GiveWell's Moral Weights Underweight the Value of Transfers to the Poor

By Ethan Ligon and Trevor Woolley

University of California, Berkeley

October 31, 2022

1 Summary

GiveWell bases much of their cost-effectiveness analysis on the value of doubling consumption. Since increasing consumption expenditures is the primary effect of the GiveDirectly cash transfers program, GiveWell uses the effectiveness (value generated per dollar) of cash transfers as a metric for evaluating the effectiveness of all other programs. However, by valuing "doubling consumption", GiveWell has assumed the functional form of utility over "real consumption" x to be *log(x)* and the functional form of marginal utility over consumption to be $1/x$ (since this is the derivative of *log(x)).* This is a valid utility function in the sense that it is one of many functions that satisfies the conditions of rationality, but there is strong evidence that it is *not* a good representation of the preferences of the Kenyan beneficiaries of the GiveDirectly experiment.

The purpose of this article is to explain why GiveWell should reconsider using "doubling consumption" as the basis for assessing the value of consumption (or income) changes and instead value "halving marginal utility of expenditure"—what we think GiveWell actually *intends* to value.[1] Using data from GiveDirectly's cash transfers program in Kenya (Haushofer and Shapiro 2016), we provide empirical evidence that rejects the use of any function that implies homothetic preferences (including marginal utility of $1/x$). We then empirically estimate the true marginal utility over consumption (λ) as revealed by Kenyan beneficiaries of GiveDirectly's cash transfers program and show how **the value per dollar of cash transfers is actually 2.6 times GiveWell's current number** (from 0.0034 to 0.009). This is because $1/x$ is quickly dwarfed by revealed marginal utility, λ, at low levels of

[1] Higher marginal utility of expenditure implies greater "need" for consumption. "Doubling consumption" is a special case of "halving marginal utility of expenditure" (or "halving need") when utility is logarithmic.

consumption. Therefore, valuing "doubling consumption" underweights the value of cash transfers to the very poor if we let them "speak" for themselves.[2]

The remainder of this article is as follows: In Section 2 "The Argument", we explain why "doubling consumption" should not be used to measure the value produced by consumption (or income) changes; In Section 3 "The Alternative", we give a simplified overview of how to estimate the actual value generated by consumption (or income) changes; In Section 4 "The GiveDirectly Experiment", we use data from the GiveDirectly program in Kenya (Haushofer and Shapiro 2016) to demonstrate how homotheticity does not hold in Kenya and estimate the revealed marginal utilities of consumption for this population; And in Section 5 "Recommendation & Conclusion", we discuss the implications of our results for GiveWell's cost-effectiveness analysis.

2 The Arguments

In this article, we make two arguments. The first we will call the "Deconstructive Argument" because it describes why GiveWell should not use "doubling consumption" to represent the value of GiveDirectly's cash transfers. The second we will call the "Constructive Argument" as it describes the logic behind the measure that we recommend GiveWell use instead. This second argument will be outlined in the next section, Section 3 "The Alternative".

The Deconstructive Argument is as follows:

1. **Premise 1:** Utility functions that rely only on "real consumption" measures imply homothetic preferences. *Log(x)* is one such function.
2. **Premise 2:** Recipients of cash transfers do not exhibit homothetic preferences
3. **Conclusion:** Any homothetic function will not accurately describe the value of cash transfers to recipients, especially for those in extreme poverty.

2.1 Premise 1: "Real consumption" ⇒ Homotheticity

2.1.1 **Real Consumption.** GiveWell uses what they call "consumption" as an important input to its calculation of moral weights. Examination of programs that actually estimate changes in consumption seem to indicate that GiveWell's usage corresponds to what a careful economist might call "total real non-durable household expenditures." Conceptually this measure matches up reasonably well with the construction of some other welfare measures; e.g., the World Bank's preferred method for constructing measures of poverty also relies on real non-durable expenditures.

So, how does one go about constructing a measure of "real consumption" for households in low-income countries? The process begins by conducting a survey of households which elicits reports on recent expenditures on a wide variety of different items. To take a typical example, the Kenya Life Panel Survey used to support the analysis GiveWell relies on for valuing deworming projects asks about recent purchases of nearly 100 different sorts of expenditures, about half of which are different sorts of food (Miguel, Kremer, and Hamory 2021). These are added up, delivering a measure of total non-durable expenditures.

[2] It is important to note that our results do not rely on any type of inequality aversion or "prioritarian" preferences of the planner but are purely utilitarian in the sense that they reflect the wellbeing of individual consumers regardless of their relative position to others.

But note that these are in nominal terms, in whatever the local currency is. To put these into "real" terms one divides by a price index. In the case of the KLPS this index converts the value of Kenyan shillings in different years into 2017 US dollars (combining a PPP calculation which expresses the value of KSH into USD with changes in a consumer price index over time; see this KPLS-4 Report for GiveWell). Other measures of real consumption generally adopt very similar strategies, including the calculations delivered by the World Bank for measuring poverty (Development Research Group of the World Bank) or by other projects funded by GiveWell such as the GiveDirectly cash transfers in Kenya documented by Haushofer and Shapiro 2016, who collect data on 40 expenditure items at different points in time and construct an index to convert these to constant 2012 USD.

2.1.2 **Price Indices**. Whether one uses something like the US consumer price index (CPI) to compare dollar values over time, a "purchasing power parity" (PPP) index to compare values in different currencies, or some combination of the two, the goal is to produce something we might call "real consumption" to serve as the basis of welfare calculations. The use of such indices is very standard within economics. But though widely used it's also an open secret that any such construction of "real consumption" by using a price index has real problems if we use it to evaluate consumer welfare. Though the secret is open and has long been known (google "price index problems"), the very ubiquity of such indices may lull many practical economists into a sense that the problem is not of very great importance.

The whole point of constructing measures of total real consumption is to characterize consumers' budgets. Call the sum of total non-durable expenditures the budget x, denominated in the local currency, in nominal terms. Let p be a vector of all of the prices consumers face, and let $\pi(p)$ be an index constructed from these prices.

2.1.3 **Indirect Utility**. We adopt the standard view that consumers choose to allocate their budget so as to maximize a utility function U which maps a vector of consumption goods into the real line. Thus, people value a larger budget not for its own sake, but because it gives them more options when they choose what to consume.

Of course, the value of one's budget (say x) depends not only on the size of the budget, but also on the prices the consumer faces; call these prices a vector p. The utility a consumer facing these prices can achieve from a budget x is given by the so-called "indirect" utility function

$$V(x,p) \;=\; max_{c}\, U(c) \qquad such\ that\ \ p'c \leq x.$$

Note that this indirect utility depends on not only the budget but also on the vector of all nominal prices p. If welfare is assumed to depend only on the *real* value of the budget, $x\,\pi(p)$, then this is an assertion that the indirect utility function can be expressed in the form $V(x,p) \;=\; V(x\,\pi(p))$.

But this assertion has profound consequences for the structure of consumer demand and preferences. To see this, we exploit a connection between the indirect utility function and consumer demand: Roy's Identity[3],

$$-\frac{\partial V/\partial p_i}{\partial V/\partial x}(x,p) \;=\; c_i(x,p),$$

[3] Note that Roy's Identity holds for all indirect utility functions; not just in the case of homothetic preferences

where $c_i(x, p)$ is the consumer's demand for the good indexed by *i.* In the case in which only "real consumption" matters for welfare, this implies

$$-\frac{\partial V/\partial p_i}{\partial V/\partial x}(x, p) \;=\; x\,\pi_i(p) \;=\; c_i(x, p),$$

Where $\pi_i(p)$ is the partial derivative of the price index with respect to the price of the *i*th good. The key take-away: if real consumption is adequate for measuring welfare, then demand for any good will be proportional to the budget x.

2.1.4 **Homotheticity**. This property of "expenditure proportionality" is equivalent to the utility function having the property of "homotheticity". The moral weights used by GiveWell rely on taking the logarithm of a measure of real consumption, which is an example of a homothetic utility function.[4]

2.2 Premise 2: Real-world utility functions are not homothetic

2.2.1. **Engel Curves.** When the study of economics was young, observations of human behavior yielded statements of capitalized "Laws". We were given the "Law of Malthus", the "Law of Demand", and "Engel's Law", among many others. For the most part these "Laws" have not had a particularly good run. Of the three mentioned the Law of Malthus has abjectly failed to describe the trajectory of the human population[5] and the Law of Demand has important caveats[6]; only Engel's Law has held up well in the face of modern data collection and developments in theory. But Engel's law is perhaps the least-well known of these three. A rough statement of the law is that "as the household's material well-being improves, the share of food in total expenditures falls" (Stigler 1954).

2.2.2. **Engel's Law Implies that "Real" Consumption Measures Welfare for the Wealthy**. Since food shares do empirically fall with total expenditures (see Section 4.2), then Engel's Law violates expenditure proportionality (homotheticity) and is in *direct* conflict with the use of "real consumption" as measure of household well-being. Compared to wealthier households, poorer households spend proportionally more on food than they do on less necessary goods like clothing or entertainment, so an increase in relative food prices will affect poorer households more. Relatedly, when food prices are high an income transfer to poor households will have a greater effect on their well-being than it will when food prices are low.

The implication of using a single price index to construct "real consumption" is not only that household welfare is mis-measured, but that it is mis-measured in a way that favors the wealthy. In practice, indices are typically constructed to weight the price changes of different goods in a way that disproportionately reflects the consumption portfolios of

[4] A more mathematical definition of homothetic utility function is one that is a monotonic transformation of a homogeneous function. A homogeneous function *f* is one where *f(a*x) = a*f(x).* Since *f(x) = x* is itself homogenous, any monotonic transformation of *x* is homothetic. *Log(x)* is therefore homothetic.

[5] Human rates of fertility have reliably fallen with improved economic circumstances.

[6] Demand curves slope downward, except sometimes for "inferior" goods (Jensen and Miller, 2008)

wealthier households (see Section 4.2). For this reason, indices such as those used to construct real consumption in various GiveWell programs, or the US Consumer Price Index (CPI) are sometimes called examples of "plutocratic" indices.

3 The Alternative

Our Constructive Argument is as follows:

1. **Premise 1:** A utilitarian social planner (GiveWell) is interested in allocating resources where the marginal benefit of allocation is highest.
2. **Premise 2:** People spend their money as if maximizing a utility function that maps consumption bundles to the real number line. Their financial need can be measured by the increase in their utility which results from an increase in their budget. This rate of increase is their marginal utility of expenditures, "λ", which is revealed by household consumption behavior.[7]
3. **Conclusion:** A utilitarian social planner (GiveWell) should weight household consumption changes by the household's λ (i.e., they should value halving λ).

3.1 Premise 1: Maximize U ⇒ Minimize MU

At present GiveWell assigns a "moral weight" to changes in the log of total "real expenditures." This is expressed as the "value assigned to doubling consumption for one person for one year". As argued above, this formulation does not adequately consider the fact that differently positioned households are differently affected by changes in prices.

Assigning a value to "doubling consumption" is equivalent to defining a utilitarian social welfare function over the allocation of goods in which the utility of consumption is logarithmic. With logarithmic utility, an equivalent formulation is "the value assigned to *halving* the marginal utility of expenditures". In the log case, the marginal utility for a household with "real" consumption x is precisely $1/x$.

As argued above, the formulation with real consumption is only valid with homothetic utility functions (log is a special case). But the goal of "halving the marginal utility of expenditures" is much more general and can allow for all sorts of complicated dependence on prices. It seems to us that this is indeed what GiveWell is aiming to capture, though it fails to accurately do so by using $log(x)$.

3.2 Premise 2: Estimate revealed marginal utility

3.2.1 **Consumer Theory.** Returning to our formulation of the consumer's problem, suppose that within a period the household has budget x and faces a vector of prices p, solving

$$V(x,p) = \max_c U(c) \quad such\ that\ p'c \leq x.$$

This can equivalently be expressed in Karush-Kuhn-Tucker form as

[7] Caveat: For any λ consistent with demand behavior there's actually a class of others that are also consistent. But there's a one-to-one mapping between these, and we've chosen the simplest member of the class that's analogous to GiveWell's current use of $log(x)$ as a welfare measure.

$$V(x,p) = max_c U(c) + \lambda(x - p'c),$$

where λ is the multiplier associated with the budget constraint. Then in this general case (from the Envelope theorem) the marginal utility of an additional dollar (or Kenyan Shilling) is given by

$$\frac{\partial V}{\partial x}(x,p) = \lambda(x,p)$$

Note that in this general case λ can be regarded as an unrestricted function of budget x and prices p. How should we interpret $\lambda(x,p)$? It is exactly the rate at which utility will increase as the household's budget increases; that is, the marginal utility of an extra dollar.

3.2.2 **Measuring λ.** We've argued that λ is the utilitarian return on investing a dollar in the budget of the household. Utility functions are famously not observable; one cannot directly infer them from observed behavior. However, *marginal* utilities can be inferred from behavior (and this happens to be what GiveWell needs for its cost-effectiveness calculations). This follows from the first order conditions from the consumer's problem (or another "Law" that's held up well, Gossen's Second Law),

$$\frac{\partial U(c)}{\partial c_i} = p_i\lambda.$$

An equation such as this holds for *every* good i, with λ common across each equation. Ligon (2019)[8] shows how (using data on consumption expenditures) one can exploit these equations to infer the value of λ for each household without the need to assume a potentially misleading functional form for total utility like *log(x).*

Because $\lambda(x,p)$ is marginal utility over consumption expenditures (and is monotonically decreasing), a decrease in λ for a given household represents a shift to higher total utility. Because GiveWell values halving marginal utility over consumption, we transform λ by *−log()* so that a one unit increase in *−log λ(x, p)* represents a halving of λ.

4 The GiveDirectly Experiment

The GiveDirectly unconditional cash transfer program forms the foundation for GiveWell's cost-effectiveness analysis. The experiment reported an average 18% increase in total consumption expenditures for households which received these cash transfers. But for reasons given above, this 18% may not be the best welfare measure when preferences aren't homothetic. So, let's look at the data.

4.1. **Aggregate Expenditure Shares versus Mean Expenditure Shares**. The experiment in question asked about 40 different categories of expenditure. Were preferences homothetic, we would observe that households across the distribution of total expenditures would have similar expenditure shares. A simple way to test this is to see whether everyone has the same

[8] The code for handling this estimation is open source, and available via pypi at https://pypi.org/project/CFEDemands/

expenditure shares across income levels. If this is true, then the average expenditure share (across households) should be equal to the *aggregate* expenditure share (that is, the share of, say, expenditures on cereals for the entire population divided by total expenditures for the entire population). It is these "aggregate" shares that are typically employed in the construction of CPI or PPP indices.

To show this, Figure 1 below computes the statistic $\rho = log(aggregate\ share) - log(mean\ share)$. If preferences are homothetic then ρ will be equal to zero. Figure 1 offers strong evidence against the hypothesis of homothetic preferences. So just based on this evidence we can see that, say, an increase in the relative price of vegetables will have larger welfare consequences than an increase in the price of meat. Note also that goods with smaller values of ρ in Figure 1 tend to be more important in the consumption bundles of poorer households (as shown next in Figure 2).

4.2. **Relative Income Elasticities**. Figure 1 shows values of ρ at both baseline and endline points of data collection[9]. It's worth noting that while there's some stability in the value of ρ for some goods (e.g., vegetables, clothing), for others there are dramatic changes (e.g., fuel, sugar). This is another consequence of non-homothetic preferences—when relative prices change, shares will also change.

It would be nice to have another way to characterize preferences, and, more specifically, a way to estimate the rate at which demand will change with income (or total expenditures). A demand system described by Ligon (2019) provides such an estimate, with the caveat that one can only obtain *relative* income elasticities.

So, for example, from Figure 2 we see that the relative income elasticity (β) of root vegetables ("Roots") is about 0.3, similar to the elasticity for pulses. So as total expenditures increase, we'd expect expenditures on these two goods to increase at about the same rate. But the relative income elasticity for airtime is about twice as great, meaning that that airtime is twice as sensitive to income increases as root vegetables.

Figure 2 is a convenient way to depict "Engel curves" for multiple goods at once to see how demand for various goods change as household incomes increase. Normally, homotheticity requires that (absolute) income elasticities for every good be equal to one. However, since these elasticities are relative, homotheticity would simply require that all these elasticities be equal. That they are very much *not* equal implies that Kenyans' utility functions cannot be homothetic.

4.3. **Marginal Utility of Expenditures**. Figure 3 below presents the distribution of –log λ, our preferred welfare measure (if bigger is better). No account is taken of the experimental transfers in this figure; we simply display the distribution at baseline and endline. Note that even with a large proportion of these households receiving a transfer, welfare is actually lower on average in the later period.

[9] Meaning, before and after treatment for the entire population; not separated between treated and control groups. This is not a "treatment effect" in the sense that one is a control for the other.

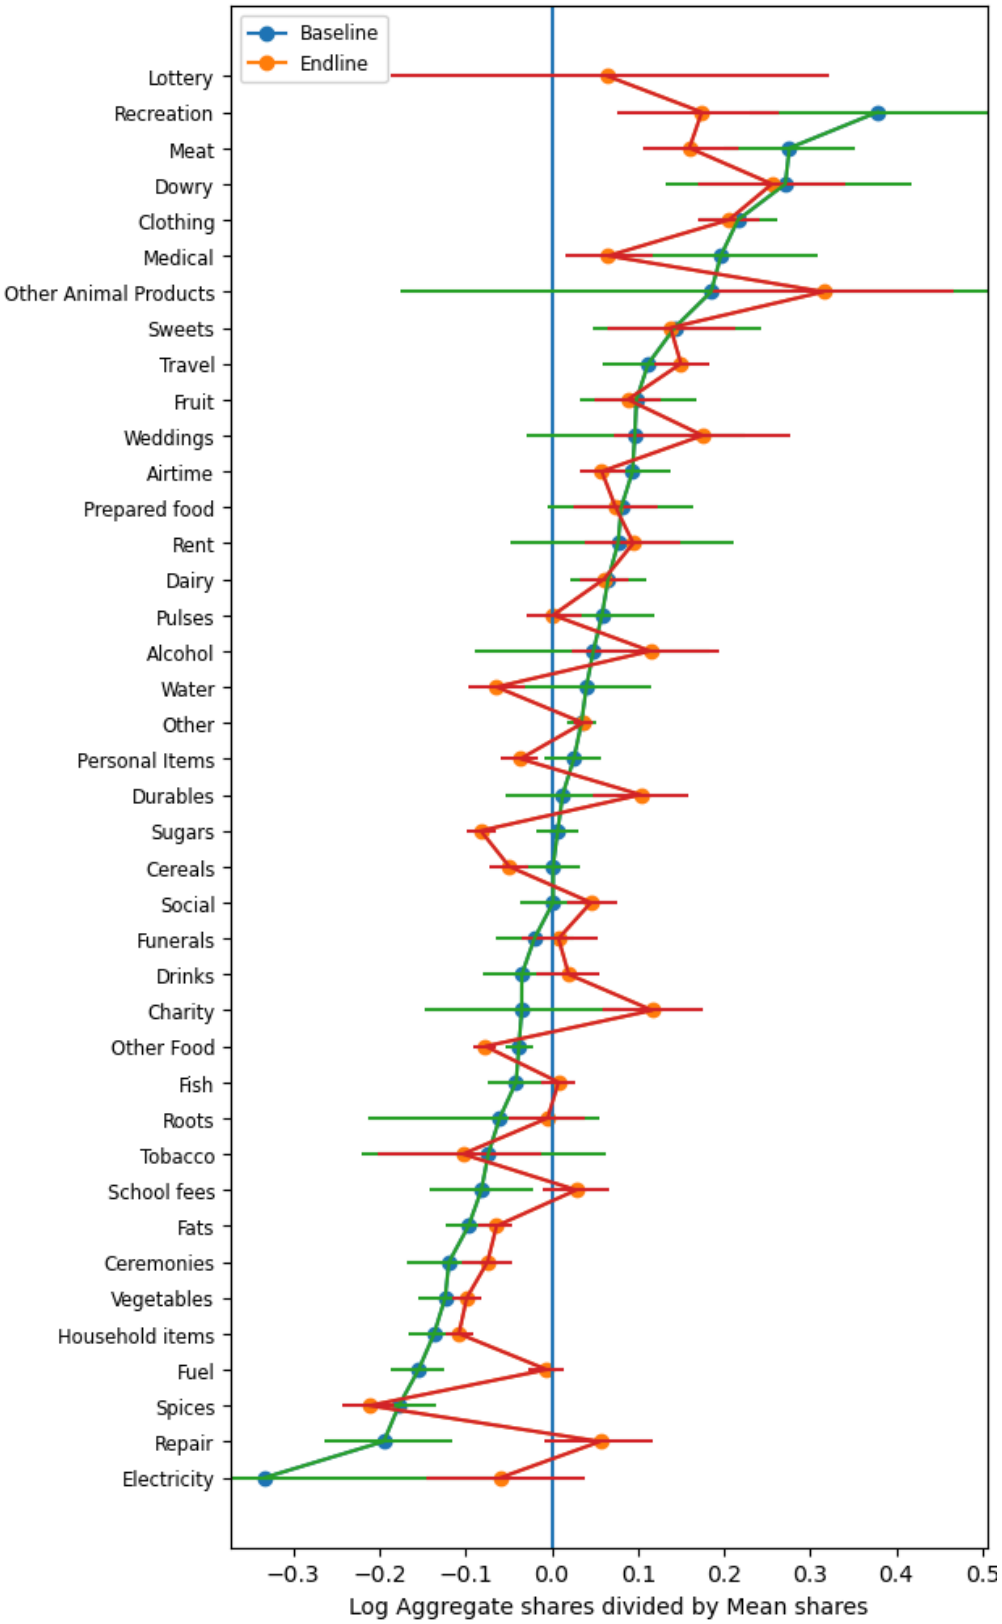

FIGURE 1: $\rho = log(aggregate\ share) - log(mean\ share)$ *for different goods Bars are 95% confidence intervals. Points farther to the right feature more prominently in the consumption bundles of wealthier households.*

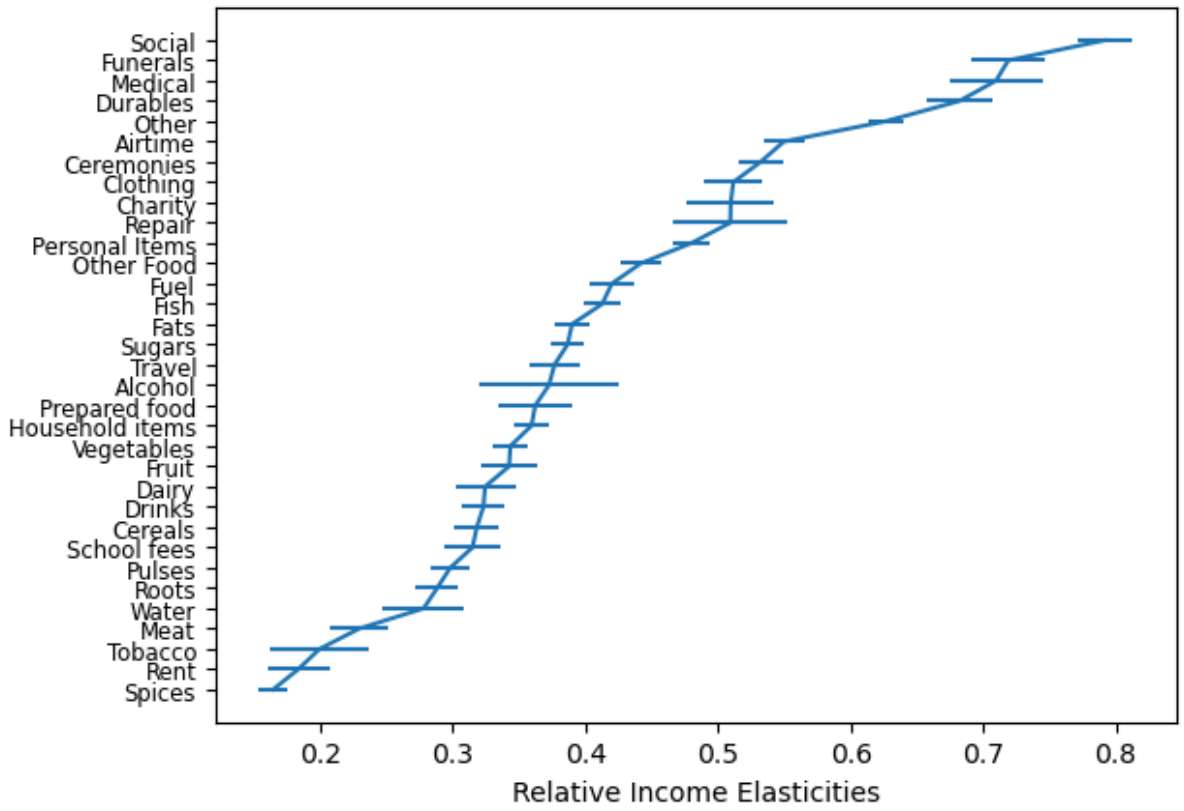

FIGURE 2: Estimates of β for each good, with 95% confidence intervals.

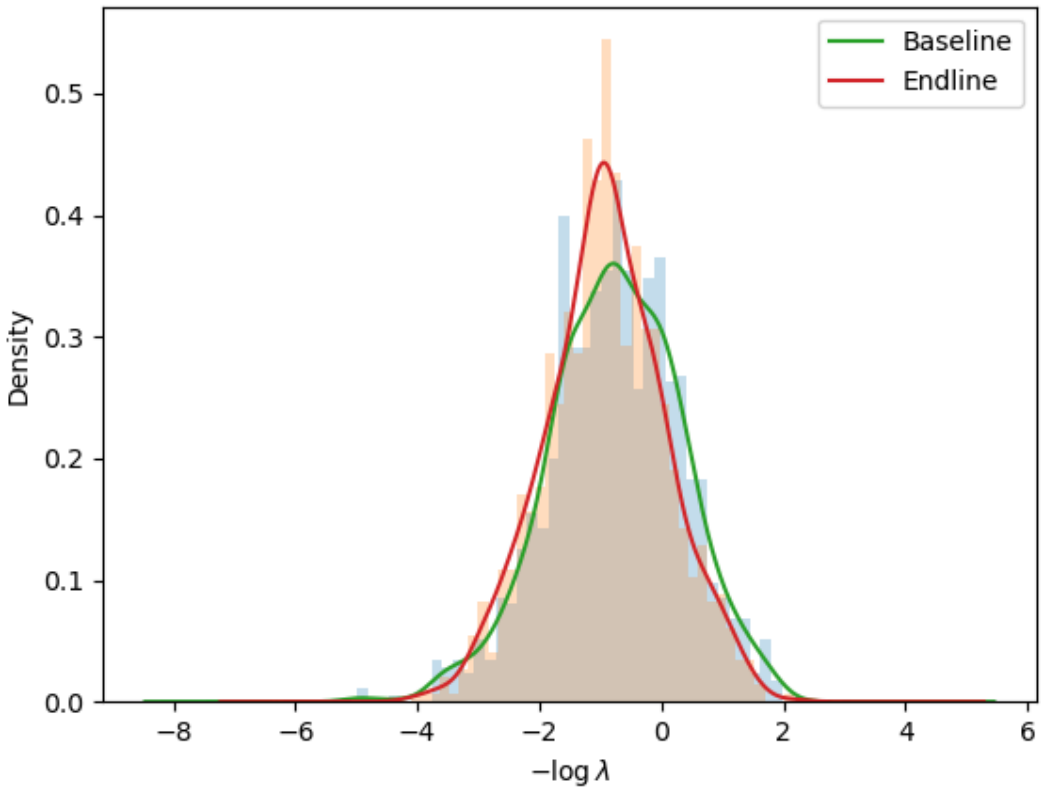

FIGURE 3: Distribution of Welfare, Baseline and Endline. Note that this is the distribution of minus log λ, so that higher levels of welfare are on the right.

4.4. **Effects of Treatment**. Now recall that the GiveWell experiment involved giving randomly selected households an unconditional gift of cash. How did these transfers affect welfare? Figure 4 tells the basic story of the effects of receiving cash transfers. The treatment distribution (in red) is unambiguously shifted to the right of the control distribution. This shift now represents changes in welfare for a given household that was given cash. The average treatment effect is about 47% with a standard error of 4%. This compares with the headline result of a roughly 18% increase in total consumption reported by Haushofer and Shapiro (2016).

Why these differences? Both sets of results are based on the same experiment and the same data. Haushofer-Shapiro's outcome of log total consumption is a close conceptual match for the way in which moral weights are assigned to consumption increases by GiveWell (and log utility is a special case of the demand system we estimate, corresponding to the case with all values of $\beta = 1$). The fact that we obtain different values of β is another way the data tells us that the households in this setting don't have logarithmic utility.. As we've argued above, assuming log utility is likely to under-weight the benefits poorer households receive from such transfers.

Our present finding is consistent with this, and suggests that the error in the construction of this moral weight has had the effect in this experiment of under-valuing cash transfers by a factor of over two and a half.

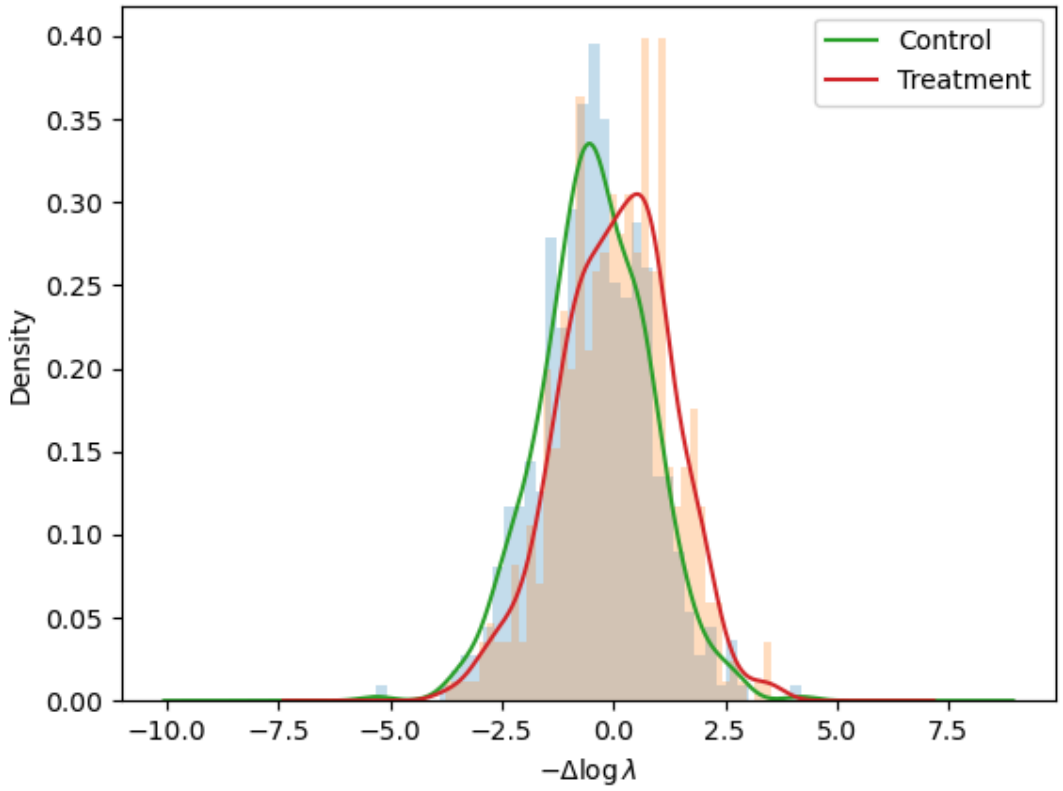

FIGURE 4: Distribution of Welfare, Control v. Treated. Note that this is the distribution of minus log λ, so that higher levels of welfare are on the right.

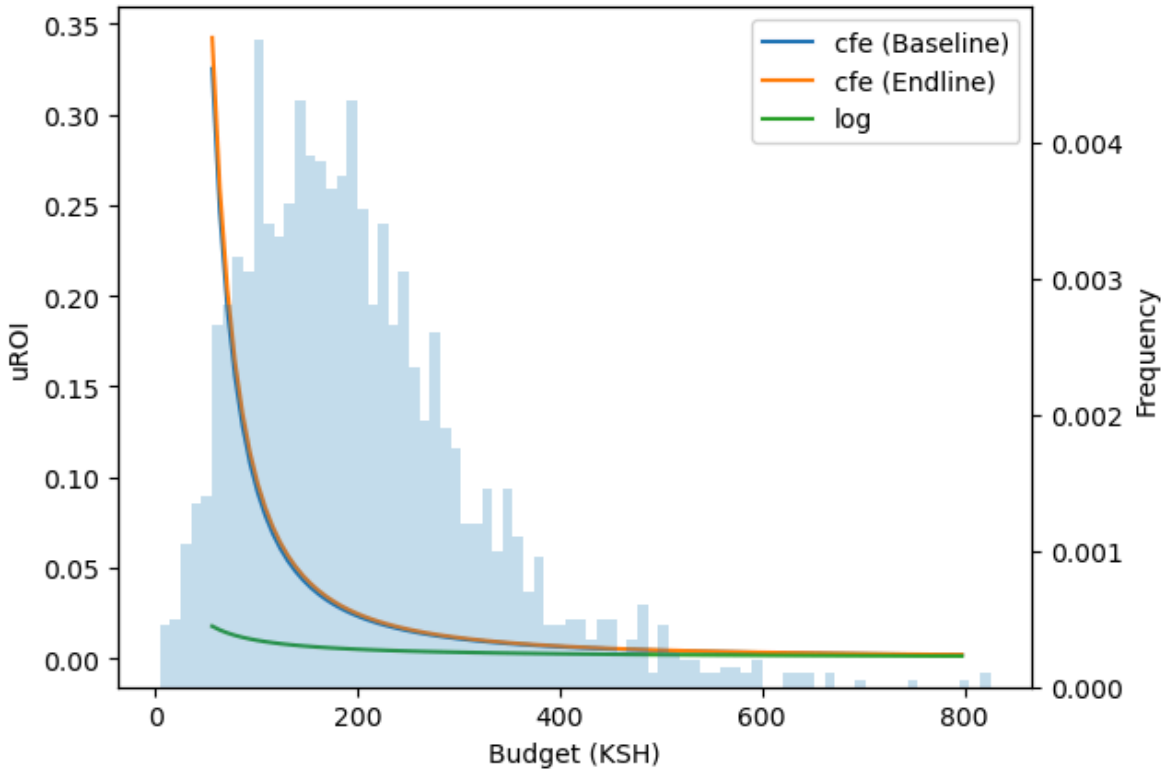

FIGURE 5: Utilitarian ROI. The curves labeled cfe use the estimated marginal utilities of expenditure λ(x, p), with p either the values at baseline, or endline. The curve labeled log uses the marginal utility of expenditures implied by a logarithmic utility function used (implicitly) by GiveWell.

4.5. **Utilitarian Return on Investment**. For every dollar given to a particular household, there's some increase in utility, which we can think of as a "utilitarian ROI". This increase depends on household characteristics (e.g., size, composition), on the household's budget (other things equal poorer households will benefit more), and on prevailing prices (e.g., if food prices are relatively high, the uROI will be relatively higher for poorer households).

For the households in the GiveDirectly experiment we trace out the uROI as a function of household budget x, using prices observed at the baseline, and "average" household characteristics. The GiveDirectly experiment distributed transfers randomly across the distribution of households pictured in Figure 5. The figure also plots marginal utilities of expenditures. The green line is the marginal utility of expenditures $\lambda(x, p)$ corresponding to log utility. At higher budgets, this is very similar to the estimated MUEs, but the large divergence at lower budgets makes it clear where the greater estimated impact comes from: poorer households benefit much more from an additional dollar than do wealthier households. This is true for log utility as well, but to a *far* lesser extent.

5 Recommendation & Conclusion

5.1 **Recommendation.** For our particular empirical example involving the GiveDirectly experiment, the logic given above suggests that GiveWell should update the value of changes in households' budgets. Specifically, wherever there's presently a *log(x)* (where x is total budget for consumption expenditures within the period) this should be replaced with the more

general $-log\ \lambda(x, p)$.[10] Note that the current use of $log(x)$ is a special case of this, since with log utility, $-log\ \lambda(x, p) = log(x)$.

In general, this calls for making predictions about future prices. In practice, we feel comfortable assuming the baseline prices for this experiment (note that Figure 5 above indicates that price changes between baseline and endline had very little effect). There are three particular cells in the GiveWell Cost Effectiveness Spreadsheet for GiveDirectly that are affected by the arguments advanced above:

Cell	Description	$\log(x)$	$-\log\lambda(x,p)$
B20	Total immediate increase in $\log(x)$ per person made possible by funds transferred	0.37	0.92
B21	Future annual increase in ln(consumption) per person — from returns on investments	0.03	0.08
B26	Present value of ln(consumption) increase in the last year of expected benefits	0.06	0.16
B31	Total units of increase in ln(consumption) per household	2.9	7.5
B38	Units of value generated per dollar spent	0.0034	0.0090

Table 1: Utilitarian ROI. Key cells from the GiveWell Cost Effectiveness Spreadsheet for GiveDirectly. Cells B20, B21, and B26 are all directly affected by the arguments in this note. B31 is a summary output.

Table 1 shows the changes that our methods would produce—an increase in the overall value of consumption increases by a factor of 7.5/2.9 = 2.6. A similar factor applies to the other cells. The impact of our recommendation on cell B38 "Units of value generated per dollar spent" is particularly significant given its role in calibrating the efficacy of other programs.

5.2 **Implications.** While we could have done this analysis for any program for which consumption or income changes are important outcomes, we chose to do it for the cash transfers program due to the centrality of its role in calibrating the entirety of GiveWell's cost-effectiveness analysis. If GiveWell cares about knowing the true value of its top charities to their beneficiaries, this exercise (Ligon 2019) should be carried out for each program that involves consumption or income effects. From the "2022 GiveWell cost-effectiveness analysis" spreadsheet, this includes AMF, Deworm the World, END Fund, SCI Foundation, Sightsavers, and Malaria Consortium.

Evaluating these other programs in a similar way is beyond the scope of this article. However, it is not unlikely that doing so would yield increases in the value of their consumption/income components that are similar in magnitude to what we estimated in Kenya (2.6 times) if their beneficiaries are similarly impoverished. As such, there are two reasons this could quite plausibly cause GiveWell to reorder its top charities. The first is that

[10] While λ is a return denominated in utils, $-log\ \lambda$ is a dimensionless rate of return.

those programs whose value relies more on income/consumption effects will now be relatively more effective (value per dollar). The second reason is that the distribution of expenditures across beneficiary populations might be different across programs. Thus, those programs that give transfers to poorer recipients will now be considered relatively even more effective.

Since GiveWell uses cash transfers as a metric by which to compare the effectiveness of other charities (especially to define "top" charities), it is also important to consider how our results would affect this relationship. Without evaluating other consumption-affecting programs as we have done for cash transfers, we cannot yet make definitive statements about how other programs that involve income/consumption effects now fare relative to cash transfers. However, it is at least safe to say that those programs which do *not* include consumption effects for the extremely poor would now be 2.6 times less effective relative to cash transfers in terms of value generated per dollar. For instance, such a program that was previously 10 times as effective as cash transfers would now only be 3.8 times as effective.

5.3 **Conclusion**. It is clear that GiveWell already cares about avoiding paternalistic evaluations of effectiveness. They have taken steps in the recent past to support research for the purpose of updating their efficiency calculations.[11] We applaud this effort and recommend that GiveWell continue this pursuit by augmenting their moral weights to reflect revealed marginal utility rather than the ad hoc functional form of $log(x)$. Doing so would not only provide an empirical foundation for GiveWell's valuation of consumption and income effects but would also incorporate the true preferences of the extremely poor. As a result, the GiveDirectly cash transfers program should be valued 2.6 times as much per dollar as GiveWell's current number (from 0.0034 to 0.009). Theoretically motivated and empirically estimated, λ is what GiveWell should use if their objective is to maximize the true value of their programs to beneficiaries.

[11] Of particular relevance is the report by IDinsight (2019) that helped GiveWell identify the value that the average poor Kenyan and Ghanaian places on doubling consumption relative to mortality. Our results do not contradict this survey but do affect its interpretation.

References

Development Research Group of the World Bank (Downloaded 2017). PovcalNet: The On-line tool for poverty measurement. url: http://iresearch.worldbank.org/PovcalNet/.

Haushofer, Johannes and Jeremy Shapiro (2016). The short-term impact of unconditional cash transfers to the poor: experimental evidence from Kenya . In: The Quarterly Journal of Economics 131.4, pp. 1973 2042.

IDinsight (2019) Beneficiary Preferences: Findings from Ghana and Kenya. url: https://www.idinsight.org/publication/measuring-peoples-preferences/.

Jensen, Robert T and Nolan H Miller (2008). Giffen behavior and subsistence consumption . In: American economic review 98.4, pp. 1553 77.

Ligon, Ethan (2019). Estimating Household Welfare from Disaggregate Expenditures . Unpublished manuscript. url: https://escholarship.org/uc/item/3ts0g5tn.

Miguel, Edward, Michael Kremer, and Joan Hamory (2021). Kenya Life Panel Survey, Round 3 (KLPS-3). Version V2. doi: 10.7910/ DVN/PXVFJD. url: https://doi.org/10.7910/DVN/PXVFJD.

Stigler, George J (1954). The Early History of Empirical Studies of Consumer Behavior . In: Journal of Political Economy 62.2, pp. 95 113.