[10]

4 Main Empirical Results

4.1 Graphical Analysis

The first results are presented in Figure 1. Panel A plots simple means of infant mortality by year across 2011-21 for two central groups of interest within the population of eligible households: births in treatment villages located in high-saturation sublocations, and births in control villages located in low-saturation sublocations. The first group comprises censused children born in villages that received the highest average intensity of cash transfers, whereas the second group encompasses censused births in villages with the lowest average intensity of transfers.

A clear pattern is evident from Panel A: while high-saturation treatment and low-saturation control villages exhibit virtually identical levels of infant mortality during the pre-period of 2011-14, a marked divergence occurs once cash transfers are distributed. Following the start of transfers, infant mortality in high-saturation treatment villages falls by 20 deaths per thousand births from 2014 to 2016, a roughly 50% decline. By contrast, low-saturation

[10]The original PAP specified a version of Equation 4 that estimated the effect of cash received in each group on mortality jointly by calculating how much cash was transferred in each group g and then estimating one regression across the sample. While in theory this could separately identify the effects of cash at different stages, a recent literature such as Callaway and Sant'Anna (2020) shows that such methods can be biased. We have opted to implement the more robust approach described here and interpret results cautiously.

control villages continue on their pre-COVID-19 trajectory of gradual improvement over time (absent a severe drought in 2017 which elevated infant mortality). Once the disbursal of transfers ends, however, infant mortality in high-saturation treatment villages swiftly returns to the rates seen in low-intensity villages.

4.2 Regression Analysis

Turning to the first regression results from the reduced-form specification, Panel B of Figure 1 displays estimates of the effect of living in a high-saturation treatment village (i.e., $\alpha_1 + \alpha_2$ from Equation 1) on infant mortality by year of birth. As can be clearly seen from Panel B, there are large impacts on infant survival during the period in which UCTs were disbursed. The coefficient estimates range from 15 to 30 infant deaths per thousand births during the three years in the UCT period, and closely mirror the simple differences in means reported in the figure. Pooling the three years of the UCT period, the effect of transfers is statistically significant at the one percent level (p-value < 0.01). Though we lack the statistical power to detect differences in coefficients across years within the transfer disbursal period, the year with the largest treatment coefficient (2017) is also the year in which a severe drought affected the region, which may suggest that UCTs reduce infant mortality most in periods of greatest economic adversity, although this is speculative. We do not detect significant treatment effects in any of the pre-period years, nor do we find significant persistent impacts in the post-UCT period.

Table 1 reports this study's main results. Column 1 displays estimates of the reduced-form effect UCTs exerted on infant (under-one) mortality for eligible households during the period in which transfers were disbursed. Infant mortality declines by 18.57 deaths per thousand births in treatment villages located within high-intensity sublocations, a result statistically significant at the one percent level. This coefficient estimate represents a 44% decline in mortality relative to the low-saturation control mean of 41.49 deaths per thousand births.[11] The total effect is driven roughly equally by the direct effect of treatment and within-village spillovers (α_1) and by the effect of cross-village spillovers (α_2). Similarly, when examining child (under-five) mortality in Column 2, there is a reduction of 17.85 deaths per thousand births during the UCT period, a 30% decline relative to the control mean.

Column 3 of Table 1 presents estimates on infant mortality from Equation 2, the instrumental variables specification which more precisely captures the full spatial dimension of spillover effects. From this specification, we find an average total effect (including direct,

[11] Distinguishing between stillbirth and neonatal deaths soon after live birth can be challenging outside of medical settings. While we include stillbirths in the main analyses, we present the main results excluding stillbirths (as determined by a verbal autopsy classification algorithm) in Appendix Table A.3 and find similar results in terms of magnitude and statistical significance.

within-village and across-village spillovers) on infant mortality in recipient households of -22.88 deaths per 1000 births across 2015-17. This represents a 55% decline in infant mortality relative to the control mean and is statistically significant at the one percent level. We estimate significant effects both from transfers within one's own village (β_1) and from other nearby villages (β_2). Column 4 reports spatial IV estimates of transfer impacts on child mortality, and estimates an average reduction of 28.98 child deaths per 1000 births during the UCT period, a 49% decline relative to control significant at the one percent level.

Both specifications yield the same striking finding: infant and child mortality declines by nearly half in this setting due to the disbursal of cash transfers. Figure 3 presents estimates for the periods outside of the time of transfer disbursal — the pre-period (2011-14) and the post-UCT period (2018-21). Reassuringly, we not find any significant differences in the pre-period, indicating balance. Furthermore, once cash transfers cease, so do the impacts on infant mortality. In the post-UCT period, we do not estimate significant differences in infant survival between treatment and control villages, though our standard errors cannot rule out moderate impacts in either direction.[12] We furthermore do not find significant changes over time in infant mortality among ineligible households who did not receive cash transfers (and were richer at baseline).

Turning to examining differences in the effect of cash by timing of exposure, we find by far the largest treatment effects among children whose household was receiving cash in the month of their birth. Figure 2 illustrates treatment effects in high-saturation sublocations for five groups — (i) children whose household received cash 9 months to three years before their birth ("pre-pregnancy"), (ii) 9 months before birth to birth ("in-utero"), (iii) birth and up to 28 days after birth ("birth month"), (iv) 28 days after birth to one year after birth ("infant $\geq$ 1 month"), and (v) one to three years after birth ("child $\geq$ 1 year").[13] Infant mortality declines by 43.77 deaths per thousand births among children receiving cash in their birth month — a finding statistically significant at the one percent level which basically implies that mortality for these infants dropped to the low levels seen in industrialized nations. We observe more modest mortality declines for infants whose households received cash in-utero (-13.95 deaths per thousand births) and little to no effects for children whose family received cash before their pregnancy had occurred (echoing the lack of persistent effects in the post-cash transfer period noted above) or after their neonatal period. This stark heterogeneity by transfer timing indicates that factors present in the neonatal period may interact with the cash to produce particularly large declines in child mortality.

[12] We do not estimate effects on under-five mortality for 2018-21 as children born after 2017 may not have reached their fifth birthday at the time of the census activity, as noted above.

[13] Figure 2 presents results using the spatial IV specification presented in Equation 4; we present results from a reduced-form specification in Appendix Figure A.2.

The mortality effects we estimate during the UCT period are broad-based across various characteristics of the child and mother. As Appendix Figure A.1 reports, the impacts on mortality are nearly identical by child gender and similar by child birth order. We find some suggestive evidence of heterogeneity by maternal age, with coefficient estimates about 75% larger for older (aged 25 and over) than younger mothers, though we cannot precisely distinguish differences between the two groups, which both exhibit statistically and economically significant declines.

4.3 Effects by Cause of Death: Verbal Autopsies

To ascertain the likely cause of death for each of the 4,696 recorded under-five deaths in the household birth census, trained field staff conducted verbal autopsies (VAs) using the World Health Organization's 2022 questionnaire. The field team was able to conduct VAs for 92% of child deaths across the period of study, and of these 82% were collected from a family member present at the time of the death, and thus is likely to be particularly knowledgeable about the circumstances. Each VA took an average of 22 minutes to collect. Following collection of the VA, a likely cause of death was assigned using the Institute for Health Metrics and Evaluation's SmartVA algorithm, which utilizes the Tariff 2.0 method for machine-learning classification of VAs and which was designed and validated with the Population Health Metrics Research Consortium Gold Standard VA database (Institute for Health Metrics and Evaluation, 2025).[14]

As pre-specified, we focus on five cause of death (COD) groups. The broad COD distribution seen in this study's VA data is reassuringly similar to other studies in western Kenya (Amek et al., 2014). The leading COD group in control villages is maternal and neonatal causes, which encompasses individual causes such as birth asphyxia, congenital malformation, and death from preterm delivery. Maternal and neonatal causes comprise 38% of deaths with non-missing causes in low-saturation control villages. The second-largest COD group, encompassing 37% of deaths with non-missing causes, is communicable and nutritional diseases such as malaria and malnutrition. Other COD groups include respiratory diseases such as pneumonia (13%), non-communicable diseases (10%), and injuries (2%). A sixth category encompasses completed VAs for which SmartVA was unable to determine a likely cause due to missing or inconsistent answers (18% of completed VAs).

Figure 4 illustrates treatment effects on infant mortality by cause of death estimated using Equation 2. The largest reduction in mortality by far is in deaths from maternal and neonatal causes: we estimate a drop of 10.68 deaths in this category per 1000 births, representing

[14]The likely cause of death is a standardized one as defined by the International Classification of Diseases, tenth edition (ICD-10).

a nearly 70% decline relative to the control mean.[15] Across other CODs, coefficients are almost always negative (the exception is non-communicable diseases, for which the point estimate is near zero), but mortality reductions across all other CODs combined amount to just half the decline seen within maternal and neonatal causes alone. We do not find any evidence that VAs were differentially likely to be undetermined or absent in treatment villages, conditional on a death (Appendix Table A.4). Yet given the sharp decline in overall mortality we do find a drop in the overall death rates for undetermined or absent causes in treatment households, accounting for the remainder of the total mortality reduction.

4.4 Comparison with Non-Experimental Variation

The existence of birth census data over the period 2011-2023 enables us to benchmark the experimental cash transfer treatment effect estimates against several dimensions of non-experimental variation in economic circumstances.

One form of variation is cross-sectional: transfer-eligible households were notably poorer on average than ineligible households (as noted above), as measured by either consumption or assets. As illustrated by Figure 3, the birth census indicates that infants born to ineligible mothers in control villages were substantially less likely to die in their first year of life than children from eligible households in these same villages — but that in the 2015-17 period of transfer disbursal, the treatment effect from transfers is sufficient to close the infant mortality gap between eligible and ineligible households. Across 2015-17, under-five mortality among control ineligible births was 12.82 deaths per thousand lower than among control eligibles, a difference statistically significant at the five percent level. This socioeconomic gradient in mortality also holds even when controlling for birth characteristics such as maternal age and gender, and a similar pattern is evident when examining under-one mortality or all the years across the birth census as opposed to the main cash transfer period. Of course, there are many other interpretations of this pattern and many potential omitted variables or confounders — which is precisely why the current study exploits experimental variation to estimate causal impacts — but we note that these cross-sectional observational patterns remain broadly consistent with the large experimental estimates described above.

Child mortality displays a similarly high degree of sensitivity to inter-temporal changes in economic conditions. The information on exact month of birth for most children in the birth census allows us to estimate month of birth effects for both infant and child mortality (controlling for several other characteristics such as birth year, gender, and maternal age).

[15] This result is consistent with the finding that neonatal mortality, i.e., in the first 30 days of life, declines by 13.24 deaths per thousand births in the census data, a nearly 60% reduction relative to the control mean (see Appendix Table A.2). We find similar results for both the VA and neonatal mortality analyses when utilizing Equation 1.

Burke et al. (2019) define the pre-harvest "lean season" in rural Kenya as encompassing the months of April through August, after which the main grain harvest occurs. The analysis of the birth census during the pre-COVID period of 2011-19 reveals a meaningful degree of seasonality in infant mortality: rates of mortality for infants born in August (often the very peak of the lean season just before the harvest) are 21.86 deaths per thousand births higher than for infants born in the next month (Appendix Figure A.3). This difference is statistically significant at the one percent level, and forms part of a larger pattern of mortality rising as the lean season progresses: from May through August, infant mortality rates rise by an average of 6.44 deaths per thousand each month. Under-five survival rates display similar patterns (Appendix Figure A.4).

Mortality in the birth census also responds to major aggregate economic shocks: as is evident in Panel A of Figure 1, infant mortality rises across the board in 2017, a year in which a severe drought affected rural Kenya (BBC, 2017). Furthermore, the COVID-19 pandemic struck Kenya in early 2020, leading to a large recession as well as access to many public services. The birth census data indicates that infant mortality nearly doubled from 2019 to 2020, with the largest jump occurring precisely in April 2020, the first full month of pandemic related lockdowns. Mortality remains elevated in 2021, a year in which the COVID-19 pandemic continued to affect Kenya and another severe drought impacted the region (World Bank, 2022).

In short, across multiple sources of variation, child survival appears to respond substantially to economic conditions in rural Kenya, with non-experimental differences exhibiting the same sign and similar magnitudes as the main cash transfer experimental results.

5 Mechanisms and Behavioral Change

The main finding of this study is that the disbursal of cash transfers leads to the halving of infant mortality, and that this sharp reduction in mortality is concentrated in birth complications and neonatal deaths. In this section, we turn to exploring potential drivers of these mortality declines. We do so by taking advantage of the numerous waves of detailed household surveys the study team has conducted in the study region across over a decade, ranging from baseline data in 2014 to the third endline survey (collected during 2023-2025). Here we document that multiple channels and behavioral mechanisms appeared to contribute to the overall effect, while noting that it is challenging to decompose exactly how much of the total child mortality reduction can be attributed to any single mechanism, absent strong or untestable assumptions.

5.1 Healthcare Access

Here we present multiple analyses that indicate that improved access to healthcare — in particular, to birth delivery services — playing an important role in accounting for the observed gains in child survival.

The first two pieces of evidence were previously presented above. First, when we examine treatment effects by the timing of cash transfer exposure, the impacts on mortality are largely concentrated among infants whose families were receiving a transfer in the month when the child was delivered or when the child was a neonate (i.e., the first 28 days of life). Second, when investigating changes in mortality by cause of death using the verbal autopsy data, by far the largest reduction is within the neonatal and maternal cause category (e.g., death from preterm delivery, birth asphyxia, other delivery complications, congenital malformation, etc.). These analyses suggest that conditions in the immediate antenatal and neonatal period play a critical role, as neonatal and maternal causes of death in particular are largely preventable if quality healthcare is utilized (World Health Organization, 2015, 2022b).

We next examine direct evidence for these mechanisms collected through the endline 3 long-form household survey, which asks a representative sample of households detailed questions about their antenatal, delivery, and postnatal healthcare utilization during pregnancies which took place between 2015-17.[16] Figure 5 reports estimated effects of the transfers on five pre-specified, World Health Organization (WHO)-recommended metrics of healthcare utilization calculated using Equation 2. The results indicate that mothers in eligible households were 30 percentage points more likely to give birth in a hospital if they received the UCT in a high-saturation sublocation, about a 70% increase over the control mean of 44%. These findings are consistent with past research which has argued that cost is a substantial barrier to institutional delivery in Kenya (Njuguna et al., 2017). Point estimates also suggest that the rate of pregnancies where the WHO-recommended number of at least four antenatal visits occurred may have increased, although this result is not statistically significant at traditional levels.

To further examine the potential role of healthcare access, we assemble detailed data on travel times to health facilities across the study region. The coordinates of every registered health facility nationwide were obtained from the Kenya Master Health Facility List (KMHFL), a database developed by the Kenya Ministry of Health and the United States Agency for International Development (USAID). We then surveyed the facilities to collect

[16]Note to readers: the second and final wave of this survey is still in the field; overall data collection is 70% complete for a representative household sample and we present results from the data we have so far. Future versions will have complete household behavioral and health spending data.

information about the care that they offer, including whether the facility is staffed by a physician. We augment this data using surveys of health facilities that we collected among facilities that household respondents had reported visiting but were not present in the KMHFL data. Walking travel times are then estimated using geospatial data from OpenStreetMap and an estimation procedure previously utilized in the study region (Ouko et al., 2019). We focus on walking times as Ouko et al. (2019) find that the majority of women walk to receive care in the study region, and validate these estimated travel times using directly-measured travel times for a subset of roads and paths collected by sending enumerators equipped with speedometers to traverse the study region. These data enable the estimation of Equations 1 and 2 interacted with bins of travel time.

We focus on distance to a physician-staffed facility, rather than to a clinic or hospital per se, because we find that this measure is more reliably measured. For instance, during surveys of facilities, we found that staff respondents reliably knew if a physician was employed at the facility, but they were often confused by the tiered numbering system that Kenya uses to classify facilities. And these official tiers are only imperfectly correlated with the reality on the ground: while all level 4 facilities should be hospitals with a physician employed, our surveys find that only 70% reported having a doctor and 5% reported not offering any birth delivery services. We therefore focus on having a doctor employed at the facility as a proxy for quality delivery care. This approach is supported by the findings of Okeke (2023), who shows that this is a key determinant of care quality in another rural Sub-Saharan African health facility setting.

The results of this analysis indicate that the reductions in infant mortality observed during the disbursal of cash transfers were largely concentrated in households located close to a physician-staffed health facility. Table 2 reports a decline in infant mortality of 35.51 deaths per thousand births among households with below median walk times to a physician-staffed facility. By contrast, households with above median walking time, which must travel an average of 81 minutes (about 50 minutes more than those with below-median travel times), exhibit a substantially smaller and significantly insignificant reduction of 12.07 deaths per thousand births. As a placebo check, note that we observe no such heterogeneity in the pre-cash transfer period. We also observe similar patterns when using alternate travel time measures such as travel time by automobile or motorcycle, or when using the study-measured travel times. The results hold when selecting controls using double partial-out LASSO to account for place-based factors — such as population density, distance to towns and roads, baseline village wealth, malaria suitability, and rainfall, among others — and importantly, the main interaction effect between cash treatment and below median travel time is robust to including the LASSO-selected variables interacted with treatment, and remains statistically

significant at the 10% level.

While some alternative interpretations remain possible, since the location of health facilities staffed with doctors is not randomly allocated, across multiple analyses a pattern clearly emerges: cash matters most for infant survival when delivered in the neonatal period to households near established health infrastructure with good delivery services. Rates of hospital delivery rise, and rates of death from neonatal complications readily prevented by appropriate obstetric care fall. These results are consistent with prior literature including Okeke (2023), which finds in an experiment that access to a physician in rural Nigeria significantly reduces neonatal mortality. In the present study, access widens when cash transfers are disbursed: a study contemporaneous with the UCT program under analysis found that the average cost of an institutional delivery in Kenya was KSH 8,812 (USD 106), which can be cost-prohibitive for the relatively poor households eligible for the UCT (Institute for Health Metrics and Evaluation, 2014). Yet even having ample cash in hand may not be enough if hospitals and doctors are difficult to reach, given the challenges of accessing on-demand rapid transportation in Kenyan villages (Floyd et al., 2020).

5.2 Fertility Patterns

In this section we explore impacts on fertility. We find modest increases in general fertility but no evidence that the characteristics of households and mothers giving birth changed. Specifically, Table 3 reports estimates from Equation 2 indicating that among recipient households the share of women giving birth rose by 13% relative to the low-intensity village mean.[17] Fertility patterns are indistinguishable by treatment status in the 2011-14 pre-period and again in the 2018-21 period following the disbursal of transfers. Similar results are obtained in the reduced-form using Equation 1.

Several approaches indicate that the characteristics of families giving birth did not change as a result of the transfers in 2015-17. To analyze family characteristics, we turn to earlier long-form household surveys conducted with a representative sample of households as the birth census itself was parsimonious. First, using the sample of transfer-eligible households surveyed at baseline, households with a birth over 2015-17 do not significantly differ by treatment status across six key baseline socioeconomic characteristics: education, age, marital status, income, assets, and household size (Appendix Table A.5).

Next, we use these six baseline household characteristics to predict the probability of giving birth over 2015-17 for adult females censused in EL3, and then examine whether treatment status is associated with the predicted probability of birth among women who

[17]The population of women for which we calculate birth rates is all censused women who were adults in 2023. We did not collect age data for women who did not give birth, and hence are unable to create age-adjust fertility rates.

gave birth during the key 2015-17 period. We find in Column 4 of Table 3 that the predicted probability of birth among actual mothers does not significantly differ by treatment status in the transfer disbursal period. We further find in Column 5 that treatment status is not associated with predicted birth probability among actual mothers in the post-transfer period of 2018-21. In both periods, the coefficient estimate on treatment is in fact slightly negative in sign though not significant. Predicting birth probability using LASSO from a set of 261 baseline household- and village-level characteristics, as we do in Appendix Table A.6, yields similar results indicating an absence of treatment impacts on the characteristics of women who gave birth (in terms of their individual predicted birth probabilities).

Last, we run Equation 2 augmented with additional baseline controls to test whether the inclusion of household baseline characteristics affects our results (Appendix Table A.7). These regressions are performed only on households surveyed in the detailed baseline survey. The first column only includes pre-specified controls (e.g., year of birth fixed effects, birth gender, and mother age group), results in this sample of households also indicate sharp infant mortality declines due to treatment. In Column 2, we then include controls for the six baseline socioeconomic characteristics previously considered (education, age, marital status, income, assets, and household size) along with a full set of interactions between these variables. The infant mortality results are again largely unaffected. In Column 3, we augment Equation 2 to include controls selected by PDS LASSO, with a total of 261 baseline household- and village-level characteristics available for selection, and the results are also unaffected, with point estimates if anything slightly increasing in size.

A natural conclusion from the above analyses is that any changes in the characteristics of the mothers giving birth (at least on the basis of observables) in the treatment group were at most modest and cannot readily explain the sharp infant mortality declines that we document.

5.3 Maternal Labor Supply

We documented in a previous paper (Egger et al., 2022) positively signed but imprecise increases in labor supply from the cash transfers among recipient households. Families with a late-stage pregnancy or newborn at the time transfers are disbursed, however, may exhibit a different pattern, particularly with respect to female labor supply. Increased wealth from the transfers may enable women to work less at critical periods for fetal and newborn health, enabling them to reallocate time to rest or other health-producing activities. This channel may be particularly relevant in Kenya as female labor force participation, at 72% of adult women, nearly equals rates of participation by men (World Bank, 2025).

Detailed data on hours worked by gender combined with the birth census enable analysis

of heterogeneity in the cash transfer's effects on labor supply by gender and the presence of a third-trimester pregnancy in the household. Table 4 reports results separately by gender from Equation 1 augmented with an interaction between treatment status and an indicator for a third-trimester pregnancy present during transfer disbursal. Estimates encompass only those households selected for the long-form surveys which inquired about labor supply; as such, the sample size here is reduced compared to the full birth census.

We find quantitatively substantial heterogeneity in labor supply effects among women by whether a late-stage pregnancy is present within the household. Women in recipient households in high-intensity villages (i.e., treated villages within high-saturation sublocations) without a third-trimester pregnancy exhibit no change in hours of labor supplied as a result of the transfers. When a late-stage pregnancy is present, however, female labor supply in recipient households falls by 14.60 hours a week — a 57% decline over the control mean that is statistically significant at the ten percent level. By contrast, we find considerably more muted and statistically insignificant interactions between household pregnancy status and treatment status among men in recipient households, suggesting that it is largely women in particular who are able to temporarily exit the labor force when a pregnancy is present due to the cash transfers.[18]

These results provide suggestive evidence that the cash transfers may have enabled some women to reallocate time from potentially strenuous labor to rest or other activities more beneficial for fetal and child health. Reductions in labor supply during pregnancy as a channel lowering infant mortality would be consistent with prior work highlighting the importance of parental time in the production of child health and development (Mayer, 1997; Kruger, 2007; Miller and Urdinola, 2010), and accord with studies in other LMIC settings indicating an association between cash transfers and reductions in maternal labor supply (Novella et al., 2012; Amarante et al., 2016; Garganta et al., 2017).

5.4 Nutrition

Another important consideration is the role of nutrition. The first endline survey conducted in 2016-17 indicates that food security for children significantly rose for recipient households. A pre-specified index of child food security, which encompasses household survey questions on whether children skipped meals, went to bed hungry, and went entire days

[18]We lack the statistical precision to precisely detect interactions between the presence of a newborn less than three months of age and treatment status, but note that the coefficient estimate would also represent a 20% decline in female labor supply relative to the control mean, a pattern that is again not seen among men. Further heterogeneity analysis based on the gender of the cash recipient within the household could be valuable to further understand these dynamics. However, we observe only the name of the individual that the cellphone used to receive cash transfers is registered to. Cellphones are often shared by households, so we cannot reliably determine if the cash recipient was male or female in most cases.

without food over the past week, increased by 0.17 standard deviation units on average among recipients as estimated using Equation 2.[19] This result, reported in Table B.6 of Egger et al. (2022), is statistically significant at the five percent level, and similar results are found using Equation 1.

5.5 Cost-Effectiveness Implications

The prior discussion of potential mechanisms prompts two important questions. First, what do the estimated child mortality benefits imply about the welfare gains from UCTs versus other forms of aid that donors may invest in? Second, are cash transfers a cost-effective tool for reducing child mortality rates, and can the mechanism analyses reported above help guide efforts to target transfers to those most likely to benefit?

It is straightforward to estimate the number of child deaths averted due to the UCT program under study. First, we focus on recipient households and births during the 2015-17 period of transfer disbursal, as estimated effects after 2017 for this group are near zero and so are estimated impacts among non-recipients across all years. Second, we estimate the number of child deaths averted among recipient households during the transfer disbursal period using the following "back-of-the-envelope" calculation:[20]

$$\begin{aligned}&(\text{Estimated average treatment effect on recipient child mortality})\\&\times(\text{Number of births among recipient treatment households during 2015-17})\,.\end{aligned} \tag{5}$$

As from Equation 2 we find that the average treatment effect for child mortality is -26.63 deaths per thousand births (Table 1, Column 4) and there were 3,397 births in treatment villages across 2015-17, we estimate that approximately 90 child deaths were averted due to the UCTs. This is a substantial reduction: based on the low-saturation control village mean of 56.82 deaths per thousand births, we would have expected approximately 193 child deaths to have occurred across this period, again indicating that the cash transfer treatment led to a drop of nearly half in the number of child deaths.

The leading approach to estimate the welfare gains of these mortality reductions is to value lives saved by recipients' value of a statistical life (VSL). Revealed preference estimates of VSL, or consumer demand for mortality risk reductions, tend to be low among populations with income levels similar to this study. In fact, most studies we are aware of find values

[19] As noted earlier, Equation 2 in this paper departs slightly from Equation 2 in Egger et al. (2022) as the econometric literature has since advanced, but results are similar regardless of the econometric specification used.

[20] We focus primarily on child mortality as opposed to infant mortality due to its relevance for policymakers and foundations. For example, the United Nations Sustainable Development Goals refer specifically to under-five mortality but not to under-one mortality (United Nations, 2015). As previously noted, the estimates for child and infant mortality are very similar.

below $5,000 (Killeen, 2025; Kremer et al., 2011; Berry et al., 2020). This would imply welfare gains from the mortality reductions of less than $500,000. However, GiveWell, a non-profit charitable giving advisor applies "moral weights" which (based on our understanding) value averting an under-5 death at 116 times doubling consumption, implying a value of $87,935 in this setting. This would value welfare gains from the child mortality reductions at approximately USD PPP 7.96 million, a sizable portion of the expenditure on UCTs (USD PPP 25.75 million/USD 10.75 million nominal).

To account for the wide range of VSL estimates among the population, we report the estimated benefits of a $1,000 investment in UCTs or a leading health intervention, malaria medication, by VSL in Figure 6.[21] We also plot an estimate of the posterior distribution of VSL in this population obtained using Bayesian hierarchical meta analysis of the most relevant revealed preference estimates as well as Redfern et al. (2019), which estimates a VSL of over $55,000 using a stated preference approach and informed GiveWell's moral weights. We focus on revealed preference estimates since they are based on choices with real stakes and less prone to social desirability bias, but we include the estimate from Redfern et al. (2019) due to its policy importance and to capture the fact that stated preference estimates are typically much higher (Hamory et al., 2021). The plot reveals that the estimated welfare gains of the UCT program dominate Malaria medicine for about 75% of the distribution of VSLs, although the mean estimate of gains from malaria medicine are about $300 higher. This holds because the Redfern et al. (2019) induces a substantial rightward skew in the distribution. UCTs yield higher estimated welfare for the full 95% confidence interval of VSLs if the study is excluded. Details of the estimation of the distribution of VSLs and the welfare analysis are presented in Appendix B.

The welfare gains of the UCT program are mainly driven by consumption gains, and there is a proportionally small difference in welfare when accounting for mortality reductions across most VSLs in the distribution. This is true because the UCTs were not targeted to pregnant women, so the cost per life saved is high. We therefore also plot estimated welfare gains if UCTs were targeted to pregnant women, assuming a consumption multiplier of 1 rather than 2.5 since general equilibrium effects of cash transfers would likely be smaller. The welfare gains from mortality reductions are much larger in this scenario, and targeted cash transfers yield higher estimated welfare than untargeted transfers for VSL values above about $50,000 in this scenario. In cases where UCTs can be distributed broadly enough to produce the general equilibrium effects documented in Egger et al. (2022), we therefore estimate that UCTs are among the socially efficient forms of aid for about 75% of the distribution of

[21] We use GiveWell's estimate of the cost per life saved of $4,304. We selected this program because it was GiveWell's top listed charity in March, 2025.

posterior VSL values. But when cash transfers cannot produce these general equilibrium effects, the results suggest that UCTs targeted to pregnant women or health interventions may yield larger welfare gains.[22]

A second question related to the welfare implications of the child mortality reductions is how cost effective UCTs are compared to other health interventions. As documented in Killeen (2025), the prior welfare analysis guides optimal decision making if donors are deciding between various programs to benefit a population such as one in the study, but if funds are earmarked for safety, economic theory does not support the use of recipients' VSL. We therefore benchmark the cost per death averted to other health interventions in Sub-Saharan Africa considered cost-effective by the World Health Organization and other health experts. In total, the UCT program under study disbursed USD PPP 25.75 million, and given that we estimate the program averted 90 child deaths, this implies a cost of USD PPP 285,000 per death averted. However, this calculation is likely highly conservative when thinking about cost-effectiveness for at least two reasons. First, unlike most health interventions, the UCTs were intended to affect non-health outcomes such as raising household consumption. Second, the UCT was not targeted towards pregnant women or households with small children. Given the detailed household data we possess, it would be possible to target transfers more narrowly based on factors such as whether a woman is pregnant.

As a first pass, we consider targeting transfers to households with women in the third trimester of pregnancy.[23] Disbursing UCTs to these households would cost a total of USD PPP 1.6 million (USD 650,000 nominal). Calculating the number of deaths averted under this scenario is challenging due to at least two opposing factors. On one hand, restricting transfers to a subset of households reduces the impact of spillovers from other treated households. On another hand, targeting cash to particular subpopulations may result in larger treatment effects among those high-impact groups. We found earlier in Figure 2, for example, that mortality was virtually eliminated for children whose households were receiving the UCT in the month when they were born, in contrast to other children who saw more modest effects. For simplicity and to take a middle-ground approach, we simply utilize the average treatment effect for all recipients (-26.63 deaths per thousand births) and apply this to the back-of-the-envelope calculations for neonatal targeting.

Targeting UCTs to women in the third trimester of pregnancy under these assumptions would cost USD PPP 38,016 per child death averted, net of household consumption gains.

[22] Appendix Figure A.5 reports the results of a decision theoretic model which yields similar results. Namely, UCTs minimize median regret, but Malaria medicine narrowly minimizes Bayesian regret when Redfern et al. (2019) is included in posterior estimates.

[23] We consider treatment effect estimates for those exposed to cash in their birth month for this analysis, since almost all of these households received a cash transfer shortly before the child's birth month.

Without accounting for direct consumption increases, the cost per death averted would be USD PPP 86,470. We can then benchmark these calculations to 11 other cost-effectiveness estimates for health interventions in Sub-Saharan Africa sourced from a recent review article (Zeng et al., 2018) and World Bank report (Jamison et al., 2006). We restrict attention to studies that contained estimates of cost per child death averted (as opposed to cost per DALY averted or alternative metrics) and which contained geographic data that enabled PPP conversion. These estimates, along with ours, are presented in Figure 7. In 2023 PPP USD, the cost per death averted estimates range from $900 to $119,776, and all these interventions were deemed "cost-effective" by their authors. Net of direct consumption impacts, targeting UCTs to neonates costs approximately as much as interventions such as newborn home visits and strengthening community health systems. Furthermore, even without accounting for consumption increases UCTs are not an outlier among these interventions.[24]

A final point is that the previous analysis of mechanisms, particularly with respect to health facility access, may be able to inform the targeting of UCTs to the children that could most benefit. Targeting to households where a women is about to give birth located relatively close to physician-staffed health facilities, for example, would result in a cost per death averted estimate of USD PPP 31,246. This is far below the USD PPP 65,000 under-five death VSL value in the GiveWell-commissioned study of preferences in Kenya and Ghana, and below the average cost per death averted among the Sub-Saharan African health interventions we analyze. In short, under rough but plausible assumptions, targeted UCTs may be comparable in cost to a number of health interventions considered cost-effective in the global health literature, and this cost-effectiveness may be amplified if UCTs are targeted near places with already-established health infrastructure.

6 Conclusion

A large-scale unconditional cash transfer program in rural Kenya led to a sharp drop of one half in infant mortality. The largest mortality reductions were observed in households located near physician-staffed health facilities and those receiving the cash in the months around a child's birth. Concomitant with large mortality reductions among these households, we find that cash leads to higher rates of hospital deliveries for households who live near

[24]Figure 7 also displays the average cost per death averted (USD PPP 12,841) across 37 WHO-recommended maternal and child health interventions in East Africa as estimated by Stenberg et al. (2021). Across interventions and scenarios, the cost per death averted for the WHO-recommended interventions ranges from USD PPP 27 to USD PPP 233,952. UCTs targeted to women about to give birth and with neonates living near health infrastructure would fall in approximately the 90th percentile of cost per death averted across these interventions (i.e., in the top 10% of most cost-effective interventions), which include extremely low-cost and high impact programs such as vitamin A supplements, family planning, oral rehydration salts, and vaccination.

health facilities. A rough calculation suggests that transfers targeted to these households are similarly cost-effective to a number of child health interventions recommended by the World Health Organization, and represent benefits beyond the direct and spillover household consumption and wealth gains already documented by Egger et al. (2022).

The large magnitude of the child survival gains documented here underscores the fact that infant and child mortality appears extremely sensitive to economic conditions in the rural Kenyan setting such as the one under study. The socioeconomic gradient in mortality appears quite steep in rural Kenya. Not only do UCTs halve infant and under-5 child mortality, but we also document that mortality rates vary dramatically by household baseline wealth in the cross-section, as well as inter-temporally by the agricultural harvest season. Furthermore, once the UCTs cease, the mortality reductions do not persist, indicating that contemporaneous income is critical. Interventions such as UCTs can result in substantial child survival gains, but they likely will need to be sustained over time rather than in a one-time program to generate persistent child survival gains. Of course, that does not mean that even one-time cash transfers are worthless, as indicated by the large number of approximately 90 children who we estimate survived in the study sample who otherwise would not have.

Another lesson from this analysis is that cash alone, even if sustained, may not be enough to produce the reductions in infant mortality needed to fully realize public health targets. Health infrastructure appears to complement cash in the present study, with substantially larger reductions in mortality displayed among households living nearer to health facilities with a doctor present. While giving birth in a hospital is expensive for poor households, even with ample cash in hand it may be arduous or near impossible to deliver in a facility if there is no hospital nearby. In the short-run, one option for policymakers who wish to maximize the impact of cash transfers on mortality could be to target cash transfers to places with established health infrastructure. However, this may be undesirable for equity and political economy reasons. In the longer-term, investments in improved health infrastructure, particularly child delivery services, combined with cash transfers may be an attractive approach to ensure the equitable achievement of child survival goals.

References

Abadie, Alberto, Jiaying Gu, and Shu Shen, "Instrumental variable estimation with first-stage heterogeneity," *Journal of Econometrics*, 3 2023, p. 105425.

Amarante, Verónica, Marco Manacorda, Edward Miguel, and Andrea Vigorito, "Do Cash Transfers Improve Birth Outcomes? Evidence from Matched Vital Statistics, Program, and Social Security Data," *American Economic Journal: Economic Policy*, 2016, *8* (2), 1–43.

Amek, Nyaguara O., Annemieke Van Eijk, Kim A. Lindblade, Mary Hamel, Nabie Bayoh, John Gimnig, Kayla F. Laserson, Laurence Slutsker, Thomas Smith, and Penelope Vounatsou, "Infant and child mortality in relation to malaria transmission in KEMRI/CDC HDSS, Western Kenya: validation of verbal autopsy," *Malaria Journal*, January 2018, *17* (1), 37.

_ , Frank O. Odhiambo, Sammy Khagayi, Hellen Moige, Gordon Orwa, Mary J. Hamel, Annemieke Van Eijk, John Vulule, Laurence Slutsker, and Kayla F. Laserson, "Childhood cause-specific mortality in rural Western Kenya: application of the InterVA-4 model," *Global Health Action*, December 2014, *7* (1), 25581. Publisher: Taylor & Francis _eprint: https://doi.org/10.3402/gha.v7.25581.

Apouey, Benedicte and Andrew E. Clark, "Winning big but feeling no better? The effect of lottery prizes on physical and mental health," *Health Economics*, 2015, *24* (5), 516–538.

Baird, Sarah, Craig McIntosh, and Berk Özler, "When the money runs out: Do cash transfers have sustained effects on human capital accumulation?," *Journal of Development Economics*, September 2019, *140*, 169–185.

Barham, Tania, "A healthier start: The effect of conditional cash transfers on neonatal and infant mortality in rural Mexico," *Journal of Development Economics*, January 2011, *94* (1), 74–85.

BBC, "Kenya's Uhuru Kenyatta Declares Drought a National Disaster," 2017. Accessed: 2025-02-23.

Berry, James, Greg Fischer, and Raymond P. Guiteras, "Eliciting and Utilizing Willingness to Pay: Evidence from Field Trials in Northern Ghana," *Journal of Political Economy*, April 2020, *128* (4), 1436–1472.

Blattman, Christopher and Paul Niehaus, "Show Them the Money: Why Giving Cash Helps Alleviate Poverty," *Foreign Affairs*, 2014, *93* (3), 117–126. Accessed: 2025-02-23.

_ , Nathan Fiala, and Sebastian Martinez, "The Long-Term Impacts of Grants on Poverty: Nine-Year Evidence from Uganda's Youth Opportunities Program," *American Economic Review: Insights*, September 2020, *2* (3), 287–304.

Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel, "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets," *Quarterly Journal of Economics*, 2019, *134* (2), 785–842.

Burstein, Roy, Nathaniel J. Henry, Michael L. Collison, Laurie B. Marczak, Amber Sligar, Stefanie Watson, Neal Marquez, Mahdieh Abbasalizad-Farhangi, Masoumeh Abbasi, Foad Abd-Allah et al., "Mapping 123 million neonatal, infant and child deaths between 2000 and 2017," *Nature*, October 2019, *574* (7778), 353–358.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, December 2020.

Cesarini, David, Erik Lindqvist, Robert Östling, and Björn Wallace, "Wealth, health, and child development: Evidence from administrative data on Swedish lottery players," *The Quarterly Journal of Economics*, 2016, *131* (2), 687–738.

Conley, Timothy G., "Spatial Econometrics," in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, second edition ed., Vol. 7, Houndsmills: Palgrave Macmillan, 2008, pp. 741–47.

Crosta, Tommaso, Dean Karlan, Finley Ong, Julius Rüschenpöhler, and Christopher R Udry, "Unconditional Cash Transfers: A Bayesian Meta-Analysis of Randomized Evaluations in Low and Middle Income Countries," Working Paper 32779, National Bureau of Economic Research August 2024.

Cutler, David, Adriana Lleras-Muney, and Tom Vogl, "Socioeconomic Status and Health: Dimensions and Mechanisms," in "The Oxford Handbook of Health Economics," New York: Oxford University Press, 2012.

Deaton, Angus S. and Christina Paxson, "Mortality, Income, and Income Inequality over Time in Britain and the United States," in David A. Wise, ed., *Perspectives on the Economics of Aging*, University of Chicago Press, 2004, pp. 247–286.

Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker, "General Equilibrium Effects of Cash Transfers: Experimental Evidence From Kenya," *Econometrica*, 2022, *90* (6), 2603–2643. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA17945.

Evans, David and Katrina Kosec, "Cash Transfers and Health: It Matters When You Measure and It Matters How Many Health Care Workers," https://blogs.worldbank.org/en/impactevaluations/cash-transfers-and-health-it-matters-when-you-measure-and-it-matters-how-many-health- 2016. Accessed: 2025-02-23.

Floyd, Jessica R., Joseph Ogola, Eric M. Fèvre, Nicola Wardrop, Andrew J. Tatem, and Nick W. Ruktanonchai, "Activity-Specific Mobility of Adults in a Rural Region of Western Kenya," *PeerJ*, 2020, *8*, e8798.

Gacheri, Susan, Hillary Kipruto, Evans Amukoye, Jane Ong, Ellen M H Mitchell, Joseph Sitienei, Richard Kiplimo, and Charles Muturi, "Performance of clinicians in identifying tuberculosis as cause of death using verbal autopsy questionnaires in Siaya County, Kenya," *African Journal of Health Sciences*, 2014, *27*, 232–238.

Garganta, Santiago, Leonardo Gasparini, and Mariana Marchionni, "Cash Transfers and Female Labor Force Participation: The Case of AUH in Argentina," *IZA Journal of Labor Policy*, 2017, *6*, 1–22.

Guldi, Melanie, Amelia Hawkins, Jeffrey Hemmeter, and Lucie Schmidt, "Supplemental Security Income for Children, Maternal Labor Supply, and Family Well-Being: Evidence from Birth Weight Eligibility Cutoffs," *Journal of Human Resources*, 2024, *59* (4), 975–1010.

Hamory, Joan, Edward Miguel, Michael Walker, Michael Kremer, and Sarah Baird, "Twenty-year economic impacts of deworming," *Proceedings of the National Academy of Sciences of the United States of America*, 4 2021, *118*.

Haushofer, Johannes and Jeremy Shapiro, "The long-term impact of unconditinoal cash transfers: experimental evidence from Kenya," 2018.

Institute for Health Metrics and Evaluation, *Health Service Provision in Kenya: Assessing Facility Capacity, Costs of Care, and Patient Perspectives*, Seattle, WA: Institute for Health Metrics and Evaluation (IHME), 2014.

_ , "Verbal Autopsy Tool," https://www.healthdata.org/data-tools-practices/verbal-autopsy 2025. Accessed: 2025-02-22.

Jamison, Dean T, Joel G Breman, Anthony R Measham, George Alleyne, Mariam Claeson, David B Evans, Prabhat Jha, Anne Mills, and editors Philip Musgrove, *Disease Control Priorities in Developing Countries, 2nd edition*, The International Bank for Reconstruction and Development / The World Bank, 2006.

Killeen, Grady, "A New Experimental Method for Estimating Demand for Non-market Goods: With an Application to the Value of a Statistical Life," November 2025.

Kim, Seonghoon and Kanghyock Koh, "The effects of income on health: Evidence from lottery wins in Singapore," *Journal of Health Economics*, 2021, *76*, 102414.

Kremer, Michael, Jessica Leino, Edward Miguel, and Alix Peterson Zwane, "Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions," *The Quarterly Journal of Economics*, 2011, *126* (1), 145–205. Accessed: 2025-02-23.

Kruger, Daniel, "Coffee Production Effects on Child Labor and Schooling in Rural Brazil," *Journal of Development Economics*, 2007, *82* (2), 448–463.

León, Gianmarco and Edward Miguel, "Risky Transportation Choices and the Value of a Statistical Life," *American Economic Journal: Applied Economics*, January 2017, *9* (1), 202–28.

Lleras-Muney, Adriana, Hannes Schwandt, and Laura Wherry, "Poverty and Health," Working Paper 32866, National Bureau of Economic Research August 2024.

Lyons-Amos, Mark and Timothy Stones, "Trends in Demographic and Health Survey data quality: an analysis of age heaping over time in 34 countries in Sub-Saharan Africa between 1987 and 2015," *BMC Research Notes*, December 2017, *10* (1), 760.

Mayer, Susan, *What Money Can't Buy: Family Income and Children's Life Chances*, Cambridge, MA: Harvard University Press, 1997.

Miller, Grant and B. Piedad Urdinola, "Cyclicality, Mortality, and the Value of Time: The Case of Coffee Price Fluctuations and Child Survival in Colombia," *Journal of Political Economy*, 2010, *118* (1), 113–155.

Nareeba, T., F. Dzabeng, N. Alam, G. A. Biks, S. M. Thysen, J. Akuze, H. Blencowe, S. Helleringer, J. E. Lawn, K. Mahmud, T. A. Yitayew, A. B. Fisker, and the Every Newborn-INDEPTH Study Collaborative Group, "Neonatal and child mortality data in retrospective population-based surveys compared with prospective demographic surveillance: EN-INDEPTH study," *Population Health Metrics*, February 2021, *19* (Suppl 1), 7.

Njuguna, John, Njoroge Kamau, and Charles Muruka, "Impact of free delivery policy on utilization of maternal health services in county referral hospitals in Kenya," *BMC Health Services Research*, 6 2017, *17*, 1–6.

Novella, Rafael, Laura Ripani, Guillermo Cruces, and María Laura Alzúa, "Conditional Cash Transfers, Female Bargaining Power and Parental Labour Supply," Technical

Report IDB-WP-368, Inter-American Development Bank November 2012.

Okeke, Edward N., "When a Doctor Falls from the Sky: The Impact of Easing Doctor Supply Constraints on Mortality," *American Economic Review*, March 2023, *113* (3), 585–627.

Ouko, Jacob Joseph Ochieng, Moses Karoki Gachari, Arthur Wafula Sichangi, and Victor Alegana, "Geographic information system-based evaluation of spatial accessibility to maternal health facilities in Siaya County, Kenya," *Geographical Research*, 8 2019, *57*, 286–298.

Preston, Samuel H., "The changing relation between mortality and level of economic development," *Population Studies*, 1975, *29* (2), 231–248.

Prieto, J. Romero, A. Verhulst, and M. Guillot, "Estimating the infant mortality rate from DHS birth histories in the presence of age heaping," *PLoS One*, November 2021, *16* (11), e0259304.

Rao, M. R., R. J. Levine, N. K. Wasif, and J. D. Clemens, "Reliability of maternal recall and reporting of child births and deaths in rural Egypt," *Paediatric and Perinatal Epidemiology*, April 2003, *17* (2), 125–131.

Redfern, Alice, Martin Gould, Maryanne Chege, Sindy Li, Felipe Acero Garay, and William Slotznick, "Beneficiary preferences: Findings from Kenya and Ghana," Technical Report, IDinsight 2019.

Richterman, Aaron, Christophe Millien, Elizabeth F. Bair, Gregory Jerome, Jean Christophe Dimitri Suffrin, Jere R. Behrman, and Harsha Thirumurthy, "The effects of cash transfers on adult and child mortality in low- and middle-income countries," *Nature*, June 2023, *618* (7965), 575–582. Number: 7965 Publisher: Nature Publishing Group.

Romano, Joseph P. and Michael Wolf, "Stepwise Multiple Testing as Formalized Data Snooping," *Econometrica*, 7 2005, *73*, 1237–1282.

Serina, Peter, Ian Riley, Andrea Stewart, Spencer L. James, Abraham D. Flaxman, Rafael Lozano, Bernardo Hernandez, Meghan D. Mooney, Richard Luning, Robert Black, Ramesh Ahuja, Nurul Alam, Sayed Saidul Alam, Said Mohammed Ali, Charles Atkinson, Abdulla H. Baqui, Hafizur R. Chowdhury, Lalit Dandona, Rakhi Dandona, Emily Dantzer, Gary L. Darmstadt, Vinita Das, Usha Dhingra, Arup Dutta, Wafaie Fawzi, Michael Freeman, Sara Gomez, Hebe N. Gouda, Rohina Joshi, Henry D. Kalter, Aarti Kumar, Vishwajeet Kumar, Marilla Lucero, Seri Maraga, Saurabh Mehta, Bruce Neal, Summer Lockett Ohno, David Phillips, Kelsey Pierce, Rajendra Prasad, Devarsatee Praveen, Zul Premji, Dolores Ramirez-Villalobos, Patricia Rarau, Hazel Remolador, Minerva Romero, Mwanaidi Said, Diozele Sanvictores, Sunil Sazawal, Peter K. Streatfield, Veronica Tallo, Alireza Vadhatpour, Miriam Vano, Christopher J. L. Murray, and Alan D. Lopez, "Improving performance of the Tariff Method for assigning causes of death to verbal autopsies," *BMC Medicine*, December 2015, *13* (1), 291.

Stedman, Nancy, "Cash Transfer Programs Are Growing More Common in the U.S. as Studies Show They Improve Peo-

ple's Health," https://ldi.upenn.edu/our-work/research-updates/cash-transfer-programs-are-growing-more-common-in-the-u-s-as-studies-show-\they-improve-peoples-health/ 2023. Accessed: 2025-02-23.

Stenberg, K., R. Watts, M. Y. Bertram, K. Engesveen, B. Maliqi, L. Say, and R. Hutubessy, "Cost-Effectiveness of Interventions to Improve Maternal, Newborn and Child Health Outcomes: A WHO-CHOICE Analysis for Eastern Sub-Saharan Africa and South-East Asia," *International Journal of Health Policy and Management*, November 2021, *10* (11), 706–723.

United Nations, "Goal 3: Ensure Healthy Lives and Promote Well-Being for All at All Ages," 2015. Accessed: 2025-02-23.

World Bank, "Monitoring COVID-19 Impact on Households and Firms in Kenya," 2022. Accessed: 2025-02-23.

_ , "In Kenya, the labor force participation rate among females is 72.2% and among males is 75.3% for 2023," https://genderdata.worldbank.org/en/economies/kenya 2025. Accessed: 2025-02-22.

World Health Organization, *WHO Recommendations on Interventions to Improve Preterm Birth Outcomes*, Geneva, Switzerland: World Health Organization, 2015. Accessed: 2025-02-22.

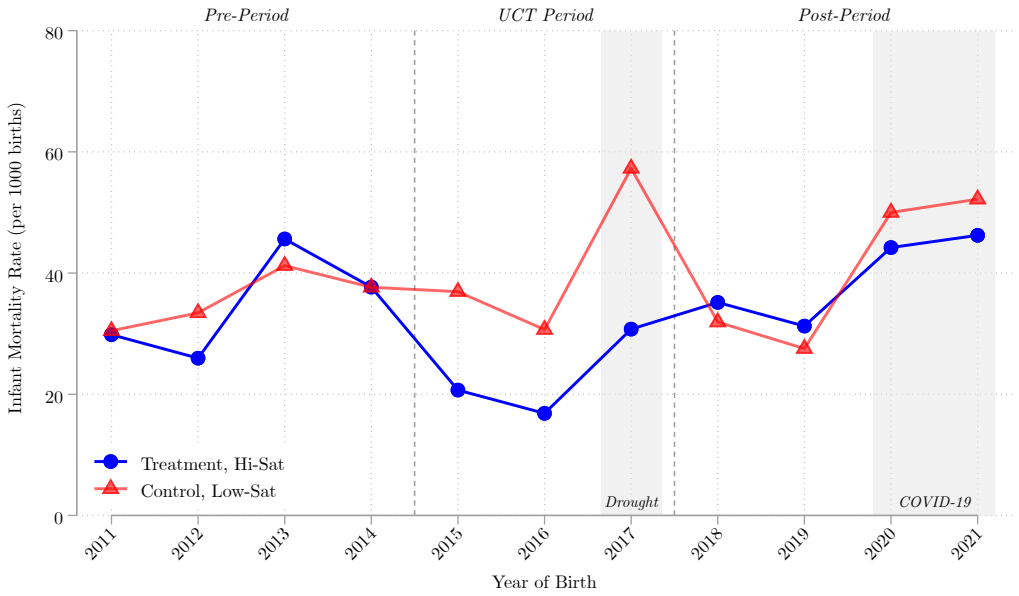
_ , "2022 WHO Verbal Autopsy Instrument," Technical Report, World Health Organization 2022.

_ , "Perinatal Asphyxia," https://www.who.int/teams/maternal-newborn-child-adolescent-health-and-ageing/newborn-health/perinatal-asphyxia 2022. Accessed: 2025-02-22.

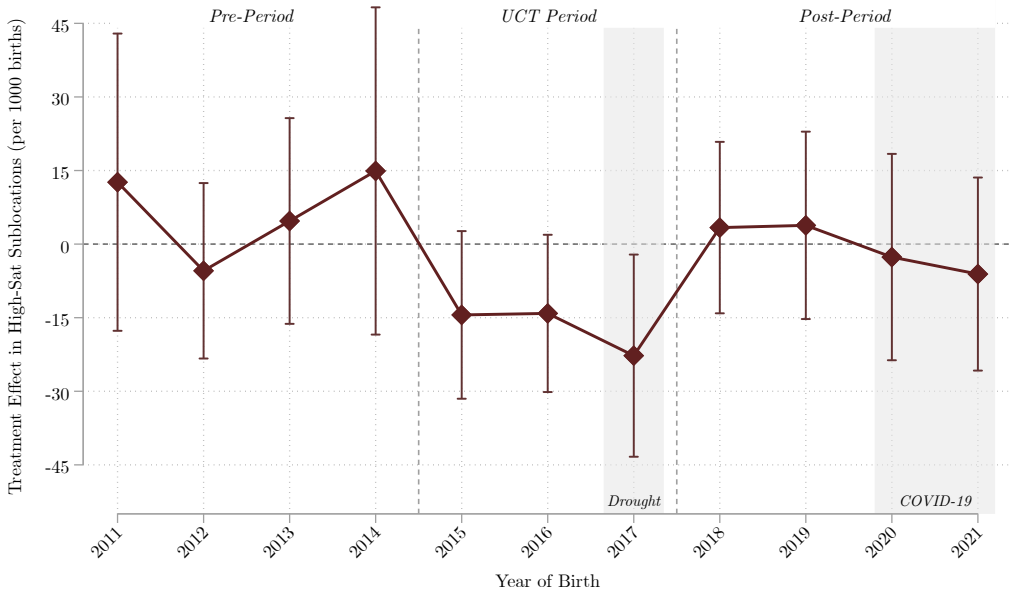
Wyndham-Douds, Kiara and Sarah K. Cowan, "Estimating the effect of a universal cash transfer on birth outcomes," *American Sociological Review*, 2024, *89* (5), 789–819.

Zeng, Wu, Guohong Li, Haksoon Ahn, Ha Thi Hong Nguyen, Donald S Shepard, and Dinesh Nair, "Cost-effectiveness of health systems strengthening interventions in improving maternal and child health in low- and middle-income countries: a systematic review," *Health Policy and Planning*, 2018, *33*, 283–297.

Figure 1: Unconditional Cash Transfers and Infant Mortality By Year



Panel A: Infant Mortality By Year



Panel B: Reduced-Form Impacts by Year

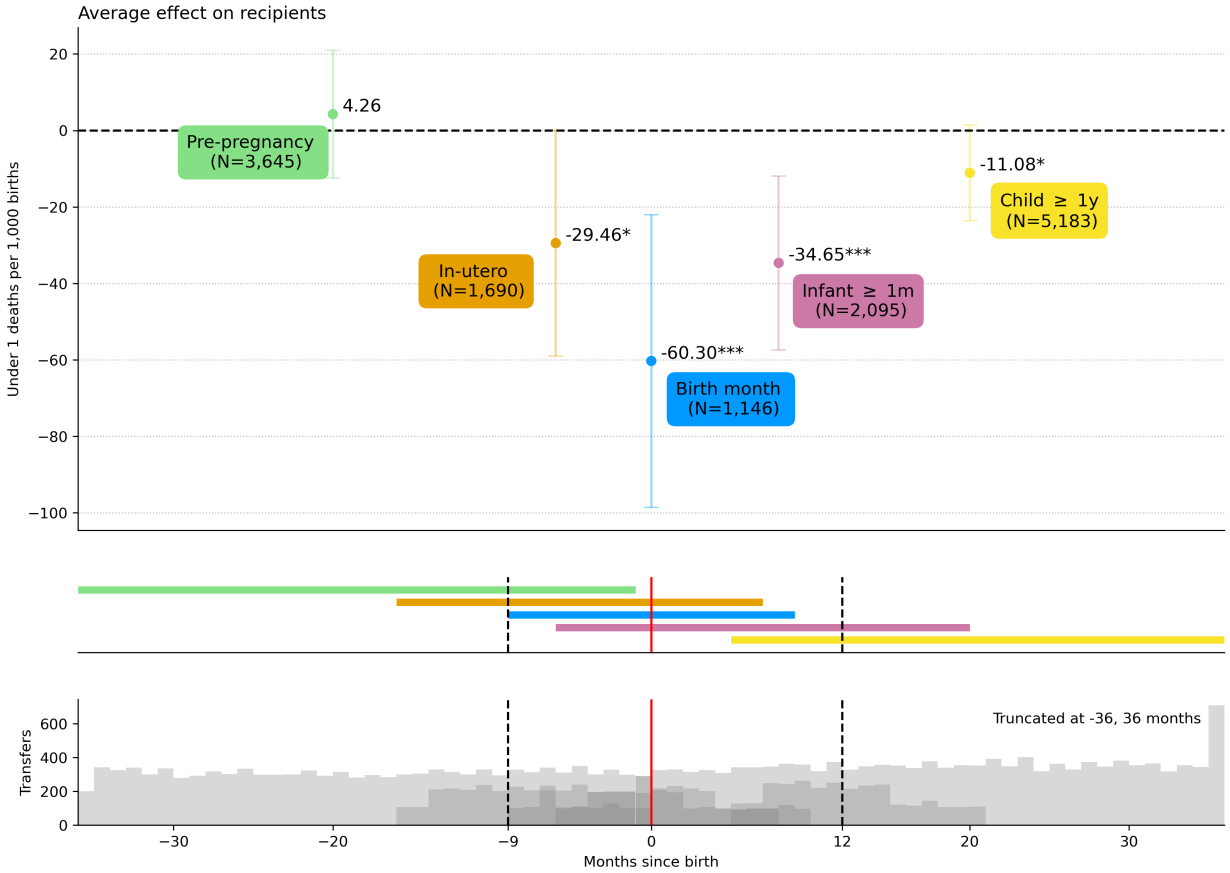
Notes: Panel A reports the mean infant mortality rate by year among eligible households for treatment villages in high-saturation areas and control villages in low-saturation areas. Panel B reports the reduced-form estimates of infant mortality impacts for treatment villages in high-saturation areas. Pre-period births refer to births occurring in the period 2011-14, whereas the unconditional cash transfer (UCT) period refers to 2015-17 and the post-UCT period denotes 2018-21. The COVID-19 pandemic spans 2020-21 and a severe drought affected Kenya in late 2016 and 2017. * $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the sublocation level.

Table 1: Unconditional Cash Transfers and Mortality, 2015-17

	Reduced-Form		Spatial IV	
	(1) Infant Mortality	(2) Child Mortality	(3) Infant Mortality	(4) Child Mortality
Own village	-7.96	-13.77**	-8.86*	-13.30**
	(5.43)	(6.04)	(4.66)	(5.49)
MHT adjusted p-value	[0.124]	[0.064]		
High-saturation spillovers	-9.17*	-2.44	-12.23*	-13.33*
	(4.76)	(6.55)	(6.42)	(7.80)
ATE in high-intensity villages	-17.13***	-16.21***	-21.08***	-26.63***
	(4.99)	(6.08)	(6.25)	(8.25)
MHT adjusted p-value			[0.012]	[0.012]
Control Mean	39.82	56.82	39.82	56.82
Observations	6,293	6,294	6,293	6,294

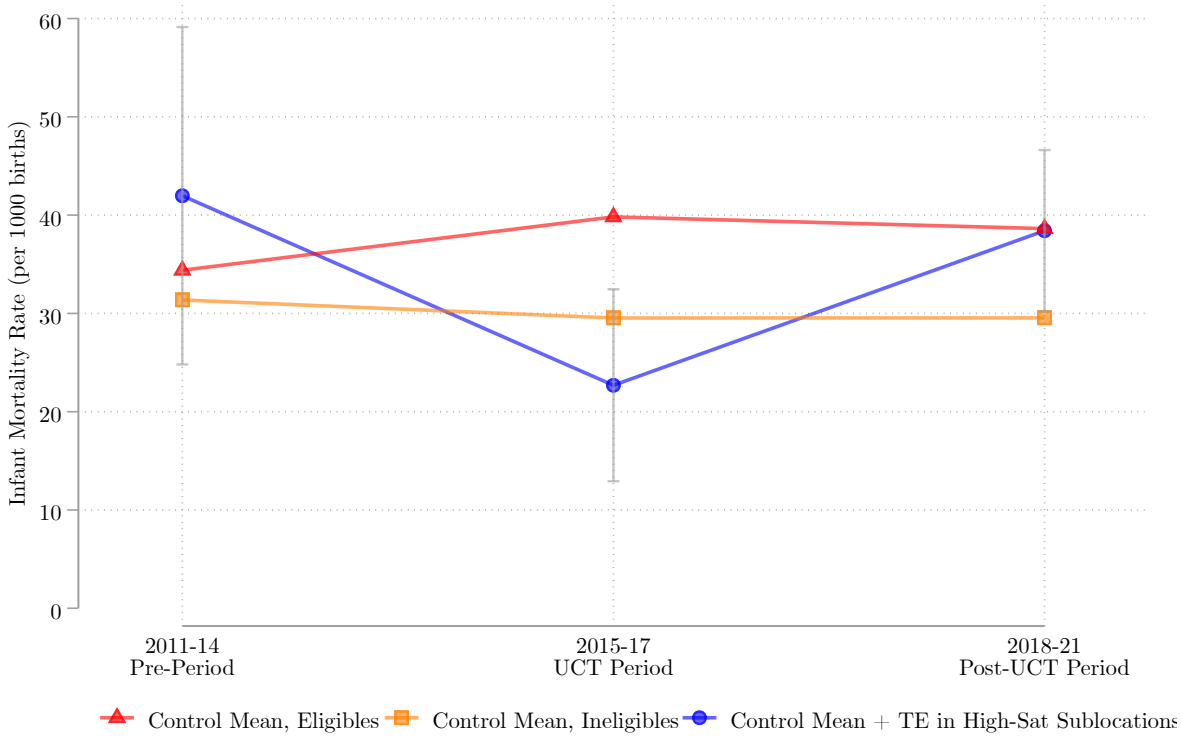
Notes: * $p < .10$, ** $p < .05$, *** $p < .01$. Reduced form standard errors are clustered at the sublocation level. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates. The ATE in high-intensity villages equals the average total effect of own-village estimates and spillovers in high-saturation sublocations. Columns (1) - (2) report estimates from Equation 1 and columns (3) - (4) report estimation from Equation 2. MHT corrected p-values in brackets for outcomes that were pre-specified calculated using a Romano and Wolf (2005) step-down correction based on randomization inference with 500 iterations.

Figure 2: Transfer Effects on Infant Mortality by Timing of Cash Disbursal



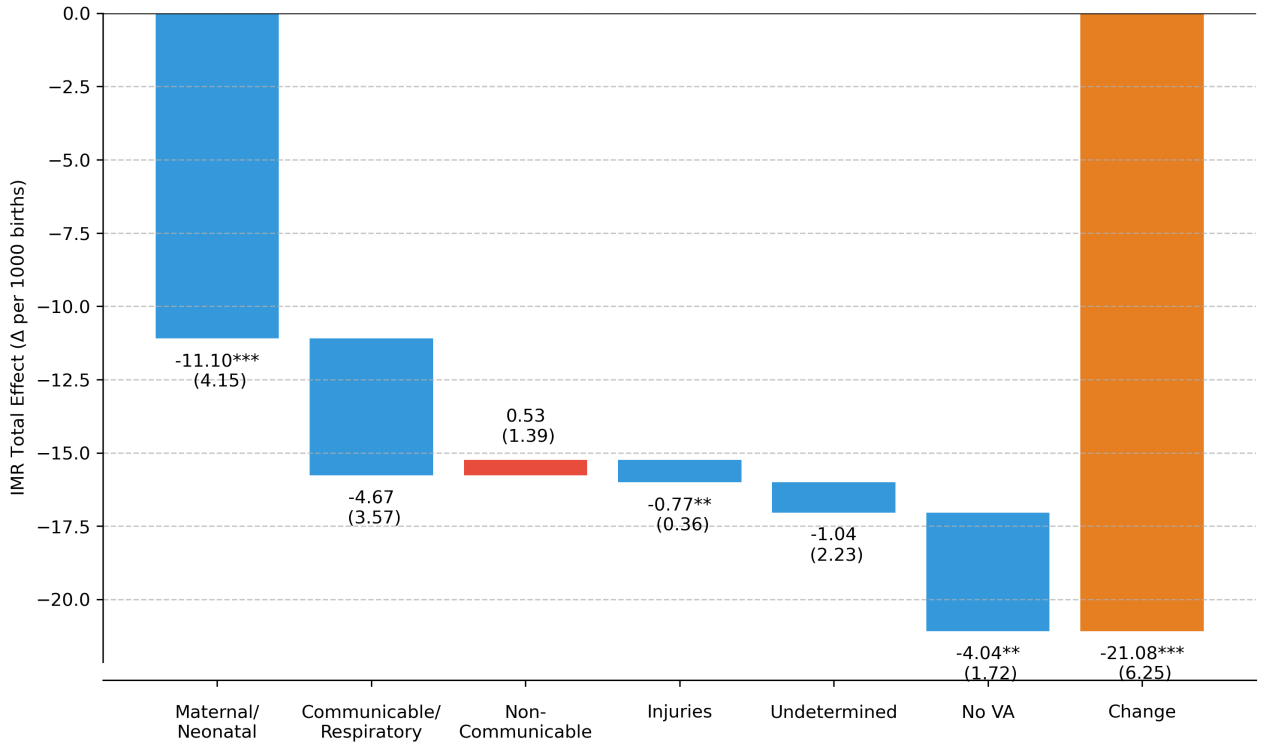
Notes: This figure plots estimates of dynamics based on the time period of exposure to cash during the child's life. The transfer timing is defined relative to the "experimental start date", as this is well-defined for both treatment and control villages. Since the transfers were distributed in a series of three payments, some households in the 1-2 years post-transfer bin were still receiving cash transfers during this period.

Figure 3: Unconditional Cash Transfers and Infant Mortality in Baseline Households



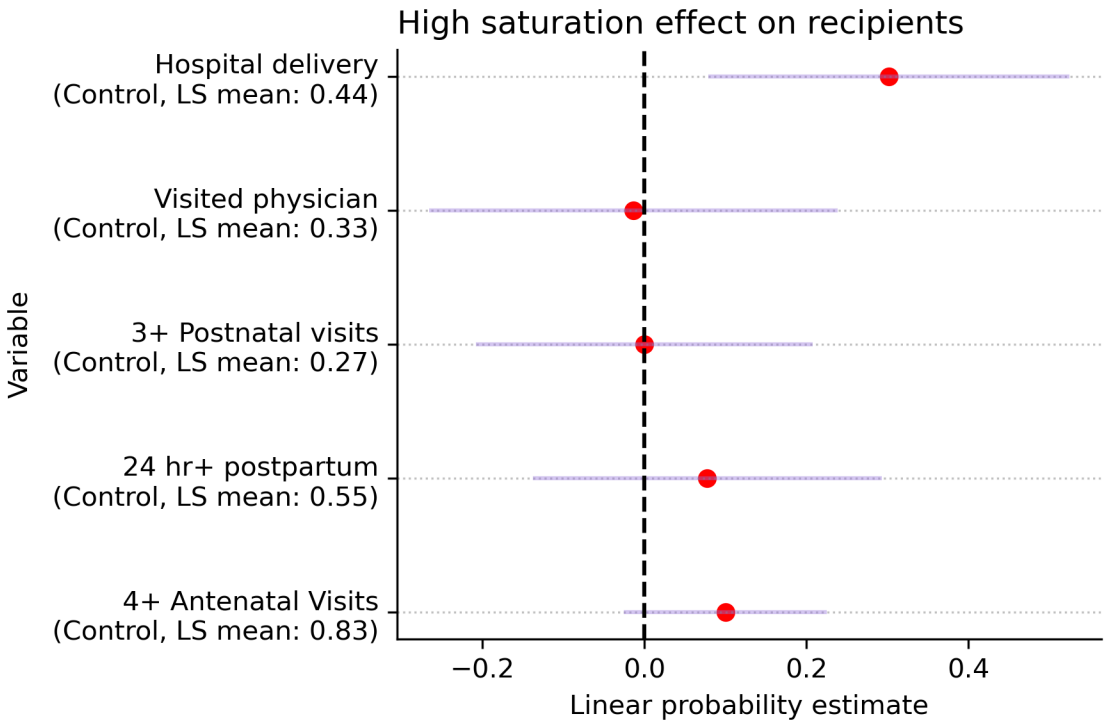
Notes: This figure plots the infant (under-1) mortality rate per 1000 births (y-axis) for three different periods of child births: pre-transfers (child birth years of 2011-14), the transfer period (child birth years 2015-17), and the post-transfer period (child birth years 2018-21). All data is restricted to households present at the time of the baseline census. The red line reports average rates for transfer-eligible households in control, low saturation villages. The blue line adds in the estimated treatment effects for treatment villages in high saturation areas by period (from Equation (1), reported in Table 1), with 95% confidence intervals shown. The yellow line reports infant mortality rates for transfer-ineligible households residing in control, low saturation villages.

Figure 4: Transfer Effects on Infant Mortality by Likely Cause of Death, 2015-17



Notes: This figure reports the estimated reduction in infant mortality (per 1000 births) based on cause of death determinations from verbal autopsies. Treatment effects are estimated using Equation (2).

Figure 5: Unconditional Cash Transfers and Healthcare Utilization, 2015-17



Notes: This figure reports treatment effects on indicators for visiting the indicated health service during a pregnancy associated with a birth between 2015 and 2017. Treatment effects are estimated using Equation (2). 95% confidence intervals are constructed using spatial HAC standard errors with a 10km cutoff (Conley, 2008). The data considers a representative subset of the Endline 3 household survey containing about 70% of anticipated observations.

Table 2: Heterogeneity: Complementarity with Health Services (Infant Mortality)

	Open Street Maps Estimates			**Study Measurements**
	(1) Walk time (OSM)	(2) Motorcycle time (OSM)	(3) Drive time (OSM)	(4) Measured time (GPS)
Total Effect IV	-35.51***	-35.87***	-39.17***	-30.02***
(Time: Below median)	(9.51)	(9.67)	(10.80)	(10.28)
Total Effect	23.24*	23.39	28.18*	14.10
x Above median	(13.43)	(14.82)	(15.75)	(13.97)
Double-partial out LASSO controls, no treatment interactions				
Above median:	21.04	21.77	28.44*	15.82
DP LASSO Controls	(14.07)	(15.68)	(15.85)	(15.14)
Double-partial out LASSO controls, with treatment interactions				
Above median:	30.29*	30.72*	37.98**	27.94*
DP LASSO x Treat Controls	(16.74)	(17.52)	(16.88)	(16.51)
Control Mean	43.52	44.06	50.69	43.56
Observations	6,254	6,254	6,254	6,254

Notes: * $p < .10$, ** $p < .05$, *** $p < .01$. Spatial HAC standard errors in parenthesis.

Estimates in columns (1) - (3) use estimates of the time to a physician-staffed health facility estimated using data from Open Street Maps, accounting for faster speeds on roads and slower walking speeds through difficult terrain. Column (4) uses GPS measurements of travel speeds obtained via the study team during travel throughout the study area. We measured physician-staffed facilities via clinic surveys in 2024. If the facility reported it was open in 2014 and had a physician employed during the survey the facility is included. Rows under "Double-partial out LASSO controls, no treatment interactions" include covariates selected by double-partial out LASSO. The possible covariates includes malaria suitability, rainfall, baseline village income and assets, proximity to a road, population, distance to a town, and proximity to a water source. Results under "Double-partial out LASSO controls, with treatment interactions" are similar, but include covariate times treatment interactions in the list of possible covariates.

Table 3: Unconditional Cash Transfers and Fertility

	Actual Fertility Among Adult Women			*Predicted Fertility Among Actual Mothers*	
	(1) 2011-14	(2) 2015-17	(3) 2018-21	(4) 2015-17	(5) 2018-21
Total Effect IV	0.25	0.98***	0.20	-0.18	-0.64
	(0.24)	(0.35)	(0.27)	(0.61)	(0.61)
Control Mean	6.985	7.357	8.095	8.009	10.014
Observations	23527	23527	23527	930	930

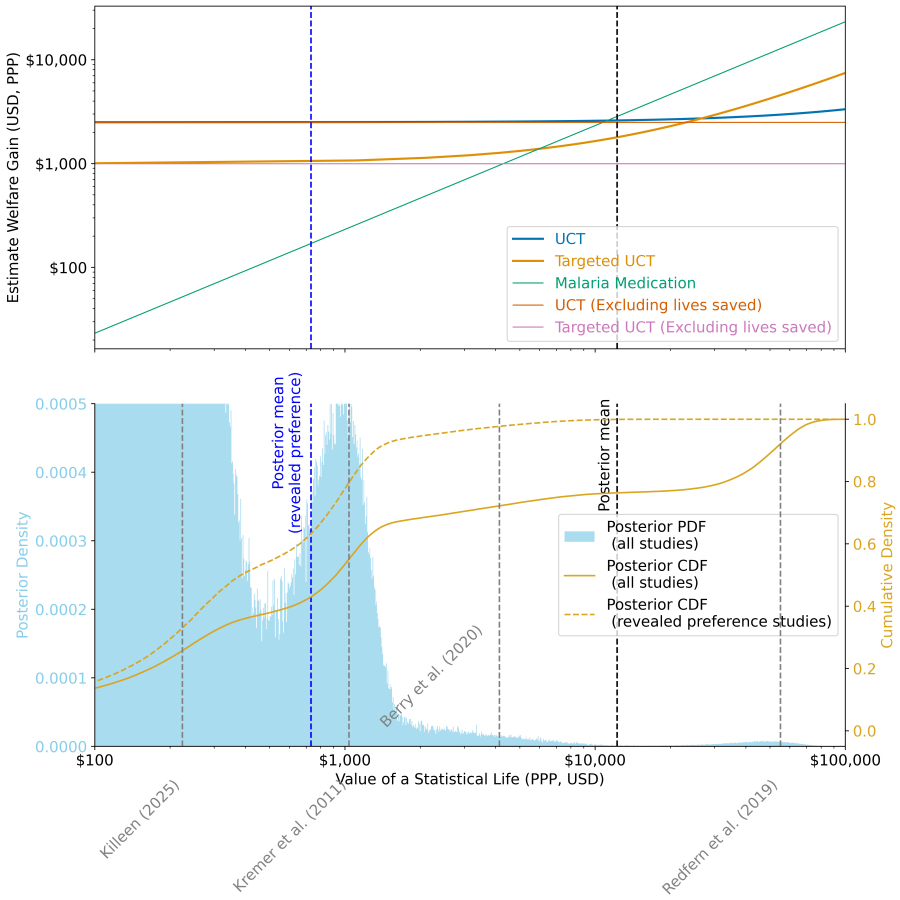
Notes: Data comprise women from the EL3 census (columns 1-3), pregnancies from the EL2 survey (columns 4-5), and women from the EL3 census from households surveyed at baseline (columns 6-7). Fertility outcomes represent the annual probability a woman gives birth to at least one child. The last two columns report predicted fertility among women actually giving birth in a given period using a random forest trained on women in control, low-saturation villages with household baseline income, assets, education, marital status, age, and household size as predictor variables. * $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the sublocation level.

Table 4: Transfers and Labor Supply by Gender and Third Trimester Pregnancy

	(1) Female Labor Supply Hours Last Week	(2) Male Labor Supply Hours Last Week
Treatment Village	-1.378	0.444
	(1.653)	(1.929)
High-Saturation Sublocation	2.519	5.385***
	(1.553)	(2.038)
Treat x 3rd Trimester In-Utero	-7.570	-4.794
	(8.813)	(8.385)
Hi Sat x 3rd Trimester In-Utero	-7.026	-3.120
	(6.656)	(8.399)
Treatment Villages in	1.14	5.83***
High-Saturation Sublocations	(2.18)	(1.96)
(Treat x Heterogeneity Var) +	-14.60***	-7.91
(Hi Sat x Heterogeneity Var)	(8.78)	(8.93)
Percentage Change	-57.15%	-28.62%
Control Mean	25.540	27.651
Observations	2716	2716

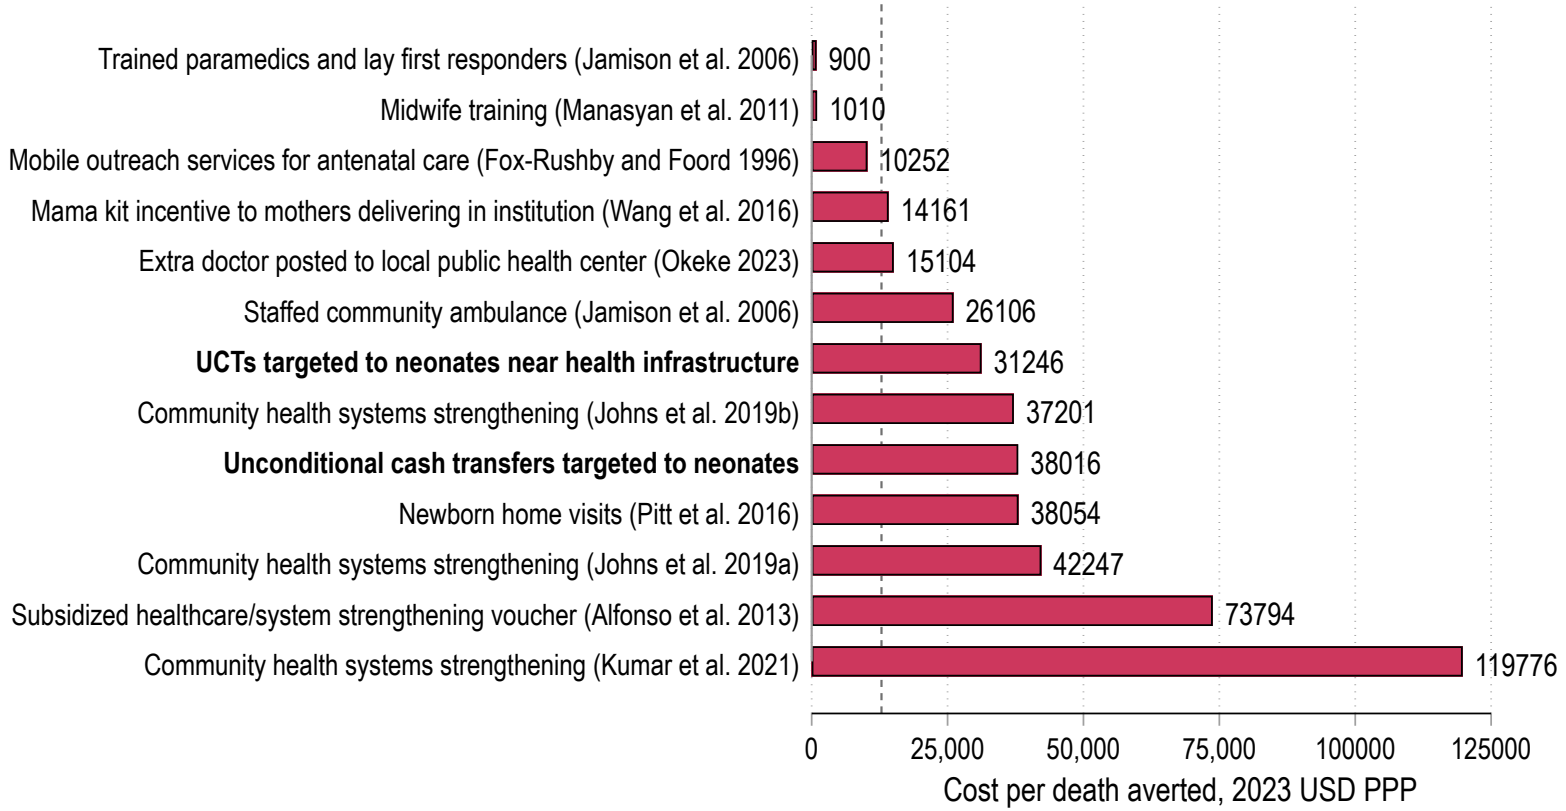
Notes: Labor supply encompasses hours worked (in agricultural employment, non-agricultural self-employment, and wage employment) as well as hours spent searching for work. * $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at sublocation level.

Figure 6: Welfare Gains from UCTs or Malaria Medicine by Value of a Statistical Life



Notes: The top panel of this figure plots the estimated welfare gains from a $1,000 investment in UCTs or GiveWell's top recommended program as of March 2025, malaria medicine, based on the value of a statistical life (VSL). We plot UCT benefits assuming the multiplier of 2.5 reported in Egger et al. (2022) with and without including the benefits of the child mortality reductions estimated in this paper, obtained by multiplying VSL by 1,000 over the cost per life saved. We also consider a targeted transfer with a consumption multiplier of 1 without child mortality benefits and with child mortality benefits if targeted to women in the third trimester of pregnancy. Malaria medicine benefits are estimated using the cost per life saved reported by GiveWell of $4,304. The bottom panel reports a posterior distribution of VSL estimates in this sample obtained from Killeen (2025), Kremer et al. (2011), Berry et al. (2020), and Redfern et al. (2019) using Bayesian hierarchical meta analysis with a log-normal prior. This final estimate is not obtained via revealed preference, so we also report the cumulative density function without it (labeled by revealed preference studies).

Figure 7: Comparative Cost-Effectiveness with Other Interventions in Sub-Saharan Africa



Notes: This figure displays the estimated cost per death averted for 11 health interventions in Sub-Saharan Africa sourced from Zeng et al. (2018) and Jamison et al. (2006), as well as the average cost per death averted across 37 WHO-recommended maternal and child health interventions in East Africa as estimated by Stenberg et al. (2021). Cost per death averted estimates of UCTs targeted to neonates net out three years of estimated consumption gains for recipient households. All estimates are converted to 2023 USD PPP.

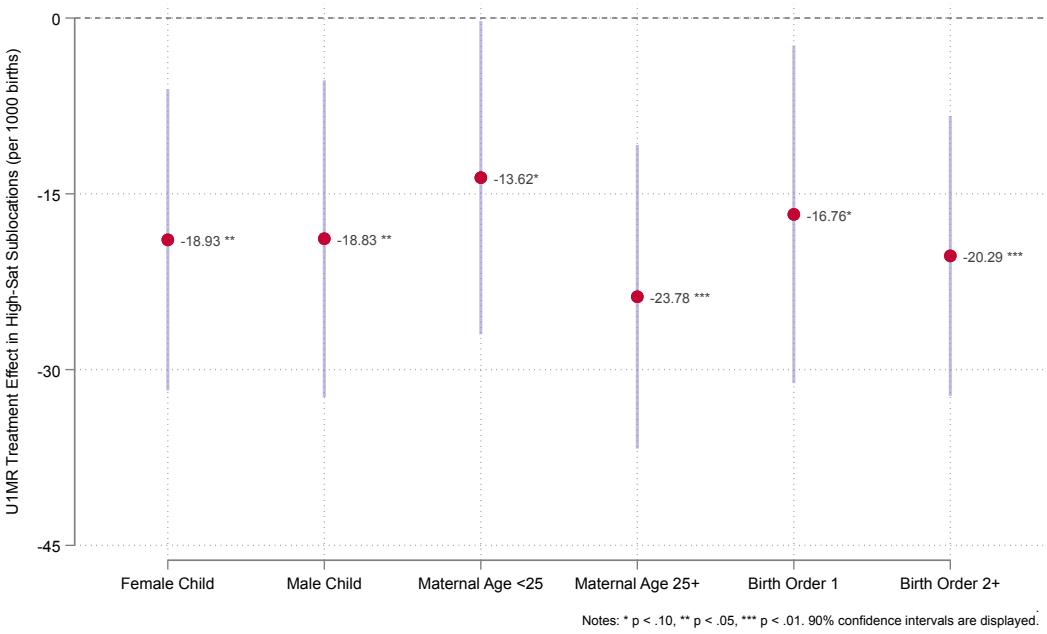
A Additional Exhibits

Table A.1: Balance on Tracking, Study Area Presence, and Census Completion

	(1) Household Tracked	(2) Present in Study Area	(3) Completed Census
Treatment Village	-0.000	-0.001	0.003
	(0.006)	(0.007)	(0.006)
High-Saturation Sublocation	-0.005	-0.003	-0.008
	(0.006)	(0.007)	(0.006)
Treatment Villages in High-Saturation Sublocations	-0.01	-0.00	-0.01
	(0.01)	(0.01)	(0.01)
Observations	62199	59394	52327

Notes: This table reports the share of households present at baseline tracked in the third endline census (EL3), present in the study area at EL3, and completed the EL3 census, as well as balance by treatment status. The share tracked (column 1) is reported as a share of baseline households non-deceased in the second endline census (EL2). The share present in study area (column 2) is reported as a share of baseline households confirmed as non-deceased in EL2 and EL3. The share completing the census (column 3) is reported as a share of baseline households present in the study area. A total of 65,383 households were censused at baseline. * $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the village level.

Figure A.1: Heterogeneity by Gender, Maternal Age, and Birth Order



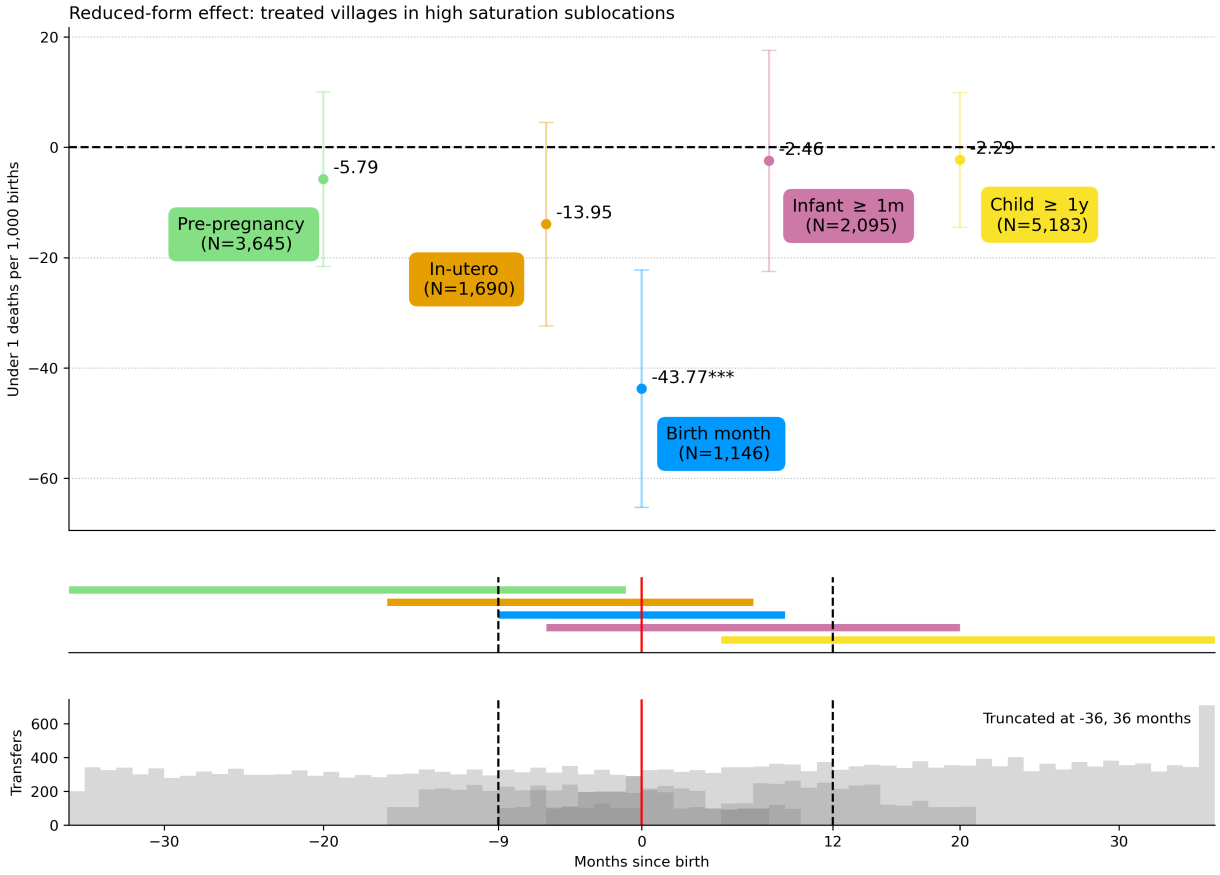
Notes: This figure plots the annualized birth rate per 100 females by period, using data on female household members in transfer-eligible households in the study area since baseline. The red line reports the control, low saturation group mean, while the blue line adds in treatment effects for treatment villages in high saturation areas from a regression of Equation (1).

Table A.2: Transfers and Mortality Including Secondary Outcomes, 2015-17

	Primary outcomes		Secondary outcomes	
	(1) Infant Mortality 2015-17	(2) Under-5 Mortality 2015-17	(3) Neonatal Mortality 2015-17	(4) Days survived under-5 2015-17
Panel A: Reduced-form results				
Treatment village	-8.55	-14.65**	-3.55	17.49*
	(5.59)	(6.18)	(4.01)	(10.23)
	[0.108]	[0.068]		
High-saturation sublocation	-10.33**	-3.65	-6.00	11.63
	(4.91)	(6.79)	(3.75)	(10.49)
Treatment village in high-saturation sublocation	-18.88***	-18.30***	-9.56**	29.12***
	(5.02)	(6.03)	(3.80)	(9.74)
Panel B: Spatial IV results				
Spatial IV:	-22.07***	-28.41***	-13.24***	41.43***
ATE on eligibles	(5.97)	(7.73)	(4.89)	(12.09)
	[0.01]	[0.014]		
Control Mean	41.79	59.63	23.04	1,735.21
Observations	6,260	6,261	6,241	6,261

Notes: Reduced form standard errors are clustered at the sublocation level. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates. MHT corrected p-values in brackets.

Figure A.2: Transfer Effects on Infant Mortality by Timing of Cash (Reduced Form)



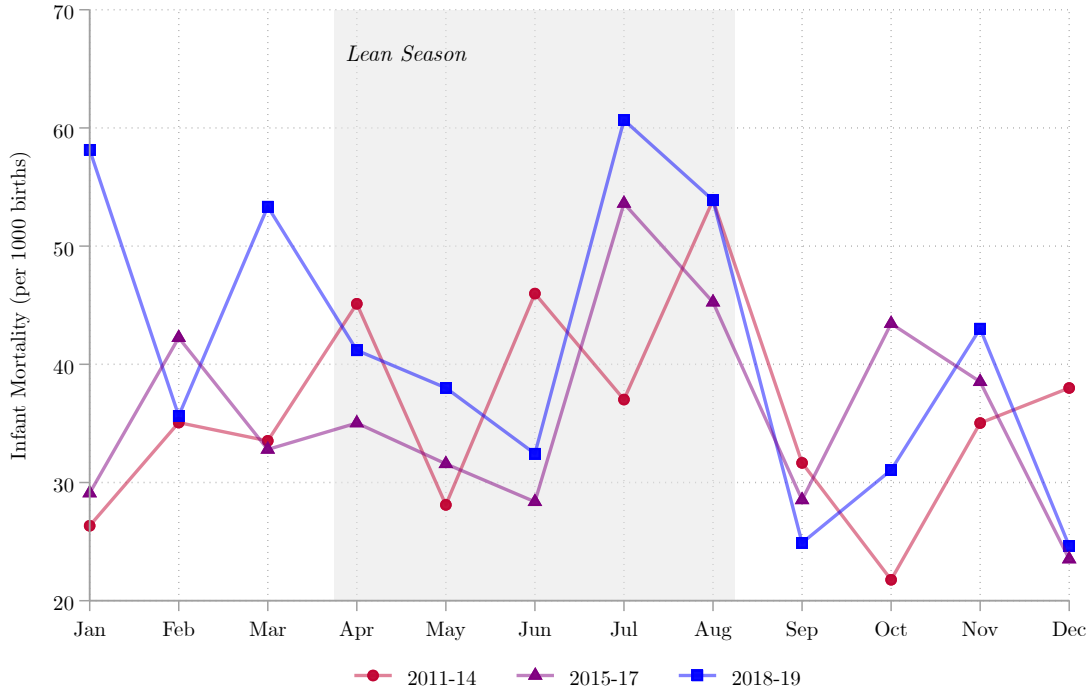
Notes: This figure plots estimates of dynamics based on the time period of exposure to cash during the child's life. The transfer timing is defined relative to the "experimental start date", as this is well-defined for both treatment and control villages. Since the transfers were distributed in a series of three payments, some households in the 1-2 years post-transfer bin were still receiving cash transfers during this period.

Table A.3: Transfers and Mortality Excluding Stillbirths, 2015-17

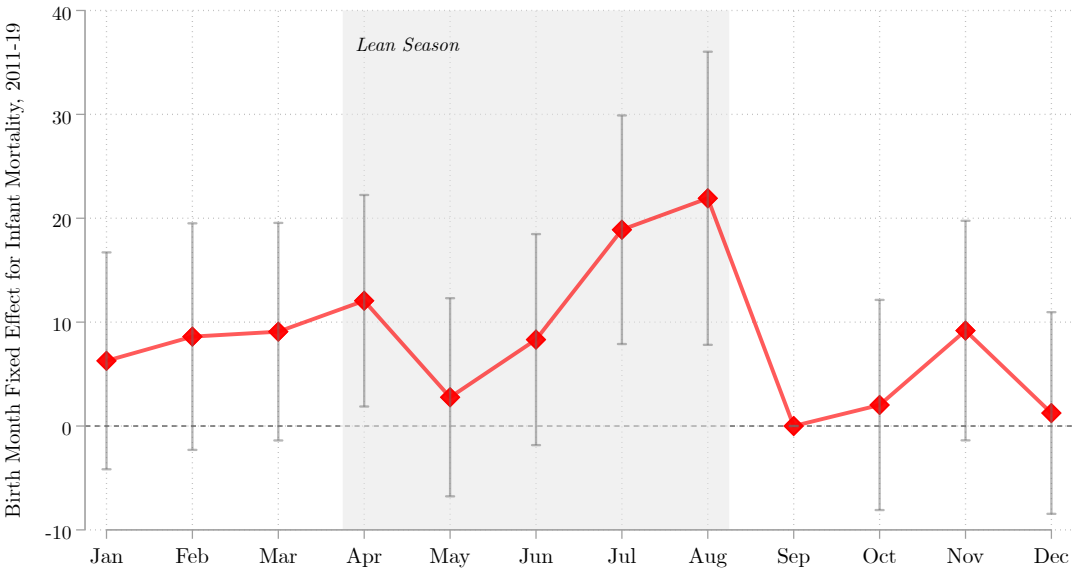
	Reduced-form		Spatial IV	
	(1) Infant Mortality 2015-17	(2) Under-5 Mortality 2015-17	(3) Infant Mortality 2015-17	(4) Under-5 Mortality 2015-17
Own village	-3.41	-7.68	-5.76	-8.45
	(5.24)	(6.05)	(4.56)	(5.17)
High-saturation spillovers	-12.04**	-6.46	-15.86**	-20.07***
	(4.55)	(6.17)	(6.24)	(7.43)
ATE in high-saturation sublocations	-15.45***	-14.14***	-21.62***	-28.52***
	(4.34)	(5.47)	(6.16)	(7.91)
Control Mean	34.54	48.97	34.54	48.97
Observations	6,203	6,204	6,203	6,204

Notes: * $p < .10$, ** $p < .05$, *** $p < .01$. Reduced form standard errors are clustered at the sublocation level. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates.

Figure A.3: Seasonality in Infant Mortality in the Birth Census



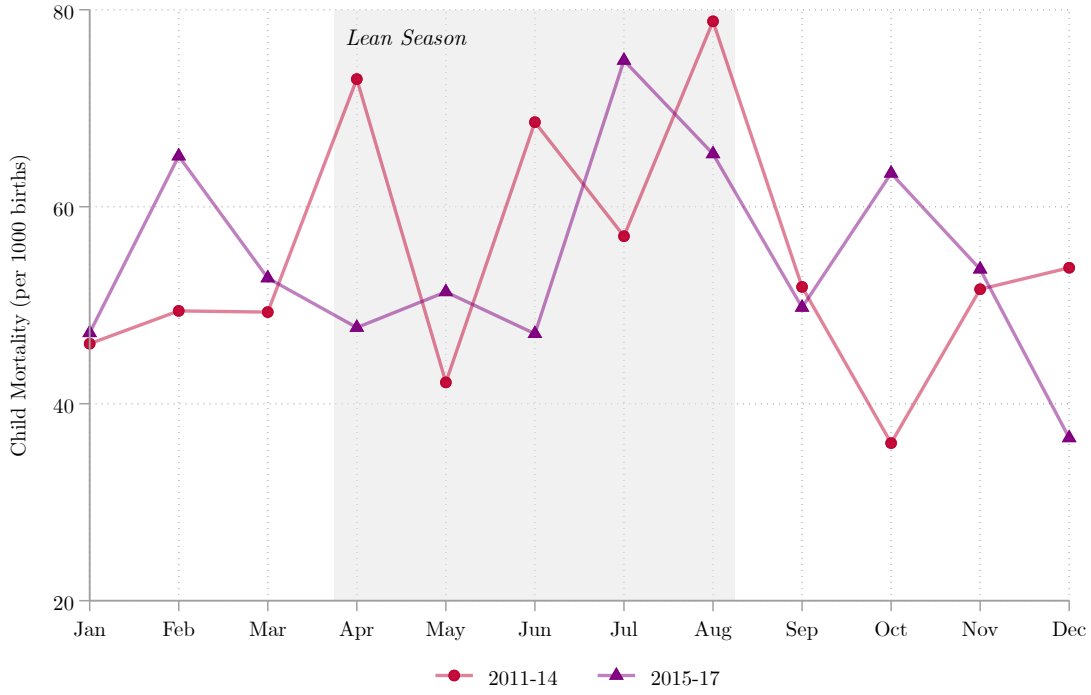
Panel A: Infant Mortality Means By Month



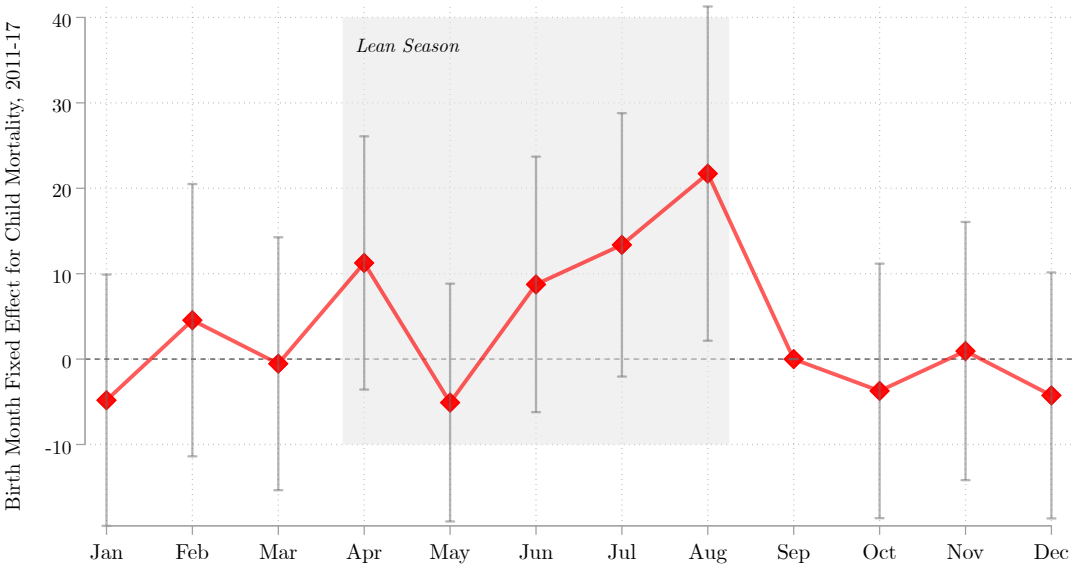
Panel B: Birth Month Fixed Effects

Notes: Panel A reports the mean infant mortality rate by month across 2011-14, 2015-17, and 2018-19. Panel B reports birth month fixed effects controlling for birth year fixed effects, with September as the omitted month. The lean season is defined as April-August using the definition from Burke et al. (2019), whose study area comprised a nearby region in rural Kenya. The years 2020 and 2021 are omitted as COVID-19 pandemic impacted the region, though patterns appear similar if those years are included. 95% confidence intervals are reported in Panel B.

Figure A.4: Seasonality in Child Mortality in the Birth Census



Panel A: Child Mortality Means By Month



Panel B: Birth Month Fixed Effects

Notes: Panel A reports the mean child mortality rate by month across 2011-14 and 2015-17. Panel B reports birth month fixed effects controlling for birth year fixed effects, with September as the omitted month. The lean season is defined as April-August using the definition from Burke et al. (2019), whose study area comprised a nearby region in rural Kenya. The years 2018-21 are omitted as children in those group had not yet reached the age of five at the time of the birth census. 95% confidence intervals are reported in Panel B.

Table A.4: Unclassified Verbal Autopsies by Treatment Status

	(1) Undetermined VA (2011-14)	(2) Undetermined VA (2015-17)	(3) Undetermined VA (2018-21)
Panel A: Reduced-Form Estimates			
Treatment Village	0.053	-0.091	-0.026
	(0.048)	(0.060)	(0.049)
High-Saturation Sublocation	0.023	0.088	0.024
	(0.048)	(0.066)	(0.049)
Treatment Villages in	0.08	-0.00	-0.00
High-Saturation Sublocations	(0.06)	(0.08)	(0.05)
Panel B: Spatial IV Estimates			
Total Effect IV	0.10	0.04	0.06
	(0.08)	(0.07)	(0.06)
Control Mean	0.124	0.194	0.169
Observations	266	200	328

Notes: Standard errors in parentheses. Sample encompasses verbal autopsies (VAs) for deceased infants from eligible households surveyed at baseline. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.5: Association Between Treatment Status and Baseline Characteristics For Eligible Households with 2015-17 Births

	(1) Education	(2) Age	(3) Marital Status	(4) Income	(5) Assets	(6) Household Size
Total Effect IV	-0.08	-1.60	0.07	-4056.17	-5412.13	0.32
	(0.30)	(1.27)	(0.04)	(2772.12)	(4949.98)	(0.23)
Control Mean	8.588	34.499	0.718	14895.334	32446.214	4.403
Observations	1159	1158	1159	1159	1159	1159

Notes: This table reports regressions exhibiting the association between treatment status and six baseline household socioeconomic characteristics (maximum years of education of household members, primary respondent age, number of adults, number of children, baseline household assets, and baseline household income) for transfer-eligible households present at baseline with births over 2015-17 that were sampled in the baseline survey. * $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the village level.

Table A.6: Association Between Treatment Status and Predicted Probability of Giving Birth Among Mothers Giving Birth Over 2015-17

	(1) Probability of Birth Among Actual Mothers Socioeconomic Variables	(2) Probability of Birth Among Actual Mothers LASSO Variables
Total Effect IV	-0.18	0.27
	(0.61)	(0.21)
Control Mean	8.009	7.921
Observations	929	929

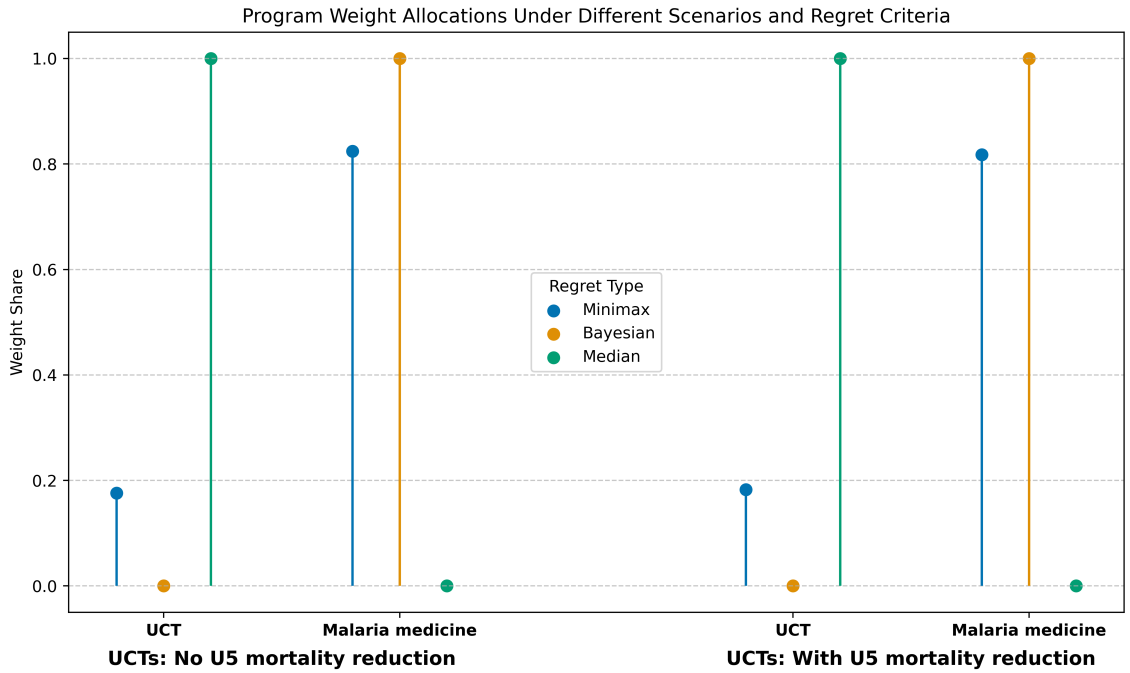
Notes: This table reports regressions exhibiting the association between treatment status and the probability of birth among mothers who actually gave birth over 2015-17. The sample is comprised of mothers in transfer-eligible households present at baseline sampled in the baseline survey. In column 1, the probability of birth is predicted using six baseline household socioeconomic characteristics (maximum years of education of household members, primary respondent age, marital status of primary respondent, income, assets, and household size) using a random forest trained on women in control, low-saturation villages. In column 2, the probability of birth is predicted linearly using variables selected by LASSO. In both cases, the training sample is not included in the final regression. * $p < .10$, ** $p < .05$, *** $p < .01$. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates.

Table A.7: Impacts on Infant Mortality in Baseline Sample After Including Controls

	(1) No Controls BL Survey Sample	(2) Socioeconomic Controls	(3) PDS LASSO-Selected Controls
Total Effect IV	-42.44***	-34.89**	-43.55***
	(16.13)	(16.12)	(15.72)
Control Mean	60.810	60.810	60.810
Observations	1432	1432	1432

Notes: This table reports effects estimated by Equation 2 of the cash transfer on infant mortality rates for transfer-eligible households present at baseline and surveyed in the baseline survey. Mortality data are sourced from the endline 3 census whereas household- and village-level control variables are sourced from the baseline survey. Column 1 presents infant mortality results in the baseline survey sample with no additional controls beyond those pre-specified (e.g., year of birth fixed effects, birth gender, mother age group). Column 2 includes controls for six baseline household socioeconomic characteristics (maximum years of education of household members, primary respondent age, marital status of primary respondent, income, assets, and household size) along with a full set of interactions between these variables. Column 3 includes controls selected by PDS LASSO, with a total of 261 baseline variables available for selection. * $p < .10$, ** $p < .05$, *** $p < .01$. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates.

Figure A.5: Share of Funds Allocated to UCTs and Malaria Medicine by Regret Type



Notes: This figure reports the share of funds allocated to UCTs or Malaria medicine by an agent under minimax, Bayesian (mean), and median regret excluding and including the mortality reduction benefits of UCTs. The agent faces uncertainty about the value of a statistical life (VSL), and they minimize the given regret function over the posterior distribution of VSL. We treat other parameters as known. We use GiveWell's estimate of the cost per life saved from Malaria medicine to estimate its returns. The posterior distribution of VSL estimates in this sample obtained from Killeen (2025), Kremer et al. (2011), Berry et al. (2020), and Redfern et al. (2019) using Bayesian hierarchical meta analysis with a log-normal prior.

B Details of benefit-cost calculations

This appendix details the estimates presented in Section 5.5. We first describe how the estimated posterior distribution of VSLs presented in Figure 6 is constructed and detail the benefit curves in the figure. We then describe the econometric decision theory problem producing the weights plotted in Figure A.5.

The posterior distribution of VSLs is constructed using Bayesian hierarchical meta analysis. We begin by positing that VSLs in countries with similar income levels follow a log-normal distribution with weakly informative priors. The log-normal distribution imposes that VSL is never negative and captures the possibility of a large right tail. We assume that the location parameter of the distribution is normally distributed with a mean of 10 and standard deviation of 10, and the scale parameter is 5 times a half Cauchy distribution with parameter 0.5. Estimates of the posteriors are not sensitive to these decisions, consistent with the fact that the prior is only weakly informative.

We then use Markov Chain Monte Carlo methods in the Python package "pymc" to estimate the posterior distribution. We do so by matching the likelihood to estimates from Killeen (2025), Kremer et al. (2011), Berry et al. (2020), and Redfern et al. (2019). We utilize these estimates because they were obtained from populations with similar income levels. For instance, we exclude León and Miguel (2017) because the population studied has a much higher income level, and Killeen (2025) shows that this predicts a VSL orders of magnitude higher. We also restricted our focus to revealed preference estimates, other than Redfern et al. (2019), since they reflect real world choices rather than hypothetical questions that may be prone to social desirability bias.

Figure 6 then plots welfare gains of UCTs and other programs by VSL. For UCTs, we plot estimates without child mortality reductions priced in imposing the 2.5 consumption multiplier from Egger et al. (2022), the same with our estimates of the cost per life saved of the UCTs, and targeted transfers to pregnant women that have a consumption multiplier of 1 with and without lives saved priced in. We also plot the benefit curve by VSL of GiveWell's top recommendation, the Malaria Consortium, since this is one of the most effective health interventions and GiveWell reports estimates of the cost per life saved. We assume no direct consumption benefit and therefore all benefits are (1,000/cost per life saved) x VSL.

Figure A.5 plots the share of funds that a decision maker would invest in UCTs or the Malaria Consortium by regret type in the following econometric decision problem.

A social planner allocates a budget across J programs, each characterized by its consumption gain and expected lives saved per dollar spent. The planner's objective is to minimize regret given uncertainty about the true value of a statistical life (VSL), which we will denote by $\theta \in \Theta$.

Each program $j \in \{1, \ldots, J\}$ has two attributes: g_j, the consumption gain per dollar spent, and S_j, the expected number of lives saved per dollar spent.

The planner chooses a weight vector $\mathbf{w} = (w_1, \ldots, w_J)$ such that $\sum_{j=1}^{J} w_j = 1$, where w_j represents the share of the budget allocated to program j.

The true VSL, denoted θ, is uncertain and follows a distribution $\theta \sim F(\theta)$. We use the estimated posterior distribution of θ in calculations, but in minimax regret impose that it falls within the 95% confidence interval since one would otherwise mechanically invest only in the most effective health program since VSL has no upper bound in the log-normal prior

assumed. For simplicity, and since we were unable to find confidence intervals on the cost per life saved of the Malaria Consortium program, we take point estimates of consumption multiplier and cost per life saved as given.

For a given realization of θ, the optimal allocation $\mathbf{w}^*(\theta)$ maximizes total benefit:

$$\mathbf{w}^*(\theta, \mathbf{S}) = \arg\max_{\mathbf{w}} \sum_{j=1}^{J} w_j \left(S_j\theta + g_j\right).$$

The regret for a given $(\theta, \mathbf{S})$ and choice $\mathbf{w}$ is:

$$R(\mathbf{w}, \theta) = \max_{\mathbf{w}'} \sum_{j=1}^{J} w'_j \left(S_j\theta + g_j\right) - \sum_{j=1}^{J} w_j \left(S_j\theta + g_j\right).$$

We consider three regret criterion. First, the minimax regret criterion selects $\mathbf{w}$ to minimize the worst-case regret:

$$\min_{\mathbf{w}} \sup_{\theta} R(\mathbf{w}, \theta).$$

where we use the 5th and 95th percentile of VSL values estimated in the posterior distribution to define the respective worst case draws of θ for each program. This case reflects the optimal allocation for an extremely risk averse decision maker that wishes to maximize welfare in the worst possible case.

Second, the Bayesian (average) regret criterion minimizes the expected regret given the posterior distribution of VSL:

$$\min_{\mathbf{w}} \mathbb{E}_{\theta}\left[R(\mathbf{w}, \theta)\right].$$

In practice, we use 1,000 posterior draws from the posterior to estimate this value. This value would be optimal for a risk neutral decision maker that wishes to maximize expected welfare gains, but is heavily influenced by the outlying right tail of VSL values. In practice, this results in the Malaria Consortium's program receiving 100% of investment in all cases, even though it is dominated by UCTs for 75% of VSL values simulated from the posterior distribution.

Third, we consider the allocation that minimizes the median regret over θ:

$$\min_{\mathbf{w}} \text{median}_{\theta}\, R(\mathbf{w}, \theta).$$

This provides a regret measure that is less sensitive to outliers in VSL.

We construct these estimates including and excluding the benefits of mortality reductions from UCTs. In practice, this has little effect on results since the cost per life saved is lower under the Malaria Consortium than UCTs.