A Additional Exhibits

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

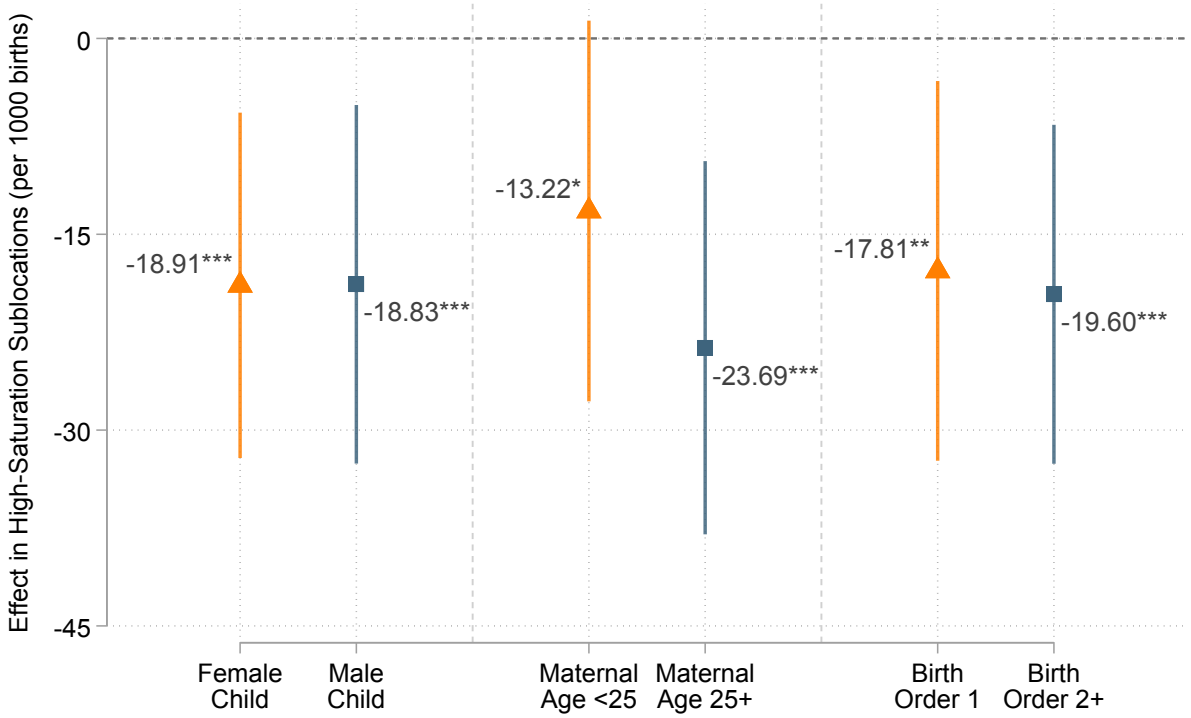

Notes: This figure is based on the main EL3 birth census sample, and plots estimated effects of cash transfers on infant mortality in treatment villages within high-saturation sublocations by gender, maternal age, and birth order. Maternal age at birth is analysed using pre-specified five-year age groups (under 20, 20-24, 25-29, 30-34, and 35 and above); this analysis splits on the median five-year group, with 46% of births having a mother below the age of 25. Child birth order is calculated using the order of births across the 2011-23 period and does not take into account births to mothers prior to 2011 as information on those births was not obtained in the census. Effects are estimated using equation (1) augmented with an interaction with the characteristic of interest. The p-values for the difference across groups (the effect on the interaction term in treatment villages in high-saturation sublocations) is as follows: 0.993 (gender), 0.337 (maternal age), and 0.856 (birth order). The results shown in this figure are additionally presented in Appendix Table A.6. Standard errors are clustered at the sublocation level. 95% confidence intervals are shown. * $p < .10$, ** $p < .05$, *** $p < .01$.

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

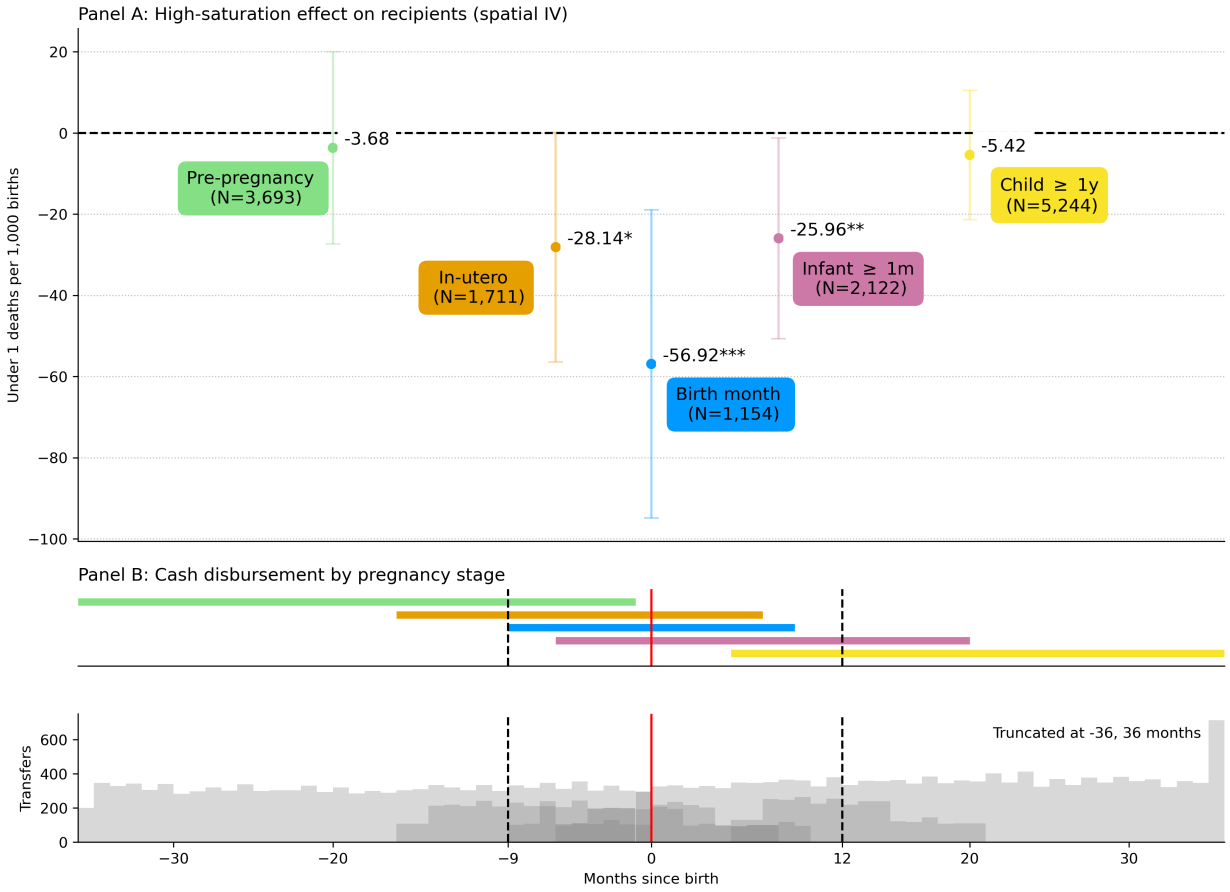

Notes: This figure is based on the main EL3 birth census sample. Panel A plots estimates of dynamics based on the time period of exposure to cash during the child's life. The transfer timing is defined relative to the "experimental start date", as this is well-defined for both treatment and control villages. "Pre-pregnancy" includes household exposure 3 years to 10 months before birth. In-utero is 9 to 1 month before birth. Birth month includes cash within the first month of life. Infant includes 1 month to 12 months. Child includes 1 to 3 years. Estimates are constructed using the same approach as equation (4), but using the spatial IV. We estimate equation (2) after restricting the sample to those exposed to cash at a particular time in life. Observations appear in multiple groups since cash transfers were distributed over 8 months. Panel B plots the range of experimental start dates, relative to birth month, included in each estimate and a histogram of transfers by month in each bin. 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. The spatial IV version of this figure (which is estimated using equation 1) is presented as Figure 3. 95% confidence intervals are shown. * $p < .10$, ** $p < .05$, *** $p < .01$.

Figure A.3: Non-Experimental Variation in Child Mortality

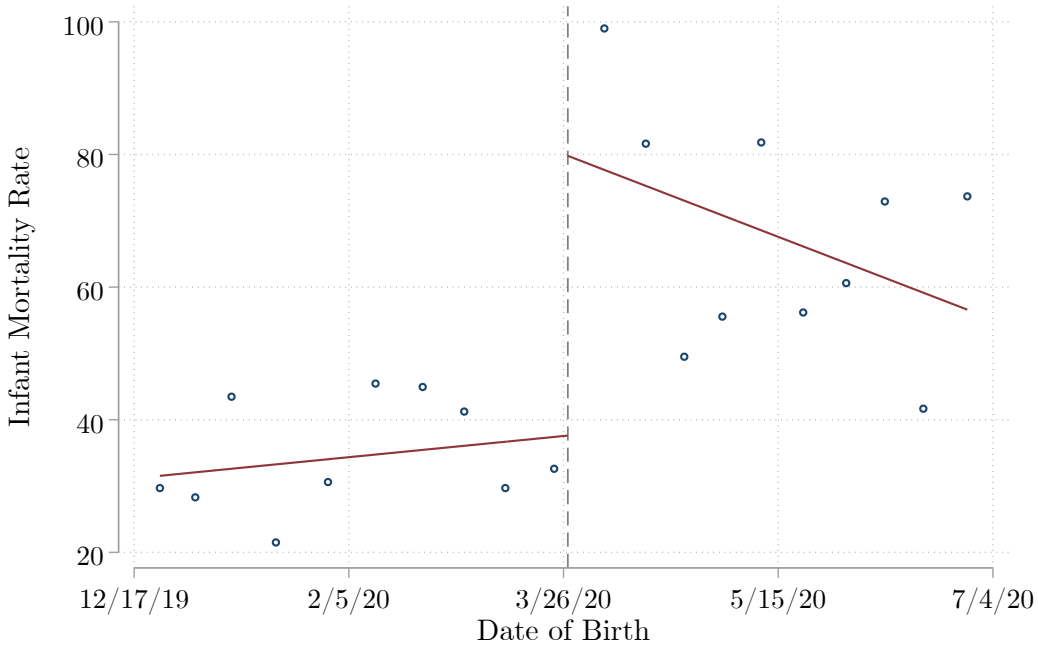

Panel A: COVID-19 Lockdown and Infant Mortality in Study Region

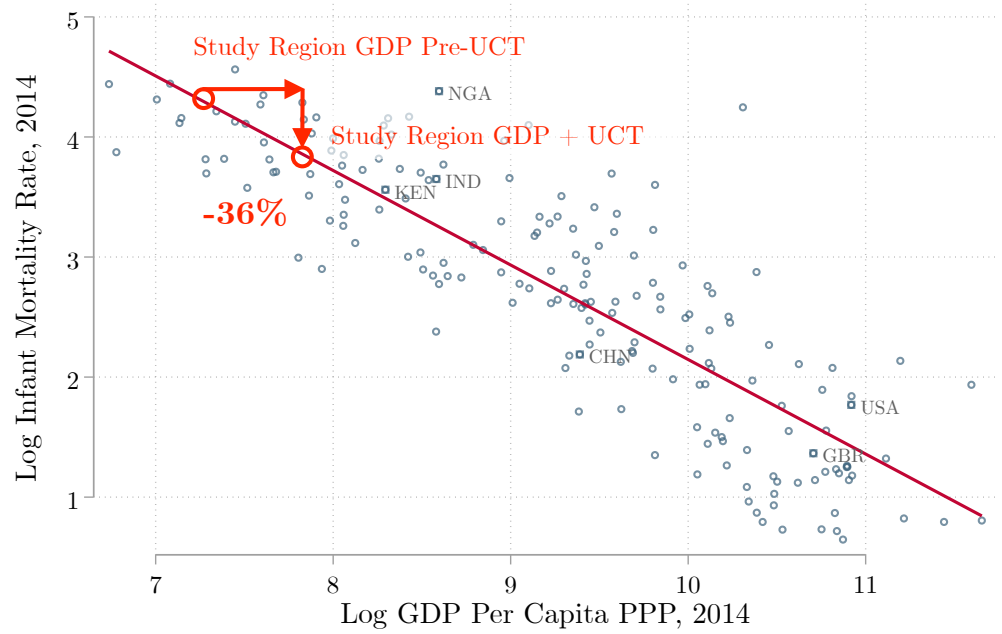

Panel B: Cross-Country GDP - Infant Mortality Relationship

Notes: This figure is based on the main EL3 birth census sample. Panel A reports infant mortality in the study region before and after Kenya's March 2020 COVID-19 lockdown using exact date of birth data from this project's birth census, utilizing all censused births with available data. The data in Panel B are shown as a binned scatter plot. Panel B reports the cross-country association between the natural log of the infant mortality rate (per 1000 births) as reported by the United Nations (2024) and the natural log of GDP per capita (in USD PPP) in 2014 as reported by the World Bank (2024). The "Study Region GDP Pre-UCT" dot highlights the study region's per capita GDP in 2014 as estimated by the Kenya National Bureau of Statistics (2019) as well as predicted infant mortality based on the cross-country relationship. The "Study Region GDP + UCT" dot illustrates the study region's per capita GDP adding a UCT of the same fraction of GDP as the transfer disbursed as a fraction of household consumption (75%) as part of this project (and – consistent with the evidence – assuming it is spent within a year), and the predicted change in infant mortality at that level of income based on the cross-country relationship. Note that the Kenyan national average ("KEN") lies nearly on the regression line and at a higher per capita income level than the rural study region, since the national average also include far richer urban areas like Nairobi.

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

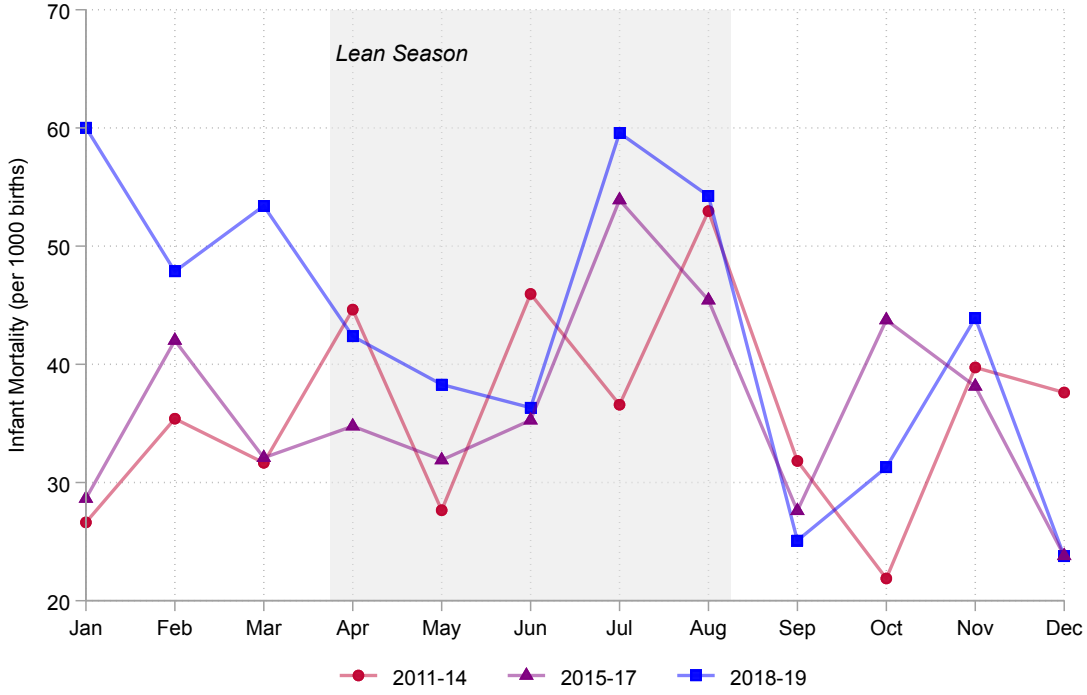

Panel A: Infant Mortality Means By Month

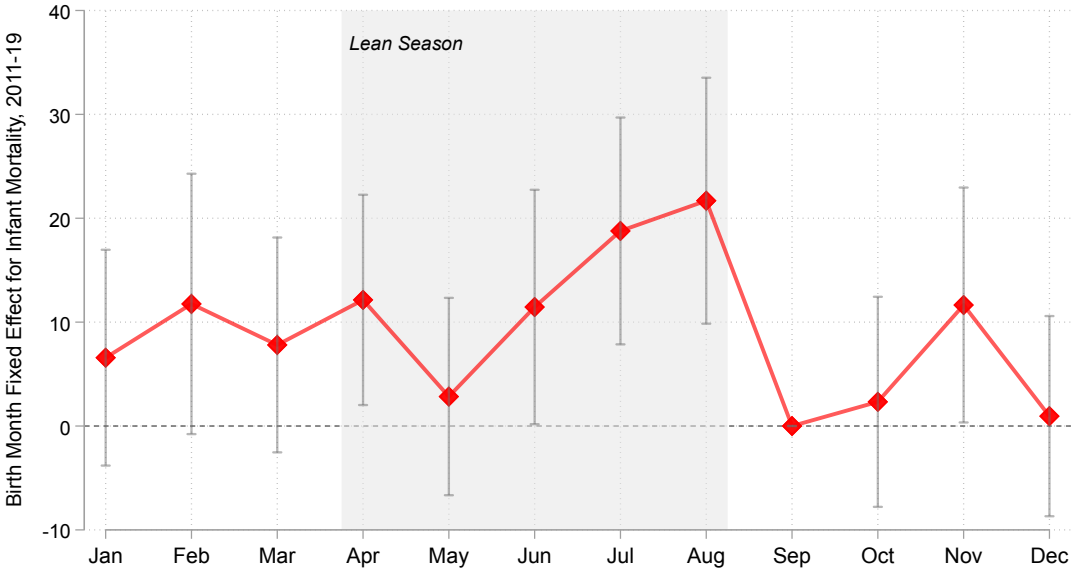

Panel B: Birth Month Fixed Effects

Notes: This figure is based on the main EL3 birth census sample. 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.5: Share of Funds Allocated to UCTs and Malaria Medicine by Regret Type

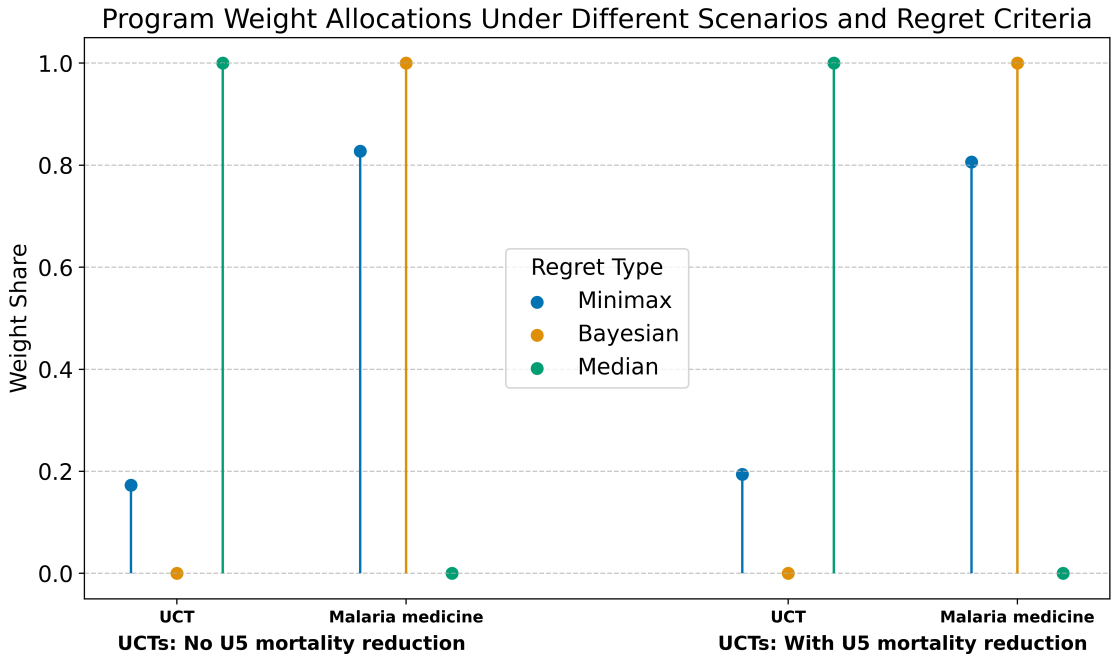

Notes: This figure reports the share of funds allocated to UCTs or Malaria medicine by an agent under minimax, Bayesian (mean), and median regret excluding and including the mortality reduction benefits of UCTs. The agent faces uncertainty about the value of a statistical life (VSL), and they minimize 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.

Table A.1: Descriptive Statistics for the Census of Births

	(1) Infant Mortality	(2) Child Mortality	(3) Neonatal Mortality	(4) Maternal Age at Birth
Births to All Households (2011-23)				
Mean	33.71	46.69	15.11	25.65
Observations	101,336	101,348	100,498	95,353
Births to Households Present at Baseline (2011-23)				
Mean	33.61	46.40	18.50	25.72
Observations	77,847	77,849	77,229	73,247
Births to Eligible Baseline Households, UCT Period (2015-17)				
Mean	34.54	50.68	19.23	25.55
Observations	6,334	6,335	6,240	6,003
Births to Eligible Baseline Households, Pre-UCT Period (2011-14)				
Mean	35.43	54.18	14.11	24.73
Observations	8,634	8,635	8,504	8,208
Births to Eligible Baseline Households, Post-UCT Period (2018-21)				
Mean	37.89	51.54	20.09	26.35
Observations	9,577	9,577	9,457	9,129
Births to Eligible Baseline Households (2011-21)				
Mean	36.15	52.25	17.77	25.57
Observations	24,545	24,547	24,201	23,340
Births to Ineligible Baseline Households (2011-21)				
Mean	31.28	45.44	17.39	25.60
Observations	40,872	40,872	40,712	38,205

Notes: This table reports means for descriptive characteristics and observation counts for groups in the EL3 census of births. The four descriptive characteristics reported are the infant under-one mortality rate (in deaths per thousand births), the child under-five mortality rate (deaths per thousand births), the neonatal mortality rate (deaths per thousand births), and the mean maternal age at birth. The five groups in the census for which these data are presented are as follows: all births recorded in the census (across 2021-23), all births to households present at baseline (2011-23), births to transfer-eligible households present at baseline across 2015-17 (our main sample under study), births to eligible households present at baseline across 2011-14, births to eligible households present at baseline across 2018-21, births to eligible baseline households across 2011-21, and births to ineligible baseline households across 2011-21.

Table A.2: Balance on Tracking, Study Area Presence, and Survey Completion

	(1) Household Tracked (Census)	(2) Present in Study Area (Census)	(3) Completed Census (Census)	(4) Household Tracked (Survey)	(5) Completed Survey (Survey)
Treatment Village	0.001	-0.006	0.006	0.005	0.008
	(0.006)	(0.004)	(0.004)	(0.005)	(0.007)
High-Saturation Sublocation	-0.006	0.008*	-0.005	-0.003	0.004
	(0.008)	(0.004)	(0.006)	(0.004)	(0.007)
Treatment Villages in	-0.005	0.002	0.001	0.002	0.012
High-Saturation Sublocations	(0.009)	(0.004)	(0.007)	(0.005)	(0.009)
Control Mean	0.897	0.915	0.905	0.965	0.893
Observations	62742	52582	48125	8372	8075

Notes: This table reports the share of households present at baseline (2014-15) tracked in the third endline (EL3) census, present in the study area at EL3, completed the EL3 census, tracked in the EL3 survey, and completed the EL3 survey, as well as balance by treatment status. The share tracked in the census (column 1) is reported as a share of baseline households non-deceased in the second endline census (EL3). The share present in study area (column 2) is reported as a share of non-deceased baseline households tracked in EL3. The share completing the EL3 census (column 3) is reported as a share of baseline households present in the study area, excluding holiday homes which do not comprise a household's primary residence. The share tracked in the EL3 survey (column 4) is reported as a share of baseline households sampled for the EL3 survey. The share completing the EL3 survey (column 5) is reported as a share of baseline households tracked in the survey. Observations encompass split households, which were created since 2014-15 from the 65,383 parent households originally present at baseline. Standard errors are clustered at the sublocation level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.3: 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	-5.74	-11.96*	-3.56	17.63*
	(5.85)	(6.38)	(4.00)	(10.22)
MHT adjusted p-value	[0.234]	[0.110]		
High-saturation sublocation	-12.13**	-5.68	-6.00	11.78
	(5.04)	(6.66)	(3.75)	(10.50)
Treatment village in high-saturation sublocation	-17.87***	-17.64***	-9.56**	29.41***
	(4.94)	(5.86)	(3.80)	(9.75)
Panel B: Spatial IV results				
Spatial IV: Own village effect	-7.98*	-12.72**	-3.56	17.83*
	(4.82)	(5.55)	(3.62)	(9.55)
Spillover effect	-11.49*	-12.91	-11.00*	26.94**
	(6.84)	(8.12)	(6.16)	(13.12)
ATE on eligibles	-19.46***	-25.63***	-14.57**	44.78***
	(6.94)	(8.54)	(5.75)	(13.47)
MHT adjusted p-value	[0.044]	[0.036]		
Control Mean	40.21	57.37	23.05	1,735.17
Observations	6,317	6,318	6,237	6,257

Notes: This table is based on the main EL3 birth census sample and includes estimated effects on pre-specified secondary outcomes. 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. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.4: 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	-6.85	-13.10**	-9.02**	-13.79***
	(5.53)	(6.13)	(4.44)	(5.26)
High-saturation spillovers	-9.63**	-3.18	-8.35	-9.84
	(4.84)	(6.52)	(5.47)	(7.33)
ATE in high-saturation sublocations	-16.48***	-16.28***	-17.38***	-23.63***
	(4.64)	(5.60)	(5.73)	(8.11)
Percent reduction in HS sublocations	52.04%	33.24%	54.88%	48.24%
Control Mean	31.66	48.98	31.66	48.98
Observations	6,266	6,267	6,266	6,267

Notes: This table is based on the main EL3 birth census sample and estimates effects on infant and child mortality excluding stillbirths. The sample excludes 51 stillbirths that were classified as such by the SmartVA verbal autopsy algorithm as opposed to directly reported by the household as a stillbirth. 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. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.5: Mortality Effects of Cash Transfer Spillovers on Non-Recipients

	(1) Infant Mortality 2015-17	(2) Under-5 Mortality 2015-17	(3) Neonatal Mortality 2015-17	(4) Days survived under-5 2015-17
Total effect IV	0.74	2.20	-4.04	-1.15
	(3.84)	(4.96)	(3.01)	(7.27)
Control Mean	31.82	46.84	19.21	1,753.45
Observations	13,191	13,192	13,102	13,154

Notes: This table reports estimates of spillover effects of the cash transfers on non-recipients of cash (both ineligible households and eligible households in control villages) on mortality for births across 2015-2017. Estimates are constructed using a spatial instrumental variable approach similar to Equation 2. Spatial HAC standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.6: Heterogeneity by Gender, Maternal Age, and Birth Order

	Infant Mortality			**Child Mortality**		
	(1)	(2)	(3)	(4)	(5)	(6)
	Male	Maternal Age 25+	Birth Order 2+	Male	Maternal Age 25+	Birth Order 2+
Treatment Village	-12.80*	-5.80	-14.34*	-12.92*	-12.92*	-22.77***
	(7.02)	(7.26)	(8.39)	(7.09)	(7.09)	(8.27)
High-Saturation Sublocation	-6.11	-7.42	-2.34	1.35	1.35	2.74
	(6.98)	(6.06)	(7.62)	(7.56)	(7.56)	(9.32)
Treatment Village x Heterogeneity	8.53	-4.57	9.98	-3.06	-3.06	13.63
	(9.04)	(9.57)	(9.26)	(10.32)	(10.32)	(10.54)
High Saturation Sublocation x Heterogeneity	-8.45	-5.90	-13.68	-10.36	-10.36	-11.34
	(9.09)	(9.71)	(10.64)	(11.00)	(11.00)	(11.51)
Treatment Villages in	-18.91***	-13.22*	-16.68**	-11.56	-11.56	-20.03**
High-Saturation Sublocations	(6.75)	(7.44)	(7.52)	(8.52)	(8.52)	(8.96)
Treatment in High-Saturation Sublocations	-18.83***	-23.69***	-20.39***	-24.99***	-24.99***	-17.74**
Including Heterogeneity Characteristic	(7.01)	(7.29)	(6.70)	(8.48)	(8.48)	(7.62)
Control Mean	39.95	39.95	39.95	57.00	57.00	57.00
Observations	6256	6331	6256	6332	6332	6257

Notes: This table is based on the main EL3 birth census sample. Effects are estimated using Equation (1) augmented with an interaction with the characteristic of interest. Maternal age at birth is analysed using pre-specified five-year age groups (under 20, 20-24, 25-29, 30-34, and 35 and above); this analysis splits on the median five-year group, with 46% of births having a mother below the age of 25. Child birth order is calculated using the order of births by mother across the 2011-23 period and does not take into account births to mothers prior to 2011 as information on those births was not obtained in the census. Standard errors clustered at the sublocation level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.7: Unclassified Verbal Autopsies by Treatment Status

	(1) Undetermined VA (2011-14)	(2) Undetermined VA (2015-17)	(3) Undetermined VA (2018-21)
Treatment Village	0.06	-0.10*	-0.03
	(0.05)	(0.05)	(0.05)
High-Saturation Sublocation	0.03	0.08	0.03
	(0.05)	(0.06)	(0.05)
Treatment Villages in	0.09	-0.02	0.00
High-Saturation Sublocations	(0.06)	(0.07)	(0.05)
Control Mean	0.11	0.21	0.16
Observations	263	201	325

Notes: This table reports regressions exhibiting the association between treatment status and the probability that a completed verbal autopsy (VA) had an undetermined cause by the SmartVA classification algorithm for transfer-eligible under-5 deaths across three periods: 2011-14 (pre-UCT), 2015-17 (the period contemporaneous with UCT disbursal), and 2018-21 (post-UCT). Standard errors are clustered at the sublocation level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.8: Cross-Sectional Relationship Between Socio-Demographic Status and Infant Mortality in Control Eligible Households

	(1) Under-1 Mortality	(2) Under-1 Mortality	(3) Under-1 Mortality	(4) Under-1 Mortality
Transfer-Eligible	10.54**	10.58**		
	(5.25)	(5.31)		
Below-Median Baseline Assets			50.88**	60.51**
			(23.74)	(25.34)
Mean in Low SES Group	39.95	39.95	81.38	81.38
Mean in High SES Group	29.40	29.40	30.50	30.50
Observations	5202	5202	489	489
Birth Characteristic Controls	NO	YES	NO	YES

Notes: This table reports regressions exhibiting the cross-sectional association between infant mortality and markers of socio-demographic status for births to eligible households in control low-saturation villages across 2015-17. Columns 1-2 present the association between eligibility for the unconditional cash transfers (determined by having a thatched roof) and infant mortality. Columns 3-4 present the association between holding a level of assets at baseline below the median level and infant mortality within the sample of transfer-eligible households surveyed in the representative baseline survey sample. Columns 1 and 3 present bivariate relationships, whereas Columns 2 and 4 include pre-specified controls for birth and maternal characteristics (i.e., gender, year of birth, and maternal age). Robust standard errors are utilized. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.9: Heterogeneity: Complementarity with Health Services (Delivery in indicated facility type)

	(1) Physician-staffed facility	(2) Hospital
Panel A: Spatial IV, no controls		
Total Effect IV (Time: Above median)	-0.04	0.11
	(0.07)	(0.09)
Total Effect x Below median time to facility	0.07	0.27**
	(0.12)	(0.14)
Panel B: Spatial IV, double-partial out LASSO controls, with treatment interactions		
Total Effect x Below median time to facility	0.04	0.33**
	(0.14)	(0.13)
Average time to facility (above median, minutes)	55.5	52.6
Average time to facility (below median, minutes)	22.5	18.8
Control Mean, above median time	0.26	0.39
Control Mean, below median time	0.30	0.45
Observations	1,075	1,110

Notes: The outcome is an indicator for delivery in a physician-staffed facility in column (1) and a hospital in column (2). Estimates in column (1) use the time to a physician-staffed health facility estimated using GPS measurements of travel speeds obtained via the study team during travel throughout the study area. Physician-staffed facilities were measured via clinic surveys in 2024. Estimates in column (2) are similar, but consider time to a hospital. If the facility reported it was open in 2014 and had a physician employed during the survey the facility is included in column 1. Column 2 includes level 4 and higher facilities surveyed, plus those categorized level 4 or higher on the Kenya Master Health Facility List that were unsurveyed. Rows under "Double-partial out LASSO controls, with 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. Covariate times treatment interactions are also included. Spatial HAC standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.10: UCTs and Household Labor Supply by Gender Around the Time of Birth

	(1) Women Hours Worked Last Week	(2) Men Hours Worked Last Week
Panel A: First 6 Months In-Utero		
Total Effect IV	-7.06	1.94
	(11.14)	(18.56)
Panel B: 3rd Trimester & First 3 Months Postpartum		
Total Effect IV	-20.79**	1.40
	(8.64)	(11.61)
Panel C: Next 6 Months Postpartum		
Total Effect IV	-8.48	-3.15
	(10.82)	(16.01)
Control Mean	40.79	46.50
Observations	876	659

Notes: This table encompasses households surveyed in the first endline survey (2016-17) with a recorded birth in the EL3 census during 2015-17. Labor supply encompasses hours worked across the household (in agricultural employment, non-agricultural self-employment, and wage employment), as well as hours spent searching for work, in the week prior to the first endline survey. The table displays the total estimated effect of the cash transfers on labor supply for three groups: households with a woman in the first 6 months of pregnancy when surveyed, households with a woman in the third trimester of pregnancy or who gave birth in the past 3 months when surveyed, and households with a woman who gave birth 3-9 months ago when surveyed. Treatment effects are estimated separately for each group using Equation (2), augmented with interactions with the group of interest as well as indicators for each group. Spatial HAC standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.11: 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	0.05	-5.07	0.05	-1642.99	-7399.24	0.09
	(0.33)	(4.18)	(0.04)	(2452.42)	(4573.36)	(0.28)
Control Mean	8.50	26.69	0.72	12326.22	33837.74	4.34
Observations	1354	1354	1354	1354	1354	1354

Notes: This table encompasses transfer-eligible households which reported a birth from 2015-17 in the EL3 census and were surveyed at baseline. Regressions are reported exhibiting the association between treatment status and six baseline household socio-demographic characteristics (maximum years of education of household members, average household age, number of adults, number of children, baseline household assets, and baseline household income). Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.12: 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	0.35	0.11
	(0.96)	(0.10)
Control Mean	8.72	7.10
Observations	1354	1354

Notes: This table encompasses mothers in transfer-eligible households present at baseline sampled in the baseline survey, and reports regressions exhibiting the association between treatment status and the probability of birth among mothers who actually gave birth over 2015-17. In column 1, the probability of birth is predicted using six baseline household socio-demographic characteristics (maximum years of education of household members, average household member age, marital status of primary respondent, baseline household income, baseline household 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 LASSO from a set of 243 baseline household- and village-level characteristics. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table A.13: Impacts on Infant Mortality in Household Survey Sample

	(1) Pre-Specified Controls	(2) Socioeconomic Controls	(3) PDS LASSO-Selected Controls
Panel A: Spatial IV Estimates			
Total Effect IV	-36.16*	-35.41*	-37.83**
	(18.81)	(19.16)	(18.70)
Panel B: Reduced-Form Estimates			
Treatment Village	-25.76*	-26.13*	-26.48**
	(12.97)	(13.24)	(12.94)
High-Saturation Sublocation	-10.97	-8.10	-11.96
	(11.81)	(11.83)	(11.60)
Treatment Villages in High-Saturation Sublocations	-36.73**	-34.23*	-38.44
	(17.91)	(17.90)	(17.95)
Control Mean	65.83	65.83	65.83
Observations	1486	1486	1486

Notes: This table is based on births reported in the EL3 birth census to transfer-eligible households present at baseline and surveyed in the baseline household survey. Effects estimated by Equation (2) (Panel A) and Equation (1) (Panel B) of the cash transfer on infant mortality rates are reported. 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 only pre-specified (e.g., year of birth fixed effects, birth gender, mother age group) covariates. Column 2 additionally 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 the pre-specified controls. Column 3 includes controls selected by PDS LASSO, with a total of 243 baseline variables available for selection, along with pre-specified controls. For consistency across the two equations, we use the PDS LASSO covariates selected for Equation (1) for Equation (2) as well. Spatial HAC standard errors (Conley, 2008) with a cutoff of 10km are reported for IV estimates, and standard errors clustered at the sublocation level are reported for reduced-form estimates.* $p < .10$, ** $p < .05$, *** $p < .01$.

B 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. 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 reported above.

B.1 Cross-Sectional Variation

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 Columns 1-2 of Appendix Table A.8, 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 treated eligible and ineligible households. Across 2015-17, under-one mortality among control eligible births was 10.5 deaths per thousand higher than among control ineligibles, a 36% 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. Furthermore, even within eligible households, births to households with below-median assets at baseline are substantially less likely to survive their first year in control villages: Columns 3-4 of Appendix Table A.8 show that infant mortality among households with below-median assets is 2.7 times higher (81.4 deaths per thousand births) than among households with above-median assets (30.5 deaths per thousand). 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.

B.2 Seasonality in Infant Mortality

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). Burke et al. (2019) define the pre-harvest "lean season" in rural western 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 21.9 deaths per thousand births higher than for infants born in the next month (Appendix Figure A.4). This difference – which is similar in magnitude to the cash transfer treatment effect we estimate – 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 virtually identical patterns.

B.3 Aggregate Economic Shocks

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 reduced 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).

As data were collected on exact date of birth for the majority of children censused, we can more precisely examine the impact of the COVID-19 shock on child health outcomes. On March 27, 2020, the Kenyan government imposed a strict lockdown intended to combat the spread of COVID-19, including a nationwide curfew during evening and night hours (Achuka and Gisesa, 2020). Individuals who left their homes during the curfew were subject to a maximum of three months imprisonment, a Ksh. 1000 (USD PPP 23.5) fine, or both, and local media reported numerous instances of police brutality meted out on Kenyans breaking curfew (Mireri, 2020).

A regression discontinuity analysis indicates that the infant mortality rate sharply increased immediately after the lockdown. Panel B of Figure A.3 illustrates that infant mortality doubled in the study region the week after the lockdown, a result statistically significant at the five percent level. These effects are likely caused by changes to health care access rather than COVID-19 deaths per se: reported COVID-19 deaths exhibit no such discontinuity at the lockdown date, and mortality from COVID-19 is largely concentrated among the elderly as opposed to children (Ma et al., 2021). Furthermore, the jump in infant mortality is particularly high in areas near a physician-staffed health facility, suggesting that a decline in access to care may have contributed to part of the deterioration in child health outcomes. The sharp rise in infant mortality in the study setting during the COVID-19 lockdown is consistent with non-experimental evidence from India suggesting that mortality among infants at six months rose 44% during India's lockdown, with the authors positing that reduced healthcare access and utilization may have driven this decline in child survival (Asker et al., 2025).

B.4 Cross-Country Association Between Income and Health

We last examine the cross-country relationship between per capita GDP and infant mortality. As previously documented by others, income and health are closely associated with each other (Preston, 1975; Cutler et al., 2012). Data on per capita GDP is obtained from the World Bank (2024), whereas infant mortality estimates are sourced from the United Nations (2024).[37] In 2014, the year in which cash transfers in our study region commenced,

[37] We also utilize an estimate of the study region's GDP, which is obtained from the Kenya National Bureau of Statistics (2019).

each log point increase in GDP per capita was associated with a 0.79 log point reduction in infant mortality, a relationship statistically significant at the one percent level and with an R-squared of 0.74. As seen in Panel A of Figure A.3, augmenting the study region's GDP by the same proportion as the cash transfers did for treated households would result in a 36% decline in infant mortality as predicted by the global cross-country GDP-mortality relationship. While the cross-country association is correlational rather than causal in nature, the magnitude of this predicted reduction is notably similar to our experimental estimates.

In summary, child mortality appears to be highly sensitive to economic conditions in this study's context, irrespective of the source of variation. Non-experimental approaches recover similar magnitudes as the large experimental effects estimated in this study.

C Details of Benefit-Cost Calculations

This appendix details the estimates presented in Section 7. We first describe how the estimated posterior distribution of VSLs presented in Figure 8 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 location parameter 0 and scale parameter 0.5. A half-Cauchy distribution with a scale parameter of 0 and scale of 0.5 is a recommended best practice for imposing weakly informative priors (Reis et al., 2022). We multiply these draws by 5 to ensure priors do not favor VSL values near 0. 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 8 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 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.