

BENCHMARKING A CHILD NUTRITION PROGRAM AGAINST CASH: EVIDENCE FROM RWANDA*

Craig McIntosh[†] and Andrew Zeitlin[‡]

December 18, 2020

Abstract

We present the results of a study that benchmarks a multi-dimensional child nutrition intervention against the same donor program cost being disbursed directly to beneficiary households to spend as they see fit. Randomized variation in cash transfer amounts allows us to estimate impacts of cash transfers at expenditure levels equivalent to the in-kind program, as well as to estimate the return to increasing cash transfer values. Focusing on poor households with children or pregnant mothers, we find that neither the in-kind program nor a cost-equivalent transfer of \$134 per household moves core child outcomes within a year. By contrast, a much larger cash transfer of more than \$500 improves consumption and investment outcomes, as well as driving modest improvements in dietary diversity and child growth. We find little evidence of heterogeneity that would allow overall benefits to be improved by selectively targeting modalities, whether by pre-selected characteristics or machine learning. We find one-time cash transfers to generate as much benefit as a flow of transfers across 12 months, and little additional benefit from allowing households to choose their cash modality. Our methodology allows us to contrast a comparative cost effectiveness approach with one that looks at differential impacts across programs at a cost equivalent level.

Keywords: Experimental Design, Cash Transfers, Malnutrition

JEL Codes: O12, C93, I15

*We are grateful to DIV, Google.org, and USAID Rwanda for funding, and to USAID, CRS, GiveDirectly, CEGA/DIL, and IPA for their close collaboration. We thank Leodimir Mfura and Marius Chabi for overseeing the fieldwork, and Richard Appell, Sarait Cardenas-Rodriguez, Chris Gray, Ali Hamza, and Bastien Koch for research assistance. This project is covered by Rwanda National Ethics Committee IRB 143/RNEC/2017 and IPA IRB 13730, and the study is pre-registered with the AEA as trial AEARCTR-0002559. This study is made possible by the support of the American People through the United States Agency for International Development (USAID.) The contents of this study are the sole responsibility of the authors and do not necessarily reflect the views of USAID or the United States Government.

[†]Corresponding author: University of California, San Diego. 9500 Gilman Dr., La Jolla CA, 92093-0519. Ph: 858-822-1125, Fx: 858-534-3939. ctmcintosh@ucsd.edu

[‡]Georgetown University, andrew.zeitlin@georgetown.edu

I. INTRODUCTION

Should governments and aid agencies provide assistance in kind, or would beneficiaries be better off if they just received the money? There are several reasons to compare impacts of in-kind programs to those generated when its cost is distributed directly as a cash transfer. From a practical perspective, cash transfers generate well-documented benefits for households (Aizer et al., 2016; Haushofer and Shapiro, 2016), and their costs are falling across the developing world due to the penetration of mobile money (Suri, 2017). Cash programs are flexible and can be rescaled in terms of targeting or expense, which allows them to be paired at cost-equivalent levels against more complex interventions.¹ By contrast, comparisons across programs in different contexts and at different costs rely on strong assumptions about external validity (Vivalt, 2015) and linearity (both in cost-effectiveness ratios and the preferences of policymakers across individuals). More deeply, decisions made over the use of unconditional transfers may reveal information about the preferences of the beneficiary households that is important in interpreting the welfare effects of more complex programs (Finkelstein and Hendren, 2020). While missing markets, externalities, or other failures can justify the provision of in-kind programs, benchmarking the impact of such programs against cost-equivalent cash transfers forces policymakers to be explicit about the circumstances that would merit a more multidimensional, overhead-heavy approach. For these reasons, recent years have seen increasing calls for head-to-head studies that use cash as a benchmark for more complex forms of international development assistance.²

In this study we provide such a benchmark, comparing a nutritional and maternal health intervention to a cash transfer in Kayonza and Nyabihu districts of Rwanda. This is an ideal setting and intervention for such a comparison. Nutritional deprivation is a pressing issue in both districts, which feature under-five stunting levels of 34.8 and 44.9%, respectively (DHS, 2016). The in-kind program takes a multi-dimensional approach widely used by USAID and other donors to combat malnutrition, one that has been broadly advocated in the public health literature (Ruel et

¹In medical research, new interventions are typically benchmarked against the best current treatment, referred to as the ‘standard of care’. Perhaps the best-validated anti-poverty tool in the development literature is BRAC’s Targeting the Ultra-Poor program (Banerjee et al., 2015); however given that this program is expensive (\$3,700 per household), lengthy in duration (a year of intensive intervention), and targeted at very specific households (the ultra-poor), it represents a somewhat inflexible benchmark for the broad range of potential development interventions.

²In particular Blattman and Niehaus (2014) argue for the use of cash transfers as the ‘index funds’ of international development, providing a reference rate of return that could be used to hold donors accountable.

al., 2013) and found to have positive impacts in neighboring Burundi (Leroy et al., 2016). The complex and overhead-intensive nature of such a holistic, bundled intervention makes a stark contrast with the pared-down nature of the cash-transfer alternative. And yet there is reason to believe that cash transfers can improve child growth outcomes (Aguero et al., 2006; Baird et al., 2019; Seidenfeld et al., 2014), and that gains achieved in early childhood might have lasting effects through the accrual of human capital (Currie and Almond, 2011; Hendren and Sprung-Keyser, 2020).

To form this comparison, we conduct a cluster-randomized trial across 248 Rwandan villages.³ We enroll poor households containing children or pregnant women, with an emphasis on the 1,000-day window of opportunity from pregnancy until a child’s second birthday (Currie and Almond, 2011). In the in-kind arm, eligible beneficiaries are offered Catholic Relief Service’s *Gikuriro* program.⁴ *Gikuriro* aims to improve child nutrition through superior information, direct transfer of productive assets, and improvements in household diet and sanitation. It consists of four components targeted directly at beneficiary households: a Village Nutritional School, Farmer Field Learning Schools, Savings and Internal Lending Communities, and a Water, Sanitation, and Hygiene intervention, as well as a Behavior Change Communication intervention implemented at the village level.

Gikuriro is compared to a control arm as well as to a study arm that received unconditional cash transfers via mobile money, implemented by the U.S. non-profit GiveDirectly. Such transfers have been found to improve consumption and/or dietary diversity in multiple contexts across Sub-Saharan Africa (Aker et al., 2016; Haushofer and Shapiro, 2016). At the village level, transfer amounts were randomized in a range around the expected cost of *Gikuriro*, with a subset of villages receiving a much larger transfer. Regression adjustment then allows us to calculate an exact cost-equivalent comparison of cash versus in-kind programming across *modalities*, while the larger cash transfer allows us to compare benefits across potential *scales* of intervention. Within these GiveDirectly villages, we implemented a household-level experiment, whereby beneficiary households were randomized to receive transfers in a lump-sum or as a 12-month flow, or to have their choice between these options. The ability to understand the role of cash modality preferences gives us a window into the extent to which self- and other-control problems drive the use of unconditional cash.

³The village-level study design was motivated both by the clustered nature of the *Gikuriro* intervention, and helps to allay concerns about the potential for spillovers of cash transfers.

⁴*Gikuriro* means ‘well-growing child’ in Kinyarwanda.

Thirteen months after baseline, we measure impacts on five primary outcomes: household consumption, household dietary diversity, child and maternal anemia, child growth, and household non-land net wealth. These outcomes are documented as core drivers of improved long-run outcomes for children (Hoddinott et al., 2013; Maluccio et al., 2009), as well as being plausibly impacted by both intervention types. In addition, we examine a set of secondary outcomes including borrowing and savings, fertility, health knowledge and sanitation practices, diseases and mortality, household assets, and the quality of housing stock.⁵ Prior to randomization the survey firm classified all households in study villages as ‘eligible’ (identifiable using administrative data sources as containing underweight children, or households in the bottom two income categories with children 5 years old or younger or with pregnant or lactating women), or ‘ineligible’ (everyone else). We can therefore measure impacts both on the mutually agreed-upon intended target population as well as on the study villages as a whole, even in the presence of potential differences in actual targeting across implementers. Using the eligible sample we can estimate experimental intention-to-treat effects, and using the full sample (population weighted) we estimate total causal effects on the average household in study villages.

This design, like that of the follow-on study that benchmarks cash against a youth employment program (McIntosh and Zeitlin, 2020), allows us to assess tradeoffs across *outcomes*, across *scales*, and across *sub-populations*. Across outcomes, our results provide a nuanced view of the relative impact of a highly-tailored child malnutrition program and cost-equivalent cash transfers. Impacts of both cash and in-kind programs are multidimensional. At a cost of \$142, Gikuriro was successful at delivering gains in savings, a domain that was the target of one core program component. It did not lead to improvements in consumption, dietary diversity, wealth, child anthropometrics, or anemia within the thirteen-month period of the study. A cost-equivalent cash transfer allows households to pay down debt, and generates increased investment in productive and consumption assets. Here, then, is a first type of tradeoff that this design can reveal: should policymakers prefer an increase in savings versus the decrease in debt that households choose to engage in when constraints are relaxed?

Across scales, our results show that a much larger cash transfer costing \$567 led to across-the-

⁵All primary and secondary outcomes were registered prior to receipt of endline data on the American Economic Association RCT Registry, with ID AEARCTR-0002559.

board improvements in consumption-based welfare measures, a substantial improvement in dietary diversity, a drop in child mortality, and modest improvements of about 0.1 standard deviation in the anthropometric indicators of height-for-age, weight-for-age, and mid-upper arm circumference (all significant at 10 percent or above). The large cash transfer delivers benefits even on outcomes specifically targeted by the other program. While it is unsurprising that very large amounts of money show up in consumption and productive assets, the improvements in diet and particularly child anthropometrics over such a short period of time are impressive. Further, while it may be unsurprising that the impact of cash transfers scales with amount spent in the way found here, the same may not be true of other types of development intervention that would quickly hit diminishing marginal returns once certain core objectives were achieved. For a given target population, policymakers must therefore decide how to trade off between relatively intensive interventions that can only be provided to a small subset, or lower-cost interventions that can reach a larger proportion of this target group.

Across subpopulations, our analysis speaks to tradeoffs in three distinct ways. First, policymakers face decisions about the extent to which they value outcomes among a narrowly targeted population, or the broader society in which they reside. Our study design allows us to calculate not only benefits per dollar cost among the program-eligible population, but also to compare the total causal effect of each intervention on the full population of study villages in a cost-equivalent manner and at comparable levels of treatment saturation. Despite 80 and 84 percent of the eligible households in treatment villages receiving Gikuriro and GiveDirectly, respectively (for the villages as a whole the treatment rates were 19 and 18 percent), neither treatment resulted in sufficiently widespread benefits as to be detectable at cost-equivalent levels in the general population, with the exception of an improvement in health knowledge and vaccination rates in Gikuriro villages and vaccination rates in GD large villages. Second, on the other hand, policymakers need not place equal welfare weights on all subgroups within the targeted beneficiary population, and the relative effects of cash versus in-kind transfers may differ across these groups. We test for heterogeneity in relative returns along observed subgroups of a priori interest—using consumption and child-growth levels at baseline—and find limited evidence of such heterogeneity in this instance. Third, since policymakers might also be able to fine-target intervention modalities based on beneficiary characteristics, we estimate the extent of heterogeneity in *relative* treatment response by application

of a generalized random forest (Athey et al., 2019; Wager and Athey, 2018). We find that while there is substantial heterogeneity in impacts of cash versus in-kind treatments for any one outcome dimension at a time, estimated subpopulation treatment effects are only weakly correlated across outcome dimensions. This suggests that fine targeting can exploit heterogeneity with tailored treatments to improve outcomes for policymakers with narrow objective functions, but that—at least in this case—such personalized treatments deliver only modest gains when policymaker objectives are broadly defined.

Finally, we build a sub-experiment into the cash arm of the study to shed light on the role that choice and self-control play in the comparative impact of cash transfers. Both the conventional wisdom and the behavioral literature have been concerned with the idea that liquid resources with no strings attached will be wasted or spent in ways that the households later would regret (Ashraf et al., 2006; Laibson, 1997). We speak to this issue using an incentivized choice experiment where households were asked to decide between receiving all their money in a lump-sum versus getting it in a flow over twelve months. While demand for the more rigid monthly cash transfers is substantial, we find no evidence of ‘sophisticated’ choice reflecting essential heterogeneity, and the benefits of flow or in-kind transfers are not larger for households that report self- or other-control problems at baseline. The striking absence of behavioral heterogeneity suggests that the higher-cost and more rigid programs do not deliver better outcomes even for those with behavioral challenges.

Our study extends a small number of studies using cash transfers as one arm in multi-armed randomized, controlled trials. The largest extant literature is based on the comparison of cash aid to food aid (Ahmed et al., 2016; Cunha et al., forthcoming; Hidrobo et al., 2014; Hoddinott et al., 2014; Leroy et al., 2010; Schwab et al., 2013), uncovering a fairly consistent result that food aid leads to a larger change in total calories while cash aid leads to an improvement in the diversity of foods consumed.⁶ Efforts to compare more complex, multi-dimensional programs against cash include BRAC’s Targeting the Ultra-Poor program (Chowdhury et al., 2017), microfranchising (Brudevold-Newman et al., 2017), agricultural inputs (Brudevold-Newman et al., 2017), and sustainable livelihoods (Sedlmayr et al., 2017). These studies have typically struggled with the question of how to anticipate costs and compliance well enough to realize an exact cost-equivalent comparison

⁶An interesting contribution to this food versus cash literature is (?), who find that poor and geographically marginalized households may prefer food transfers because of a lack of consumer markets during times of shortfall.

after the fact. Our study provides a methodology incorporating randomization of transfer amounts and ex-post, regression-based cost adjustment to allow cost-equivalent benchmarking comparisons to be made in a general way.

Beyond this, we contribute to a growing number of numerous studies that show meaningful impacts of cash transfers on important life outcomes in the short term, including child nutrition (Aguero et al., 2006; Seidenfeld et al., 2014), schooling (Skoufias et al., 2001), mental health (Baird et al., 2013; Samuels and Stavropoulou, 2016), teen pregnancy and HIV (Baird et al., 2011), microenterprise outcomes (De Mel et al., 2012), consumer durables (Haushofer and Shapiro, 2016), and productive assets (Gertler et al., 2012). The evidence on the long-term impacts of cash transfers is more mixed, but some studies have found substantial impacts (Aizer et al., 2016; Barham et al., 2014; Fernald et al., 2009; Hoynes et al., 2016).⁷

The rest of the paper is organized as follows. The next section lays out the study design, including a detailed description of the interventions, sampling routine, costing principles, experimental structure, and primary and secondary outcomes. Section III presents the core empirical results of the benchmarking exercise, as well as the results of sub-experiments on cash transfer modalities. Section IV estimates the extent of heterogeneity, both using pre-specified attributes as well as conducting a machine learning exercise to speak to optimal targeting rules. Section V discusses the choice experiment and evidence of the role of rigidity. Section VI concludes with a discussion of how results from such cash-benchmarking studies can be used to bound the preferences over benefit/cost ratios across outcomes required to justify each program.

II. STUDY DESIGN

In this section, we describe four dimensions of the study design: the interventions, the study outcomes, the allocation of villages and individuals to treatment arms, and the approach to using experimental variation in cash transfer values to make cost-equivalent comparisons across modalities.

⁷For examples of studies that find dissipating long-term benefits, see Baird et al. (2016) and Araujo et al. (2017). Evidence from systematic reviews of cash transfers on schooling (Molina-Millan et al., 2016) and child health (Manley et al., 2013; Pega et al., 2014) has been similarly uneven.

II.A. Description of Interventions

The Gikuriro program was developed by USAID, Catholic Relief Services (CRS), and the Netherlands Development Organization (SNV) to combat food insecurity among pregnant women and children, particularly during the critical first 1,000 days of life that play such a dominant role in later-life outcomes and cognition (Figlio et al., 2014). The resulting multi-faceted program brings together several components in order to attack this problem from multiple directions at once, and is a central pillar of the Government of Rwanda’s approach to combatting malnutrition in rural Rwanda.⁸ Gikuriro combines an integrated nutrition program with a standard WASH curriculum (water, sanitation, and hygiene), and seeks to build the capacity of the health infrastructure providing services to mothers and newborns, particularly Community Health Workers (CHWs). The program also seeks to improve livelihoods by providing additional assistance to eligible households, including (a) Village Nutrition Schools (VNS); (b) Farmer field learning schools (FFLS), which potentially includes distribution of small livestock, fortified seed, etc.; (c) Savings and Internal Lending Communities (SILCs); and (d) the Government of Rwanda’s Community-Based Environmental Health Promotion Program (CBEHPP). In addition, Gikuriro provided a program of Behavioral Change Communication (BCC), supporting participation in all components of the program including savings, agriculture, and nutrition, as well as hygiene. This comprehensive approach seeks to build supply and demand for child health services simultaneously, and is fairly typical of the kinds of multi-sectoral child health programs implemented by USAID in many parts of the developing world.⁹

To benchmark the impact of this program against cash we worked with GiveDirectly, a US-based 501(c)3 Non-Profit organization. GiveDirectly specializes in sending mobile money transfers directly to the mobile phones of beneficiary households to provide large-scale household grants in developing countries including Kenya, Uganda, and Rwanda. GiveDirectly’s typical model has involved targeting households using mass-scale proxy targeting criteria such as roof quality. GiveDirectly builds an in-country infrastructure that allows them to enroll and make transfers to households

⁸USAID’s Global Health and Nutrition Strategy explicitly calls for multi-sectoral interventions that incorporate agriculture, WASH, education, and outreach to mothers in the first 1,000 days through the public health system. The agency reports reaching 27 million children worldwide under the age of 5 in 2016 alone through such programs, which represent the prescribed USAID modality for Scaling up Nutrition (SUN) countries.

⁹Examples of similar integrated WASH/agriculture/child nutrition programs funded by USAID include SPRING in Bangladesh, RING in Ghana, Yaajende in Senegal, and ENGINE in Ethiopia.

while simultaneously validating (via calls from a phone bank) that transfers have been received by the correct people and in a timely manner. Their typical transfers are large and lump-sum, on the order of \$1,000, and the organization provides a programatically relevant counterfactual to standard development aid programs because it has a scalable business model that would in fact be capable of providing transfers to the tens of thousands of households that are served by the Gikuriro program. Because of the nutritional focus of the Gikuriro intervention, GiveDirectly incorporated a ‘nudge’ into the way the program was introduced (Benhassine et al., 2015), distributing a flyer at the time of the intervention that emphasized the importance of child nutrition. An English translation of this flyer is included in the Online Appendix. Given observed impacts of cash transfers on other goods, e.g., productive assets and housing value, it is evident that households felt at liberty to spend the grants on items not directly related to child nutrition.

Rwanda may be a particularly interesting environment in which to pose the benchmarking question for several reasons. First, child malnutrition rates overall are high—the prevalence of stunting among children under age five in the 2014-15 Demographic and Health Survey was 37.9 percent, underweight 9.3 percent and wasted 2.2 percent—though this represents an improvement in recent years (DHS, 2016). Second, Rwanda is a country notable in Africa for its bureaucratic capacity and the public health infrastructure has been successful in delivering substantial improvements in child and maternal health outcomes (NISR, 2015) through schemes such as Pay-for-Performance (Basinga et al., 2011). Hence, it may provide a relatively strong case in terms of interventions like Gikuriro that are led through the public health system and lean heavily on Community Health Workers (CHWs). Third, the Government of Rwanda has been experimenting extensively with cash transfer programs over the past few years, such as the inclusion of cash in the flagship *Umurenge* poverty reduction program (Gahamanyi and Kettlewell, 2015), the \$50 million ‘Cash-to-poor’ program supported by the World Bank, as well as a number of efforts to transition the support systems for the country’s large population of refugees to cash transfers (such as a World Food Programme (WFP) program that is now supporting 15,000 refugees in Gihembe Camp using cash rather than traditional in-kind aid mechanisms (Taylor et al., 2016)). Hence there should be the bureaucratic capacity to implement Gikuriro well, and there is both experience with and interest in cash transfers as a safety net modality in the country.¹⁰

¹⁰Given the framing provided by GiveDirectly and the unusually strong degree of social control exerted by local

II.B. Study outcomes and eligibility

We pre-committed to five primary outcomes and seven secondary outcomes for the study. The primary outcomes are consumption, dietary diversity, child growth (height-for-age and weight-for-age, as well as mid-upper arm circumference), and household non-land wealth. The secondary outcomes are borrowing and saving, pregnancy and live birth rates, health knowledge, mortality, health-seeking behaviors, productive assets, and housing quality. A more detailed description of all outcomes is provided in the Online Appendix. These data were collected through two survey waves; the baseline was conducted in August and September of 2016, and the endline was conducted in September and October of 2017. The baseline included a comprehensive household survey module as well as anthropometric measurement of all children five and under in eligible households.¹¹ The endline consisted of these same measures plus blood-based anemia testing for the study children and for all women of childbearing age.¹²

Both implementers made contact with the study subjects and began enrollment immediately after baseline. GD began implementation shortly after the baseline meaning that at endline individuals in that arm had experienced about 12 months of the household grants treatment (running up through the month before endline). Gikuriro was slower than the cash program to begin implementation on the ground; in that arm households had typically experienced 8-9 months of household-level implementation at the time of the endline.¹³ The duration of the RCT component of the study was limited by the fact that local governments wanted to hit targets for the broader, national rollout of nutritional and WASH programming, and hence we were not able to maintain the control groups for more than one year. We cannot therefore speak to the long-term impacts of the interventions. Anticipating this issue, we took two approaches to measurement. One of them was to try capture the stocks of intertemporal assets that would be the obvious conduits to future consumption benefits for the households. The second was to emphasize outcomes such as dietary diversity and

officials in the Rwandan context, it is certainly possible that our ‘unconditional’ transfers have been more forcibly devoted to child consumption than they would have been in a different context.

¹¹We weighed all children younger than 6 years once using a Seca 385 scale. We measured length for children under two years with Seca 417 length boards and height for older children to the nearest 0.1 cm with Seca 213 stadiometers. These were measured twice and we took the average of the two measurements unless they differed by more than 0.7 cm in which case we took a third measurement and averaged the two closest measurements.

¹²Using the guideline for anemia testing in population-based surveys, we used HemoCue Hb 301 system.

¹³Since both programs had six months of notice that they would be implementing in the study sample in these two districts and began national-level implementation at the same time, this differential delay likely reflects a real difference in the relative ramp-up speeds of cash versus more complex programming.

anemia that have the potential to respond quickly to changes in consumption patterns, while also retaining the more standard metrics of child malnutrition such as height for age (HAZ), weight for age (WAZ), and mid-upper arm circumference (MUAC).¹⁴ Further, a number of recent RCTs have shown that programs can have meaningful impacts on biometric outcomes over timeframes similar to that analyzed in this study, such as Desai et al. (2015), Leroy et al. (2016), Fink et al. (2017), and Null et al. (2018).

A major practical issue with this kind of comparative study is how to define the the target group, given that each implementer would naturally do this differently. We collectively resolved this by agreeing to a set of readily observable characteristics that both implementers deemed eligible for their funding and desirable to target, and then tasking the survey firm with listing all households in study villages and allocating them to the ‘eligible’ or ‘ineligible’ stratum. Households were considered eligible if they contained a child known to the government growth monitoring system to be malnourished, or else were in the poorest two government poverty classifications (Ubudehe 1 or 2) and contained either children under the age of 5 or a pregnant/lactating woman. All other households in study villages are classified as ineligible. The survey firm passed the names of all eligible households to the implementing partners in an identical way, and we then randomly sampled up to 8 eligible and 4 ineligible households per village for the study. The identity of the sampled households was not revealed to either implementer. This structure allows us to use survey weights to estimate impacts that are representative either for eligible households (which we call the Intention to Treat) or for all households in study villages (the Total Causal Effect). Impacts among ineligibles may arise either because the implementers treated some households outside of the IPA-defined eligible group, or because of spillovers from beneficiary to non-beneficiary households.

Both implementers concurred quite closely with our definition of eligibility on the ground, and compliance was high: we have 80 percent of the eligibles treated by Gikuriro and 84 percent by GiveDirectly.¹⁵ The IPA-driven process of eligibility definition however encountered an average of 13.9 eligible households per village, much lower than the rate of 30 anticipated. When Gikuriro

¹⁴Dietary diversity is an immediate indicator of improvements in consumption, and the clinical literature has shown that anemia tests respond within 3 months of improvements in diet (Habicht and Pelletier, 1990).

¹⁵Because eligibility status was determined from records, in some cases the Ubudehe status of the households proved to be incorrect or unverifiable when they were approached by GD for treatment, and hence were not offered the program. Gikuriro implemented a consultative process with community members that formed the basis of their targeting.

began their actual program implementation, they continued their standard consultative process for beneficiary identification, which included the use of soft targeting information not available to the survey firm. Using this additional, richer information set to target, they identified and treated an average of 25.8 households in study villages in Kayonza district, and 26.97 in Nyabihu. Since our first tranche of GiveDirectly treatments were only among IPA-defined eligible households, we found ourselves with a substantial discrepancy in the intensity of treatment across implementers in the ineligible group. We responded to this asymmetry by drawing in an additional sample of the poorest ineligible households in GiveDirectly villages to receive cash grants so as to maintain parity in village-level treatment intensity. Therefore, while the targeting of ineligible households differs across implementers, the treatment intensities across the two arms are identical by construction.

II.C. Design of the Experiment

Randomization occurred at the village level across 248 villages, using a blocked randomization where the blocks were formed by the combination of districts and village-level poverty scores within district, creating a total of 22 blocks with between 10 and 13 villages per block. Fixed effects for these blocks are included in all analysis. A computer was used by the researchers to conduct the randomization based on a frame of villages agreed to by CRS and government officials.

Table I presents a schematic of the design of the study. 74 villages were assigned to the Gikuriro intervention, 74 were assigned to the control group (no intervention), and 100 were assigned to GiveDirectly household grants. The GiveDirectly villages were further split into four transfer amounts, randomized at the village level. Three treatment amount arms, with 22 villages in each, received transfer amounts in a range around the anticipated cost of Gikuriro. A final 34 villages were assigned to the ‘large’ GiveDirectly transfer amount which was selected by GiveDirectly as the amount anticipated to maximize the cost effectiveness of cash. The transfers actually received by households in the GD ‘main’ arms were \$41.32, \$83.63, and \$116.91. Then, the large GD arm actually transferred \$532 to households. All transfer amounts were translated into Rwandan Francs at an exchange rate of 790 RwF per US dollar, and were rounded to the nearest hundred.¹⁶ Figure

¹⁶GiveDirectly believed that most cost-effective use of these funds would be to attempt to equalize the amount transferred per household member, rather than to have households of very different sizes receiving the same transfer amount. To accomplish this, we scaled the transfer amounts within a village by household size, such that larger households received larger transfers, but leaving the mean transfer amount at the village level unaffected. This formula first calculated the per-capita transfer for a village using household sizes and the desired average household

B.1 provides a box and whisker plot of the randomly assigned mean transfer amount per village relative to the actual amount received per household observed in the GD institutional data, and shows that the two correspond closely.

Within the GD arm we conducted an additional experiment that elicited choice over cash transfer modality in an incentivized way. We wished to test the hypothesis that a regular monthly flow of transfers is likely to be a more effective way of delivering the kinds of nutrition and health outcomes central to Gikuriro (Haushofer and Shapiro, 2016), and to understand how much of this heterogeneity in benefits is understood by the recipients themselves. Large Essential Heterogeneity in the sense of Heckman et al. (2006) would suggest that program choice could deliver superior overall impacts. To do this, all respondents were given a menu illustrating the choice between a single lump-sum transfer delivered in any of the 12 months from August 2016 to July 2017 or a flow of monthly payments totaling the same amount. They were informed that they would then be assigned to receive one or the other, and with 1/12th probability would receive what they chose. We then assigned 5/8ths of the cash sample to the flow arm (evenly divided over 12 monthly payments), the 1/4th to receive an up-front lump sum cash transfer, and the remainder to receive their choice. While we do not use the choice arm as a treatment in itself, being able to observe the choice for all cash subjects allows to understand whether those who ‘got what they wanted’ experience superior treatment effects.

II.D. Cost Equivalence, Before and After the Fact

A high-quality estimate of the costs of each program is a critical building block of cost-adjusted comparisons. The design-based approach to cost-equivalent comparisons across programs is hampered by the fact that we can only measure costs well when the programs have been fully implemented, but the head-to-head comparison would need to know these costs at the design phase. Anticipating this issue, we costed both programs in detail prior to and after the intervention period, following Levin and McEwan (2001). The ex-ante costing exercise arrived at an ex-ante cost of \$119 per eligible household. We then randomized transfer amounts at the village level to bracket this anticipated cost. Three smaller cash transfer costing \$77, \$119, and \$152 (with beneficiaries actually receiving \$41, \$84, and \$117, respectively) are used to form the cost-adjusted comparison.

transfer value. Second, it scaled household-level transfer amounts with household size, applying a minimum of 3 members and a maximum of 8 members, so as to achieve the intended mean transfer amount per household per village.

The larger cash transfer cost \$567 and transferred \$517. We pre-committed to a regression structure that linearly adjusts the cash cost to form the exact comparison to the final, ex-post costing figure.

The ex-post costing was conducted using the ‘ingredients method’, valuing inputs at their opportunity costs (see Dhaliwal and Tulloch, 2012; Levin and McEwan, 2001, for more discussion). The costing question is asked from the perspective of the donor (in this case, USAID), meaning that we consider the total money spent per beneficiary to achieve the benefits observed. Overhead expenditures in the implementation chain are an inherent part of these costs, and so the lower transactions costs in getting mobile money to the beneficiary play an important role in their potential attractiveness.¹⁷

Since the Gikuriro program covers eight districts (e.g. much larger than the study population only) and many of the startup and administrative costs must be amortized over this whole beneficiary pool, we calculate cost per beneficiary in the full national program. Because we did not want differences in scale to drive differential costs per beneficiary, we asked GiveDirectly to artificially scale up their operations and provide us with numbers reflecting the costs per beneficiary if they were running a national-scale program across eight districts, including 56,127 beneficiary households like Gikuriro. We costed each GD arm separately, asking what the overhead rate would have been if GD had run a national program at the scale of Gikuriro giving only transfers of that amount.¹⁸ This allows us to conduct the benefit/cost comparisons ‘at scale’, rather than having the artificial, multi-amount environment of the study contaminate the costing exercise across arms.

The ex-post costing exercise arrived at a cost per beneficiary household for Gikuriro of \$141.84. Actual GD costs were \$66, \$111, and \$145, and \$567, meaning realized Gikuriro costs were within the range over which we randomized but 28 percent higher than the average of the three smaller arms. Cost per beneficiary is not the natural counterpart to the Intention to Treat in cost/benefit terms, however, because the ITT ignores non-compliance and gives average outcomes in the entire eligible sample. Hence, to calculate cost per eligible household (which is what was actually spent to generate

¹⁷Costs are inclusive of all direct costs, all indirect in-country management costs including transport, real estate, utilities, and the staffing required to manage the program, and all international overhead costs entailed in managing the Gikuriro program. All administrative costs, including the appropriate share of the costs of maintaining international headquarters infrastructure, were included in the costing. Beneficiary identification costs, incurred by the survey firm and identical across all arms of the study, are excluded from the cost-benefit calculation. Monitoring and Evaluation costs, similarly, were excluded so as to be costing only the implementation component.

¹⁸Overhead costs as a percentage of the amount transferred decline sharply with transfer amount for GD because fixed costs represent a large share of their total overhead.

the outcome observed), we need to adjust for non-compliance in costing. We define ‘averted’ costs, as those which are not spent on a household if they do not comply with treatment; and ‘non-averted’ costs which will be expended whether or not the household complies. For GiveDirectly all costs are averted, while for Gikuriro all variable costs except for the village-level behavior change component are averted. Gikuriro’s non-averted costs are those conducted at the village level, namely (WASH and BCC), which drive 40 percent of the total cost of the program. The averted costs are those arising from direct household treatment; for these we multiply the cost per beneficiary times the compliance rate to calculate the cost per eligible and cost per household overall. Using this approach we can recover a cost-equivalent comparison even if the compliance rates are different across arms.

Table II provides the exact costing numbers arrived at by the ex-post exercise. Gikuriro treatment rates are 80 percent among eligibles and 19 percent in the population as a whole. Given an actual cost to USAID of \$141.84 per beneficiary and the role of averted costs, this gives a cost of \$124.49 per eligible household and \$28.02 per household in the village. GD compliance rates among eligibles range from 81-91 percent across arms, and among ineligibles are 18-19 percent. Given this, GD costs per eligible household are \$54, \$96, \$121, and \$517, and costs per village household overall are \$12, \$21, \$27, and \$100.

We use these estimates to regression-adjust outcomes for the differential cost from Gikuriro, providing a way of testing the differential impact of the programs at identical cost to the donor (approach described in detail in the next section). To compare cost-effectiveness across scale, we conduct F-tests for differences in the impact/cost ratio across all the arms of the study, each of which operated at a different cost. We present comparative impacts across a wide range of outcomes, providing nuance as to the specific outcomes chosen by beneficiaries in the cash arm relative to those that the in-kind intervention drives. We provide a number of ways to consider targeting heterogeneity, including standard interaction analysis, the use of machine learning matching algorithms to recover individual treatment effects, and analysis of the choice experiment conducted within the cash arm that speaks to the extent to which the beneficiaries themselves understood the heterogeneity in treatment effects.

III. COMPARATIVE IMPACT ANALYSIS

III.A. Attrition and Balance of the Experiment

Endline outcome measurement is subject to a number of distinct forms of attrition; we start our empirical analysis by considering each in turn. The most straightforward of these is standard household-level attrition, meaning that a household sampled into the baseline survey attrited from the endline survey. In Table A.1, we see that overall rates of attrition at the household level were low, around 3.3 percent in the control. We see the pattern typical in RCT studies where attrition is somewhat lower in the treatment groups (where both ongoing contact and a sense of reciprocity may keep individuals in the endline), but these differentials are small, from 0.89 percentage point in the GD ‘small’ arm to 1.7 percentage point in the GD ‘large’ arm; only the latter is significant, and only at the 10 percent level. Looking at the other covariates of attrition in column 2 we see that attriters and non-attriting households are similar. Hence we conclude that household-level attrition is unlikely to be a source of bias in the study.

When we turn to the analysis of individual-level outcomes in Columns (3)-(9) the picture is more complex because many of the primary and secondary outcomes are only measured for certain types of individuals (anthropometrics for children, birth outcomes only for those pregnant). We analyze four types of individual missingness that may occur. First, we compare the attrition of all household members from the roster in the household survey; both the rates and the differentials here are very similar to the household attrition problem suggesting that there has been little additional differential attrition of individuals. Next we examine the anthropometric panel, whereby all children under 6 at baseline who were given anthropometrics at the baseline should have been followed up with at endline. Here the absolute rates of attrition are a little more than double what they are for individuals overall, presumably because of the greater difficulty of finding and measuring children for this exercise. More concerningly, the decline in attrition in the treatment groups now becomes strongly significant, particularly for Gikuriro villages (which perhaps is evidence of the superior monitoring of malnourished children taking place in those villages). Given this significance, we follow our Pre-Analysis Plan in also presenting results for the anthropometric impacts that are corrected by inverse propensity weights to correct for the observable determinants of selection.

Third, we examine whether individuals who should have been anemia tested in the followup were; here we see no evidence of differential attrition across arms. Finally, we examine the likelihood that a new household member appears (typically due to births subsequent to baseline), and find no significant differences. Overall, then, differential selection across treatment arms is not a major problem for study outcomes other than anthropometrics. We return to the issue of unequal attrition in anthropometrics in the following section.

Next we present the baseline comparison of primary and secondary outcomes as well as household-level control variables, using the attrited panel sample that will be the basis of the endline evaluation. The regressions used here mimic as closely as possible the impact regressions, using fixed effects for randomization blocks, including a battery of baseline control variables, weighting to make the sample representative of all eligibles, and clustering standard errors at the village level to account for the design effect. All tables also present p-values adjusted for False Discovery Rates within outcome families (primary or secondary) using the technique of Anderson (2008); stars in the tables are based on clustered standard errors. Looking first at balance on study outcomes, Tables A.2, B.1 show balance for primary and secondary outcomes, respectively. Overall balance on secondary outcomes is excellent, but we do find evidence that we have superior outcomes on child anthropometrics in the GD Large arm. HAZ and WAZ are about .2 standard deviations higher than the control in this arm, and the latter difference is significant at the 10% level even after adjusting for multiple inference. Given this imbalance, we present all core results on anthropometrics using an ANCOVA structure (controlling for baseline values), focusing on impacts on children born prior to baseline. Table B.2 examines balance over baseline covariates and finds all treatment arms to be comparable to the control across all covariates.

Compliance with Gikuriro is complex because of the multi-dimensional nature of the program. Table B.3 examines the determinants of participation with specific sub-components of Gikuriro, based on self-reports within the eligible population. We have five forms of participation that can be ascribed directly to the program: participation in three types of training (nutritional, cooking/hygiene, and agricultural extension), whether households have themselves harvested the nutrient-rich household vegetable plots that they were trained by the Farmer Field Schools to grow, and whether they received livestock directly from Gikuriro. Slightly more than half of eligible households report receiving nutrition or agricultural training, half receive hygiene training and harvest

FFLS gardens, and a third receive livestock. The two driving determinants that emerge from this table are that Gikuriro was successfully targeted at the poorer households even within the (relatively poor) eligible group, and that they met more success with households headed by younger individuals. Conditional on this, however, they were not differentially successful at reaching households with children, or female-headed households.

III.B. Basic Experimental Results on Eligibles

Tables III and IV present the core intention-to-treat results of the study on primary and secondary outcomes for the eligible population. The analysis is conducted as an ANCOVA, including the lagged dependent variable where available, fixed effects for the randomization blocks, and for each outcome including the baseline controls selected for it by LASSO as described in the Online Appendix. For parsimony, our main tables include a dummy for the Gikuriro treatment, another for the Large GD treatment, and then a single dummy that indicates the three smaller GD transfers. Standard errors are clustered at the village level to reflect the design effect and observations are weighted to make the analysis representative of all eligible study households. Panel A of each table analyzes household-level outcomes, and Panel B individual-level outcomes.

Taking the Gikuriro treatment first, we see household savings increase by a 109 percent, consistent with the creation of SILCs. Dietary diversity, anthropometrics, and maternal anemia all move in the right direction but none of these changes is significant. No consistent impacts appear for consumption and wealth outcomes or for health knowledge and sanitation practices. Hence the program has been successful in moving an indicator closely related to one of its sub-components, but at least within the time-frame of the study these changes in savings have not translated into significant improvements in the anthropometric child outcomes.

We turn next to the impact of the three smaller GiveDirectly arms whose average cost per eligible is \$90.28, 72 percent of the cost of Gikuriro. Here, we see quite a different set of outcomes move. Instead of increasing savings, small GD transfers lead to a 77 percent paydown of debt, and an increase in the value of productive and consumption assets, by 26 percent and 35 percent respectively. Thus far, then, the comparison of Gikuriro to cash breaks down into two distinct dimensions of improvement, each of which has a different and entirely plausible pathway to long-term improvements: savings (Gikuriro), or debt reduction and asset investment (GiveDirectly).

When we examine the third column, however, a more transformative impact arising from the Large cash transfer is clearly apparent. Not only do omnibus measures of consumption and wealth go up across the board, but metrics of consumption closely linked to child health improve. The dietary diversity score increases by .55 food groups, or by 12 percent off a base of 4.77. Figure II displays the fraction of each arm consuming each food group, and shows the treatment effects of the large arm to be most pronounced in fish, fruits, oils and fats, spices and condiments, and cereals. Productive assets increase by 80 percent, consumer assets almost double in value, and home value increases substantially. In the individual outcomes the benefits of this surge in consumption are evident as well; within the course of one year we see a 0.09 SD improvement in HAZ, a 0.07 SD improvement in WAZ, and a 0.13 SD improvement in MUAC, all significant at least at the 10 percent level, prior to adjustment for multiple inference.²⁰ The ANCOVA specification should be particularly important in the analysis of the anthropometric indicators that showed signs of imbalance at baseline; indeed if we examine these outcomes in post-treatment levels we see substantially stronger apparent treatment effects. Anemia falls slightly (not significant) and there is a substantial decrease in child mortality of almost 1 percentage point (or 70 percent off of the baseline value). To contextualize these effects using unweighted numbers, the control group eligibles saw 13 cases of child mortality out of 2,596 children (0.5 percent) while the GD Large arm saw 2 cases out of 1,200 children (0.16 percent). Hence the GD Large arm saw meaningful improvements in consumption and child health.

In the final columns of each table we provide t-statistics on F-tests that the ratio of the benefits across arms differs from the ratio of their costs. These statistics ask whether we can reject that the impacts scale in a manner similar to the costs. For the comparison across the two GD arms we find two significant differences: only in the case of debt reduction (where small transfers have a big effect and big transfers do not) and house quality (where small transfers have a negative and large transfers a positive effect) can we reject cost-symmetric benefit scaling for cash transfers. In comparing Gikuriro to the GD large arm we see more differences, with Gikuriro being more cost-

¹⁹Table A.3 shows the experimental results with each of the three treatment cells within the ‘main’ treatment broken out individually. Because these cells are both small and feature treatment amounts that are similar, this more granular analysis does not turn up evidence of meaningful heterogeneity across the three smaller GD transfer amounts.

²⁰These improvements should be viewed against the backdrop of a sharp deterioration in anthropometrics subsequent to birth that typically occurs in LDCs, leaving rural African children often two full SDs below the international norm by age 3 (Shrimpton et al., 2001; Victora et al., 2010).

effective in generating savings and GD Large superior in consumption, as well as productive and consumption assets.

One of the most fundamental ideas in development economics is that poor households should have a single ‘shadow value’ of cash which pulls down investment in all capital-hungry endeavors in a symmetric way. The above findings are generally consistent with this view of the world, as an intervention that relaxes credit constraints leads to shifts in consumption patterns that are very broadly spread across domains. This property means that small cash transfers are hard to detect because they move too many outcomes by too small an amount to be significant, while large cash transfers result in a broad-based increase in consumption in many dimensions.

Our pre-analysis plan states that for any outcomes where we find differential attrition, we would estimate a propensity to remain in the sample incorporating covariates, dummies for treatments, and their interactions on the right-hand side, and then re-weight the analysis by the product of the inverse of this probability and the standard sampling weight. This procedure corrects the impacts for the observable determinants of attrition, and uses regression weighting to attempt to make the treatment and control samples comparable on important covariates even after attrition. Because we primarily found significantly differential attrition for the anthropometric outcomes, in Table B.4 we present the results of this correction. We interact with each treatment dummy the same right-hand side covariates we use the same controls in the anthropometric regressions: gender, a linear, quadratic, and cubic for age in months, baseline household wealth, and a dummy for membership in Ubudehe poverty category 1. The results in this table are virtually identical to Table III, indicating that the types of children who attrited from the study are similar across arms and hence differential attrition is unlikely to be driving our impacts.

III.C. Cost-Equivalent Benchmarking

The core purpose of this experiment is comparative; namely to exploit the randomized variation in transfer amounts to conduct an exact cost-equivalent benchmarking exercise. Using the costing numbers emerging from the ex-post exercise, we use the observed costs, overhead rates, and compliance rates to calculate the donor cost per eligible household in each arm of the study.

To present our cost adjustment strategy more formally, let the subscript i indicate the individual (or household depending on unit of analysis), c the cluster (village), and b the randomization block.

E_{icb} is an indicator for eligibility status, defined at the household level. Outcomes are observed both at baseline (Y_{icb1}) and at endline (Y_{icb2}). First, begin with the total GD donor cost per eligible within each transfer amount arm, denoted by t_c . Subtract from this number the benchmarked Gikuriro cost per eligible household C described above, and denote the difference $t_c - C = \tau_c$; this is the deviation (positive or negative) of each GD arm from the benchmarked Gikuriro cost. Set τ_c to zero in the control and Gikuriro arms. We can then run an ANCOVA impact regression as above, but controlling for a linear term in τ_c , a dummy for either treatment, and a dummy for receiving Gikuriro. Because τ absorbs the deviation of the GD arm from the benchmarked Gikuriro cost, the dummy coefficient on Gikuriro treatment will serve as an intercept measuring the impact of Gikuriro benchmarked an exactly donor cost-equivalent cash transfer. So, we have:

$$Y_{icb2} = \alpha_b + \delta^T T_c + \delta^{GK} T_c^{GK} + \beta X_{icb1} + \rho Y_{icb1} + \gamma_1 \tau_c + \epsilon_{icb2} \quad \forall \quad I : E_{icb} = 1 \quad (1)$$

In this specification T_c is a dummy variable indicating any treatment (Gikuriro or GD). Subject to the assumption of linear transfer amount effects, the slope coefficient τ_c captures impacts arising from deviations in GD cost from Gikuriro cost, the coefficient δ^T effectively gives the impact of GD at the cost of Gikuriro), and the dummy variable δ^{GK} provides a direct benchmarking test: the differential impact of Gikuriro over GD at the same cost per eligible. Per the pre-analysis plan, we impose the simple linear functional form to preserve as much statistical power as possible for the core cost-equivalent benchmarking comparison, although it is straightforward to make this more flexible.²¹

The Total Causal Effect can also be benchmarked at a cost-equivalent level. The methodology is very similar; we redefine τ_c as the deviation of GD spending per household in the overall village, relative to the Gikuriro cost defined in the same way. By then cost-adjusting an estimate of the TCE weighted to be representative of all households in the village, we can measure how the overall village-level average impact of each program differs when spending per household is the same.

A graphical representation of our strategy is provided in the left-hand panel of Figure I, which plots the average outcome on the y-axis for all four GD treatment amounts (colored circles), for GK (gray diamond), and the control (white circle). The line represents the fitted average savings

²¹Because in this study we have three very similar small transfer amounts and one much larger amount we have limited ability to measure non-linear impacts of the transfer amount.

by GD transfer amount. By predicting the outcome on this line at the exact cost of Gikuriro (gray triangle), the benchmarked differential is then the vertical difference between the Gikuriro impact and the projected cost-equivalent GD impact.

The results of the cost equivalent analysis for primary and secondary outcomes are presented in Tables V and VI, respectively. Starting with the third column, we have a direct estimate of the marginal effect of an additional 100 dollars in donor cost on primary and secondary outcomes. As could be inferred from previous tables, this coefficient is strongly significant across a wide range of outcomes, particularly those most related to household consumption. An extra 100 dollars leads to a 5 percent increase in consumption, a 9 percent increase in dietary diversity, a 19 percent increase in savings, an 13 percent and 14 percent increase in productive and consumption assets, respectively, and leads housing value to improve by 5 percent and the index of housing quality to increase by 0.1 SD. In terms of anthropometrics, the change in value of transfer is positive but small in absolute magnitude for HAZ, WAZ, and MUAC, and does not survive correction for multiple inference. An extra \$100 per beneficiary household—with eligible households containing an average of 2.7 children under the age of six—increases HAZ by 0.022 standard deviations. Beyond this, none of the other individual outcomes respond to transfer amount in a manner that we can reject at 95 percent significance.

With the third column estimating the linear heterogeneity in impacts by transfer size around the cost-equivalent transfer, the second column (dummy for ‘any treatment’) becomes an intercept term that estimates the impact of cash transfers at a cost equivalent to Gikuriro, although this exact amount was not included in the experiment. Given that the mean transfer amount in the ‘small’ arm is only slightly lower than the GK cost, this estimate looks generally similar to the second column of the ITT tables (the simple average experimental effect across the ‘small’ transfer amounts). At the exact cost of Gikuriro, we estimate that cash transfers would have led to a significant 73 percent decrease in the stock of debt, and a 30 percent and 40 percent increase in productive and consumption assets, respectively.

The first column of these tables contains the heart of the comparative benchmarking exercise. Compared to an exactly cost-equivalent cash transfer, Gikuriro is less effective at driving consumption, the paydown of debt, and the accumulation of assets. On the other hand, it is significantly more effective than cash at creating savings. The differential effect of the programs on savings and

borrowing is interesting, and suggests that while both interventions serve to improve the net stock of liquid wealth (savings net of borrowing), the focused push on SILC groups in Gikuriro drives this more strongly through the vehicle of new savings while households making their own choices are more strongly disposed to reduce debt instead. Which of these strategies makes more sense? A simple comparison of interest rates is revealing. Gikuriro SILCs were free to set their own interest rates, but typically paid about 5 percent per annum nominal. Credit interest rates, by comparison, vary from an average of 22 percent in the MFI sector to upwards of 60 percent in informal credit markets. Given that 32 percent of eligible households reported having both borrowing and savings at baseline (and 79 percent had either borrowing or saving) it seems that the desire to pay down debt might be warranted.

III.D. Cost Equivalence versus Cost Effectiveness

This study is designed to make a specific form of cost-equivalent comparison, namely the impact of a cash transfer intervention assessed at the exact cost of an in-kind intervention. This comparison fixes the amount to be spent per beneficiary and asks which intervention is more effective. A different but related question is that of cost effectiveness, where we compare programs that operate at different costs and ask which generates the greatest benefit per dollar spent.

The comparison between these two approaches is represented graphically in Figure I for the primary study outcomes. In the left-hand panel we plot comparative cost-equivalence; here the focus is on the value on the X-axis that is the ex-post cost of Gikuriro, and we are interested in seeing which is greater, the observed benefit of Gikuriro (represented by a black diamond) or the predicted benefit of cash at this cost (the hollow circle). In the right-hand panel we connect the shaded circle that represents the outcome in the control group (at zero cost) with the outcome in each arm; because the plot represents outcomes in benefit/cost space, the intervention that features the steepest slope in this graph has the highest cost effectiveness. Interestingly, this graphic illustrates that while in general there are not substantial differences between Gikuriro and cash at benchmarked cost, because the smallest cash arm is so inexpensive while producing outcomes that are generally as good as (or better than) more expensive treatments, for four out of the five outcomes represented the smallest cash transfer has the highest cost effectiveness. The difference in cost effectiveness across arms is tested statistically for all primary and secondary outcomes in the final columns of

Tables III and IV.

Table VII provides a statistical analysis of the benefit-cost slope terms represented in the right-hand panel of Figure I, pooling the three smaller GD arms together as done in the rest of the paper. It presents the ITT coefficient from the prior tables divided by the cost of the arm, and so gives the household improvement generated per \$100 spent through that modality, as well as the corresponding standard error on this BCR.²² The final three columns of this Table VII provide statistical tests of the difference in cost-effectiveness slopes and show how difficult it is to power a study to reject these; despite the relatively large sample size of this study, it is only for the consumption outcome that we are able to reject equality of CBRs across interventions, with the cash interventions outperforming Gikuriro. Between the two cash arms we are unable to reject equal BCRs for any outcome, indicating benefits that scale linearly with cash transfer amounts.

III.E. Total Causal Effects

Because we randomly sampled from among ineligible as well as eligible households, we can also conduct an analysis representative of the population of study villages by pooling the strata together and using population sampling weights. The average weight in the ineligible sample is 24.4 and in the eligible sample it is 2, meaning that while the unweighted eligible sample is larger, it is the ineligibles who will dominate the weighted sample. Recall that the treatment effects on ineligible households may arise either from the treatment and targeting of the two interventions among ineligible households (with saturations set to be the same at just over 18 percent on average in both arms, but with targeting differing), or from spillovers between eligible and ineligible households. With 11.4 percent of all households being defined as eligible, the treatment rate in the ineligible sample is 8.4 percent. This means that the large majority of the additional sample included in the TCE analysis only receive impacts through spillover effects to untreated households.

These impacts are presented in Table VIII and A.4 for primary and secondary outcomes, respectively. Here, the overall picture is very different from the impact among eligibles. For Gikuriro, we see improvements at the 99 percent significance level in the index of health knowledge, a core component of the program and one which was broadly targeted at the village population by the

²²Given that individual-level outcomes in this table represent average benefits for amount spent on the household, if we want the BCRs at the individual level we need to scale up the coefficients in this table by the average number of individuals per household for each outcome given the notes to the table.

program (as reflected by our accounting of these costs as ‘non-averted’). Vaccinations, presumably provided by government health facilities but not a central focus of Gikuriro, also improve significantly. So there is some real evidence of holistic benefits in health-related domains for the population of Gikuriro villages. These changes, it is true, do not translate into observable improvements in health outcomes for children or adults within the timeframe of the study, but still suggest that Gikuriro implementers have been successful in driving community-level health knowledge.

With cash transfers, on the other hand, improvements appear to be more narrowly limited to the beneficiaries of the transfers. The ‘small’ transfers do not move any village-level outcome at the 5 percent significance level, even before adjustment for multiple inference. The ‘large’ transfers, so positive among beneficiaries, in general see negative signs across the consumption indicators, lead to a significant drop in savings at the village level, and are only positively associated with vaccination rates. On net there is little evidence that the widespread benefits observed in the eligibles carry over to the broader population of the village when the transfers are targeted at a relatively small fraction of the households.

A comparison of the ITT and TCE, particularly for the GD Large arm, suggests that for a number of outcomes (such as dietary diversity and savings) we see a swing from positive effects on the eligibles to negative point estimates for the village as a whole. This raises the question as to whether the GD treatment may be having negative spillover effects on ineligibles. We are able to analyze spillovers experimentally within a subset of ineligible households. Because the cash transfers were provided only to ineligible households who were in Ubudehe 1 and 2 and had 3 or more members (and thus might contain children), wealthier or childless households were never treated by GD. Therefore, by comparing ineligibles either in Ubudehe 3 and 4 or with fewer than three members between the GD and control villages, we have a clean window into the spillovers effects of cash.²³ The results of this analysis are presented in Tables A.5 and A.6. The GD Small transfer has no detectable spillover effects on any outcome, and while the GD large does show negative spillovers on savings and positive spillovers on vaccination rates, no systematic picture emerges. So, recognizing that this test is relatively low-powered, we conclude that in this case cash transfers do not appear to have resulted in meaningful spillovers on smaller, wealthier households in treatment

²³Individuals in all Ubudehe categories report receiving some benefits from Gikuriro, so we have no comparable ability to separate spillovers from treatment of ineligibles in that arm.

villages.

III.F. Benchmarked Total Causal Effects

We can perform a similar cost-equivalent benchmarking exercise for the village-level TCE, adjusting now by cost per household in the overall village from Table II rather than the cost per eligible household. This allows us to ask how the two programs differ in their impact on the village population when the same amount is spent by each program per household in the village. This analysis is presented in tables and B.5 and B.6. In general we lack the power to make many clear comparative statements across interventions in terms of TCEs. Before multiple inference correction Gikuriro is significantly better at generating health knowledge in the population, but neither this result nor even any of the expenditure slope terms retain their significance once we look at the sharpened Q-values. Clearly, the implication from this result is not that these outcomes would be unmoved by large cash transfers to the whole population (as shown in (Egger et al., 2019)), but rather that no arm in this study actually moved average village-level costs by a large enough sum to drive detectable population-level outcomes.

IV. HETEROGENEOUS COMPARATIVE EFFECTS AND TRADEOFFS ACROSS SUB-POPULATIONS

The average comparative impacts of cash transfers relative to Gikuriro’s in-kind programming may mask substantial heterogeneity in relative impacts across groups. This has (at least) two consequences. First, it implies that decisions over intervention modality entail tradeoffs between population subgroups: there can be both winners and losers from the switch from kind to cash. If policymakers place particular weight on certain subgroups, then these tradeoffs may change their preferences over modalities. Second, to the extent that this heterogeneity is predictable based on observed characteristics, policymakers may be able to improve on the average effects of both programs by finely targeting benefit types to individuals or groups on the basis of predicted relative impacts. Here, we test for heterogeneity on pre-specified dimensions of heterogeneity anticipated to matter for relative efficacy, as well as by using machine learning methods to identify subgroups for whom cash is advantageous over in-kind programming, and vice-versa. We find that there is little

heterogeneity on attributes selected ex ante, and that while machine-learning methods do deliver heterogeneity in impacts for each group, these bear little correlation across outcome dimensions, suggesting that targeting modalities on the basis of expected benefits is a promising approach for policymakers with narrow objective functions, but—at least in this instance—that even fine targeting offers only modest gains for broad welfare metrics.

Our pre-analysis plan highlighted two forms of heterogeneity that we anticipated would be important at the design phase; namely how baseline malnutrition and child age may moderate the impact of nutritional interventions. Given that we have children who start the study outside of the first 1,000 days (those 2–5 years old at baseline), we might expect that the impact of the program on these more fully developed children would be smaller. Similarly, we might expect that both of these interventions would be most effective for children who began the intervention most malnourished.

Our analysis generally reveals a lack of heterogeneity, in that impacts are not larger for children most malnourished at baseline (Table B.7), or for children exposed to the treatments at younger ages (Table B.8). Figures B.2 and B.3 provide fan regressions of impacts by child age and there is a suggestion that children exposed to large cash transfers in utero realize the largest benefits.²⁴ In general, however, our results are not suggestive of strong age- or malnutrition-driven heterogeneity of these interventions.²⁵

While these pre-specified dimensions are not drivers of differential outcomes, it remains possible that there be substantial heterogeneity in the impacts of cash transfers relative to in-kind transfers. This heterogeneity may simply be poorly predicted by these pre-specified covariates. To the extent that it *is* predictable on the basis of observable attributes, however, then predictions of relative impacts that differ by subgroup might be used to target aid modalities to improve their impacts.

Modern methods for using machine learning techniques to develop predictive models of heterogeneous causal effects are well suited to this problem. In this section, we deploy the *causal forest* approach of Wager and Athey (2018) to estimate the relative effects of cash versus in-kind programming. We first characterize the extent of heterogeneity across observable subgroups in each of the core outcome dimensions. Though there is substantial heterogeneity, we then demonstrate

²⁴This pattern is similar to the medium-term results in Baird et al. (2016), who find unconditional transfers in Malawi to have the largest effect on children exposed in utero.

²⁵The longer-term literature has typically found impacts of large cash transfer programs on HAZ in the range of 0.2–0.45 standard deviations (Aguero et al., 2006; Barham et al., 2014); comparison to the broader literature suggests that these impacts may grow over time.

that it is limited in its correlation across outcome dimensions. Turning to welfare, we show that for a simple metric of welfare gains, tradeoffs in impacts across outcome dimensions mean that even policymakers who are able to individually assign cash versus kind are limited in their ability to achieve broadly defined welfare gains by ‘personalized’ choice of benefit type in this setting.

We estimate the extent of heterogeneity in treatment response in the sample by generalized random forest (Athey and Wager, 2019; Athey et al., 2019). To ensure that treatment effect estimates are centered on point estimates for the full sample as were reported in Section III.B, we begin by residualizing both the outcomes studied and an indicator for the GiveDirectly ‘Main’ treatment on the basis of all covariates used in the ITT analysis of that section. We then restrict attention only to the Gikuriro and GiveDirectly Main treatment arm; a regression of primary outcomes on the residualized GD-Main indicator alone in this sample replicates the point estimates of Table III. In this estimating sample, we allow the same candidate covariates used in the Post-Double Lasso to act as candidate effect moderators.

Appendix Figure A.1 displays estimates of treatment effect heterogeneity in groups defined by this approach. Moderate point estimates often mask substantial heterogeneity across observed characteristics. These estimates suggest that even where average effects are small, a policymaker who values a single dimension of these outcomes in particular can achieve substantial welfare gains by optimizing choice of modality.

However, welfare is surely not one-dimensional, and typically policymakers will value outcomes across multiple dimensions. If this is the case, then the ability to optimize effectively will depend on the extent to which heterogeneity in the relative benefits of cash versus kind is correlated across dimensions. If individuals who benefit from cash in one outcome dimension also tend to benefit from cash in other dimensions, then it will be possible to assign treatments in ways that benefit them across the board. On the other hand, there may be no such free lunch: if individuals’ treatment responses differ across outcome dimensions, then there may be no assignment choice that uniformly raises these outcomes.

In Table IX we illustrate that the extent of these cross-outcome correlations in the sign of treatment responses are low. None of the correlations have an absolute value greater than .3, indicating that no single set of determinants predict benefits across a wide range of outcomes. Perhaps unsurprisingly households that see larger wealth benefits see smaller consumption benefits (since transfers

can either be saved or spent). Interestingly there does seem to be a positive correspondence between the types of households effective in generating wealth and those that see child growth improve.²⁶ This is suggestive of a shared investment in long-term outcomes in response to cash transfers occurring through both the accumulation of physical and human capital in a subset of households. The limited concordance between treatment effect heterogeneity across outcome dimensions suggests that policymakers who value a multi-dimensional welfare metric will be constrained in their ability to target cash relative to in-kind transfers. To illustrate this, consider a policymaker who values an additive social welfare metric combining consumption, wealth, dietary diversity, and child growth, with each measure standardized to have a variance of one.

For policymakers who value only one outcome, optimizing targeting based on estimates like the ones above does deliver gains. Table X reports the impacts in standard deviation terms from optimal targeting. These are modest in the case of targeting for consumption, but rise to 5.7, 8.6, and 10.3 percent of a standard deviations of the outcomes wealth, dietary diversity, and child growth, when targeting is based on these outcomes and they are the sole metric of benefit. If optimal policies were perfectly aligned across these outcomes, the gains to personalized treatment assignment for a welfare metric that is the average of these outcomes would be equal to the average of the outcome-specific benefits. But as Table X shows, it is instead substantially smaller: at 2.9 percent of a standard deviation of the composite welfare metric, this is approximately 44 percent of that theoretical maximum. Tradeoffs across outcomes limit the extent to which policymakers can optimize cash versus kind regimes based on heterogeneous responses by households.

V. PROGRAM RIGIDITY, CHOICE, AND BENEFICIARY WELFARE

A central argument against cash transfers, and particularly lump-sum unconditional transfers, is that they provide too large a temptation for misuse of resources. From this behavioral perspective, the costs required to build additional rigidity into programs (by giving cash in smaller increments over time, or by giving benefits in-kind) can be justified if they deliver superior outcomes. Even better, if those with behavioral problems are sophisticated about their own actions, then offering

²⁶The determinants of impact heterogeneity across the three biometric measures of child growth concur very closely as shown in Online Appendix Figure B.4; consequently we have collapsed the three child growth metrics down into a single variable for simplicity of presentation.

choice over program rigidity can permit the right individuals to ‘bind themselves to the mast’ while allowing those without such problems to choose cheaper, lump-sum transfers (Ashraf et al., 2006). We shed light on this behavioral angle by combining incentivized measures of present bias with an experiment on beneficiary choices over transfer timing embedded in the cash-transfer arm of the study.

At baseline, we implemented an incentivized choice experiment over transfer modality within the cash arm. All cash transfer subjects were shown a graphic representing the upcoming twelve months of the study, and for each month were asked to choose whether they would prefer to receive a lump sum cash transfer in that month, or flow transfers across all 12 months, with the payout amounts being equal. Subjects were told that a small fraction of the sample (1 out of 12) would be randomly selected, a month would be randomly selected, and they would receive the actual choice that they had made in that month. The remainder of the sample would then be randomly assigned to receive lump sum or flow transfers. The purpose of this exercise was not to be able to evaluate the choice arm per se (it is not adequately powered to be studied as an independent treatment), but rather to have a highly incentivized way of understanding beneficiary preferences across the two modalities. This exercise can then be used either to test the impact of being assigned to lump sum versus flow transfers or as a tool for revealing preferences to understand whether allowing choice over transfer modality would improve outcomes.

We combine the choice experiment, the direct randomization of transfer modality, and measures of behavioral heterogeneity to examine three core questions. First, despite greater administrative complexity, are flow transfers more effective than lump-sum transfers at delivering child health? This can be answered using the simple experimental comparison between those directly assigned to the flow and lump-sum arms. Second, are the beneficiaries sufficiently sophisticated as to their own relative benefits to be able to choose the modality from which they would derive the most benefit? This can be examined by seeing whether benefits are larger when beneficiaries are assigned to the arm that they selected in the choice exercise. And third, are the comparative impacts of alternative modalities—cash versus kind, or lump-sum versus flow—heterogeneous according to behavioral measures? Specifically, we examine whether more rigid interventions (such as in-kind over cash or flow transfers over lump sum) are disproportionately beneficial for individuals with self- or other-control problems, and whether such individuals who chose the flow arm are hurt when

they are instead assigned to receive a lump sum transfer.

V.A. Lump-sum versus flow transfers

Beginning with the question of the impact of lump sum versus flow transfers, we can form a simple randomized comparison by dropping the households assigned to the choice arm and comparing those experimentally assigned to lump sum or flow to the control group. Because of the large difference between the three smaller transfers and the large one, we analyze this experiment with four sets of dummies: one for the main transfer cells, a second that measures the additional impact of lump-sum transfers in these cells, then a dummy for large transfers, and a dummy that measure the additional impact of large lump sum in the large arm. We provide F-tests that give the p-values on the total effect of small lump sum and large lump sum.

This analysis, presented in Tables A.7 and A.8, provides little support for the idea that flow transfers generate superior child health outcomes. Looking first at the individual outcomes we see if anything lump-sum transfers generate slightly better child health outcomes (this is true for both WAZ and MUAC, while small flow transfers are modestly more effective at delivering HAZ). Turning to the household analysis, we find evidence of interesting differences in the use of intertemporal assets; only when the smaller transfers are given as a flow do they lead to a reduction in debt loads. Conversely, lump sum transfers are significantly better for generating savings. Put together these suggest that households under the small transfer only avoid building debt when transfers come as a flow, but only save when transfers come as a lump sum. For both small and large transfers there is a moderately greater increase in consumption assets when the transfers are lump-sum. Taken as a whole, this analysis is indicative of muted overall differences between cash transfer modalities, with a slight edge for lump-sum transfers even on child health outcomes. Given the greater ease and lower costs of making one-time transfers, the takeaway should be that we find no clear reason to incur the costs of flow transfers.

V.B. Transfer modality choice

Next we move to analysis of the choice experiment. The choices made by households reveal an overall preference for lump-sum transfers. In every one of the 12 months over which individuals made choices, roughly 65% of the subjects chose lump sum. Contrary to what we might expect if a

large share of the sample were highly impatient or time inconsistent, we do not see any systematic decrease in the fraction choosing lump sum as we move from the immediate choice (month 1) to more distant months. 56% of the subjects chose lump sum in every month and 74% ever chose lump sum in any month.

The differential impacts of ‘getting what you wanted’ can be thought of as a test of *essential heterogeneity* (Heckman et al., 2006), in that if impacts were heterogeneous across modalities and this was well understood by beneficiaries, then we would expect to see systematically superior outcomes when the assignment was consonant with the choice. Improved outcomes when subjects get what they choose would provide a powerful argument in favor of beneficiary-driven flexibility in modalities.

This analysis must account both for the (endogenous) choice and the (randomized) treatment status. The regression includes a dummy for choosing lump sum, a dummy for being assigned to the lump sum treatment, and then a dummy for ‘got what I wanted’.²⁷ The outcome of this analysis, presented in tables A.9 and A.10, is very clear (and is the same as seen if we instead partition the analysis to look at those who originally chose LS and flow separately). The results indicate that the receipt of the desired treatment modality has no detectable benefit in terms of the main study outcomes, and indeed appears to result in significantly worse HAZ. This result is potentially quite important, suggesting that no meaningful essential heterogeneity exists across transfer modalities in this sample. The two potential candidate explanations for this result would be either that no heterogeneity exists, or that it exists but is not understood by the participants. Given the lack of strong differences between lump sum and flow transfers when randomly assigned from the prior section, the former explanation may be the most reasonable.

V.C. Does rigidity help those with self- or other-control problems?

We added features to the design of the experiment that let us speak to behavioral dimensions of the cash benchmarking question. One of the most powerful motivations for in-kind programs is the idea that beneficiaries lack the control to manage large sums of cash effectively. Whether because of behavioral self-control problems (Ashraf et al., 2006; Thaler and Benartzi, 2004) or intra-household

²⁷Given this structure, the lump sum dummy gives the impact of being assigned to LS and not wanting it, the dummy for LS choice is the impact of choosing LS and not getting it, and the last dummy gives the average difference between getting and not getting one’s choice.

other-control problems (Anderson and Baland, 2002; Ashraf, 2009), any such failure will lead to the inability to translate one-shot lump sum transfers into the type of steady, sustained increases in consumption that are likely to deliver improvements in child nutrition and health. In this sense, changing the cash transfer modality away from lump sum transfers and towards a steady flow of smaller monthly payments can be seen as analogous to the types of intentional rigidity that underlie in-kind programs. This logic would suggest that program rigidity would prove helpful those with control issues. If this use of flow transfers as a way of ‘binding oneself to the mast’ is understood by sophisticated beneficiaries, then we should see the choice to receive flow transfers, and the benefit of receiving flow over lump sum, to be particularly pronounced among individuals who are both time inconsistent and sophisticated about this fact.

To measure the relevant behavioral heterogeneity, we conducted an incentivized Convex Time Budget (CTB) exercise (Andreoni and Sprenger, 2012; Giné et al., 2018) that provides an individual-level measure of discounting (δ) and time inconsistency (β) for each individual. This exercise asks participants to allocate a sum of money between being received immediately and received 30 days in the future, varying both the interest rate and the starting point for the 30 days (today or 90 days from now). To measure self-control, we define a dummy variable for ‘impatient’ and a dummy variable for ‘time inconsistent’ based on a straightforward use of the data from the Convex Time Budget exercise.²⁸ To measure other-control, we use the responses to a set of baseline questions that measure the extent to which an individual reports financial threats from others around them.²⁹

We have three ways to address this question. First, do individuals with self- or other-control problems benefit from the rigidity of receiving a program in kind rather than in cash? This can be asked by comparing the interaction effects of these attributes with indicators for the Gikuriro and the smaller GD arms (omitting the large GD arm for comparability). Second, within the cash transfer arm do they benefit from the comparative rigidity of receiving transfers as flows rather

²⁸Specifically, we count an individual as impatient if they elect to receive money today rather than in 30 days despite being paid double to wait, and as inconsistent if they allocate more money to ‘soon’ when they are faced with the same choices in 90 versus 120 days. We estimated structural utility models for the CTB using three different methods: Logit, Tobit, and Non-Linear Least Squares. The estimates from these different models correlate poorly with each other and with survey-based measures of impatience and time inconsistency, so we have elected to use the CTB questions in the simplest way possible. Using these measures 25% of the sample is impatient, and 67% is inconsistent. Results are similar if we instead use a survey-based measure of self-control problems.

²⁹An individual was counted as lacking other control if they: (a) did not want to keep money around the house because of a lack of trust of others, because of thieves, or relatives who might steal from them, or (b) reported that they had a spouse with whom they had conflict about money, or who was ‘wasteful’ or ‘irresponsible’. By this metric 27% of our sample lacks other control.

than as a single lump sum (similar analysis, now comparing lump sum to flow within the GD arm). Finally, we can focus on the most concentrated case for rigidity as a potential commitment device: are individuals with control problems most damaged by being having rigidity removed when they had indicated they need it? We can pose this question by looking among those who chose flow transfers for heterogeneity in whether they were then assigned to lump sum or flow.

In the end, despite these numerous ways of posing questions about the benefits of rigidity, we find little evidence of behavioral heterogeneity. In Appendix Tables B.9 and B.10 we present the result of first two analyses described above. Focusing on the F-statistics at the bottom that test for differential impacts across the two interventions on those with self- or other-control problems, in both the comparison of Gikuriro to cash as well as the comparison of Lump Sum versus Flow we never find the behavioral attributes to have significantly differential effects for the more rigid program relative to the more flexible one. In B.11 we conduct the third analysis described above; even in the group we might most expect to be signalling a desire for rigidity (chose flow) we find no harm from being moved to a less rigid modality (lump sum). Further, the simple determinants of choice do not support the idea that this decision is driven by behavioral parameters; if we construct a measure of ‘sophistication’ as being those individuals who are time inconsistent per the CTB exercise and also report in the survey having concerns about their own control over money, these individuals are no more likely to choose flow transfers than individuals who are time inconsistent and naive.³⁰ Those who report other control problems also choose lump sum at virtually identical rates to those who do not. The takeaway is that at least by these metrics of individual heterogeneity, we see no evidence to support the ‘rigidity as commitment’ idea.

VI. CONCLUSION

This study uses a large-scale randomized experiment to pose a number of questions in comparative cost-effectiveness. Most centrally, we establish an approach to ask whether it is better to run complex multi-dimensional programs or simply to provide cash grants of equal ex-post cost. Our

³⁰We consider an individual to be sophisticated if they are time inconsistent by the above measure and respond yes to any of the three questions at baseline: doesn’t keep cash on hand because doesn’t trust self with money, regrets a purchase made in the past month, or reports has problems with self wasting money on things the house doesn’t need. Just over 10% of time inconsistent individuals are sophisticated by this measure; about 66% of both groups choose Flow.

application of this approach to the critical battle against child malnutrition highlights an even more basic lesson: evidently, regardless of modality, it simply costs more than \$140 per household to deliver clinically relevant impacts on child malnutrition outcomes within a year.³¹ Even at this relatively low cost, however, the programs do trigger meaningfully different responses in the household use of intertemporal assets, with Gikuriro households moving to save more, while cost equivalent cash instead leads to the pay-down of debt and growth in asset investment. Our cost-equivalent estimates allow us to make the trade-off across outcomes very explicit. For example, a policy preference for Gikuriro is equivalent to asserting that it is worth foregoing the increase in productive assets in order to obtain the improvement in savings, or that the relative benefit from percent increases in savings versus productive assets exceeds 4.06, the ratio of the differential treatment effects.³² This kind of precision can help policymakers be much more exact as to the types of trade-offs required to justify one type of intervention over another.

Given the modest impacts of both programs at costs equivalent to Gikuriro, comparisons across program scale highlight a key tradeoff. When we compare the cost-equivalent programs to a cash transfer of almost five times the amount, we see that larger sums of money can powerfully improve not only overall consumption and dietary diversity, but lead to modest improvements in child growth. The transfer amounts used in this study are not well-powered to test the linearity assumption (three small transfers of similar size, one much larger), however we generally find outcomes scaling with transfer amount in a simple way. Randomized variation in transfer values lets us form a number of interesting counterfactuals; for example the smallest transfer at which the benefit of cash would exceed Gikuriro; for savings this number is \$694, and for HAZ it is \$277.³³ Policymakers seeking to move child growth outcomes face difficult tradeoffs between the depth and breadth of their interventions; the comparative ease with which the resource intensity of cash transfers can be adapted makes this counterfactual modality particularly capable of revealing these tradeoffs.

³¹Gikuriro's costs were substantially lower than those of a related program in neighboring Burundi.

³²When compared with a cost-equivalent cash transfer, Gikuriro has relatively larger effects on savings: a differential impact of 118 percent, as shown in Table VI. On the other hand, a cost-equivalent cash transfer delivers a 79% larger decrease in debt loads. Expressed as dollar changes over the control group mean, this means that a policy preference for Gikuriro in these two variables implies willingness to gain \$18.26 in new savings at the cost of \$27.68 in foregone debt reduction.

³³Evidence from the sister study in Rwanda (McIntosh and Zeitlin, 2020), which featured more intermediate treatment amounts, did in fact find evidence of a non-linear effect, with outcomes involving investment and time use peaking at transfers a little over \$400. Nonetheless, in this study for no primary outcome are we able to reject a constant benefit/cost ratio across different cash amounts.

Finally, we are able to test how the program impacts differ according to the sub-populations targeted. Overall this analysis provides evidence of surprisingly homogeneous returns from these interventions, whether looking across pre-identified study subgroups, using machine learning to identify optimal targeting rules that span multiple outcomes, or by introducing self-targeting through beneficiary choice. The general lack of heterogeneity or scale effects, the lack of evidence that flow transfers are superior, absence of essential heterogeneity, and the uncorrelated individual-level benefits across different dimensions of impact all suggest that it is reasonable for policymakers to use simple lump-sum cash transfers and to target them using preference weights based on individual attributes or inequality.

There are some important limitations to what can be learned from the study. First, the cross-village experiment lets us measure the impact of those parts of the Gikuriro intervention that are implemented at the village or household level only. Village-level assignment means that some of our tests involving transfer amounts are not highly powered statistically. Most importantly, the time frame of the study (13 months, but only 8–9 months after treatment with Gikuriro) means that we are measuring endline outcomes more quickly than would be ideal, particularly for anthropometrics that may respond slowly to improvements in nutrition. Gikuriro made substantial investments in local capacity around health and sanitation and this infrastructure may drive future benefits in a way not captured in our study. However, given that the targeting criteria for these programs will cause children to age out of eligibility within a few years, it remains important to show that such programs can generate benefits quickly for a given set of vulnerable children, as do the larger cash transfers we study.

Given the nuance of our findings, it is hard to square them with any simple idea of cash transfers as a kind of uni-dimensional ‘index fund’. While business investment may have a single, cardinal objective—financial profit—development policy is undertaken with many goals in mind, and a perfect reconciling of these competing benefits would require a clear statement of trade-offs in this multi-dimensional space. Perhaps a clearer way of expressing the counterfactual provided by unconditional cash is that it gives us a statement of the priorities that the beneficiaries themselves hold when credit constraints are relaxed, and thereby motivates us to be clearer about the logic underlying paternalistic development programs. While beneficiary decisions may not be ‘optimal’ in terms of long-term social welfare (for example, due to high discount rates, to self- or other-

control problem, or to resource and information constraints), the impact of unconditional cash is nonetheless a powerful statement of the outcomes that the beneficiaries themselves want changed. For us to argue that a program is justified in using resources to drive outcomes different from the ones the beneficiaries would choose, we should have a clear reason why they fail to arrive at the welfare-maximizing outcome themselves. This is a view of benchmarking that is not about picking a winner but rather about quantifying trade-offs.

Author Affiliations:

Craig McIntosh, School of Global Policy and Strategy, University of California San Diego. 9500 Gilman Drive, La Jolla CA, 92093-0519. ctmcintosh@ucsd.edu.

Andrew Zeitlin, McCourt School of Public Policy, Georgetown University. Old North, Suite 100, 37th and O Streets, N.W., Washington, D.C., 20057. andrew.zeitlin@georgetown.edu

REFERENCES

- Aguero, Jorge, Michael Carter, and Ingrid Woolard**, “The impact of unconditional cash transfers on nutrition: The South African Child Support Grant,” 2006.
- Ahmed, AU, JF Hoddinott, S Roy, E Sraboni, WR Quabili, and A Margolies**, “Which Kinds of Social Safety Net Transfers Work Best for the Ultra Poor in Bangladesh,” *Operation and Impacts of the Transfer Modality Research Initiative*, 2016.
- Aizer, Anna, Shari Eli, Joseph Ferrie, and Adriana Lleras-Muney**, “The long-run impact of cash transfers to poor families,” *The American Economic Review*, 2016, 106 (4), 935–971.
- Aker, Jenny C, Rachid Boumnijel, Amanda McClelland, and Niall Tierney**, “Payment mechanisms and antipoverty programs: Evidence from a mobile money cash transfer experiment in Niger,” *Economic Development and Cultural Change*, 2016, 65 (1), 1–37.
- Anderson, Michael L**, “Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American statistical Association*, 2008, 103 (484), 1481–1495.
- Anderson, Siwan and Jean-Marie Baland**, “The economics of roscas and intrahousehold resource allocation,” *The Quarterly Journal of Economics*, 2002, 117 (3), 963–995.
- Andreoni, James and Charles Sprenger**, “Estimating time preferences from convex budgets,” *American Economic Review*, 2012, 102 (7), 3333–3356.
- Araujo, M Caridad, Mariano Bosch, and Norbert Schady**, “Can Cash Transfers Help Households Escape an Inter-Generational Poverty Trap?,” in “The Economics of Poverty Traps,” University of Chicago Press, 2017.
- Ashraf, Nava**, “Spousal control and intra-household decision making: An experimental study in the Philippines,” *American Economic Review*, 2009, 99 (4), 1245–77.
- , **Dean Karlan, and Wesley Yin**, “Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines,” *The Quarterly Journal of Economics*, 2006, 121 (2), 635–672.
- Athey, Susan and Stefan Wager**, “Estimating treatment effects with causal forests: An application,” *Observationl Studies*, 2019, 5, 21–35.
- , **Julie Tibshirani, and Stefan Wager**, “Generalized random forests,” *The Annals of Statistics*, 2019, 47 (2), 1148–1178.
- Baird, Sarah, Craig McIntosh, and Berk Özler**, “Cash or condition? Evidence from a cash transfer experiment,” *The Quarterly Journal of Economics*, 2011, 126 (4), 1709–1753.
- , – , and – , “When the Money Runs Out: Do Cash Transfers Have Sustained Effects on Human Capital Accumulation?,” *World Bank Policy Research Working Paper*, 2016, 7901.
- , – , and **Berk Özler**, “When the money runs out: Do cash transfers have sustained effects on human capital accumulation?,” *Journal of Development Economics*, 2019, 140, 169 – 185.
- , **Jacobus De Hoop, and Berk Özler**, “Income shocks and adolescent mental health,” *Journal of Human Resources*, 2013, 48 (2), 370–403.

- Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry**, “A multifaceted program causes lasting progress for the very poor: Evidence from six countries,” *Science*, 2015, *348* (6236).
- Barham, Tania, Karen Macours, and John A Maluccio**, “Schooling, Learning, and Earnings: Effects of a 3-Year Conditional Cash Transfer Program in Nicaragua After 10 Years,” *La Plata, Argentina: Centro para los Estudios Distributivos, de Trabajo y Sociales. Disponible en: http://www.cedlas-er.org/sites/default/files/aux_files/barham-maluccio-macours_.pdf*, 2014.
- Basinga, Paulin, Paul J Gertler, Agnes Binagwaho, Agnes LB Soucat, Jennifer Sturdy, and Christel MJ Vermeersch**, “Effect on maternal and child health services in Rwanda of payment to primary health-care providers for performance: an impact evaluation,” *The Lancet*, 2011, *377* (9775), 1421–1428.
- Belloni, A, D Chen, V Chernozhukov, and C Hansen**, “Sparse models and methods for optimal instruments with an application to eminent domain,” *Econometrica*, November 2012, *80* (6), 2369–2429.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen**, “High-dimensional methods and inference on structural and treatment effects,” *The Journal of Economic Perspectives*, 2014, *28* (2), 29–50.
- , – , and – , “Inference on treatment effects after selection among high-dimensional controls,” *Review of Economic Studies*, 2014, *81*, 608–650.
- Benhassine, Najj, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen**, “Turning a shove into a nudge? A labeled cash transfer for education,” *American Economic Journal: Economic Policy*, 2015, *7* (3), 86–125.
- Blattman, Christopher and Paul Niehaus**, “Show them the money: why giving cash helps alleviate poverty,” *Foreign Aff.*, 2014, *93*, 117.
- Brudevold-Newman, Andrew Peter, Maddalena Honorati, Pamela Jakiela, and Owen W Ozier**, “A firm of one’s own: experimental evidence on credit constraints and occupational choice,” 2017.
- Card, David and Alan B Krueger**, “Time-series minimum-wage studies: a meta-analysis,” *American Economic Review*, May 1995, *85* (2), 238–243.
- Casey, Katherine, Rachel Glennerster, and Edward Miguel**, “Reshaping institutions: Evidence on aid impacts using a preanalysis Plan,” *Quarterly Journal of Economics*, 2012, *127* (4), 1755–1812.
- Chowdhury, Reajul, Elliott Collins, Ethan Ligon, and Munshi Sulaiman**, “Valuing Assets Provided to Low-Income Households in South Sudan,” 2017.
- Cunha, Jesse M, Giacomo De Giorgi, and Seema Jaychandran**, “The price effects of cash versus in-kind transfers,” *Review of Economic Studies*, forthcoming.
- Currie, Janet and Douglas Almond**, “Human capital development before age five,” *Handbook of labor economics*, 2011, *4*, 1315–1486.

- Desai, Amy, Laura E Smith, Mduduzi NN Mbuya, Ancikaria Chigumira, Dadirai Fundira, Naume V Tavengwa, Thokozile R Malaba, Florence D Majo, Jean H Humphrey, and Rebecca J Stoltzfus**, “The SHINE Trial infant feeding intervention: pilot study of effects on maternal learning and infant diet quality in rural Zimbabwe,” *Clinical Infectious Diseases*, 2015, *61* (suppl_7), S710–S715.
- Dhaliwal, Iqbal and Caitlin Tulloch**, “From research to policy: using evidence from impact evaluations to inform development policy,” *Journal of Development Effectiveness*, 2012, *4* (4), 515–536.
- DHS**, “Rwanda Demographic and Health Survey, 2014–15,” Technical Report, National Institute of Statistics of Rwanda, Ministry of Finance and Economic Planning, Ministry of Health, and the DHS Program, ICF International March 2016.
- Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael W Walker**, “General equilibrium effects of cash transfers: experimental evidence from Kenya,” Technical Report, National Bureau of Economic Research 2019.
- Fernald, Lia CH, Paul J Gertler, and Lynnette M Neufeld**, “10-year effect of Oportunidades, Mexico’s conditional cash transfer programme, on child growth, cognition, language, and behaviour: a longitudinal follow-up study,” *The Lancet*, 2009, *374* (9706), 1997–2005.
- Figlio, David, Jonathan Guryan, Krzysztof Karbownik, and Jeffrey Roth**, “The effects of poor neonatal health on children’s cognitive development,” *The American Economic Review*, 2014, *104* (12), 3921–3955.
- Fink, Günther, Rachel Levenson, Sarah Tembo, and Peter C Rockers**, “Home-and community-based growth monitoring to reduce early life growth faltering: an open-label, cluster-randomized controlled trial,” *The American journal of clinical nutrition*, 2017, *106* (4), 1070–1077.
- Finkelstein, Amy and Nathaniel Hendren**, “Welfare Analysis Meets Causal Inference,” *Journal of Economic Perspectives*, 2020, *34* (4), 146–167.
- Gahamanyi, Vincent and Andrew Kettlewell**, “Evaluating graduation: insights from the Vision 2020 Umurenge Programme in Rwanda,” *IDS Bulletin*, 2015, *46* (2), 48–63.
- Gertler, Paul J, Sebastian W Martinez, and Marta Rubio-Codina**, “Investing cash transfers to raise long-term living standards,” *American Economic Journal: Applied Economics*, 2012, *4* (1), 164–192.
- Giné, Xavier, Jessica Goldberg, Dan Silverman, and Dean Yang**, “Revising commitments: Field evidence on the adjustment of prior choices,” *Economic Journal*, 2018, *128* (608), 159–188.
- Habicht, Jean-Pierre and David L Pelletier**, “The importance of context in choosing nutritional indicators,” *The Journal of nutrition*, 1990, *120* (11 Suppl), 1519–1524.
- Haushofer, Johannes and Jeremy Shapiro**, “The short-term impact of unconditional cash transfers to the poor: Experimental Evidence from Kenya,” *The Quarterly Journal of Economics*, 2016, *131* (4), 1973–2042.
- Heckman, James J, Sergio Urzua, and Edward Vytlacil**, “Understanding instrumental variables in models with essential heterogeneity,” *The Review of Economics and Statistics*, 2006, *88* (3), 389–432.

- Hendren, Nathaniel and Ben Sprung-Keyser**, “A unified welfare analysis of government policies,” *The Quarterly Journal of Economics*, 2020, *125* (3), 1209–1318.
- Hidrobo, Melissa, John Hoddinott, Amber Peterman, Amy Margolies, and Vanessa Moreira**, “Cash, food, or vouchers? Evidence from a randomized experiment in northern Ecuador,” *Journal of Development Economics*, 2014, *107*, 144–156.
- Hoddinott, John, Harold Alderman, Jere R Behrman, Lawrence Haddad, and Susan Horton**, “The economic rationale for investing in stunting reduction,” *Maternal & Child Nutrition*, 2013, *9* (S2), 69–82.
- , **Susanna Sandström, and Joanna Upton**, “The impact of cash and food transfers: Evidence from a randomized intervention in Niger,” 2014.
- Hoynes, Hilary, Diane Whitmore Schanzenbach, and Douglas Almond**, “Long-run impacts of childhood access to the safety net,” *The American Economic Review*, 2016, *106* (4), 903–934.
- Laibson, David**, “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, May 1997, *112* (2), 443–478.
- Leroy, Jef L, Deanna Olney, and Marie Ruel**, “Tubaramure, a Food-Assisted Integrated Health and Nutrition Program in Burundi, Increases Maternal and Child Hemoglobin Concentrations and Reduces Anemia: A Theory-Based Cluster-Randomized Controlled Intervention Trial–3,” *The Journal of nutrition*, 2016, *146* (8), 1601–1608.
- , **Paola Gadsden, Sonia Rodríguez-Ramírez, and Teresa González De Cossío**, “Cash and in-kind transfers in poor rural communities in Mexico increase household fruit, vegetable, and micronutrient consumption but also lead to excess energy consumption,” *The Journal of nutrition*, 2010, *140* (3), 612–617.
- Levin, Henry M and Patrick J McEwan**, *Cost-effectiveness analysis: Methods and applications*, Vol. 4, Sage, 2001.
- Maluccio, John A, John Hoddinott, Jere R Behrman, Reynaldo Martorell, Agnes R Quisumbing, and Aryeh D Stein**, “The impact of improving nutrition during early childhood on education among Guatemalan adults,” *The Economic Journal*, 2009, *119* (537), 734–763.
- Manley, James, Seth Gitter, and Vanya Slavchevska**, “How effective are cash transfers at improving nutritional status?,” *World development*, 2013, *48*, 133–155.
- McIntosh, Craig and Andrew Zeitlin**, “Benchmarking a Workforce Readiness Program Against Cash: Experimental Evidence from Rwanda,” Technical Report, CECA Working Paper 2020.
- Mel, Suresh De, David McKenzie, and Christopher Woodruff**, “One-time transfers of cash or capital have long-lasting effects on microenterprises in Sri Lanka,” *Science*, 2012, *335* (6071), 962–966.
- Molina-Millan, Teresa, Tania Barham, Karen Macours, John A Maluccio, and Marco Stampini**, “Long-Term Impacts of Conditional Cash Transfers in Latin America: Review of the Evidence,” Technical Report, Inter-American Development Bank 2016.
- NISR**, “Rwanda Demographic and Health Survey 2014-15. Rockvill,” Technical Report, National Institute of Statistics of Rwanda (NISR) [Rwanda], Ministry of Health (MOH) [Rwanda] Rockville, Maryland, USA. 2015.

- Null, Clair, Christine P Stewart, Amy J Pickering, Holly N Dentz, Benjamin F Arnold, Charles D Arnold, Jade Benjamin-Chung, Thomas Clasen, Kathryn G Dewey, Lia CH Fernald et al.**, “Effects of water quality, sanitation, handwashing, and nutritional interventions on diarrhoea and child growth in rural Kenya: a cluster-randomised controlled trial,” *The Lancet Global Health*, 2018, 6 (3), e316–e329.
- Pega, Frank, Stefan Walter, Sze Yan Liu, Roman Pabayo, Stefan K Lhachimi, and Ruhi Saith**, “Unconditional cash transfers for reducing poverty and vulnerabilities: effect on use of health services and health outcomes in low-and middle-income countries,” *The Cochrane Library*, 2014.
- Ruel, Marie T, Harold Alderman, Maternal, Child Nutrition Study Group et al.**, “Nutrition-sensitive interventions and programmes: how can they help to accelerate progress in improving maternal and child nutrition?,” *The Lancet*, 2013, 382 (9891), 536–551.
- Samuels, Fiona and Maria Stavropoulou**, “Being Able to Breathe Again: The Effects of Cash Transfer Programmes on Psychosocial Wellbeing,” *The Journal of Development Studies*, 2016, 52 (8), 1099–1114.
- Schwab, Benjamin et al.**, “In the form of bread? A randomized comparison of cash and food transfers in Yemen,” in “Agricultural & Applied Economics Association 2013 AAEA & CAES Joint Annual Meeting” 2013, pp. 4–6.
- Sedlmayr, Richard, Anuj Shah, Munshi Sulaiman et al.**, “Cash-Plus: Cash Transfer Extensions and Recipient Agency,” Technical Report, Centre for the Study of African Economies, University of Oxford 2017.
- Seidenfeld, David, Sudhanshu Handa, Gelson Tembo, Stanfield Michelo, Charlotte Harland Scott, and Leah Prencipe**, “The impact of an unconditional cash transfer on food security and nutrition: the Zambia child grant programme,” 2014.
- Shrimpton, Roger, Cesar G Victora, Mercedes de Onis, Rosângela Costa Lima, Monika Blössner, and Graeme Clugston**, “Worldwide timing of growth faltering: implications for nutritional interventions,” *Pediatrics*, 2001, 107 (5), e75–e75.
- Skoufias, Emmanuel, Susan W Parker, Jere R Behrman, and Carola Pessino**, “Conditional cash transfers and their impact on child work and schooling: Evidence from the programa program in Mexico [with comments],” *Economia*, 2001, 2 (1), 45–96.
- Suri, Tavneet**, “Mobile money,” *Annual Review of Economics*, 2017, 9, 497–520.
- Taylor, J Edward, Mateusz J Filipski, Mohamad Alloush, Anubhab Gupta, Ruben Irvin Rojas Valdes, and Ernesto Gonzalez-Estrada**, “Economic impact of refugees,” *Proceedings of the National Academy of Sciences*, 2016, 113 (27), 7449–7453.
- Thaler, Richard H and Shlomo Benartzi**, “Save more tomorrow™: Using behavioral economics to increase employee saving,” *Journal of political Economy*, 2004, 112 (S1), S164–S187.
- Victora, Cesar Gomes, Mercedes de Onis, Pedro Curi Hallal, Monika Blössner, and Roger Shrimpton**, “Worldwide timing of growth faltering: revisiting implications for interventions,” *Pediatrics*, 2010, pp. peds–2009.

Vivalt, Eva, “How much can we generalize from impact evaluations?,” *Journal of the European Economic Association*, 2015.

Wager, Stefan and Susan Athey, “Estimation and inference of heterogeneous treatment effects using random forests,” *Journal of the American Statistical Association*, 2018, *113* (523), 1228–1242.

TABLES AND FIGURES

TABLE I: RESEARCH DESIGN

	Control	Gikuriro	GiveDirectly				Total
			Lower transfer (\$66)	Middle transfer (\$111)	Upper transfer (\$145)	Large transfer (\$566)	
<i>Panel A. Village-level randomization</i>							
Villages	74	74	22	22	22	34	248
Ineligible households	298	297	88	87	88	137	995
Eligible households	521	541	165	154	167	246	1,794
<i>Panel B. Household-level randomization of cash payment modality among eligibles</i>							
Flow transfers			83	87	104	147	421
Lump-sum transfers			51	50	41	68	210
Choice			31	17	22	31	101

NOTES—Table gives the number of observations across the arms of the study. The first row is the number of villages, the unit of assignment, and the remaining rows are number of households.

TABLE II: INTERVENTION COSTS AND COMPLIANCE RATES

	Gikuriro	GiveDirectly			
		Lower	Middle	Upper	Large
Cost to USAID per beneficiary	\$141.84	\$66.02	\$111.09	\$145.43	\$566.55
Share averted if untreated	0.60	1.00	1.00	1.00	1.00
Compliance rate among eligibles	0.80	0.81	0.86	0.83	0.91
Cost to USAID per eligible household	\$124.49	\$53.58	\$95.86	\$121.24	\$517.44
Compliance rate in population	0.19	0.18	0.19	0.18	0.18
Cost to USAID per village household	\$28.02	\$12.20	\$20.83	\$26.69	\$99.56

NOTES—Table gives costs for each of the study arms. The first row is the cost per beneficiary that emerged from the ex-post costing exercise. The second row gives the share of program costs averted among non-compliers, and the third row the compliance rate. Using these the fourth row provides the average amount spent per study eligible household. The fifth and sixth rows provide cost numbers for the average household in the village population.

TABLE III: ITT ESTIMATES: PRIMARY OUTCOMES

	Gikuriro	GiveDirectly		Control Mean	Obs.	R^2	p -values: B/C ratios	
		Main	Large				GD=GDL	GK=GDL
<i>Panel A. Household outcomes</i>								
Consumption [†]	−0.11 (0.10) [0.41]	0.06 (0.09) [0.57]	0.30*** (0.11) [0.02]	10.69	1750	0.14	0.93	0.04
Household dietary diversity score	0.19 (0.12) [0.38]	0.17 (0.15) [0.41]	0.55*** (0.13) [0.00]	4.77	1751	0.18	0.60	0.59
Household non-land wealth [†]	0.01 (0.18) [0.60]	0.00 (0.21) [0.60]	0.40 (0.28) [0.38]	13.04	1751	0.22	0.74	0.60
<i>Panel B. Individual outcomes</i>								
Height-for-age	0.05 (0.04) [1.00]	−0.02 (0.04) [1.00]	0.09** (0.05) [0.78]	−1.97	2125	0.71	0.33	0.47
Weight-for-age	0.04 (0.04) [1.00]	0.01 (0.03) [1.00]	0.07* (0.04) [0.78]	−1.04	2104	0.68	0.97	0.54
Mid-upper arm circumference	0.02 (0.06) [1.00]	−0.01 (0.07) [1.00]	0.13* (0.08) [0.78]	−0.59	1629	0.51	0.63	0.85
Child anemia	0.00 (0.02) [1.00]	0.02 (0.02) [1.00]	−0.01 (0.04) [1.00]	0.18	2372	0.07	0.30	0.77
Maternal anemia	−0.02 (0.03) [1.00]	−0.00 (0.03) [1.00]	−0.03 (0.03) [1.00]	0.12	1581	0.11	0.91	0.48

NOTES—Table presents Intention to Treat impacts on primary outcomes, with the three study arms presented in rows and the three smaller GiveDirectly transfers pooled into the ‘Main’ treatment. Regressions include but do not report the lagged dependent variable, fixed effects for randomization blocks, and a set of LASSO-selected baseline covariates, and are weighted to be representative of the eligible population. Standard errors (in parentheses) are clustered at the household level to reflect the design effect. Asterisks denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a † are in inverse hyperbolic sines. Reported p -values in final two columns derived from F -tests of hypotheses that cost-benefit ratios are equal between GD Main and Large transfer amounts (GD=GDL), and between Gikuriro and GD Large (GK=GDL).

TABLE IV: ITT ESTIMATES: SECONDARY OUTCOMES

	Gikuriro	GiveDirectly		Control		R^2	p -values: B/C ratios	
		Main	Large	Mean	Obs.		GD=GDL	GK=GDL
<i>Panel A. Household outcomes</i>								
Stock of borrowing [†]	0.067 (0.350) [0.71]	-0.765** (0.316) [0.05]	-0.341 (0.397) [0.64]	7.39	1751	0.12	0.01	0.63
Stock of saving [†]	1.115*** (0.324) [0.01]	-0.128 (0.345) [0.71]	0.656** (0.329) [0.14]	5.88	1751	0.16	0.46	0.00
Health knowledge index	-0.076 (0.368) [0.71]	0.153 (0.321) [0.71]	0.075 (0.468) [0.71]	2.89	1751	0.04	0.64	0.78
Sanitation practices index	-0.280 (0.219) [0.37]	0.146 (0.227) [0.71]	0.087 (0.284) [0.71]	-0.68	1751	0.07	0.54	0.13
Productive assets [†]	0.020 (0.100) [0.71]	0.257** (0.100) [0.05]	0.800*** (0.116) [0.00]	11.22	1751	0.29	0.21	0.06
Consumption assets [†]	-0.367 (0.240) [0.28]	0.354* (0.206) [0.22]	0.932*** (0.243) [0.00]	8.70	1751	0.35	0.32	0.01
House value [†]	-0.023 (0.056) [0.71]	-0.029 (0.061) [0.71]	0.196*** (0.061) [0.01]	13.81	1654	0.34	0.27	0.16
Housing quality index	-0.195 (0.146) [0.36]	-0.217* (0.132) [0.24]	0.211 (0.174) [0.37]	-0.17	1751	0.10	0.05	0.08
<i>Panel B. Individual outcomes</i>								
Child Mortality	-0.006 (0.005) [1.00]	-0.004 (0.006) [1.00]	-0.009** (0.004) [1.00]	0.01	2687	0.01	0.73	0.39
Pregnancy	-0.031 (0.025) [1.00]	-0.035 (0.031) [1.00]	-0.007 (0.027) [1.00]	0.20	2552	0.08	0.25	0.18
Live Birth	0.103 (0.079) [1.00]	0.091 (0.072) [1.00]	-0.068 (0.081) [1.00]	0.68	411	0.13	0.12	0.09
Birth in Facility	-0.046 (0.059) [1.00]	0.069 (0.052) [1.00]	-0.062 (0.099) [1.00]	0.84	293	0.16	0.11	0.59
Any Vaccinations in past year	0.010 (0.034) [1.00]	-0.010 (0.032) [1.00]	-0.005 (0.039) [1.00]	0.72	1291	0.26	0.75	0.72
Completed Vaccinations	0.011 (0.039) [1.00]	-0.013 (0.035) [1.00]	0.006 (0.038) [1.00]	0.58	1291	0.17	0.67	0.79
Disease Burden	-0.020 (0.032) [1.00]	-0.031 (0.030) [1.00]	-0.018 (0.033) [1.00]	0.54	2680	0.05	0.33	0.61
Diarrheal Prevalence	-0.003 (0.015) [1.00]	-0.000 (0.016) [1.00]	-0.007 (0.015) [1.00]	0.09	2680	0.04	0.95	0.89

NOTES—Table presents Intention to Treat impacts on secondary outcomes, with the three study arms presented in rows and the three smaller GiveDirectly transfers pooled into the ‘Main’ treatment. Regressions include but do not report the lagged dependent variable, fixed effects for randomization blocks, and a set of LASSO-selected baseline covariates, and are weighted be representative of the eligible population. Standard errors (in parentheses) are clustered at the household level to reflect the design effect. Asterices denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a † are in inverse hyperbolic sines. Reported p -values in final two columns derived from F -tests of hypotheses that cost-benefit ratios are equal between GD Main and Large transfer amounts (GD=GDL), and between Gikuriro and GD Large (GK=GDL).

TABLE V: COST-EQUIVALENT COMPARISONS: PRIMARY OUTCOMES

	Gikuriro: Differential	Cost-equivalent GD impact	Transfer Cost	Control Mean	Observations	R^2
<i>A. Household outcomes</i>						
Consumption [†]	−0.19** (0.08) [0.06]	0.08 (0.09) [0.49]	0.05** (0.02) [0.06]	10.69	1750	0.14
Household dietary diversity score	−0.01 (0.13) [0.71]	0.20 (0.14) [0.28]	0.09** (0.03) [0.06]	4.77	1751	0.18
Household non-land wealth [†]	−0.03 (0.20) [0.71]	0.04 (0.21) [0.71]	0.09 (0.07) [0.30]	13.04	1751	0.22
<i>B. Individual outcomes</i>						
Height-for-age	0.06 (0.04) [0.96]	−0.01 (0.04) [1.00]	0.02** (0.01) [0.40]	−1.97	2125	0.71
Weight-for-age	0.02 (0.04) [1.00]	0.02 (0.03) [1.00]	0.01 (0.01) [0.96]	−1.04	2104	0.68
Mid-upper arm circumference	0.02 (0.06) [1.00]	0.00 (0.06) [1.00]	0.03 (0.02) [0.96]	−0.59	1629	0.51
Child anemia	−0.02 (0.03) [1.00]	0.02 (0.02) [1.00]	−0.01 (0.01) [1.00]	0.22	2372	0.07
Maternal anemia	−0.02 (0.02) [1.00]	−0.00 (0.03) [1.00]	−0.00 (0.01) [1.00]	0.12	1581	0.10

NOTES—First column is a dummy for Gikuriro treatment, giving the differential effect of Gikuriro over cash at equivalent cost. Second column is a dummy for either treatment, giving the impact of cash at the cost of Gikuriro. Third column is the cost slope, measured as the dollar-value deviation (in hundreds of dollars) of the treatment received from the cost of Gikuriro. Asterices denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a [†] are in inverse hyperbolic sines.

TABLE VI: COST-EQUIVALENT COMPARISONS: SECONDARY OUTCOMES

	Gikuriro: Differential	Cost-equivalent GD impact	Transfer Cost	Control Mean	Observations	R^2
<i>A. Household outcomes</i>						
Stock of borrowing [†]	0.793*** (0.277) [0.01]	-0.726** (0.309) [0.04]	0.094 (0.077) [0.19]	7.39	1751	0.12
Stock of saving [†]	1.184*** (0.333) [0.01]	-0.070 (0.327) [0.54]	0.188** (0.080) [0.04]	5.88	1751	0.16
Health knowledge index	-0.227 (0.329) [0.44]	0.151 (0.316) [0.54]	-0.023 (0.103) [0.54]	2.89	1751	0.04
Sanitation practices index	-0.421** (0.213) [0.06]	0.141 (0.220) [0.44]	-0.014 (0.069) [0.54]	-0.68	1751	0.07
Productive assets [†]	-0.283*** (0.097) [0.01]	0.302*** (0.096) [0.01]	0.125*** (0.026) [0.00]	11.22	1751	0.29
Consumption assets [†]	-0.764*** (0.256) [0.01]	0.397** (0.199) [0.06]	0.139** (0.058) [0.04]	8.70	1751	0.35
House value [†]	-0.017 (0.051) [0.54]	-0.007 (0.059) [0.55]	0.049*** (0.014) [0.01]	13.81	1654	0.34
Housing quality index	-0.008 (0.151) [0.56]	-0.187 (0.126) [0.12]	0.104** (0.048) [0.05]	-0.17	1751	0.10
<i>B. Individual outcomes</i>						
Child mortality	-0.002 (0.005) [1.00]	-0.004 (0.006) [1.00]	-0.001 (0.001) [1.00]	0.01	2687	0.01
Pregnancy	0.002 (0.027) [1.00]	-0.033 (0.029) [1.00]	0.007 (0.007) [1.00]	0.20	2552	0.08
Live birth	0.024 (0.067) [1.00]	0.079 (0.069) [1.00]	-0.037** (0.017) [1.00]	0.68	411	0.13
Birth in facility	-0.104** (0.052) [1.00]	0.058 (0.051) [1.00]	-0.029 (0.023) [1.00]	0.84	293	0.16
Any vaccinations in past year	0.018 (0.032) [1.00]	-0.008 (0.031) [1.00]	-0.000 (0.009) [1.00]	0.72	1291	0.26
Completed vaccination schedule	0.021 (0.036) [1.00]	-0.010 (0.033) [1.00]	0.003 (0.009) [1.00]	0.58	1291	0.17
Disease burden	0.012 (0.033) [1.00]	-0.031 (0.029) [1.00]	0.004 (0.008) [1.00]	0.54	2680	0.05
Diarrheal prevalence	-0.003 (0.014) [1.00]	-0.000 (0.015) [1.00]	-0.002 (0.004) [1.00]	0.09	2680	0.04

NOTES—First column is a dummy for Gikuriro treatment, giving the differential effect of Gikuriro over cash at equivalent cost. Second column is a dummy for either treatment, giving the impact of cash at the cost of Gikuriro. Third column is the cost slope, measured as the dollar-value deviation (in hundreds of dollars) of the treatment received from the cost of Gikuriro. Asterices denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a [†] are in inverse hyperbolic sines.

TABLE VII: BENEFIT-COST RATIOS FOR PRIMARY OUTCOMES

	Gikuriro	GiveDirectly		<i>p</i> -values		
		Main	Large	(a)	(b)	(c)
<i>Panel A. Household outcomes</i>						
Consumption [†]	−0.090 (0.077)	0.066 (0.100)	0.058 (0.021)	0.05	0.93	0.04
Household dietary diversity score	0.155 (0.098)	0.186 (0.162)	0.106 (0.024)	0.83	0.60	0.59
Household non-land wealth [†]	0.007 (0.148)	0.005 (0.235)	0.077 (0.054)	0.99	0.74	0.60
<i>Panel B. Individual outcomes</i>						
Height-for-age	0.042 (0.036)	−0.021 (0.044)	0.018 (0.009)	0.11	0.33	0.47
Weight-for-age	0.031 (0.032)	0.012 (0.037)	0.013 (0.007)	0.61	0.97	0.54
Mid-upper arm circumference	0.018 (0.045)	−0.007 (0.072)	0.026 (0.015)	0.71	0.63	0.85
Child anemia	0.003 (0.018)	0.026 (0.028)	−0.002 (0.007)	0.38	0.30	0.77
Maternal anemia	−0.019 (0.022)	−0.002 (0.032)	−0.005 (0.006)	0.46	0.91	0.48

NOTES—Benefit-cost ratios derived from ITT estimates and estimated costs per eligible household. *p*-values report tests of equal BCR between (a) Gikuriro and GD-Main; (b) GD-Main and GD-Large; and (c) Gikuriro and GD-large. Household-level BCRs; the average eligible household contains 5.2 members, 1.5 anthro-eligible children, 1.7 children eligible for anemia testing, and 1.2 adult women eligible for anemia testing. Per-person BCRs should be scaled up by these numbers.

TABLE VIII: VILLAGE-LEVEL TOTAL CAUSAL EFFECTS: PRIMARY OUTCOMES

	Gikuriro	GiveDirectly		Control		R^2	p -values: B/C ratios	
		Main	Large	Mean	Obs.		GD=GDL	GK=GDL
<i>Panel A. Household outcomes</i>								
Consumption [†]	-0.14 (0.10) [0.60]	-0.14 (0.10) [0.60]	-0.06 (0.17) [0.81]	10.55	2717	0.15	0.18	0.20
Household dietary diversity score	0.12 (0.13) [0.60]	0.00 (0.13) [0.85]	-0.28 (0.19) [0.60]	4.29	2718	0.22	0.64	0.11
Household non-land wealth [†]	-0.19 (0.23) [0.60]	-0.32 (0.28) [0.60]	-0.46 (0.30) [0.60]	14.02	2718	0.29	0.40	0.79
<i>Panel B. Individual outcomes</i>								
Height-for-Age	-0.01 (0.05) [1.00]	-0.01 (0.05) [1.00]	0.05 (0.06) [1.00]	-1.75	2618	0.73	0.66	0.62
Weight-for-Age	-0.06 (0.05) [1.00]	-0.01 (0.04) [1.00]	0.02 (0.05) [1.00]	-0.87	2594	0.73	0.62	0.17
Mid-Upper Arm Circ	-0.07 (0.07) [1.00]	-0.02 (0.07) [1.00]	0.06 (0.09) [1.00]	-0.61	1981	0.57	0.57	0.17

NOTES—Analysis pools eligible and ineligible households and is weighted to be representative of the population in study villages. Asterics denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a † are in inverse hyperbolic sines. Final two columns present p -values on tests of the equality of benefit-cost ratios between arms.

TABLE IX: CORRELATION BETWEEN HOUSEHOLD-LEVEL IMPACT DOMAINS

	Consumption	Wealth	Dietary diversity	Child growth
Consumption				
Wealth	-0.29			
Dietary diversity	0.07	0.07		
Child growth	0.04	0.28	0.15	

NOTES—Table shows correlation coefficients between estimated subgroup mean treatment effects across outcome domains. Treatment effects estimated by causal forest. Child growth is a standardized index of HAZ, WAZ, and MUAC.

TABLE X: WELFARE GAINS FROM TARGETING MODALITY

Target	Gains (standard deviation of outcome)
Consumption	0.016
Wealth	0.057
Dietary diversity	0.086
Child growth	0.103
Composite welfare metric	0.029

NOTES—Table reports gains relative to control from offering households the treatment (Gikuriro or GD Main) that is expected to deliver the bigger impact, based on causal forest estimates. All outcomes are standardized to have mean of zero and variance of one in the control group at endline. The first four rows represent welfare gains if treatment assignment is optimized for that outcome only; the final row represents average gains across the four dimension if households are assigned the treatment that is expected to deliver the largest average impact across these dimensions.

FIGURE I: COST EQUIVALENCE VERSUS COST EFFECTIVENESS

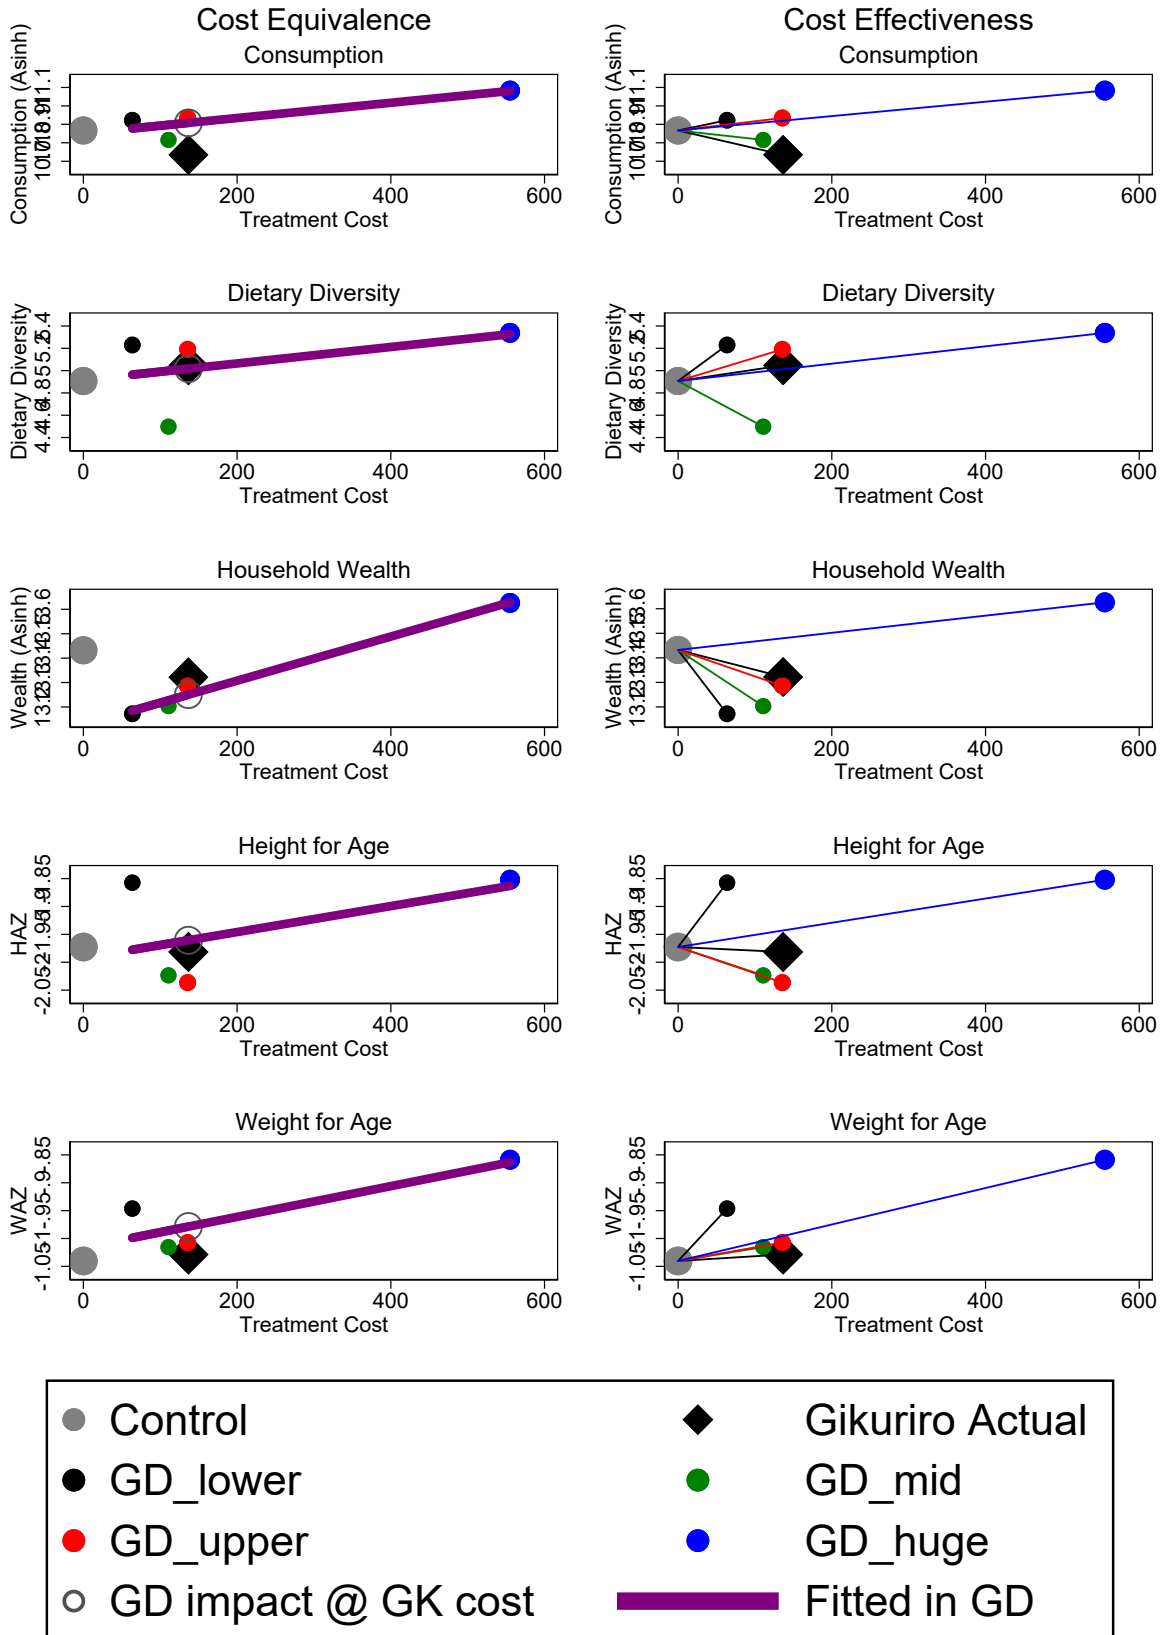
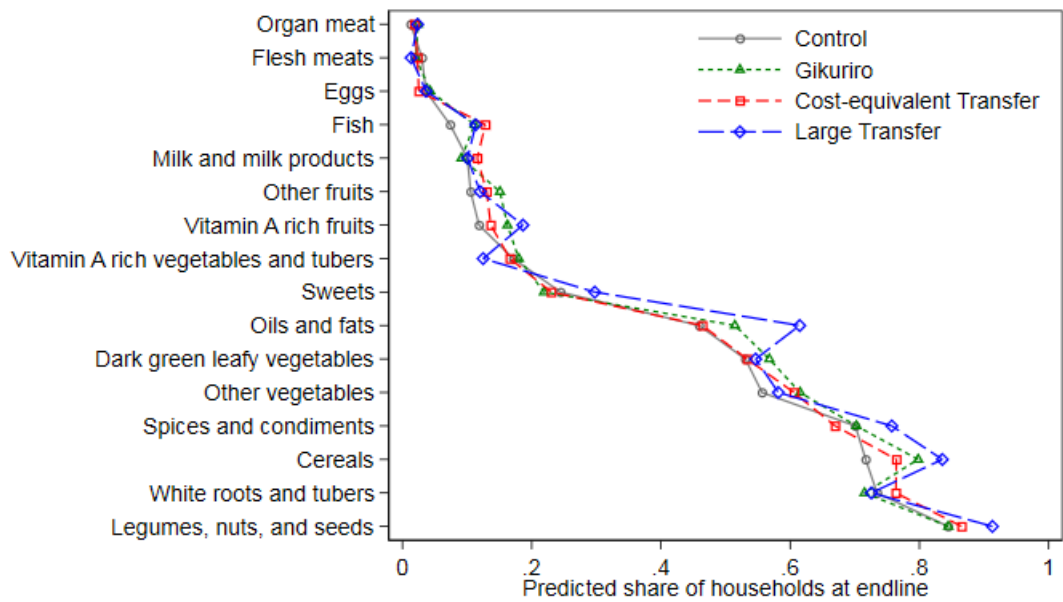


FIGURE II: IMPACTS ON DIETARY DIVERSITY



APPENDIX A. SUPPLEMENTARY TABLES AND FIGURES

TABLE A.1: ATTRITION

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Household	Household	Roster	Roster	Anthro	Anthro	Anemia	Anemia	New Member
	b/se	b/se	b/se	b/se	b/se	b/se	b/se	b/se	b/se
Gikurio	-0.013 (0.0092)	-0.011 (0.0085)	-0.010 (0.0083)	-0.0086 (0.0079)	-0.037*** (0.013)	-0.038*** (0.012)	-0.0044 (0.021)	-0.0046 (0.016)	-0.00050 (0.0081)
GD Main	-0.0089 (0.010)	-0.0097 (0.010)	-0.0077 (0.0093)	-0.0076 (0.0097)	-0.024* (0.013)	-0.025** (0.013)	-0.0069 (0.019)	-0.0013 (0.016)	-0.0027 (0.0088)
GD Large	-0.017* (0.0097)	-0.015 (0.0093)	-0.015* (0.0086)	-0.013 (0.0088)	-0.022 (0.016)	-0.020 (0.015)	0.027 (0.026)	0.0056 (0.023)	-0.0051 (0.0096)
Agricultural		-0.015 (0.012)		-0.020 (0.013)		-0.037* (0.020)		0.023 (0.023)	
Wage Worker		0.0076 (0.0093)		0.016 (0.0099)		0.019 (0.015)		0.0038 (0.016)	
Microenterprise		-0.0083 (0.0075)		-0.0085 (0.0063)		-0.00060 (0.017)		-0.039* (0.022)	
Savings Group		0.00040 (0.0085)		-0.0051 (0.0071)		-0.015 (0.012)		-0.0020 (0.017)	
Village Eligibility Ratio		0.021 (0.053)		0.015 (0.048)		0.032 (0.050)		-0.0030 (0.060)	
Age of Head		-0.00034 (0.00022)		-0.00025 (0.00019)		0.00013 (0.00034)		-0.0017*** (0.00052)	
Schooling of Head		0.11 (0.14)		0.21 (0.21)		0.053 (0.11)		0.13*** (0.025)	
Dependency Ratio		0.035 (0.030)		0.027 (0.032)		0.058 (0.056)		0.029 (0.063)	
Household Size		-0.0036 (0.0023)		-0.0031 (0.0023)		-0.0060 (0.0040)		-0.040*** (0.0058)	
Poorest Category		0.0097 (0.011)		-0.0055 (0.0098)		0.0065 (0.018)		-0.017 (0.021)	
Next Poorest Category		0.0061 (0.0085)		-0.0048 (0.0094)		0.0059 (0.014)		0.0064 (0.017)	
female				0.0011 (0.0026)		0.0012 (0.0088)		-0.045*** (0.014)	
Control group mean	0.033	0.033	0.030	0.030	0.071	0.071	0.80	0.80	0.057
Observations	1793	1793	9189	9189	2265	2265	5192	5192	9509
R ²	0.0019	0.023	0.0014	0.028	0.0049	0.025	0.00068	0.068	0.000054

NOTES—First two columns analyze household-level attrition from the baseline sample of households. (3) and (4) examine individual-level attrition from the roster of household members. (5) and (6) examine attrition from the baseline sample of child anthropometrics. (7) and (8) analyze attrition among individuals who should have received anemia testing, and (9) analyzes the probability of new household members appearing in rosters.

TABLE A.2: BALANCE ON PRIMARY OUTCOMES

	Gikuriro Village	GD Main Village	GD Large Village	Control Mean	Observations	R^2
<i>A. Household outcomes</i>						
Consumption [†]	0.053 (0.116) [1.00]	0.047 (0.122) [1.00]	-0.103 (0.130) [1.00]	10.39	1751	0.05
Household dietary diversity score	-0.064 (0.137) [1.00]	-0.071 (0.139) [1.00]	-0.058 (0.172) [1.00]	4.16	1751	0.10
Household non-land wealth [†]	0.027 (0.227) [1.00]	-0.042 (0.214) [1.00]	-0.288 (0.254) [1.00]	12.94	1751	0.06
<i>B. Individual outcomes</i>						
Height-for-Age	0.046 (0.086) [1.00]	0.080 (0.097) [1.00]	0.215** (0.098) [0.13]	-1.93	2187	0.02
Weight-for-Age	0.023 (0.069) [1.00]	0.041 (0.070) [1.00]	0.187*** (0.069) [0.07]	-1.06	2180	0.02
Mid-Upper Arm Circ	0.015 (0.068) [1.00]	0.025 (0.068) [1.00]	0.070 (0.081) [1.00]	-0.72	1987	0.04

NOTES—Columns present coefficients and standard errors from a regression of each baseline outcome on treatment indicators, with fixed effects for blocks. Asterics denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a † are in inverse hyperbolic sines.

TABLE A.3: REGRESSIONS WITH GRANULAR GD TREATMENT CELLS

	Gikuro	GiveDirectly				Control Mean	Obs.	R^2
		Small	Mid	Upper	Large			
<i>Panel A. Household outcomes</i>								
Consumption [†]	−0.11 (0.10) [0.48]	0.10 (0.11) [0.66]	−0.02 (0.12) [0.96]	0.10 (0.12) [0.66]	0.30*** (0.11) [0.04]	10.69	1750	0.14
Household dietary diversity score	0.20 (0.12) [0.29]	0.40** (0.19) [0.12]	−0.25 (0.22) [0.48]	0.36** (0.15) [0.10]	0.56*** (0.13) [0.00]	4.77	1751	0.19
Household non-land wealth [†]	0.01 (0.19) [0.96]	0.06 (0.40) [0.96]	0.04 (0.26) [0.96]	−0.09 (0.26) [0.96]	0.40 (0.28) [0.41]	13.04	1751	0.22
<i>Panel B. Individual outcomes</i>								
Height-for-age	0.05 (0.04) [1.00]	0.04 (0.05) [1.00]	−0.04 (0.06) [1.00]	−0.05 (0.05) [1.00]	0.09** (0.05) [1.00]	−1.97	2125	0.71
Weight-for-age	0.04 (0.04) [1.00]	0.01 (0.05) [1.00]	−0.01 (0.05) [1.00]	0.03 (0.04) [1.00]	0.07* (0.04) [1.00]	−1.04	2104	0.68
Mid-upper arm circumference	0.02 (0.06) [1.00]	−0.04 (0.08) [1.00]	0.07 (0.12) [1.00]	−0.06 (0.07) [1.00]	0.13* (0.08) [1.00]	−0.59	1629	0.51
Child anemia	0.00 (0.02) [1.00]	0.05 (0.06) [1.00]	0.01 (0.03) [1.00]	0.02 (0.02) [1.00]	−0.01 (0.04) [1.00]	0.18	2372	0.07
Maternal anemia	−0.02 (0.03) [1.00]	−0.03 (0.03) [1.00]	−0.00 (0.03) [1.00]	0.03 (0.05) [1.00]	−0.03 (0.03) [1.00]	0.12	1581	0.11

NOTES—Analysis includes dummies for each of the four GD transfer amounts separately. Asterisks denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a † are in inverse hyperbolic sines.

TABLE A.4: VILLAGE-LEVEL TOTAL CAUSAL EFFECTS: SECONDARY OUTCOMES

	Gikuriro	GiveDirectly		Control		R^2	p -values: B/C ratios	
		Main	Large	Mean	Obs.		GD=GDL	GK=GDL
<i>Panel A. Household outcomes</i>								
Stock of borrowing [†]	0.141 (0.408) [1.00]	-0.392 (0.394) [1.00]	-0.262 (0.428) [1.00]	5.35	2715	0.11	0.35	0.56
Stock of saving [†]	-0.222 (0.370) [1.00]	-0.427 (0.395) [1.00]	-0.755** (0.380) [1.00]	6.03	2718	0.15	0.46	0.98
Health knowledge index	1.448*** (0.392) [0.01]	0.692* (0.405) [1.00]	-0.388 (0.545) [1.00]	-0.12	2718	0.05	0.05	0.00
Sanitation practices index	0.165 (0.202) [1.00]	-0.269 (0.232) [1.00]	0.304 (0.345) [1.00]	0.29	2718	0.07	0.14	0.68
Productive assets [†]	-0.081 (0.141) [1.00]	-0.190 (0.145) [1.00]	0.222 (0.173) [1.00]	12.14	2718	0.30	0.08	0.26
Consumption assets [†]	0.102 (0.223) [1.00]	-0.028 (0.238) [1.00]	0.354 (0.361) [1.00]	9.85	2718	0.32	0.66	0.99
House value [†]	-0.003 (0.072) [1.00]	0.009 (0.065) [1.00]	0.062 (0.074) [1.00]	13.96	2531	0.39	0.96	0.76
Housing quality index	-0.172 (0.174) [1.00]	-0.007 (0.155) [1.00]	-0.019 (0.192) [1.00]	0.20	2718	0.16	0.98	0.28
<i>Panel B. Individual outcomes</i>								
Child Mortality	0.004 (0.006) [1.00]	-0.001 (0.002) [1.00]	-0.003 (0.002) [1.00]	0.01	3373	0.02	0.69	0.38
Pregnancy	0.022 (0.022) [1.00]	0.003 (0.019) [1.00]	-0.003 (0.022) [1.00]	0.12	4137	0.11	0.83	0.25
Live Birth	0.006 (0.078) [1.00]	-0.047 (0.092) [1.00]	-0.041 (0.093) [1.00]	0.70	594	0.13	0.66	0.81
Birth in Facility	-0.029 (0.069) [1.00]	0.053 (0.061) [1.00]	-0.007 (0.088) [1.00]	0.90	416	0.16	0.30	0.66
Any Vaccinations in past year	0.115* (0.060) [0.41]	0.076 (0.056) [1.00]	0.159*** (0.062) [0.08]	0.73	1479	0.31	0.37	0.16
Completed Vaccinations	0.199*** (0.064) [0.03]	0.080 (0.064) [1.00]	0.210*** (0.064) [0.03]	0.48	1479	0.17	0.51	0.01
Disease Burden	0.018 (0.052) [1.00]	-0.027 (0.047) [1.00]	-0.002 (0.055) [1.00]	0.50	3366	0.06	0.53	0.69
Diarrheal Prevalence	0.007 (0.018) [1.00]	0.034 (0.022) [1.00]	0.030 (0.025) [1.00]	0.07	3366	0.05	0.19	0.93

NOTES—Analysis pools eligible and ineligible households and is weighted to be representative of the population in study villages. Asterics denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a [†] are in inverse hyperbolic sines. Final two columns present p -values on tests of the equality of benefit-cost ratios between arms.

TABLE A.5: ESTIMATED SPILLOVER EFFECTS OF THE GD INTERVENTION ON PRIMARY OUTCOMES

	GD Main Village	GD Large Village	Control Mean	Observations	R^2
<i>A. Household outcomes</i>					
Consumption [†]	-0.112 (0.140) [0.91]	-0.158 (0.204) [0.91]	10.95	477	0.19
Household dietary diversity score	-0.004 (0.186) [0.97]	-0.454 (0.279) [0.91]	4.77	477	0.24
Household non-land wealth [†]	-0.463 (0.391) [0.91]	-0.635 (0.479) [0.91]	14.09	477	0.23
<i>B. Individual outcomes</i>					
Height-for-Age	0.040 (0.083) [1.00]	0.073 (0.082) [1.00]	-1.58	267	0.77
Weight-for-Age	-0.005 (0.065) [1.00]	-0.031 (0.082) [1.00]	-0.69	267	0.75
Mid-Upper Arm Circ	-0.004 (0.115) [1.00]	0.086 (0.107) [1.00]	-0.22	188	0.62

NOTES—Analysis of spillover effects of cash within the strata (non-poor or no kids) not eligible to be treated by GD Weighted regression, using product of sampling weight and probability that an observation is *not* treated, if in control village. Households treated by GD in villages assigned to that arm receive a weight of zero. p -values for test of equality of ratio between effect sizes and costs per eligible individual in the village.

TABLE A.6: ESTIMATED SPILLOVER EFFECTS OF THE GD INTERVENTION ON SECONDARY OUTCOMES

	GD Main Village	GD Large Village	Control Mean	Observations	R^2
<i>A. Household outcomes</i>					
Stock of borrowing [†]	0.071 (0.508) [1.00]	-0.143 (0.538) [1.00]	6.73	475	0.18
Stock of saving [†]	-0.878 (0.536) [1.00]	-1.631*** (0.564) [0.08]	7.04	477	0.17
Health knowledge index	0.608 (0.559) [1.00]	0.232 (0.733) [1.00]	2.12	477	0.07
Sanitation practices index	-0.366 (0.280) [1.00]	0.082 (0.410) [1.00]	0.06	477	0.12
Productive assets [†]	-0.325 (0.218) [1.00]	-0.048 (0.220) [1.00]	11.97	477	0.32
Consumption assets [†]	-0.404 (0.298) [1.00]	0.053 (0.447) [1.00]	10.08	477	0.34
House value [†]	0.087 (0.086) [1.00]	0.079 (0.097) [1.00]	14.13	425	0.39
Housing quality index	-0.037 (0.181) [1.00]	-0.135 (0.250) [1.00]	0.11	477	0.25
<i>B. Individual outcomes</i>					
Pregnancy	-0.007 (0.028) [1.00]	-0.001 (0.032) [1.00]	0.12	745	0.13
Live Birth	-0.136 (0.148) [1.00]	0.023 (0.158) [1.00]	0.68	86	0.48
Birth in Facility	-0.039 (0.097) [1.00]	-0.042 (0.130) [1.00]	0.93	59	0.45
Any Vaccinations in past year	-0.041 (0.063) [1.00]	0.053 (0.070) [1.00]	0.70	290	0.42
Completed Vaccinations	-0.035 (0.064) [1.00]	0.084 (0.078) [1.00]	0.44	290	0.26
Disease Burden	-0.030 (0.073) [1.00]	0.035 (0.083) [1.00]	0.53	364	0.13
Diarrheal Prevalence	0.015 (0.034) [1.00]	0.035 (0.043) [1.00]	0.08	364	0.15

NOTES—Weighted regression, using product of sampling weight and probability that an observation is *not* treated, if in control village. Households treated by GD in villages assigned to that arm receive a weight of zero. p -values for test of equality of ratio between effect sizes and costs per eligible individual in the village.

TABLE A.7: COMPARISON OF LUMP SUM AND FLOW TRANSFERS: PRIMARY

	Main GD Treatment	Main GD Lump Sum	Large GD Treatment	Large GD Lump Sum	Control Mean	Observations	R^2
<i>A. Household outcomes</i>							
Consumption [†]	0.044 (0.097) [1.00]	0.028 (0.127) [1.00]	0.286*** (0.104) [0.04]	0.027 (0.118) [1.00]	10.69	1131	0.15
Household dietary diversity score	0.068 (0.159) [1.00]	0.014 (0.176) [1.00]	0.614*** (0.137) [0.00]	-0.268 (0.237) [1.00]	4.77	1131	0.19
Household non-land wealth [†]	0.012 (0.262) [1.00]	-0.273 (0.502) [1.00]	0.209 (0.388) [1.00]	0.631 (0.441) [1.00]	13.04	1131	0.25
<i>B. Individual outcomes</i>							
Height-for-age	-0.002 (0.045) [0.93]	-0.095* (0.057) [0.52]	0.074 (0.050) [0.52]	0.047 (0.062) [0.56]	-1.97	1380	0.74
Weight-for-age	-0.036 (0.040) [0.56]	0.125* (0.071) [0.52]	0.028 (0.041) [0.56]	0.148** (0.073) [0.52]	-1.04	1369	0.70
Mid-upper arm circumference	-0.097 (0.063) [0.52]	0.320*** (0.112) [0.11]	0.121 (0.082) [0.52]	0.150 (0.107) [0.52]	-0.59	1057	0.53
Child anemia	0.044 (0.032) [0.52]	-0.054 (0.038) [0.52]	-0.007 (0.039) [0.92]	-0.005 (0.052) [0.93]	0.22	1544	0.08
Maternal anemia	0.013 (0.031) [0.72]	-0.033 (0.035) [0.56]	-0.024 (0.037) [0.56]	0.009 (0.053) [0.92]	0.12	1025	0.12

NOTES—Analysis uses only the Control and GD arms; the first and third columns are dummies for GD transfer amount cells and the second and fourth are dummies measuring the additional effect of lump sum transfers within each cell. Asterics denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a [†] are in inverse hyperbolic sines.

TABLE A.8: COMPARISON OF LUMP SUM AND FLOW TRANSFERS: SECONDARY

	Main GD Treatment	Main GD Lump Sum	Large GD Treatment	Large GD Lump Sum	Control Mean	Observations	R^2
<i>A. Household outcomes</i>							
Stock of borrowing [†]	-1.080*** (0.323) [0.02]	1.071** (0.469) [0.16]	0.024 (0.411) [1.00]	-0.356 (0.892) [1.00]	7.39	1131	0.16
Stock of saving [†]	-0.584 (0.375) [0.51]	0.888** (0.429) [0.22]	0.573 (0.410) [0.59]	0.009 (0.723) [1.00]	5.88	1131	0.17
Health knowledge index	0.229 (0.360) [1.00]	0.115 (0.499) [1.00]	0.186 (0.554) [1.00]	0.176 (0.591) [1.00]	2.89	1131	0.05
Sanitation practices index	0.080 (0.302) [1.00]	0.406 (0.511) [1.00]	0.252 (0.290) [0.91]	0.040 (0.544) [1.00]	-0.68	1131	0.08
Productive assets [†]	0.199* (0.121) [0.47]	0.069 (0.295) [1.00]	0.792*** (0.135) [0.00]	0.148 (0.225) [1.00]	11.22	1131	0.29
Consumption assets [†]	0.336 (0.239) [0.59]	0.442 (0.352) [0.72]	0.831*** (0.300) [0.07]	0.465 (0.475) [0.90]	8.70	1131	0.38
House value [†]	0.023 (0.065) [1.00]	-0.162** (0.068) [0.16]	0.169* (0.088) [0.27]	0.012 (0.111) [1.00]	13.81	1071	0.35
Housing quality index	-0.260 (0.214) [0.72]	0.177 (0.433) [1.00]	0.206 (0.206) [0.90]	0.253 (0.235) [0.88]	-0.17	1131	0.11
<i>B. Individual outcomes</i>							
Child mortality	-0.001 (0.010) [1.00]	-0.002 (0.014) [1.00]	-0.010** (0.004) [0.60]	0.010 (0.009) [1.00]	0.01	1751	0.02
Pregnancy	-0.028 (0.032) [1.00]	-0.031 (0.034) [1.00]	-0.033 (0.029) [1.00]	0.045 (0.052) [1.00]	0.20	1646	0.09
Live birth	0.057 (0.088) [1.00]	0.099 (0.112) [1.00]	-0.087 (0.093) [1.00]	-0.123 (0.148) [1.00]	0.68	273	0.22
Birth in facility	0.077 (0.058) [1.00]	-0.013 (0.073) [1.00]	-0.095 (0.100) [1.00]	-0.033 (0.117) [1.00]	0.84	188	0.25
Any vaccinations in past year	-0.013 (0.039) [1.00]	0.038 (0.059) [1.00]	-0.022 (0.051) [1.00]	0.048 (0.101) [1.00]	0.72	838	0.27
Completed vaccination schedule	-0.006 (0.039) [1.00]	0.013 (0.070) [1.00]	-0.008 (0.050) [1.00]	0.063 (0.103) [1.00]	0.58	838	0.21
Disease burden	-0.016 (0.032) [1.00]	-0.055 (0.058) [1.00]	-0.004 (0.044) [1.00]	0.014 (0.084) [1.00]	0.54	1746	0.06
Diarrheal prevalence	-0.005 (0.018) [1.00]	0.021 (0.026) [1.00]	0.002 (0.019) [1.00]	-0.021 (0.024) [1.00]	0.09	1746	0.05

NOTES—Analysis uses only the Control and GD arms; the first and third columns are dummies for GD transfer amount cells and the second and fourth are dummies measuring the additional effect of lump sum transfers within each cell. Asterisks denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a \dagger are in inverse hyperbolic sines.

TABLE A.9: EFFECT OF TRANSFER MODALITY CHOICE: PRIMARY

	Got Ones Choice	Chose Lump Sum	Treated Lump Sum	Control Mean	Observations	R^2
<i>A. Household outcomes</i>						
Consumption [†]	0.089 (0.115) [1.00]	-0.018 (0.126) [1.00]	-0.032 (0.117) [1.00]	10.69	534	0.18
Household dietary diversity score	0.067 (0.210) [1.00]	0.062 (0.197) [1.00]	-0.204 (0.178) [1.00]	4.77	534	0.23
Household non-land wealth [†]	-0.650 (0.468) [1.00]	0.033 (0.471) [1.00]	0.172 (0.460) [1.00]	13.04	534	0.29
<i>B. Individual outcomes</i>						
Height-for-age	-0.160** (0.062) [0.21]	0.034 (0.058) [1.00]	-0.026 (0.063) [1.00]	-1.97	671	0.75
Weight-for-age	-0.042 (0.052) [1.00]	0.013 (0.045) [1.00]	0.093* (0.049) [0.36]	-1.04	668	0.67
Mid-upper arm circumference	0.070 (0.071) [1.00]	0.025 (0.065) [1.00]	0.153** (0.067) [0.21]	-0.59	520	0.57
Child anemia	0.015 (0.024) [1.00]	-0.017 (0.030) [1.00]	-0.034 (0.030) [1.00]	0.22	750	0.08
Maternal anemia	0.007 (0.034) [1.00]	0.008 (0.043) [1.00]	-0.004 (0.037) [1.00]	0.12	496	0.14

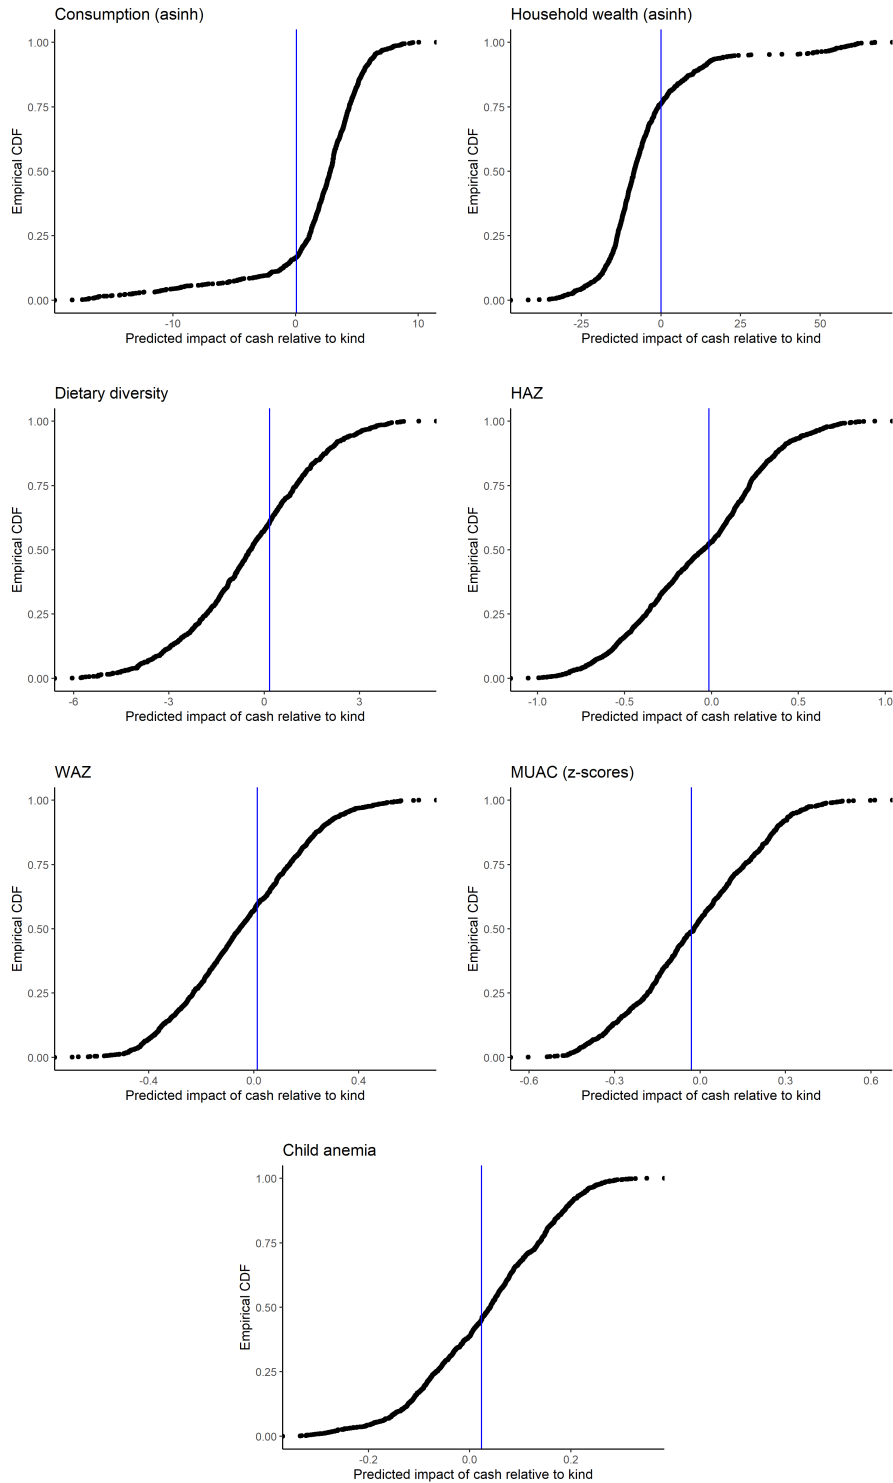
NOTES—Analysis uses only the GD arm. First column is a dummy for getting the chosen transfer modality, the second column is an (endogenous) indicator for choosing lump sum, and the third column is a dummy for actually receiving the lump sum treatment. Asterics denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a [†] are in inverse hyperbolic sines.

TABLE A.10: EFFECT OF TRANSFER MODALITY CHOICE: SECONDARY

	Got Ones Choice	Chose Lump Sum	Treated Lump Sum	Control Mean	Observations	R^2
<i>A. Household outcomes</i>						
Stock of borrowing [†]	-0.558 (0.504) [1.00]	-0.197 (0.425) [1.00]	0.956* (0.494) [0.67]	7.39	534	0.16
Stock of saving [†]	-0.222 (0.463) [1.00]	0.074 (0.491) [1.00]	0.132 (0.450) [1.00]	5.88	534	0.21
Health knowledge index	0.056 (0.534) [1.00]	0.559 (0.479) [1.00]	0.524 (0.454) [1.00]	2.89	534	0.09
Sanitation practices index	0.068 (0.286) [1.00]	0.632* (0.349) [0.67]	0.221 (0.383) [1.00]	-0.68	534	0.15
Productive assets [†]	-0.082 (0.212) [1.00]	-0.166 (0.207) [1.00]	0.001 (0.203) [1.00]	11.22	534	0.29
Consumption assets [†]	-0.189 (0.362) [1.00]	-0.917*** (0.319) [0.06]	0.952*** (0.301) [0.05]	8.70	534	0.37
House value [†]	-0.065 (0.085) [1.00]	-0.029 (0.090) [1.00]	-0.094 (0.070) [1.00]	13.81	508	0.40
Housing quality index	-0.048 (0.175) [1.00]	-0.017 (0.119) [1.00]	-0.114 (0.123) [1.00]	-0.17	534	0.21
<i>B. Individual outcomes</i>						
Child mortality	0.001 (0.010) [1.00]	-0.011 (0.014) [1.00]	0.000 (0.015) [1.00]	0.01	838	0.03
Pregnancy	0.008 (0.035) [1.00]	0.027 (0.032) [1.00]	-0.025 (0.030) [1.00]	0.20	757	0.11
Live birth	-0.064 (0.151) [1.00]	0.286* (0.162) [1.00]	-0.007 (0.120) [1.00]	0.68	129	0.32
Birth in facility	0.137 (0.127) [1.00]	0.077 (0.121) [1.00]	-0.153 (0.115) [1.00]	0.84	83	0.57
Any vaccinations in past year	0.012 (0.050) [1.00]	0.102* (0.062) [1.00]	0.016 (0.058) [1.00]	0.72	434	0.25
Completed vaccination schedule	-0.004 (0.054) [1.00]	0.060 (0.065) [1.00]	0.002 (0.064) [1.00]	0.58	434	0.19
Disease burden	0.017 (0.051) [1.00]	0.050 (0.050) [1.00]	-0.026 (0.049) [1.00]	0.54	835	0.07
Diarrheal prevalence	-0.001 (0.024) [1.00]	-0.030 (0.026) [1.00]	0.017 (0.025) [1.00]	0.09	835	0.05

NOTES—Analysis uses only the GD arm. First column is a dummy for getting the chosen transfer modality, the second column is an (endogenous) indicator for choosing lump sum, and the third column is a dummy for actually receiving the lump sum treatment. Asterices denote significance at the 10, 5, and 1 percent levels, and are based on clustered standard errors, in parentheses. Anderson (2008) sharpened q -values presented in brackets. Variables marked with a [†] are in inverse hyperbolic sines.

FIGURE A.1: HETEROGENEITY IN CASH IMPACTS RELATIVE TO IN-KIND PROGRAMMING



Notes—Panels display CDFs of subgroup average impacts of receiving GD Main transfer as opposed to Gikuriro, estimated by causal forest. Estimates presented for all primary household and child-level outcomes.