What Would it Cost to End Extreme Poverty?

Roshni Sahoo* Joshua Blumenstock† Paul Niehaus‡ Leo Selker†
Stefan Wager§

September 9, 2025

Abstract

We study poverty minimization via direct transfers, reframing this as a statistical learning problem while retaining the information constraints faced by real-world programs. Using nationally representative household consumption surveys from 17 African countries that together account for 37% of the world's poor, we estimate that reducing the poverty rate to 1% (from an as-of-survey-year baseline of 33%) costs $119B nominal per year. This is 4.6 times the corresponding reduction in the aggregate poverty gap, but only 45% of the cost of universal basic income. Extrapolated out of sample, the results correspond to a cost of (approximately) ending extreme poverty of roughly 0.31% of global GDP.

*Corresponding author (email rsahoo@stanford.edu), Department of Computer Science, Stanford University, † School of Information, UC Berkeley, ‡ Department of Economics, UC San Diego, § Graduate School of Business, Stanford University. We are grateful for helpful comments from seminar participants at the Brown University Bravo/SNSF Workshop, UChicago Machine Learning in Economics Summer Conference, Harvard Data Science Institute, University of Maryland Workshop on AI and Analytics for Social Good, Stanford Econometrics, and Stanford Statistics Seminar. We thank Anik Ashraf, Juan Manuel Monroy Barragan, S. Chandrasekhar, Tanzed Hossain, Abhiroop Mukhopadhyay, Fabrice Mukendi Mutombo, Iary Rakotondradany, David Stifel, and Jonathan Weigel for their generous assistance accessing and interpreting data; Puneet Bhaskar, Qiqi Chen, Nolan Chu, Karl Gorski, Neela Kaushik, Henry Lopez Rodriguez, Matthewos Mesfin, Francesca Vescia, and Victor Zhenyi Wang for help with key data preparation; and Aditi Acharya and Shruthi Ramesh for excellent research assistance. RS is supported by a Stanford Data Science Fellowship and Stanford DARE (Diversifying Academia, Recruiting Excellence) Fellowship. This work was partially supported by NSF grant SES-2242876 and by gifts from Luke Ding and from the Manglaben Chimanbhai Patel Family Foundation.

1 Introduction

The share of the world's population in extreme poverty declined from 43% in 1981 to 36% in 1990 to 10% in 2015, arguably one of humanity's great recent successes.[1] It might thus appear that the end of extreme poverty, the first of the Sustainable Development Goals, is within reach. Recent estimates of the global poverty gap—the cumulative amount by which everybody living below the extreme poverty line is below it—range from just $120 to $180B (nominal 2023 USD) per year.[2] This has led to calls for the creation of new multilateral funding vehicles devoted exclusively to ending extreme poverty by funding direct, targeted transfers to people in poverty (e.g., Kharas and McArthur, 2023). If it were feasible to give everyone exactly the amount they need to reach the extreme poverty line, then transfers sufficient to end extreme poverty would cost just 0.11%–0.17% of global GDP.[3]

However, this is not feasible. Existing estimates of the global poverty gap are based on household surveys that typically cover a tiny fraction, on the order of 0.001% to 0.01%, of the population of the countries in question. Thus, for almost all households in low- and middle-income countries, we do not have a direct estimate of their standard of living. Collecting that information through regular surveys would be prohibitively costly (Kilic et al., 2017) and not incentive compatible, if the amounts households were to receive depended directly on what they reported about hard-to-verify quantities like consumption. Such cost and incentive considerations are among the reasons that existing programs typically allocate transfers on the basis of proxies for living standards that are cheaper to measure and verify (Hanna and Olken, 2018). For example, in the widely-used Proxy-Means Test (PMT) approach, policymakers use survey data on the living standards (typically, consumption) of a representative but small sample of households to train a model that predicts household standard of living based on proxies available in census data (for instance, housing materials); model predictions are then used to determine transfer eligibility for all households (Grosh and Baker, 1995).

Learning to Target Cash Transfers In this paper we maintain this same information environment and ask how to design feasible policies that would end (up to a given tolerance) extreme poverty at the lowest possible cost. We formulate poverty targeting as a statistical learning problem, building on previous work that proposes to learn data-driven decision rules via loss minimization (Athey and Wager, 2021; Bertsimas

[1] According to World Bank (2021) estimates; see also Armentano et al. (2025), using the international poverty line of $2.15 (2017 PPPs) per person per day.

[2] Estimates from https://www.brookings.edu/articles/a-purpose-driven-fund-to-end-extreme-poverty-by-2030, based on Cuaresma et al. (2018), and from Our World in Data (https://ourworldindata.org/grapher/total-shortfall-from-extreme-poverty?tab=chart), converted to nominal 2023 USD.

[3] Authors' calculations using the preceding poverty gap estimates and IMF estimated 2023 nominal GDP of $106T (https://www.imf.org/external/datamapper/NGDPD@WEO/WEOWORLD, accessed 6 September 2025.) Here and throughout we report the cost of transfers; delivering these transfers would require some additional administration costs and service fees.

and Kallus, 2020; Kitagawa and Tetenov, 2018). The policymaker learns a transfer policy which maps a household's observable characteristics into a nonnegative transfer amount, seeking to minimize expected loss for a given budget (or, in the dual formulation, to control expected loss at minimum cost). We consider both unrestricted transfer policies, where each household can receive any nonnegative amount, and binary policies, where each household must receive either zero or some common positive amount. The domain of the loss function is the household's post-transfer standard of living.[4]

The choice of loss function requires some care. Because the headcount poverty rate is easily understood and widely used in public discourse about poverty, we would like to learn policies that efficiently reduce the rate.[5] But directly minimizing the rate may conflict with notions of equity: it is well-known that in the perfect information case controlling the post-transfer poverty rate can lead to inequitable policies (Atkinson, 1987; Bourguignon and Fields, 1990; Sen, 1976), and we find that the same is true under imperfect information. Intuitively, the most budget-efficient way to reduce the rate is to allocate transfers to households just below the poverty line, not to the poorest households; whereas we want the opposite, i.e. policies that are equitable in the sense that they allocate (weakly) larger transfers to poorer households. We show that using algorithms designed to minimize the poverty gap—rather than a headcount metric—enables us to avoid inequity in this sense, all while preserving formal guarantees on poverty rate reduction. Specifically, gap minimization is equivalent to minimizing the worst-case (across population subgroups) conditional poverty rate, and equivalent to rate minimization overall in a regime where rate minimization is itself not inequitable. We therefore provide algorithms for both rate and gap minimization, but emphasize in our discussion the effect of the optimal gap-minimizing policy on the (more familiar) poverty rate metric.

The Cost of Targeted Transfers We quantify the aggregate cost of reducing poverty in a sample of 17 countries, including essentially all of those for which (i) a recent, high-quality, nationally-representative living standards survey is available publicly with a consumption aggregate included, (ii) the national poverty rate exceeds 10%, and (iii) the country contains more than 1% of the world's extreme poor. The first restriction aims to ensure the reliability of our living standards measures; the second to ensure that we observe enough poor households in each survey for effective statistical learning; and the third simply to prioritize the countries that matter most.[6] Each could be relaxed in the future, as we discuss below. This

[4]To be precise, and as we discuss below, this is the standard of living the household *could* attain if it spent the entire transfer on consumption. While we abstract from dynamics here, conceptually we think of households that save or invest a portion of their transfer as revealing that they are at least weakly better-off by doing so. One can thus interpret the exercise as learning policies that raise welfare to at least the level that ending today's consumption poverty would achieve.

[5]For example, Target 1.1 of the Sustainable Development Goals is "By 2030, eradicate extreme poverty for all people everywhere, currently measured as people living on less than $1.25 a day" and the corresponding indicator 1.1.1 is the "Proportion of the population living below the international poverty line by sex, age, employment status and geographic location (urban/rural)" (UN Statistics, 2019).

[6]See Section 3 for further discussion and Appendix D for details of the data sources.

sample is very poor, with a pre-transfer population-weighted average poverty rate as of the (varied) years in which the surveys were conducted of 33% and poverty gap index of 10%, and collectively accounts for 37% of the world's extreme poor.[7]

We estimate that reducing the $2.15 2017 PPP poverty rate to 1% in all countries in this sample as of their survey year, using a gap-minimizing (and hence weakly equitable) policy, would cost $119B nominal per year. This is substantially lower than the cost of blunter, uniform policies: it is 45% of the cost of providing all individuals a basic income at the level of the poverty line and 68% of the cost of providing all individuals basic income amounts, varied by country, that achieve the same rate. This highlights the relative benefit of targeting with observable proxies of consumption compared to not targeting at all. On the other hand, it is 4.6 times the amount by which it reduces the aggregate poverty gap, or equivalently 4.6 times the cost of achieving the same gap reduction in an oracle scenario. This comparison quantifies the relative cost of targeting with proxies versus perfect information on consumption. Reducing the poverty rate implied by the newer $3.00 2021 PPP poverty line to 1%, meanwhile, costs 1.5 times as much as doing so under the older $2.15 2017 PPP line, reflecting the fact that the revision implied a meaningful real shift in the definition of poverty (Foster et al., 2025).

Real-world transfer policies are often simpler than our baseline gap-minimizing one. If we restrict attention to binary policies the cost of achieving poverty rates of 1% increases by 38%, to $164B nominal per year. This implies that, for a given poverty-reduction goal, the fiscal cost of using a simpler "eligibility rule" structure is meaningful. The equity benefits of gap-minimizing policies, on the other hand, turn out to cost little, if anything: policies learned to directly minimize the rate are if anything *more* expensive, for a given reduction in the rate, than policies learned to minimize the gap. This likely reflects the fact that rate minimization is a non-convex and computationally more difficult problem, which in turn limits the number of predictors we can use. Overall, the results suggest that gap-minimizing policies offer an attractive combination of equity and performance with respect to the familiar poverty rate metric.

We end by discussing some implications of, and caveats to, these findings. With some exceptions, the policies we learn would cost more than could plausibly be financed using purely domestic sources of revenue, and would thus require international transfers. International transfers at such scale could in turn have substantial macroeconomic effects, including exchange rate appreciation as well as stimulus to or inflation in the domestic economy. We briefly discuss pertinent recent evidence. Recent work on the general equilibrium effects of cash transfers, for example, has documented substantial expansionary effects (Egger et al., 2022;

[7]The estimate of the share of the world's poor that our sample accounts for is based on national poverty rates and populations for all countries from World Poverty Clock as of 25 August 2025. The estimates of the sample's poverty rate and poverty gap are based on available survey data from these countries (Table E.2). National poverty rates in the surveys are generally consistent with poverty rate estimates from the World Bank in the survey year.

Gerard et al., 2024) which we necessarily abstract from here. We also ask what the results suggest about the cost of (approximately) ending poverty globally, using the in-sample results to predict the potential costs of poverty reduction in others. One caveat is that most in-sample surveys were conducted prior to the COVID-19 pandemic, which increased poverty rates and for which some additional temporal extrapolation might therefore be desirable. That said, our in-sample results are consistent with an overall cost of reducing the global poverty rate to 1% of $341B per year.[8] This is 0.31% of global GDP, and 0.5% of OECD GDP—a bit less than total OECD official development assistance in 2023, prior to recent cuts to aid budgets.

Related Work Our work sits at the intersection of two long-standing traditions of thought about economic development. One has posited ambitious development goals and then sought to calculate the funding required to achieve them. From the 1960s onwards, for example, analysts at development banks used a "two-gap" framework based on the Harrod-Domar constant-returns growth model to calculate the capital investment required to achieve a target rate of economic growth (Domar, 1946; Easterly, 1999; Harrod, 1939). In the early 2000s, as attention shifted to poverty reduction and a broader set of concurrent social objectives, the UN Millennium Project estimated that its proposed Millennium Development Goals could be achieved if rich countries increased foreign aid to 0.54% of their GNI in 2015, from 0.25% in 2003 (Sachs, 2005). For the Millenium Development Goal for extreme poverty specifically, contemporary estimates were constructed by (i) calculating the growth in GDP per capita required to achieve a given level of poverty, assuming that the relative distribution of income stayed fixed, and then (ii) calculating the amount of aid needed to generate that growth using either a "two-gap" model (e.g., Devarajan et al., 2002) or estimated coefficients from cross-country aid-on-growth regressions (e.g., Anderson and Waddington, 2007). As the authors took care to point out, the causal assumptions in these analyses were heroic, and indeed in some cases counterfactual.[9] This was necessarily, in the words of Devarajan et al. (2002), "a highly speculative exercise." Ours, in contrast, is built on the more mechanical idea that giving someone $1 increases their disposable income by $1. It is in this sense less speculative (though it may for that very reason yield conservative estimates, if other strategies that reduce poverty through less direct causal mechanisms are also cheaper).

The second literature concerns the optimal design of transfer programs specifically. Theoretical work on poverty measurement with full knowledge of the income or consumption distribution (e.g., Foster et al., 1984; Greer and Thorbecke, 1986; Sen, 1976) provided one potential point of departure: a logical next

[8] Note that achieving a 1% global rate requires reducing country-specific poverty rates to 1.3%, since some countries have essentially no extremely poor people. Note also that this calculation excludes the cost of ending poverty in the following low- and middle-income countries, as their PPP conversion factors and market exchange rates were not available from the World Bank as of 25 August 2025: Taiwan, Turkmenistan, Tuvalu, Venezuela Rb, West Bank and Gaza, Yemen Rep, and Zimbabwe.

[9] For example, Devarajan et al. (2002) used the "two-gap" approach despite the fact that it was soundly rejected by the data (Easterly, 1999) and had been widely criticized, including by one of the authors himself. Anderson and Waddington (2007) used coefficients from cross-country growth regressions despite the fact that many would have agreed with Mankiw et al. (1995) even a decade earlier that "using these regressions to decide how to foster growth is ... most likely a hopeless task."

step might have been to develop methods to minimize the proposed metrics in more realistic, limited-information environments. Along these lines, Kanbur (1987) and Ravallion and Chao (1989) discussed some theoretical properties of poverty-minimizing transfer policies under specific, stylized information structures,[10] and Glewwe (1992)—in perhaps the closest antecedent to our work—proposed and illustrated learning a transfer policy as a function of a (small) set of individual-level observable covariates X_i via the plug-in principle, i.e. by minimizing a given poverty metric in-sample. In practice, however, most subsequent applied work has focused on the simpler problem of classifying households as either poor or non-poor in order to determine their eligibility for a given, fixed transfer. Potential methods for doing so include proxy-means testing (Alatas et al., 2012; Brown et al., 2018; Grosh and Baker, 1995), community-based targeting (Alatas et al., 2016; Coady et al., 2004), and more recently targeting with non-traditional data (Aiken et al., 2022; Noriega-Campero et al., 2020; Smythe and Blumenstock, 2022). Today the advent of modern, high-dimensional statistical learning techniques and the availability of far greater computational power make it a propitious time to revisit the formulation of Glewwe (1992), as we do here.

The methods we develop add to the literature on data-driven decision making (Athey and Wager, 2021; Bertsimas and Kallus, 2020; Bhattacharya and Dupas, 2012; Kallus and Zhou, 2021; Kitagawa and Tetenov, 2018; Manski, 2004). In particular, the problems we solve are examples of a "prediction-policy problem" in the sense of Kleinberg et al. (2015), meaning that we can impute the counterfactual rewards of different actions and obtain an optimal policy by minimizing a loss function that depends on these counterfactual rewards. It is distinct in this sense from recent work on learning treatment assignment policies where rewards under different actions are learned from experimental data (Bhattacharya and Dupas, 2012; Huang and Xu, 2020; Luedtke and van der Laan, 2016; Sun, 2021; Sun et al., 2021; Sverdrup et al., 2023; Wang et al., 2018), though there are many parallels—a treatment assignment policy in that literature is often a function of observable covariates and subject to constraints related to its functional form, budget, or fairness considerations, and it can be useful to reformulate budget-constrained problems as (fractional) knapsack-style problems (Sun et al., 2021; Sverdrup et al., 2023) as we do here. Björkegren et al. (2022) and Haushofer et al. (2025) are particularly relevant, as they study welfare-maximizing allocations of cash transfers (of a fixed size) in the presence of treatment effect heterogeneity.[11] Future work might combine methods like theirs with those in this paper in order to learn dynamic transfer policies that efficiently minimize poverty over time, taking advantage of the fact that some recipients have predictably higher investment propensities and returns (i.e., are more likely to "graduate") in response to a large initial transfer.

[10] Specifically, they assumed that the population consists of pre-defined, mutually-exclusive groups with known income distributions, with no further predictors are available.

[11] For a more general review of the effects of cash transfers on recipients, see Bastagli et al. (2019) and Crosta et al. (2024).

2 Targeting Transfers via Statistical Learning

Let $\mathcal{X} := \mathbb{R}^d$ denote the space of observable characteristics and let $\mathcal{Y} := \mathbb{R}_+$ denote the space of consumption values. We assume that a country has a population distribution F from which the observable characteristics and consumption of units can be drawn.

Assumption 1 (Sampling Model)**.** A country has a population distribution F from which pairs of observable characteristics and per-capita consumption (X_i, Y_i), where $X_i \in \mathcal{X}$ and $Y_i \in \mathcal{Y}$, can be drawn i.i.d. for units $i = 1, 2, \ldots n$.

Cash transfer policies are common in low- and middle-income countries and, due to the cost and incentive-compatibility considerations noted above, are typically allocated based on imperfect proxies $X_i \neq Y_i$ of living standards. Given an available (per-unit) budget $B \in \mathbb{R}_+$, define a transfer policy as a mapping from the space of observable characteristics $\mathcal{X}$ and possible budgets $\mathbb{R}_+$ to a nonnegative cash-transfer amount, i.e. $t : \mathcal{X} \times \mathbb{R}_+ \to \mathbb{R}_+$. This must satisfy the budget constraint $\mathbb{E}_F\left[t(X; B)\right] \leq B$.

We define a unit's post-transfer standard of living as $Y_i + t(X_i; B)$. In effect this means analyzing the consumption the unit *could* attain if they consumed their entire transfer and if Y_i were unaffected by transfers. This formulation is standard in the literature on optimal transfer policies (Bourguignon and Fields, 1990; Glewwe, 1992; Kanbur, 1987; Ravallion and Chao, 1989) and we view it as conservative in the sense that units who save or invest a portion of their transfer are presumed at least weakly better off by doing so. In practice, reviews by Banerjee et al. (2017) and Crosta et al. (2024) find that transfers tend to increase labor supply and earnings from other sources.[12] Note also that this formulation precludes general equilibrium effects; if transfers stimulate the economy, as Egger et al. (2022) and Gerard et al. (2021) find, then our approach is conservative in this regard.

Given this setting, there are many possible ways to learn transfer policies from data. As a point of reference we first outline a widely-used approach for allocating cash transfers based on a proxy-means test (PMT) (Alatas et al., 2012; Brown et al., 2018; Grosh and Baker, 1995; Hanna and Olken, 2018).

Definition 1 (Proxy-Means Test)**.** In a Proxy-Means Test (PMT), observable characteristics, or "proxies," X_i are used to define the conditional mean consumption $\mu(x) := \mathbb{E}_F\left[Y_i \mid X_i = x\right]$. The conditional mean consumption μ is then thresholded to determine the allocation of transfers to units the population, i.e.

$$t_{\text{PMT}}(x; B) := \begin{cases} \bar{t} & \mu(x) \leq \eta(\bar{t}, B), \\ 0 & \mu(x) > \eta(\bar{t}, B), \end{cases} \tag{1}$$

[12] There is also a distinct question whether the *rule* that determines eligibility for transfers might reduce Y_i by disincentivizing effort. Evidence on this point from low income countries is sparse, though Banerjee et al. (2020) find little evidence that eligibility rules distorted consumption choices in Indonesia.

for some choice of maximum transfer size $\bar{t} \in \mathbb{R}_+$ and eligibility threshold $\eta(\bar{t}, B)$ that is increasing in B and chosen so that $\mathbb{P}_F\left[\mu(X) \leq \eta(\bar{t}, B)\right] = B/\bar{t}$.

The PMT policy (1) is a binary policy because each unit can receive zero or some common amount $\bar{t}$. The maximum transfer size $\bar{t}$ is selected by the policymaker but not necessarily learned from data. If Y_i can be predicted perfectly from observable characteristics X_i, then the PMT policy has the desirable property that it will allocate transfers to the units with lowest consumption Y_i. However, when consumption Y_i cannot be perfectly predicted from X_i, or there is uncertainty in Y_i given X_i, then there is no obvious formal link between minimizing mean-squared error and reducing the post-transfer poverty rate, or any other poverty measures.

Given this, it may be possible to improve on a PMT's poverty-reducing performance by explicitly learning policies via minimization of a loss function, chosen to reflect the ultimate poverty reduction goal.

Definition 2 (Unrestricted Policies via Loss Minimization)**.** For any loss function $L : \mathbb{R} \to \mathbb{R}$, an optimal unrestricted transfer policy is one which minimizes the expected post-transfer loss $\mathbb{E}_F\left[L(Y_i + t(X_i))\right]$ subject to a budget constraint B, i.e. which solves

$$\min_{t(\cdot;B):\mathcal{X}\to\mathbb{R}_+} \left\{\mathbb{E}_F\left[L(Y_i + t(X_i))\right] : \mathbb{E}_F\left[t(X_i; B)\right] \leq B, \quad t(x) \geq 0 \quad \forall x \in \mathcal{X}\right\}. \tag{2}$$

Given data $(X_i, Y_i) \sim F$ i.i.d. for units $i = 1, 2, \ldots n$ from a representative survey, a policymaker can estimate the solution of (2) via empirical minimization of the loss. Note that based on the choice of L, custom algorithms may be required to obtain the transfer policy.

The transfer policies obtained in Definition 2 are unrestricted in that they can allocate transfers of any nonnegative amount. In practice many current programs allocate transfers in only a few different sizes, and sometimes only two: zero, and some fixed positive amount. To understand the impact on costs of restricting the support of transfers, we also study the problem of learning such "binary" policies.

Definition 3 (Binary Policies via Loss Minimization)**.** For any loss function $L : \mathbb{R} \to \mathbb{R}$, an optimal binary transfer policy is one which minimizes the expected loss $\mathbb{E}_F\left[L(Y_i + t(X_i))\right]$ subject to a budget constraint B and takes on only one non-zero value, i.e. which solves

$$\min_{\substack{t:\mathcal{X}\to\mathbb{R}_+,\\ \bar{t}\in[0,c]}} \left\{\mathbb{E}_F\left[L(Y_i + t(X_i))\right] : \mathbb{E}_F\left[t(X_i)\right] \leq B, \quad t(x) \in \{0, \bar{t}\} \quad \forall x \in \mathcal{X}\right\}. \tag{3}$$

Note that this approach requires learning both the optimal transfer size $\bar{t}$ and the eligibility rule t.

Loss minimization as in Definition 3, and Proxy-Means Test as in Definition 1, both yield binary policies;

the key difference is that the loss-minimization approach directly minimizes a post-transfer poverty measure. The cash-transfer policy that solves the loss minimization problem (3) for a particular post-transfer poverty measure yields the greatest possible reduction in that measure, while the PMT policy does not provide formal poverty reduction guarantees.

The appropriate loss function depends on one's definition of poverty. This is a classic topic in economics, with seminal contributions from Sen (1976), Foster et al. (1984), and Greer and Thorbecke (1986), among others. One important idea in this literature was that a measure of poverty should be sensitive to changes in the living standards of people below a given poverty line. Nevertheless, the poverty measure that is most easily understood and widely used in public discourse is one that lacks this feature: the poverty rate.

Example 1 (Poverty Rate)**.** Given a poverty line $c \in \mathbb{R}_+$, the pre-transfer poverty rate is the fraction of units in the population whose per-capita consumption falls below the threshold, i.e. $\mathbb{P}_F[Y_i < c]$. The post-transfer poverty rate is $\mathbb{P}_F[Y_i + t(X_i) < c]$.

We can also consider minimizing other measures such as the poverty gap index, which is sensitive to changes below the poverty line.

Example 2 (Poverty Gap Index)**.** Given a poverty line $c \in \mathbb{R}_+$, the pre-transfer poverty gap index is the average shortfall from the poverty line relative to the poverty line, i.e. $\mathbb{E}_F[(c - Y_i)_+/c]$. The post-transfer poverty gap index is, analogously, $\mathbb{E}_F[(c - Y_i - t(X_i))_+/c]$.

Both the poverty rate and the poverty gap index are special cases of the class of FGT-α indices studied by Foster et al. (1984) and Greer and Thorbecke (1986), defined as follows.

Example 3 (FGT Index)**.** Given a poverty line $c \in \mathbb{R}_+$ and any $\alpha \geq 0$, the pre-transfer FGT-α index is $\mathbb{E}_F[\mathbb{I}(Y_i < c) \cdot (c - Y_i)^\alpha]$. The corresponding post-transfer index is $\mathbb{E}_F[\mathbb{I}(Y_i + t(X_i) < c) \cdot (c - Y_i - t(X_i))^\alpha]$.

The FGT-0 index corresponds to the poverty rate and the FGT-1 index corresponds to the poverty gap.

2.1 Weakly Equitable Policies

Due to the emphasis on the poverty rate in public discourse, it may be tempting to learn transfer policies via minimization of the empirical poverty rate as in Example 1, as this approach directly seeks the greatest reduction in the poverty rate under a limited budget. However, in the setting where the policymaker has perfect information on consumption, i.e. $X_i = Y_i$, it is well known that rate minimization yields a policy family that prioritizes transfers to the "richest poor" (Bourguignon and Fields, 1990).[13] This is because the

[13] In particular, and under the additional condition that the budget is not large enough to lift all units above the poverty line, i.e. $0 < B < \mathbb{E}_F[(c - Y_i)_+]$, and Y_i has continuous support under F, the rate-minimizing policy is

$$t(y; B) = (c - y)_+ \cdot \mathbb{I}(\eta(B) < y < c), \tag{4}$$

post-transfer poverty rate can be minimized more efficiently by raising those very close to the poverty line over it than raising those who further away. To rule this out we introduce the following notion of equity, which prioritizes transfers to the worst-off units.

Definition 4 (Weakly Equitable Policy)**.** A family of policies $t(x; B)$ is weakly equitable with respect to a population distribution F if transfers are monotone increasing in budget, i.e. let $B, B' \in \mathbb{R}_+$, then

$$B \geq B' \implies t(x; B) \geq t(x; B') \quad \forall x \in \mathcal{X}, \tag{5}$$

and if incremental transfers are monotone decreasing in post-transfer consumption, i.e. let $x, x' \in \mathcal{X}$ and budgets B, B' such that $0 \leq B' \leq B$, then

$$F_{Y+t(X;B')|X=x} \preceq_{\text{SD}} F_{Y+t(X;B')|X=x'} \implies t(x; B) - t(x; B') \geq t(x'; B) - t(x'; B'). \tag{6}$$

By Definition 4, a weakly equitable policy family maintains or increases transfers to units as the budget expands and allocates incremental funds to the worst-off group in terms of post-transfer consumption.[14]

Clearly, rate minimization is not weakly equitable in the regime where the policymaker targets based on perfect information, i.e. $X_i = Y_i$. This raises the question of whether it is weakly equitable in the more realistic regime where policy is based on imperfect proxies, i.e. there is uncertainty in Y_i given X_i. The following assumption makes concrete this notion of imperfect proxies.

Assumption 2 (Imperfect Proxies)**.** The conditional distribution $F_{Y|X=x}$ has positive, continuously differentiable density $f_{Y|X=x}$ on a compact set $\mathcal{Y}_x \subset \mathbb{R}_+$ and $\mathbb{E}_F\left[\left|\frac{\partial}{\partial Y} \log f_{Y|X=x}(Y)\right| \mid X = x\right] < \infty$ for all $x \in \mathcal{X}$.

Theorem 1 then demonstrates the negative result that, in such an environment, we cannot guarantee that policies that minimize any nonconvex loss function are weakly equitable:

Theorem 1. *Suppose Assumption 1 holds. Consider Definition 2 with some loss function L that is decreasing, piecewise continuously differentiable, and bounded below by $C > -\infty$. If L is nonconvex, then there exists a population distribution F, satisfying Assumption 2, under which the optimal policy family $t_L(x; B)$ is not weakly equitable.*

where $\eta(B) \in (0, c)$ and $\eta(B)$ is decreasing in B. Clearly, transfers are not allocated to the worst-off units.

[14] A weakly equitable policy family will also have the property that transfers are monotone decreasing in (pre-transfer) consumption, i.e. let $x, x' \in \mathcal{X}$ and $B \in \mathbb{R}_+$

$$F_{Y|X=x} \preceq_{\text{SD}} F_{Y|X=x'} \implies t(x; B) \geq t(x'; B). \tag{7}$$

This follows from considering (6) with $B > 0$ and $B' = 0$.

This result applies in particular to the poverty rate loss function (Example 1), which is nonconvex. Fortunately, the opposite also holds for a broad class of relevant poverty measures.

Theorem 2. *Suppose Assumption 1 holds. Consider loss minimization in Definition 2. Let L be a loss function that either is (a) decreasing, strictly convex, differentiable, and bounded below by $C > -\infty$ or (b) is proportional to the loss function of an FGT index with $\alpha \geq 1$ (Example 3). Then, for any population distribution F that satisfies Assumption 2 and has $F_{Y|X=x}$ with positive support on $[0, c]$ for all $x \in \mathcal{X}$, there exists a unique optimal policy t_L induced by solving (2) and t_L is weakly equitable with respect to F.*

Theorem 2 provides the positive result that minimization of a convex poverty measure yields weakly equitable policy families. However, one might worry that this approach yields policies that are cost-ineffective at reducing the more familiar poverty rate metric. To help navigate this trade off we next show that minimizing one particular convex loss function—the poverty gap index (Example 2)—yields guarantees on both weak equity and poverty rate reduction. The first point is immediate, since the poverty gap index satisfies the conditions of Theorem 2. In terms of formal guarantees on poverty rate reduction, we first show that gap minimization solves a worst-case version of rate minimization: it is equivalent to minimization of the worst-case conditional poverty rate among covariate-defined subgroups of the population.[15]

Lemma 3. *Suppose that Assumption 1 and Assumption 2 hold. Consider gap minimization as specified in Example 2. The optimal policy that solves (2) also solves*

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \left\{ \max_{x\in\mathcal{X}} \mathbb{P}_F \left[Y + t(X) < c \mid X = x \right] \right\} : \mathbb{E}_F \left[t(X) \right] \leq B, \quad t(x) \geq 0 \quad \forall x \in \mathcal{X}. \tag{8}$$

Second, we show that gap minimization and rate minimization are equivalent in certain settings where rate minimization itself is weakly equitable.

Assumption 3 (Location Family)**.** The conditional distributions $F_{Y|X=x}$ come from a location family and have strictly increasing density on $[0, c]$ for all $x \in \mathcal{X}$.

Lemma 4. *Suppose that Assumption 1 and Assumption 3 hold. Then the optimal policy that solves (2) when the loss function is the poverty gap index (Example 2) also solves (2) when the loss function is the poverty rate (Example 1).*

Given the desirable equity and poverty rate reduction guarantees of gap minimization, our empirical analysis will focus on gap-minimizing policies, while also presenting rate-minimizing policies to examine how much more the poverty rate could be reduced if this equity condition were ignored.

[15] Lemma 3 is closely related to an observation of Kanbur (1987), who studies gap minimization in a population consisting of two groups and finds that relatively more should be spent on the group with the higher poverty rate.

3 Data Sources and Uses

Table E.2 lists the household surveys, in the indicated countries and years, for which we learn transfer policies. In most cases we use data from the survey on which the most recent estimates of the headcount poverty rate in the World Bank's Poverty and Inequality Platform (`pip.worldbank.org`) are based; in a few cases, these data are not publicly available, and we therefore use an alternative or earlier source. All surveys are nationally representative and contained detailed information on household expenditure and self-production, as well as data on various other household, location, and regional characteristics. All are part of the World Bank's Living Standards Measurement Study (LSMS), its flagship program to collect high-quality, standardized household surveys on poverty and welfare (World Bank, 2014), and include a pre-made consumption aggregate. All told, the sample includes 17 countries that collectively account for 37% of the world's extreme poor.[16] These countries vary meaningfully in their headcount poverty rates, from 8% (in Senegal) to 73% (in Malawi). The poverty rates we estimate from the survey data are generally close to the World Bank estimates for the corresponding country and year, with an average difference of 1.6 percentage points.[17]

The variables we use from the survey data are the household consumption aggregate, household size, household sampling weight, and covariates X_i, where i indexes households. We compute the household per capita consumption Y_i in 2017 PPP USD by converting the reported household consumption aggregate into 2017 PPP USD per day and dividing by household size. As is standard, this approach treats all individuals in the household as having equal consumption. We compute an analysis weight for each household by multiplying its household-level sampling weight by its size.[18]

Selecting covariates requires some judgment. Conceptually, we wish to select only characteristics that are plausibly verifiable, to limit scope for strategic behavior by potential recipients (Grosh and Baker, 1995; Björkegren et al., 2021) and make it easier to hold program staff accountable for enforcing eligibility criteria (Niehaus et al., 2013). We operationalize this idea as follows. First, we identify a collection of proxy-means tests (PMTs) actually implemented in practice in a variety of low- and middle-income countries.[19] We take the union of the predictors used in these PMTs, which yields a list spanning six broad categories: household demographics, human capital, household assets, livelihood activities, geographic indicators, and community

[16] As of 25 August 2025, as computed from World Poverty Clock poverty rate estimates.

[17] In Ethiopia we obtain a much closer match if we use the adult-equivalence scales provided in the survey data, but we do not in order to maintain consistency with the Poverty and Inequality Platform's methodology (World Bank, 2025a). This explains the difference between our estimated poverty rate of 43% and the World Bank's estimated rate of 32%.

[18] The same adjustment is used by the World Bank to construct poverty estimates for the population of individuals in a country rather than the population of households. See here for details

[19] These were Bangladesh (Kidd and Wylde, 2011), Colombia (Camacho and Conover, 2011), India (Niehaus et al., 2013; Planning Commission, 2011, 2012), Indonesia (Kidd and Wylde, 2011; Alatas et al., 2012; Fernandez and Hadiwidjaja, 2018), Malawi (Handa et al., 2022), and Peru (Hanna and Olken, 2018).

characteristics. Where appropriate we generalize from the specific variables used in past PMTs; if a PMT included a refrigerator as a predictor, for example, we would include all large household appliances on our list of eligible predictors. At the same time, we omit a few variables which, though they have been used in the past, would in our judgment typically be difficult to verify. Examples include the age of the household's primary dwelling, measures of food (in)security, indicators for use of fertilizer, or total income. Appendix D.1 provides the rubric for covariate selection. We then select only variables from this list from each living standards survey for use as predictors in our analysis. The resulting datasets are available online.[20]

We use data from each survey to learn transfer policies via empirical minimization of the poverty rate on the survey data; empirical minimization of the poverty gap; and proxy-means testing. Explicit characterizations of the empirical rate-minimizing and gap-minimizing policies and algorithms for deriving them are in Appendix A. Here we provide a brief summary of the issues that arise and the methods we use to address them.

Unrestricted gap minimization is a convex problem, and thus tractable. Optimal transfers are (we show) functions of the conditional (on covariates) quantile functions of the distribution of living standards. We learn these conditional quantile functions for a grid of quantiles using deep learning, minimizing pinball loss (Koenker and Bassett, 1978). When we impose the added constraint that policies must be binary, however, the problem becomes non-convex and requires more care. We express it as a nested optimization over (i) the size of the transfer, and then (ii) the eligibility rule used to determine which households received a transfer of that size. Learning the latter requires learning the expected change in loss associated with a transfer of the specified size to each household, effectively yielding a ranking of households; we do so using empirical risk minimization via deep learning. Nesting this algorithm within a grid search over values of the transfer yields the solution. The learning procedure for binary rate minimization is analogous to the binary gap minimization procedure.

Unrestricted rate minimization, meanwhile, is an intrinsically non-convex problem. We solve it by first showing that rate-minimizing transfers must have a specific structure: conditional on a given value x of the covariates, they must either be at the boundaries (i.e. $t(x) = 0$ or $t(x) = c$), or satisfy $f_{Y|X=x}(c-t) = \alpha$ for some common value α. Intuitively, this condition states that the number of people that would be lifted out of poverty by marginally increasing the transfer amount must be the same for any values of x at which the transfer amount is in the interior. We then show that, for any given value of α, this reduces the problem to a fractional multiple-choice knapsack problem which can be solved using estimates of the covariate distribution f_X (for which we use the empirical covariate distribution) and of the conditional distribution $f_{Y|X=x}$ (which we learn using an extension of Lindsey's method (Efron and Tibshirani, 1996)). Nesting this algorithm

[20] Our datasets are available here.

within a grid search over values of α yields the solution.

All statistical learning methods described in Appendix A are tuned for out-of-sample prediction using a validation set (Hastie et al., 2009). This implies that we do not need to limit the number of predictors we use to guard against overfitting up front and lets us use predictor sets that are often larger than those used in typical PMT exercises; our models use between $d = 43$ and $d = 424$ predictors (last column of Table E.2). That said, our intuition based on initial exploratory work is that one could use substantially smaller predictor sets without much performance degradation, something we plan to explore in future work. Finally, note that we will not be able to take full advantage of large predictor sets when learning the conditional distributions needed for unrestricted rate minimization, as the sample complexity of density estimation scales exponentially in the data dimension; we will return to this point in discussing the results below.

Implementing these learning methods requires that we specify precise training, hyperparameter selection, and evaluation procedures. To prevent iterative analysis of the same data, which could lead to "human-in-the-loop" overfitting, we prescribed these procedures in advance in a Data Use Plan (DUP). A complete copy of this plan is available online;[21] its key features are as follows.

First, we randomly partition each dataset into a 60% training sample and a 40% test sample. We do so maintaining geographic stratification at the level at which the survey has complete coverage. For example, if a survey covered all districts but not all subdistricts within districts, we would maintain stratification at least to the district level when partitioning. We evaluate subsample sizes for this purpose using simple observation counts; where possible, we ensure that this also splits the sampling weights in the same proportion.

Second, we use the training set to learn hyperparameters for the policy learning algorithms. These include, for example, the topology of the neural networks used to fit nuisance parameters for gap minimization or binary rate minimization; the number of quantiles to fit during quantile regression for unrestricted gap minimization; the dimensionality of the predictor set used for unrestricted rate minimization; and various learning parameters, among others. We learn distinct hyperparameters for each policy type we wish to learn, learning them from a 67% sub-sample of the training set and evaluating them based on the performance they induce on the complementary 33% sub-sample (the validation set). We also considered two alternative approaches: hard-coding hyperparameters based on prior intuition, or learning them from synthetic data generated from the training set via a generative adversarial network (Athey et al., 2024). When tested in our "sandbox" environment (see below), we found that all three approaches performed similarly. We therefore chose to learn hyperparameters directly from the training data as this approach is substantially simpler than using synthetic data, while providing some protection (relative to hard-coding) against "unexpected" data distributions we might subsequently encounter.

[21] Our data use plan is available here.

Third, with hyperparameters in place, we use the training set to estimate all nuisance parameters necessary for learning the transfer policies. After that, with access to the covariates from the test set, we generate the transfer schedule $\{t(X_i)\}$ for all units in the test set. Finally, we use the held-out outcomes Y_i from the test set to evaluate the transfer schedules. We refer to this paradigm as transductive learning. This is the relevant paradigm when the policymaker expects to the know the values of the covariates (though not the outcome) for the entire target population when deciding on a policy. This is the case when, for example, a government conducts a census gathering covariate information before finalizing the transfer rule for a program.

The main exception to our otherwise rigid implementation of this plan involves the 2018–2019 Integrated Household Survey in Malawi. We have used this survey as a sandbox environment in which to iteratively experiment with various approaches and guide the choices to which we subsequently committed in our Data Use Plan. Results based on this survey should therefore be interpreted with this caveat in mind. That said, our educated best guess is that any resulting over-fitting to the Malawi data is inconsequential to our main conclusions. Any other adjustments we made to the methodology over time are noted in a changelog within the DUP itself; to date the most salient of these was the modification of the analysis weights to reflect household size as well as survey sampling weights, which we had initially neglected to specify.

4 Empirical Results

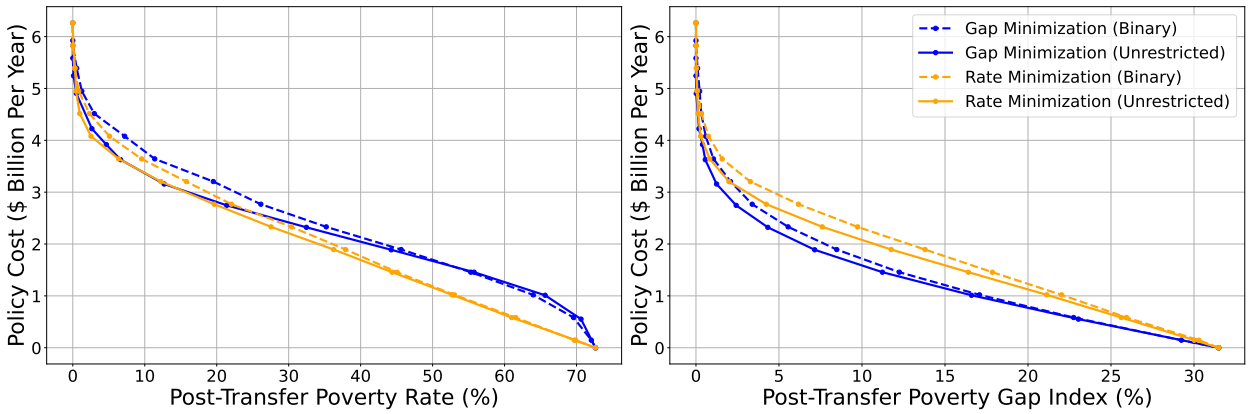
To build intuition we first present core results for a single country, Malawi, one of the world's poorest nations and the country we used as a "sandbox" for experimentation, as noted above. Next, we apply our methods to the full sample and present the aggregate policy cost of achieving uniform poverty goals.

4.1 Malawi Case Study

Figure 1 presents results based on the Fifth Malawi Integrated Household Survey 2018-2019, a survey of $n = 11434$ households that contains $d = 140$ plausibly verifiable covariates. It compares the performance of gap-minimizing and rate-minimizing transfer policies by plotting the total annual policy cost (in nominal 2023 USD) of achieving a particular poverty goal (such as a post-transfer poverty rate or post-transfer poverty gap index). Blue lines indicate methods that minimize the poverty gap; orange lines indicate methods that minimize the poverty rate. The solid and dashed lines correspond to unrestricted and binary policies, respectively.

Several patterns emerge from this example using Malawi, which in general will also hold in the other countries we analyze. First, unrestricted policies are less costly than binary ones (i.e., the solid lines fall below

Figure 1: Gap- and Rate-minimizing Policy Costs in Malawi



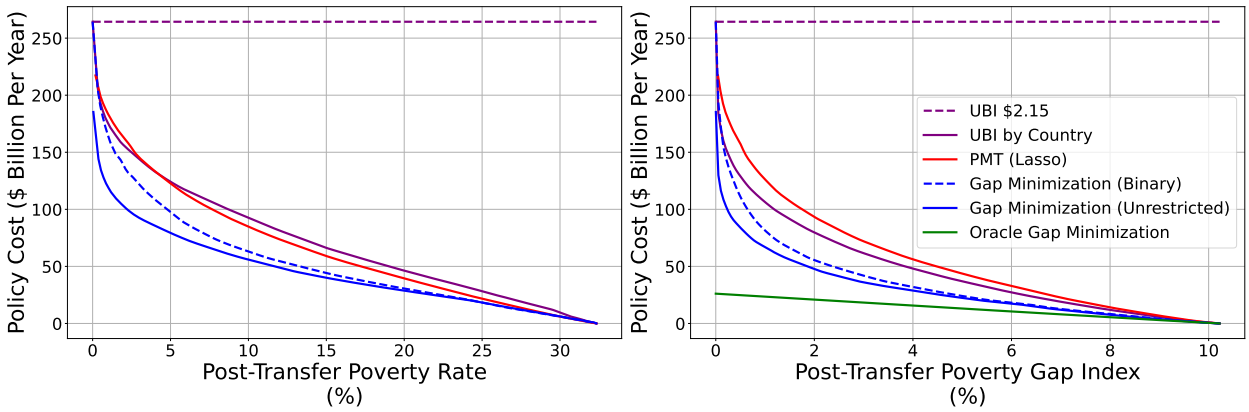
This figure reports the costs, in billions of nominal 2023 USD annually, of the indicated optimal transfer policies learned for Malawi ($n = 11434$, $d = 140$). The left-hand panel plots these against the post-transfer poverty rate, while the right-hand panel plots these against the post-transfer poverty gap index.

the dashed lines). This is expected, since unrestricted policies can allocate smaller transfers to households closer to the poverty line—but in practice the differences are fairly modest. Second, while rate minimization is more effective than gap minimization at achieving modest poverty rate goals—e.g., reducing the post-transfer poverty rate to 40%—gap minimization is equally as effective when the poverty rate goal is more ambitious, e.g., reducing the rate to 10% (Figure 1a). We hypothesize that this is because, in order to achieve more ambitious objectives, the rate-minimizing policy will have to allocate transfers to households that are well below the poverty line, reducing the difference in cost between rate minimization and gap minimization. Practically, it implies that one can reduce the poverty rate using the weakly equitable policies obtained from gap minimization at little penalty in terms of cost-effectiveness. Policies that minimize the poverty gap are analogously more effective at reducing the gap (i.e., the blue lines fall below the orange lines in Figure 1b), and this difference appears consistently across a range of poverty gap goals.

4.2 Multi-Country Analysis

We now present results aggregated across the full sample, with the underlying country-specific results available in Appendix E. We aggregate by calculating the cost of achieving a given level of a given poverty metric in *every* country in our sample, and summing these. The estimates are thus conservative for the cost of achieving the same result on *average*, reflecting as they do policies that impose a degree of equity between the countries in our sample. We may explore relaxing this requirement in future work. We present cost curves in their entirety, but also discuss in more detail the costs of achieving a 1% poverty rate in each country using policies learned via empirical poverty-gap minimization. Note that achieving this target globally would

Figure 2: Gap minimization v.s. Oracle, UBI, and PMT Benchmarks in the Full Sample



This figure reports the costs, in billions of nominal 2023 USD annually, of the indicated optimal transfer policies learned for the full sample of countries listed in Table E.2. The left-hand panel plots these against the post-transfer poverty rate, while the right-hand panel plots these against the post-transfer poverty gap index.

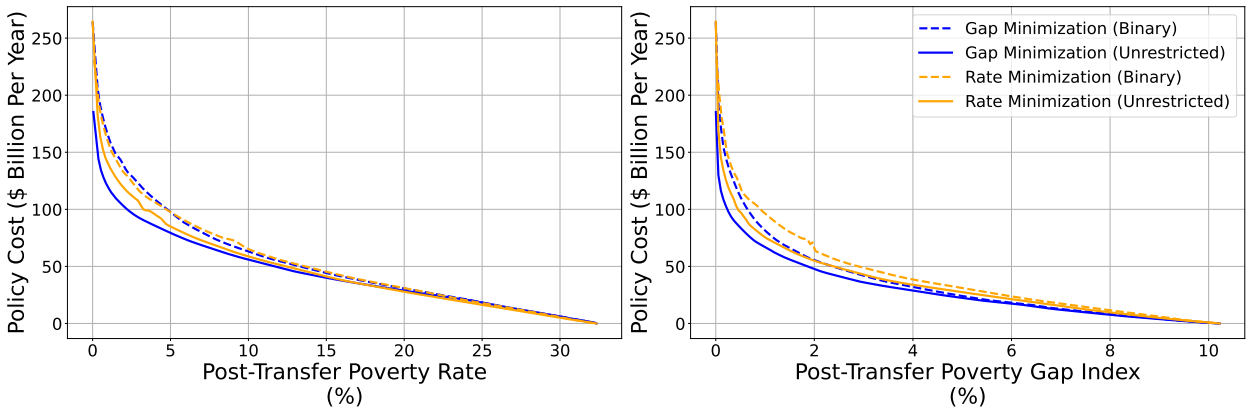
imply that the global poverty rate is 0.8%.

Figure 2 presents the core results, plotting total annual policy cost (in nominal 2023 USD) against the resulting headcount poverty rate (left-hand panel) or poverty gap index (right-hand panel) as above. The primary policy of interest here is the unrestricted gap-minimizing policy, indicated in blue. The binary policy minimizes the same objective, but adds the restriction that within each country, transfers must either be zero or some common positive amount. Other series provide relevant benchmarks. The "UBI $2.15" series indicates the (constant) cost of giving every individual $2.15 2017 PPP per day, without any targeting. The "UBI by Country" series indicates the cost of giving every individual within a given country the same transfer amount less than or equal to $2.15. The "PMT (Lasso)" series in red corresponds to allocating transfers with a proxy-means test; specifically, we fit conditional mean consumption via lasso regression (Hanna and Olken, 2018) and give individuals with mean predicted consumption below a threshold a transfer of $2.15 2017 PPP per day. Finally, the right-hand panel includes an "Oracle" series in green indicating the (hypothetical) cost of achieving a given poverty gap index with a weakly equitable policy if full information on each household's consumption were available. This is a useful benchmark given the salience of the global poverty gap in recent discourse about the cost of ending poverty (cf. Chandy et al., 2016; Kharas and McArthur, 2023; Sumner and Yusuf, 2024).[22]

Overall the results show that the lowest-cost policies cost several times more than the aggregate poverty gap, but also substantially less than the UBI and other benchmark policies. To reduce the poverty rate in all countries to 1%, for example, would cost $119B per year, which is 4.6 times the oracle cost of ($26B),

[22] We do not report oracle costs of achieving poverty rate goals as they are not a particularly useful benchmark: the only way to bring even a single household up to the poverty line while maintaining weak equity is to bring *all* households up to the line.

Figure 3: Gap- v.s. Rate-minimizing Policy Costs in the Full Sample

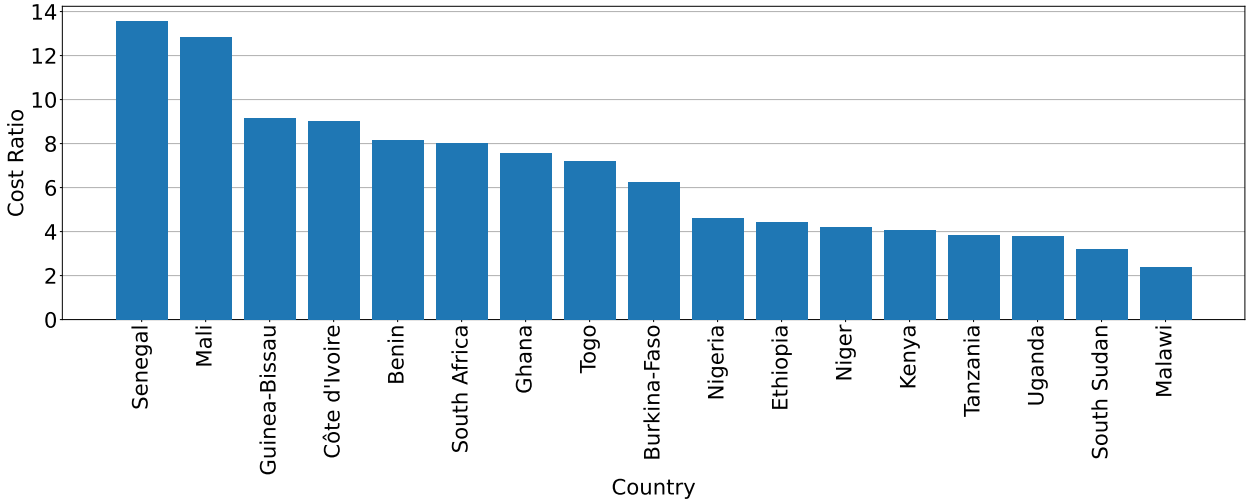


This figure reports the costs, in billions of nominal 2023 USD annually, of the indicated optimal transfer policies learned for the full sample of countries listed in Table E.2. The left-hand panel plots these against the post-transfer poverty rate, while the right-hand panel plots these against the post-transfer poverty gap.

but 45% of the cost of a $2.15 UBI, and 68% of the cost of a country-specific UBI. Similar ratios hold as we vary the desired post-transfer poverty rate. As the post-transfer poverty rate (or gap) approaches zero, however, the costs of all feasible policies increase at an increasing rate, as expected: requiring any of these algorithms to find the last needles in the population haystack becomes increasingly costly. In addition, we find that unrestricted gap minimization is 65% the cost of PMT, highlighting the performance benefits of the loss minimization approach. Lastly, the binary gap-minimizing policy costs 38% more than the unrestricted gap-minimizing policy. This implies that, if politically or logistically expedient, the transfer policy can be reduced to a simple eligibility rule without massively increasing its cost.

While none of the policies examined in Figure 2 are optimized to reduce the poverty rate, the left-hand panel shows that gap minimization still substantially reduces the poverty rate. This suggests that it would be possible to implement gap-minimizing policies, which are weakly equitable, while still communicating with broad audiences about progress in terms of the poverty rate, a more familiar concept. That said, Figure 3 directly contrasts the empirical performance of policies optimized for the rate as opposed to the gap. Axes are as in Figure 2, but here the policies include two that minimize the rate (using either unrestricted or binary-valued transfers) as well as the corresponding policies minimizing the gap. The right-hand panel shows that, as expected, policies optimized for the gap consistently reduce the gap at a lower cost than the corresponding policies optimized for the rate. While in the infinite-population regime, we expect rate minimization to obtain lower post-transfer poverty rate than gap minimization, the left-hand panel shows empirical gap minimization can actually perform somewhat *better* at reducing the rate than empirical rate minimization. A first reason that this occurs is that because under the hood, empirical rate minimization must solve a more challenging statistical problem than empirical gap minimization; empirical rate minimization relies on

Figure 4: Cost Ratio of Feasible Gap-Minimizing Policy to Aggregate Poverty Gap



This figure reports, for each country in Table E.2, the cost ratio between the gap-minimizing policy that achieves a 1% post-transfer poverty rate and the aggregate poverty gap.

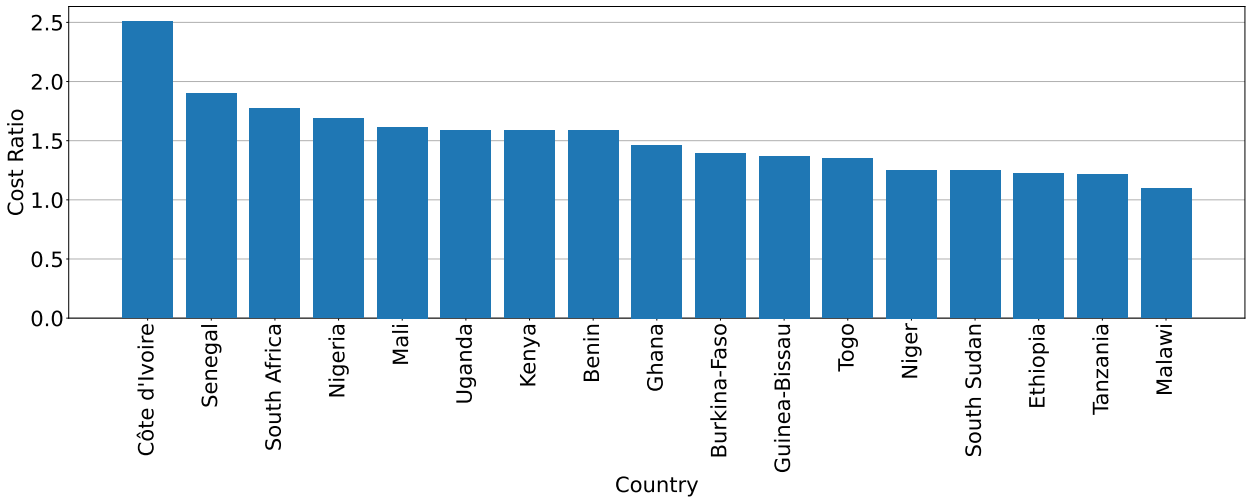
conditional density estimation, while empirical gap minimization relies on quantile estimation. This task has higher sample complexity and also requires a bespoke algorithm. Second, empirical gap minimization can leverage richer covariate information than empirical rate minimization. Since the sample complexity of density estimation scales exponentially in the data dimension, we tune the number of covariates used in our conditional density estimation algorithm. In contrast, such feature selection is not required for our quantile estimation algorithm. In terms of finite-sample performance, we find little reason to prefer rate minimization in terms of performance.

The full country-specific details underlying these results are presented in Appendix E in Figure E.1. Generally speaking, the comparisons above for the sample of countries as a whole generally hold within each country as well. Reducing the poverty rate to 1% with an optimal policy costs between 2.4 and 13.9 times the aggregate poverty gap (Figure 4), but still substantially less than the cost of doing so using UBI at the level of the international poverty line. In some of the poorest countries, however, the cost of achieving this goal using an optimally sized UBI is not much greater than the cost of doing so using variably sized transfers (Figure 5), reflecting the fact that most households in these countries are below the poverty line.

4.3 Variations and extensions

This section reports how the minimized costs of achieving poverty goals vary with parameters of the problem. For simplicity of exposition we focus discussion primarily on the cost of achieving a national poverty rate of 1% in each country, as in our summary discussion above, rather than on comparative statics for the full cost curve.

Figure 5: Cost Ratio of UBI (Variable) to Gap Minimization

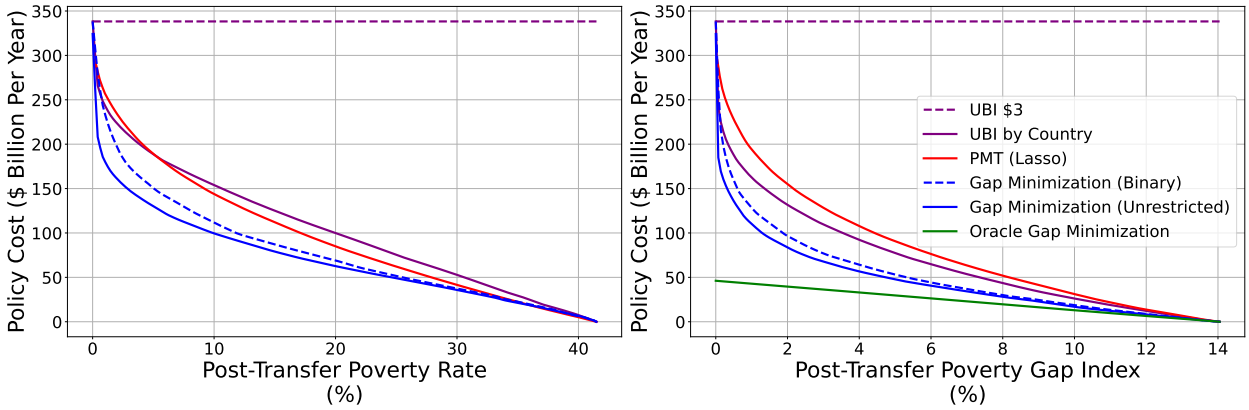


This figure reports, for each country in Table E.2, the ratio of the cost of a policy that achieves a 1% post-transfer poverty rate by giving a common transfer amount to all individuals, and the cost of a policy that does so using transfers that vary person-to-person learned via gap minimization.

We first consider an alternative poverty line. The $2.15 2017 PPP line is a well-established benchmark in both public and policy discourse, derived (via periodic inflation-indexing) from the famous concept of "dollar-a-day" poverty introduced by the World Bank in 1990 and subsequently enshrined in the Millennium Development Goals and the Sustainable Development Goals. Yet it is also a very low bar. Lant Pritchett, among others, has argued forcefully that actors working to promote economic development should view "dollar-a-day" poverty as a milestone along the way, not as an end goal (Pritchett, 2024; Pritchett and Viarengo, 2025), and in the World Bank's 2025 revision to its global poverty line it incorporated not only new information about prices, but also an upward shift in the *real* standard of living used to define the threshold. This was a meaningful change, increasing by some 125 million the number of people defined as living in extreme poverty in 2022 (Foster et al., 2025).

Our methods and data can be used flexibly to calculate the cost of achieving poverty goals with respect to any poverty line. Here we redo our analysis using the World Bank's new $3.00 2021 PPP line, which is of intrinsic interest as well as serving to illustrate the consequences of setting a higher bar more generally. Figure 6 replicates Figure 2, but using this higher line to define extreme poverty. Achieving a poverty rate of 1% in every country in our sample using policies learned via empirical poverty-gap minimization costs $181B per year in nominal 2023 dollars. This is 1.5 times the cost (again, in nominal 2023 dollars) of achieving the same rate when poverty is defined using the $2.15 2017 PPP standard. This difference quantifies the extent to which in which ending extreme poverty under the new definition is indeed a substantially more ambitious goal than under the old.

Figure 6: Policy Costs for Poverty Reduction under $3.00 (2021 PPP) Poverty Line



This figure reports the costs, in billions of nominal 2023 USD annually, of the indicated optimal transfer policies learned for the full sample of countries to reduce extreme poverty, as defined as living on less than $3.00 (2021 PPP) per day. This figure is analogous to Figure 2, which reports costs to reduce extreme poverty under the $2.15 (2017 PPP) poverty line.

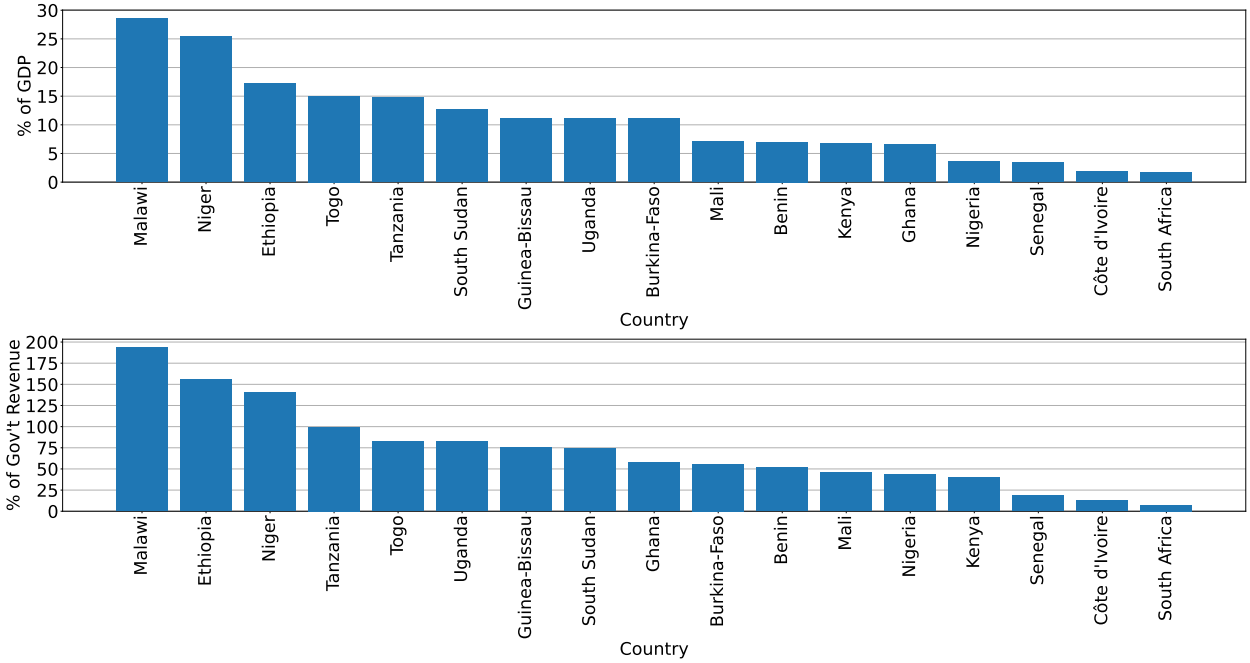
While setting more ambitious goals necessarily increases costs, other variations could help to lower them. We are examining several of these in ongoing work. First, rather than requiring a uniform poverty rate in every country, we can allow for variation in post-transfer poverty rates across countries, and thus achieve the same global rate at a (weakly) lower total cost but with greater inequity across countries. Second, we can include finer-grained regional indicators as covariates in our predictive models. This will tend to improve accuracy and thus lower transfer costs, but at the implicit cost of requiring that future living standard surveys be conducted in more regions than they were in the original data source. And third, we explore whether non-traditional covariates sourced from newer digital data sources—such as satellite imagery (Jean et al., 2016; Chi et al., 2022) or cell phone networks (Blumenstock et al., 2015; Aiken et al., 2023)—could enable accurate targeting, while also reducing (or eliminating) the cost of collecting traditional covariates through in-person surveys. We expect to share results from these exercises in future drafts.

5 Discussion & Implications

Table E.4 summarizes the estimated cost of reducing the national poverty rate to 1% in each country in our sample, along with reference information on the scale of national GDP and public revenue. Note that achieving this globally would yield a global poverty rate of 0.8%. We conclude in this section by discussing some further potential implications of redistribution on the scale implied by those figures.

The first concerns sources of funds. It does not appear (for the most part) realistic to finance a policy like this one using purely domestic sources of revenue. Figure 7 plots the amount of money spent in each country, under the learned policy which reduces the global poverty gap to 1%, as a proportion of that country's GDP

Figure 7: Policy Cost as a Percentage of GDP and of Government Revenue



This figure plots, for each indicated country, the ratio of the cost of a policy that achieves a 1% poverty rate via gap minimization to the country's GDP in the year the underlying survey was conducted (top panel) and to total government revenue in that year (bottom panel). We obtain country GDP from the World Bank (data source, accessed 13 August 2025) and government revenue percentages by country from the IMF (data source, accessed 14 July 2025).

(top panel) and total government revenue (bottom panel). For the average country, implementing such a scheme would require increasing the size of the government revenue by 84% of GDP. To put this into context, this is equivalent to roughly 70 years worth of the growth over time of tax revenue relative to GDP that Besley and Persson (2014) document in a sample of 18 countries during the 1900s (see Figure 3 in their paper). Overall, however, our assessment is similar to that reached by Hanna and Olken (forthcoming). They calculate the costs of transferring PPP $2.15 to everyone below the poverty line, taking this as a rough approximation to the true cost of eliminating poverty. They note in particular that "in principle, one could bring everyone to the poverty line for less money than this if one could give larger transfers to those further away from the poverty line" but also that "this may understate the extent of the problem, since giving transfers only to those below the poverty line assumes that one can solve targeting challenges..." After taking both of these issues into account, the policies we learn here end up costing a similarly large share of national income.

Funding such policies using transfers from abroad, on the other hand, would raise additional questions about their macroeconomic effects. One can think of these in two parts: the effects of converting foreign currency to Local Currency Units (LCUs), and the effects of recipients spending or saving those LCUs. Both

are large topics in their own right, and beyond the scope of our analysis here, but we will remark briefly on related work.

With respect to currency conversions, the amounts in question are several multiples of status quo aid flows. In our sample of countries, for example, the average country received 7% of GDP in 2023 and would require 13% of GDP to implement a policy that reduces the national poverty rate to 1% in each country (Table E.4). Recent work on the effects of currency demand shocks generally suggests that additional currency purchases on this scale would have meaningful effects on rates of exchange, at least in the short run. One approach uses (instruments for) countries' open-market operations as a source of identifying variation; Adler et al. (2019), for example, estimate that a purchase of 1 percentage point of GDP causes a depreciation of the real exchange rate of 1.4–1.7%. Another approach uses shocks to currency demand caused by changes to the composition of global indices of emerging market bonds; these are arguably more clearly exogenous, but also harder to size, as passive funds' responses are mechanical and thus estimable but those of actively-managed funds are unknown. This approach yields larger elasticities. For example, Beltran and He (2025) estimate using this approach that in a sample of countries an inflow of 0.09% of domestic GDP led to a 1% appreciation of the exchange rate in the days immediately following, with effects dissipating only partially over the subsequent 12 months (see Table 5.1 and Figure 4 in their paper). Any effects of aid inflows on exchange rates would then be mitigated to the extent that transfer recipients subsequently purchased items with imported components. That said, the broad point remains that effects on exchange rates seem likely to be large enough to matter.

With respect to domestic spending, some classic lines of thinking about the process of development imply that this might affect economic growth, through a demand channel (as in "big push" models such as that of Murphy et al. (1989)) or a credit channel (Stiglitz and Weiss, 1981), among others. Recent evidence from large-scale experiments or natural experiments broadly supports this view: Egger et al. (2022) estimate that the economies of villages in rural Kenya expanded by $2.5 for every $1 transferred into them, and Gerard et al. (2024) similarly find substantial effects of transfers on economic activity in Brazilian municipalities.[23] Any such knock-on benefits need not accrue proportionally, however—in fact, Haushofer et al. (forthcoming) find that the transfers in Egger et al. (2022) had somewhat larger effects on the income and consumption of households that would otherwise have been somewhat less poor. This is one reason, as well as for simplicity, that we have abstracted from spillovers here.

Another natural and related question is whether there is some upfront outlay that would be sufficient to achieve a given poverty goal on a lasting basis, as opposed to the flow-cost approach we have taken here.

[23] While the transfers which Gerard et al. (2024) were largely domestically financed, they study the effects of gross transfers (not net of taxation), which are the relevant ones for thinking about the impacts of externally financed transfers.

This of course inescapably requires assumptions about rates of return. One way to bound the up-front cost is to simply calculate the size of the endowment required to yield the annual flow expenditures we calculate, using any rate of return deemed plausible. But it may be possible to do better still by front-loading transfers to households themselves, as some of them may have access to much-higher-return investment opportunities than endowment managers do (see for example Haushofer et al. (forthcoming) and Hussam et al. (2022), among many others). Here the volatility of poverty would pose a challenge. Historically, households have moved back and forth across the extreme poverty line quite frequently (Armentano et al., 2025; Baulch and Hoddinott, 2000), so that the amount one would ideally wish to transfer to any given household could vary substantially from year to year (or even month to month; see Merfeld and Morduch, 2024). How well policies learned from data at a given point in time continue to approximate optimal policies in subsequent years is an open question.

While we focus on a sample of countries that meet the criteria defined above in Section 3, it is natural to wonder what the results imply for the cost of ending poverty at a global scale. We consider in particular the cost of reducing the global poverty rate to 1%. To do so it is sufficient to obtain a national poverty rate of 1.3% in all countries that currently have higher poverty rates (a target slightly higher than 1%, because there are some countries in the world with extreme poverty rates that are already essentially zero). To estimate the costs of achieving this target in the countries that are not in our current sample, we estimate the relationship between the ratio of the feasible policy cost to the poverty gap and the poverty rate within our sample, and then use this to predict the feasible policy cost to poverty gap ratio for other countries.[24] To be clear, this extrapolation is purely illustrative. In future work we plan to extend the analysis to additional countries, and it is entirely possible that the cost ratio will vary in those. Moreover, and as noted above, many of the surveys on which the input poverty rate and poverty gap index figures are based were conducted prior to the COVID-19 pandemic, and so some additional forward extrapolation would be necessary to obtain a truly current estimate. That said, the exercise may give us some sense of the scale of transfers potentially involved. It yields an estimated cost of reducing the global poverty rate to 1% equal to 0.31% of global GDP, or 0.5% of OECD GDP.[25]

In some senses this is an enormous figure. As a point of reference, major aid donors in the OECD gave 0.37% of GDP in 2023.[26] Political support for foreign aid is of course ebbing rapidly as of this writing. But in terms of sheer fiscal feasibility there is no question that wealthy countries could finance policies such as this one. And in a broader sense it implies that ending poverty is no more costly than other, less inspiring

[24] Specifically, we regress the ratio between a country's feasible policy cost and poverty gap on its poverty rate in our sample, yielding an R^2 of 0.75. We then use estimates of the poverty gap index for out-of-sample countries from the World Bank to obtain predicted feasible policy costs.

[25] Including China, the world's second-largest economy and now a major international aid donor, lowers this figure to 0.39%.

[26] From Our World in Data. Accessed 11 April 2025.

global priorities. The world spends seven times as much, for example—2.2% of global GDP—on alcoholic beverages.[27]

[27] Estimated 2024 global alcoholic beverages market size: $2,413B as per https://www.fortunebusinessinsights.com/alcoholic-beverages-market-107439, accessed 1 September 2025. Estimated 2024 GPD: $110,000B as per https://www.imf.org/external/datamapper/NGDPD@WEO/OEMDC/ADVEC/WEOWORLD, accessed 1 September 2025.

References

Adler, Gustavo, Noëmie Lisack, and Rui C. Mano, "Unveiling the effects of foreign exchange intervention: A panel approach," *Emerging Markets Review*, 2019, *40*, 100620.

Aiken, Emily L., Guadalupe Bedoya, Joshua E. Blumenstock, and Aidan Coville, "Program targeting with machine learning and mobile phone data: Evidence from an anti-poverty intervention in Afghanistan," *Journal of Development Economics*, March 2023, *161*, 103016.

Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock, "Machine learning and phone data can improve targeting of humanitarian aid," *Nature*, 2022, *603* (7903), 864–870.

Alatas, Vivi, Abhijit Banerjee, Arun G Chandrasekhar, Rema Hanna, and Benjamin A Olken, "Network structure and the aggregation of information: Theory and evidence from Indonesia," *American Economic Review*, 2016, *106* (7), 1663–1704.

_ , _ , Rema Hanna, Benjamin A Olken, and Julia Tobias, "Targeting the poor: evidence from a field experiment in Indonesia," *American Economic Review*, 2012, *102* (4), 1206–1240.

Amendola, Nicola and Giovanni Vecchi, "Durable goods and welfare measurement," *Journal of Economic Surveys*, 2022, *36* (4), 1179–1211.

Anderson, Edward and Hugh Waddington, "Aid and the Millennium Development Goal Poverty Target: How Much is Required and How Should it be Allocated?," Technical Report 1 2007.

Armentano, Vincent, Paul Niehaus, and Tom Vogl, "How Poverty Fell," Technical Report, UC San Diego February 2025.

Athey, Susan and Stefan Wager, "Policy learning with observational data," *Econometrica*, 2021, *89* (1), 133–161.

_ , Guido W. Imbens, Jonas Metzger, and Evan Munro, "Using Wasserstein Generative Adversarial Networks for the design of Monte Carlo simulations," *Journal of Econometrics*, 2024, *240* (2), 105076.

Atkinson, Anthony Barnes, "On the measurement of poverty," *Econometrica: Journal of the Econometric Society*, 1987, pp. 749–764.

Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, and Sudarno Sumarto, "The (lack of) distortionary effects of proxy-means tests: Results from a nationwide experiment in Indonesia," *Journal of Public Economics Plus*, 2020, *1*, 100001.

Banerjee, Abhijit V., Rema Hanna, Gabriel E. Kreindler, and Benjamin A. Olken, "Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs," *The World Bank Research Observer*, 08 2017, *32* (2), 155–184.

Bastagli, Francesca, Jessica Hagen-Zanker, Luke Harman, Valentina Barca, Georgina Sturge, and Tanja Schmidt, "The Impact of Cash Transfers: A Review of the Evidence from Low- and Middle-income Countries," *Journal of Social Policy*, July 2019, *48* (3), 569–594.

Baulch, Bob and John Hoddinott, "Economic mobility and poverty dynamics in developing countries," *Journal of Development Studies*, 2000, *36* (6), 1–24.

Beltran, Paula and Chang He, "Inelastic Financial Markets and Foreign Exchange Interventions," Technical Report, University of Virginia March 2025.

Bertsimas, Dimitris and Nathan Kallus, "From predictive to prescriptive analytics," *Management Science*, 2020, *66* (3), 1025–1044.

Besley, Timothy and Torsten Persson, "Why Do Developing Countries Tax So Little?," *Journal of Economic Perspectives*, November 2014, *28* (4), 99–120.

Bhattacharya, Debopam and Pascaline Dupas, "Inferring welfare maximizing treatment assignment under budget constraints," *Journal of Econometrics*, 2012, *167* (1), 168–196.

Björkegren, Daniel, Joshua E. Blumenstock, and Samsun Knight, "Manipulation-Proof Machine Learning," *Working Paper*, 2021. First version: November 30, 2018; this version: August 21, 2021.

_ , _ , and _ , "(Machine) Learning What Policies Value," June 2022. arXiv:2206.00727 [cs, econ, q-fin].

Blumenstock, Joshua Evan, Gabriel Cadamuro, and Robert On, "Predicting poverty and wealth from mobile phone metadata," *Science*, November 2015, *350* (6264), 1073–1076.

Bourguignon, Francois and Gary S Fields, "Poverty measures and anti-poverty policy," *Recherches Économiques de Louvain/Louvain Economic Review*, 1990, *56* (3-4), 409–427.

Brown, Caitlin, Martin Ravallion, and Dominique Van de Walle, "A poor means test? Econometric targeting in Africa," *Journal of Development Economics*, 2018, *134*, 109–124.

Camacho, Adriana and Emily Conover, "Manipulation of social program eligibility," *American Economic Journal: Economic Policy*, 2011, *3* (2), 41–65.

Chandy, Laurence, Lorenz Noe, and Christine Zhang, "The global poverty gap is falling," *Brookings Commentary*, January 2016.

Chi, Guanghua, Han Fang, Sourav Chatterjee, and Joshua E. Blumenstock, "Microestimates of wealth for all low- and middle-income countries," *Proceedings of the National Academy of Sciences*, January 2022, *119* (3). ISBN: 9782113658118 Publisher: National Academy of Sciences Section: Social Sciences.

Coady, David, Margaret Grosh, and John Hoddinott, "Targeting outcomes redux," *The World Bank Research Observer*, 2004, *19* (1), 61–85.

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

Cuaresma, Jesús Crespo, Wolfgang Fengler, Homi Kharas, Kunal Sen, and Mehdi Sultan, "Will the Sustainable Development Goals be fulfilled? Assessing present and future global poverty," *Palgrave Communications*, 2018, *4* (29).

Devarajan, Shantayanan, Margaret Miller, and Eric Swanson, "Goals for Development: History, Prospects, and Costs," Technical Report 2819, World Bank April 2002.

Domar, Evsey D., "Capital Expansion, Rate of Growth, and Employment," *Econometrica*, 1946, *14* (2), 137–147.

Easterly, William, "The ghost of financing gap: testing the growth model used in the international financial institutions," *Journal of Development Economics*, 1999, *60* (2), 423–438.

Efron, Bradley and Robert Tibshirani, "Using specially designed exponential families for density estimation," *The Annals of Statistics*, 1996, *24* (6), 2431–2461.

Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker, "General Equilibrium Effects of Cash Transfers: Experimental Evidence From Kenya," *Econometrica*, 2022, *90* (6), 2603–2643.

Fernandez, Luisa and Gracia Hadiwidjaja, "Do Household Socioeconomic Sstatus And Characteristics Change Over A 3 Year Period In Indonesia? Evidence From Susenas Panel 2008-2010," Technical Report 3, Jakarta:TNP2K 2018.

Foster, Elizabeth Mary, Dean Mitchell Jolliffe, Gabriel Lara Ibarra, Christoph Lakner, and Samuel Kofi Tetteh Baah, "Global Poverty Revisited Using 2021 PPPs and New Data on Consumption," Policy Research Working Paper Series 11137, The World Bank Jun 2025.

Foster, James, Joel Greer, and Erik Thorbecke, "A class of decomposable poverty measures," *Econometrica: journal of the econometric society*, 1984, pp. 761–766.

Gerard, F., J. Naritomi, and J. Silva, "Cash Transfers and the Local Economy: Evidence from Brazil," *Centre for Economic Policy Research Discussion Papers, DP16286*, 2021.

Gerard, Francois, Joana Naritomi, and Joana Silva, "Cash Transfers and the Local Economy: Evidence from Brazil," Technical Report, London School of Economics August 2024.

Glewwe, Paul, "Targeting assistance to the poor. Efficient allocation of transfers when household income is not observed," *Journal of Development Economics*, 1992, *38* (2), 297–321.

Greer, Joel and Erik Thorbecke, "A methodology for measuring food poverty applied to Kenya," *Journal of Development Economics*, 1986, *24* (1), 59–74.

Grosh, Margaret and Judy L Baker, "Proxy means tests for targeting social programs," *Living standards measurement study working paper*, 1995, *118*, 1–49.

Handa, Sudhanshu, Maxton Tsoka, and Clare Barrington, "Malawi Social Cash Transfer Programme (SCTP) Longitudinal Impact Evaluation: Baseline Report," Technical Report, University of North Carolina December 2022.

Hanna, Rema and Benjamin A Olken, "Universal basic incomes versus targeted transfers: Anti-poverty programs in developing countries," *Journal of Economic Perspectives*, 2018, *32* (4), 201–226.

_ and Benjamin Olken, "Introduction," in Rema Hanna and Benjamin Olken, eds., *The Handbook of Social Protection: Evidence to Inform Policy in Low- and Middle-Income Countries*, MIT Press, forthcoming.

Harrod, R. F., "An Essay in Dynamic Theory," *The Economic Journal*, 1939, *49* (193), 14–33.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., New York: Springer, 2009.

Haushofer, Johannes, Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael Walker, "Targeting Impact versus Deprivation," *American Economic Review*, June 2025, *115* (6), 1936–1974.

_ , _ , _ , _ , and _ , "Targeting Impact versus Deprivation," *American Economic Review*, forthcoming.

Huang, Xinyang and Jin Xu, "Estimating individualized treatment rules with risk constraint," *Biometrics*, 2020, *76* (4), 1310–1318.

Hussam, Reshmaan, Natalia Rigol, and Benjamin N. Roth, "Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field," *American Economic Review*, March 2022, *112* (3), 861–98.

IMF, "IMF Data Portal," https://data.imf.org 2025.

Jean, Neal, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon, "Combining satellite imagery and machine learning to predict poverty," *Science*, 2016, *353* (6301), 790–794.

Kallus, Nathan and Angela Zhou, "Minimax-optimal policy learning under unobserved confounding," *Management Science*, 2021, *67* (5), 2870–2890.

Kanbur, SM Ravi, "Measurement and Alleviation of Poverty: With an Application to the Effects of Macroeconomic Adjustment (Evaluation quantitative de la pauvreté et remèdes possibles: analyse des effets d'un ajustement macroéconomique)(Medición y alivio de la pobreza, con una aplicación a los efectos del ajuste macroeconómico)," *Staff Papers-International Monetary Fund*, 1987, pp. 60–85.

Kharas, Homi and John McArthur, "Updating Institutional Technologies: a Purpose-driven Fund to End Extreme Poverty," Technical Report 183, Brookings 2023.

Kidd, Stephen and Emily Wylde, "Targeting the Poorest: An assessment of the proxy means test methodology," *AusAID Research Paper, Australian Agency for International Development, Canberra, Australia*, 2011.

Kilic, Talip, Umar Serajuddin, Hiroki Uematsu, Nobuo Yoshida, Talip Kilic, Umar Serajuddin, Hiroki Uematsu, and Nobuo Yoshida, "Costing household surveys for monitoring progress toward ending extreme poverty and boosting shared prosperity," Policy Research Working Paper Series 7951, The World Bank January 2017.

Kitagawa, Toru and Aleksey Tetenov, "Who should be treated? empirical welfare maximization methods for treatment choice," *Econometrica*, 2018, *86* (2), 591–616.

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer, "Prediction policy problems," *American Economic Review*, 2015, *105* (5), 491–495.

Klenke, Achim, *Probability theory: a comprehensive course*, Springer Science & Business Media, 2013.

Koenker, Roger and Gilbert Bassett Jr, "Regression quantiles," *Econometrica: journal of the Econometric Society*, 1978, pp. 33–50.

Luedtke, Alexander R and Mark J van der Laan, "Optimal individualized treatments in resource-limited settings," *The international journal of biostatistics*, 2016, *12* (1), 283–303.

Mankiw, N. Gregory, Edmund S. Phelps, and Paul M. Romer, "The Growth of Nations," *Brookings Papers on Economic Activity*, 1995, *1995* (1), 275–326.

Manski, Charles F, "Statistical treatment rules for heterogeneous populations," *Econometrica*, 2004, *72* (4), 1221–1246.

Merfeld, Joshua and Jonathan Morduch, "Poverty at Higher Frequency," Technical Report, NYU August 2024.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny, "Industrialization and the Big Push," *Journal of Political Economy*, 1989, *97* (5), 1003–1026.

Niehaus, Paul, Antonia Attanassova, Mariane Bertrand, and Sendhil Mullainathan, "Targeting with Agents," *American Economic Journal: Economic Policy*, February 2013, *5* (1), 206–238.

Ninno, Carlo Del and Bradford Mills, *Safety nets in Africa: Effective mechanisms to reach the poor and most vulnerable*, The World Bank Group, 2015.

Noriega-Campero, Alejandro, Bernardo Garcia-Bulle, Luis Fernando Cantu, Michiel A Bakker, Luis Tejerina, and Alex Pentland, "Algorithmic targeting of social policies: fairness, accuracy, and distributed governance," in "Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency" 2020, pp. 241–251.

Planning Commission, "BPL Criteria," Technical Report, Parliament of India August 2011.

_ , "Report of the Expert Group to Recommend the Detailed Methodology for Identification of Families Living Below Poverty Line in the Urban Areas," Technical Report, Government of India December 2012.

Pritchett, Lant, "'Dollar a day' poverty was a development mileage marker, not the destination," Working Paper, London School of Economics February 2024.

_ **and Martina Viarengo**, "Raising the Bar: A Poverty Line for Global Inclusion," Growth Lab Working Paper 250, Harvard University 2025.

Ravallion, Martin and Kalvin Chao, "Targeted policies for poverty alleviation under imperfect information: algorithms and applications," *Journal of Policy modeling*, 1989, *11* (2), 213–224.

Sachs, Jeffrey, *Investing in development : a practical plan to achieve the Millennium Development Goals*, London, UK: Earthscan, 2005.

Sen, Amartya, "Poverty: an ordinal approach to measurement," *Econometrica: Journal of the Econometric Society*, 1976, pp. 219–231.

Smythe, Isabella S. and Joshua E. Blumenstock, "Geographic microtargeting of social assistance with high-resolution poverty maps," *Proceedings of the National Academy of Sciences*, August 2022, *119* (32), e2120025119. Publisher: Proceedings of the National Academy of Sciences.

Stiglitz, Joseph E and Andrew Weiss, "Credit rationing in markets with imperfect information," *The American economic review*, 1981, *71* (3), 393–410.

Sumner, Andy and Arief Anshory Yusuf, "New estimates of the cost of ending poverty and its global distribution," *WIDER Working Paper Series*, 2024. Number: wp-2024-46 Publisher: World Institute for Development Economic Research (UNU-WIDER).

Sun, Hao, Evan Munro, Georgy Kalashnov, Shuyang Du, and Stefan Wager, "Treatment allocation under uncertain costs," *arXiv preprint arXiv:2103.11066*, 2021.

Sun, Liyang, "Empirical welfare maximization with constraints," *arXiv preprint arXiv:2103.15298*, 2021.

Sverdrup, Erik, Han Wu, Susan Athey, and Stefan Wager, "Qini Curves for Multi-Armed Treatment Rules," *arXiv preprint arXiv:2306.11979*, 2023.

UN Statistics, "Global indicator framework for the sustainable development goals and targets of the 2030 agenda for sustainable development," *Developmental science and sustainable development goals for children and youth*, 2019, *439*.

Wang, Yuanjia, Haoda Fu, and Donglin Zeng, "Learning optimal personalized treatment rules in consideration of benefit and risk: with an application to treating type 2 diabetes patients with insulin therapies," *Journal of the American Statistical Association*, 2018, *113* (521), 1–13.

World Bank, "Living standards measurement study," 2014.

_ , *PovcalNet: The On-line Tool for Poverty Measurement*, Washington, DC: Development Research Group, World Bank, 2021.

_ , "World Bank Open Data," https://data.worldbank.org/ 2023.

_ , "Poverty and Inequality Platform Methodology Handbook," Technical Report, World Bank July 2025.

_ , "World Bank Poverty and Inequality Platform," https://pip.worldbank.org/ 2025.

Zemel, Eitan, "The linear multiple choice knapsack problem," *Operations Research*, 1980, *28* (6), 1412–1423.

A Algorithms

In this section, we outline algorithms for learning unrestricted policies that minimize the poverty gap index and the poverty rate, as well as binary policies. We also discuss the weakly equitable oracle policy.

A.1 Poverty Gap Minimization

We aim to solve (2) where the objective corresponds to the poverty gap index. In this case, (2) has a convex objective and is straightforward to solve using convex duality. We provide a characterization of the optimal gap-minimizing policy, a high-level overview for learning the optimal gap-minimizing policy in the population regime where we have access to the population distribution F, and finite-sample implementation details.

Lemma 5. *If $L(z) = (c - z)_+$ and $\bar{t} = c$, then the optimal policy that solves (2)*

$$t^*_{\text{gap}}(x) = (c - F^{-1}_{Y|X=x}(\lambda^*(B)))_+. \tag{9}$$

where $\lambda^*(B) \in [0, 1]$.

Theorem 5 implies that the family of optimal gap-minimizing policies, each of which minimizes the poverty gap index for a different budgets B, can be parametrized as

$$t_\lambda(x) = (c - F^{-1}_{Y|X=x}(\lambda))_+$$

for $\lambda \in [0, 1]$. Given this functional form, learning the optimal gap-minimizing policy for a particular budget B boils down to finding the appropriate choice of $\lambda^*(B)$. We can obtain the optimal gap-minimizing policy by discretizing the interval $[0, 1]$ into into a grid of quantiles $\lambda_1, \lambda_2, \dots \lambda_J$ and selecting λ that yields $\mathbb{E}_F[t_\lambda(X)] = B$.

A.1.1 Finite-Data Regime

In realistic settings, we only have access to data sampled from F. In this section, we assume the true data distribution F is unknown but we have access to a training set $\mathcal{D}_{\text{train}} = \{(X_i, Y_i)\}_{i=1}^{n_{\text{train}}}$ where $(X_i, Y_i) \sim F$ i.i.d., and we aim to optimize allocations for unlabeled samples $\mathcal{D}_{\text{test}} = \{X_i\}_{i=1}^{n_{\text{test}}}$ where $X_i \sim F$ i.i.d. We describe the algorithm for estimating the optimal gap-minimizing policy for a budget B.

First, we discretize the interval $[0, 1]$ into a grid of quantiles $\lambda_1, \lambda_2, \dots \lambda_J$. Second, estimate each $q_{\lambda_j}(\cdot)$ for $j = 1, 2 \dots J$ with minimal distributional assumptions on F. Using the training set $\mathcal{D}_{\text{train}} = \{(X_i, Y_i)\}$, we estimate each $\hat{q}_{\lambda_j}(\cdot)$ by solving a conditional quantile regression problem for quantile λ_j via deep learning. In our implementation, for each j, we train a neural network to minimize the pinball loss parametrized by the quantile λ_j (Koenker and Bassett, 1978). Recall that the pinball loss is given by

$$L_{\text{pinball}}(z; \lambda) = \lambda \cdot |z| \cdot \mathbb{I}(z \geq 0) + (1 - \lambda) \cdot |z| \cdot \mathbb{I}(z < 0),$$

and we have that

$$q_\lambda(\cdot) \in \operatorname*{argmin}_q \mathbb{E}_F\left[L_{\text{pinball}}(Y - q(X); \lambda)\right].$$

Our implementation represents q with a neural network and solves the following empirical risk minimization

problem

$$\hat{q}_{\lambda_j}(\cdot) \in \operatorname*{argmin}_q \widehat{\mathbb{E}}_F \left[L_{\text{pinball}}(Y - q(X); \lambda_j) \right].$$

After estimating the conditional quantiles, we use covariates from the test set $\mathcal{D}_{\text{test}} = \{X_i\}_{i=1}^{n_{\text{test}}}$, estimate the optimal transfer amount for each unit in the test set.

$$\hat{t}_{\lambda_j}(X_i) := (c - \hat{q}_{\lambda_j}(X_i))_+.$$

For each λ_j, we can estimate the empirical policy cost. Finally, we can select the choice of λ_j so that $\widehat{\mathbb{E}}_F \left[t_{\lambda_j}(X_i) \right] = \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \hat{t}_{\lambda_i}(X_i)$ is closest to B.

A.2 Poverty Rate Minimization

We aim to solve (2) where the objective corresponds to the poverty rate. In this case, (2) is not straightforward to solve because it has a non-convex objective. Nevertheless, we provide a characterization of the optimal deterministic rate-minimizing policy and outline an algorithm for learning a stochastic policy that is guaranteed to obtain lower or equal poverty rate than the optimal deterministic rate-minimizing policy. We also provide a high-level overview of the algorithm in the population regime where we have access to the population distribution F and finite-sample implementation details.

In this section, we make the following technical assumption that the covariate space $\mathcal{X}$ is finite-dimensional and discrete to permit t to be finite-dimensional because the existence of minimizers of non-convex functionals over infinite-dimensional spaces is not guaranteed.

Assumption 4. The covariate space $\mathcal{X}$ is finite-dimensional and discrete.

We show that when $L(z) := \mathbb{I}(z < c)$, (2) has an optimal deterministic policy, and any optimal deterministic policy has a highly-structured form; it satisfies a property called α-validity for some $\alpha > 0$.

Definition 5. For any $\alpha \geq 0$ and covariate $x \in \mathcal{X}$, the set of α-valid transfers at x is

$$\mathcal{T}_\alpha(x) := \{t \in \mathbb{R} \mid t = 0, t = c, t \text{ such that } t < c \text{ and } f_{Y|X=x}(c - t) = \alpha\}. \tag{10}$$

Using the definition of α-valid transfers, we can now define an α-valid transfer policy.

Definition 6. A transfer policy $t(\cdot)$ is α-valid if $t(x) \in \mathcal{T}_\alpha(x)$ for every $x \in \mathcal{X}$. Note that $t(x)$ may make a potentially randomized choice from the set.

Theorem 6. *Suppose that Assumption 4 holds and $F_{Y|X=x}$ has positive density on $[0, c]$ for all $x \in \mathcal{X}$. If $L(z) := \mathbb{I}(z < c)$, there exists an optimal deterministic policy that solves (2) and any optimal deterministic policy must be α-valid.*

Note that due to the non-convexity of the feasible set, the optimal deterministic policy is not necessarily unique.

To find an optimal deterministic policy, we can recast poverty rate minimization as a two-level optimiziation problem, where the outer optimization sweeps over all possible α values and the inner optimization finds an optimal deterministic α-valid policy among the class of deterministic α-valid policies.

Corollary 7. *Suppose that Assumption 4 holds. If $L(z) := \mathbb{I}(z < c)$, (2) can be re-written as*

$$\min_{\alpha \geq 0} \min_{t:\mathcal{X} \to \mathbb{R}} \{\mathbb{P}_F\left[t(X) + Y < c\right] : \mathbb{E}_F\left[t(X)\right] \leq B \text{ and } t \text{ is } \alpha\text{-valid}\}. \tag{11}$$

Note that the inner optimization problem is always feasible because $t(x) = 0$ satisfies the budget constraint and is α-valid. The inner optimization problem of (11) finds an optimal deterministic α-valid policy among the class of deterministic α-valid policies. We demonstrate that this is equivalent to solving a multiple-choice knapsack problem.

For technical convenience, we assume that for any covariate $x \in \mathcal{X}$, the set of α-valid transfers $\mathcal{T}_\alpha(x)$ contains at most K points for some $K < \infty$.

Assumption 5. There exists $K < \infty$ such that $\sup_{x \in \mathcal{X}, \alpha \in \mathcal{A}} |\mathcal{T}_\alpha(x)| \leq K$.

This assumption requires the conditional distributions $f_{Y|X}$ to have a finite number of modes.

Let $\mathcal{A}$ be the space of plausible α values. We note that $\mathcal{A}$ is bounded because $\alpha \geq 0$ and $\alpha \leq \sup_{x \in \mathcal{X}, y \in \mathcal{Y}} f_{Y|X=x}(y)$. Define $z : \mathcal{X} \times \mathcal{A} \to \mathbb{R}^K$, where $z_k(x; \alpha)$ is the k-th smallest element of $\mathcal{T}_\alpha(x)$ and $z_k(x; \alpha) = c$ for $|\mathcal{T}_\alpha(x)| < k \leq K$. Let $p : \mathcal{X} \times \mathcal{A} \to \mathbb{R}^K$, where $p_k(x; \alpha) = \mathbb{P}_F\left[z_k(X; \alpha) + Y < c \mid X = x\right]$. Note that z is a vector that captures the α-valid transfer amounts, and p is a vector that captures the conditional post-transfer poverty rates at these amounts.

Corollary 8. *Fix $\alpha > 0$. Let $\pi : \mathcal{X} \to [0,1]^K$. Under Assumption 5, the solution to the inner optimization problem in (11) is given by*

$$t(x) = \langle \pi(x), z(x; \alpha) \rangle$$

where π is the solution to a multiple-choice knapsack problem:

$$\begin{aligned} \text{minimize } & \mathbb{E}_F\left[\langle \pi(X), p(X; \alpha) \rangle\right] \\ \text{subject to } & \mathbb{E}_F\left[\langle \pi(X), z(X; \alpha) \rangle\right] \leq B \\ & \langle \pi(x), \mathbf{1} \rangle = 1 \quad \forall x \in \mathcal{X} \\ & \pi_k(x) \in \{0, 1\} \quad \forall x \in \mathcal{X}, k \in [K]. \end{aligned} \tag{12}$$

The program in (12) is a multiple-choice knapsack problem. In the multiple-choice knapsack problem, the goal is to fill a knapsack up to a capacity B by selecting exactly one item out of each class of items, where each item has an associated "loss" and "weight." In our problem, each class corresponds to a unit with covariates x and the items in the class are the α-valid transfers $\mathcal{T}_\alpha(x)$. For clarity of exposition, if $\mathcal{T}_\alpha(x)$ contains less than K elements, we pad the class of items with transfers c in the formulation in (12) until each class consists of K transfers. For item k in the class corresponding to covariates x, the loss is given by the conditional post-transfer poverty rate $p_k(x; \alpha)$ and the weight is given by the the transfer amount $z_k(x; \alpha)$.

Whereas the standard multiple-choice knapsack problem is NP-hard, its fractional relaxation can be solved using a computationally efficient algorithm. We leverage this connection to develop a practical algorithm for solving (12). In the fractional formulation, we permit the selection of fractional amounts of items, which permits fractional values $\pi_k(x) \in [0, 1]$. We emphasize that in the absence of additional assumptions, solving the fractional relaxation of (12) yields a stochastic policy that is guaranteed to obtain a poverty rate less than or equal to the optimal deterministic α-valid policy. We provide conditions under which solution to the fractional relaxation is equal to the solution to the original problem (12) in Appendix A.2.1.

Thus, our overall algorithm for solving (2) where the objective corresponds to the poverty rate relies on the representation provided in (11). We define a grid $[\alpha_1, \alpha_2, \ldots, \alpha_J]$ over $\mathcal{A}$ and solve the fractional knapsack problem for every α_j in the grid to obtain a policy $t_{\alpha_j}(\cdot)$. Finally, we return the policy t^*_{rate} that yields the lowest poverty rate over all t_α for α in the grid.

A.2.1 Algorithm

We provide an algorithm for poverty rate minimization, which leverages a computationally efficient algorithm for solving the fractional multiple-choice knapsack problem.

Input : Conditional distributions $\{f_{Y|X=x}(\cdot)\}_{x\in\mathcal{X}}$,
Covariate distribution $f_X(\cdot)$,
Budget B,
Threshold c,
Grid size m
Output: Transfer policy t^*

costs $\leftarrow \emptyset$
policies $\leftarrow \emptyset$
Define $\alpha_{\max}$ using $\{f_{Y|X=x}(\cdot)\}_{x\in\mathcal{X}}$.
for $\alpha \in [\frac{\alpha_{\max}+\delta}{m}, \frac{2\alpha_{\max}}{m}, \ldots \frac{(m-1)\cdot\alpha_{\max}}{m}, \alpha_{\max}]$ **do**
  for $x \in \mathcal{X}$ **do**
    $\mathcal{T}_\alpha(x) \leftarrow$ `ComputeAlphaValidTransfers`$(f_{Y|X=x}(\cdot), \alpha, c)$
    $\mathcal{C}_\alpha(x), \rho(x;\alpha) \leftarrow$ `ComputeLowerConvexHull`$(f_{Y|X=x}(\cdot), \mathcal{T}_\alpha(x))$
  end
  Apply Algorithm 2
  t_α, cost $\leftarrow$ `SolveFractionalMCKnapsack`$(f_X(\cdot), \{(\mathcal{C}_\alpha(x), \rho(x;\alpha))\}_{x\in\mathcal{X}}, B)$
  costs.append(cost)
  policies.append(t_α)
end
$t^* \leftarrow$ Minimum cost policy in policies (based on costs)
return t^*

Algorithm 1: Two-level optimization procedure for poverty rate minimization.

The outer loop of our algorithm is a grid search over plausible α values. Let $C := \sup_{x\in\mathcal{X}, y\in\mathcal{Y}} f_{Y|X=x}(y)$. We note that for any $\alpha > C$ the set $\mathcal{T}_\alpha(x)$ will only consist of $\{0, c\}$ for all x, meaning that the inner optimization problem is identical for these values of α. As a result, it is sufficient to define $\alpha_{\max} = C + \delta$ for some small $\delta > 0$ and restrict α to $[0, \alpha_{\max}]$. In Algorithm 1, we grid this interval into m values.

The inner loop of our algorithm finds the optimal α-valid policy among the set of stochastic α-valid policies by solving a fractional multiple-choice knapsack problem. When X has continuous support, the solution to the fractional relaxation is almost surely integer-valued, meaning that we recover the optimal deterministic α-valid policy, which solves (12). When X does not have continuous support, the solution may be stochastic. Nevertheless, the solution is guaranteed to have equal or lower policy cost than the optimal deterministic α-valid policy.

We solve the fractional multiple-choice knapsack problem using the computationally-efficient algorithm of Zemel (1980). This procedure is also used by Sverdrup et al. (2023) to solve a cost-constrained treatment allocation problem. The algorithm relies on the key observation that in the optimal solution to the fractional relaxation of (12), the only transfer amounts that are active are the ones that lie on the lower convex hull of the loss-weight plane. For any $x \in \mathcal{X}$, define the lower convex hull of $(z_k(x;\alpha), p_k(x;\alpha))$ for $k = 1, \ldots K$

to be a set $\mathcal{C}_\alpha(x)$ of points with the ordering $k_1(x), \dots k_{|\mathcal{C}_\alpha(x)|}(x)$ such that

$$\begin{aligned}
z_{k_1(x)}(x;\alpha) &< z_{k_2(x)}(x;\alpha) < \cdots < z_{k_{|\mathcal{C}_\alpha(x)|}(x)}(x;\alpha), \\
p_{k_1(x)}(x;\alpha) &> p_{k_2(x)}(x;\alpha) > \cdots > p_{k_{|\mathcal{C}_\alpha(x)|}(x)}(x;\alpha), \\
\rho_{k_1(x)}(x;\alpha) &< \rho_{k_2(x)}(x;\alpha) < \cdots < \rho_{k_{|\mathcal{C}_\alpha(x)|}}(x;\alpha) < 0,
\end{aligned}$$

where ρ is the incremental cost-loss ratio for points on this convex hull as

$$\rho_{k_j(x)}(x;\alpha) := \frac{p_{k_{j+1}(x)}(x;\alpha) - p_{k_j(x)}(x;\alpha)}{z_{k_{j+1}(x)}(x;\alpha) - z_{k_j(x)}(x;\alpha)} \quad k_{j+1}(x), k_j(x) \in \mathcal{C}_\alpha(x) \tag{13}$$

and let $\rho_k(x) := \infty$ if $k \notin \{k_1(x), \dots k_{|\mathcal{C}_\alpha(x)|}(x)\}$.

We define a thresholding rule on the incremental cost-loss ratio (13).

Definition 7 (Thresholding rule, (Sverdrup et al., 2023))**.** Given cost-loss ratios ρ as in (13), a threshold $\lambda \leq 0$, and an interpolation value $\eta \in [0, 1)$, we define a thresholding rule as

$$T_{k_j(x)}(x;\rho,\lambda,\eta) = \begin{cases} 1 & \text{if } \rho_{k_j(x)}(x) < \lambda < \rho_{k_{j+1}(x)}(x) \\ \eta & \text{if } \rho_{k_j(x)}(x) = \lambda \\ 1-\eta & \text{if } \rho_{k_{j+1}(x)}(x) = \lambda \\ 0 & \text{o.w.} \end{cases}.$$

Following from Theorem 1 of (Sverdrup et al., 2023), we have that under Assumption 5, there exists an optimal (stochastic) policy $t_\alpha(x) := \langle z(x;\alpha), \pi(x) \rangle$, where $\pi(x)$ is a thresholding rule defined in Definition 7, i.e. there exists constants $\lambda_\alpha \geq 0$, $\eta_\alpha \in [0, 1)$, such that

$$\pi(x) = T(x; \rho(\cdot;\alpha), \lambda_\alpha, \eta_\alpha).$$

If X has continuous support, then $\mathbb{P}_F[\rho_k(x) = \lambda] = 0$ for all $\lambda > 0$, so $\pi(x)$ is almost surely integer-valued (deterministic).

The fractional multiple-choice knapsack algorithm can be used to obtain the optimal (stochastic) policy. For a detailed description of the algorithm, we refer readers to Algorithm 2. At the start of the procedure, all units are assigned the transfer amount 0, which corresponds to the first point $k_1(x)$ on their convex hulls. Then, we initialize a priority queue. For each $x \in \mathcal{X}$, we add the pair of x and the second point $k_2(x)$ on its convex hull $\mathcal{C}(x)$ to the queue with priority $\rho_{k_1}(x)$. While the post-transfer poverty rate constraint has slack and the queue is nonempty, we pop points from the queue. When a unit with covariates x and point $k_j(x)$ on its convex hull is popped from the queue, the unit is assigned the corresponding transfer value $z_{k_j(x)}(x)$. If the unit has additional points on its convex hull, then the next point on its convex hull is added to the queue with priority equal to its cost-weight ratio $\rho_{k_{j+1}(x)}$. The sequence of updates is dictated by the priority queue ordered by ρ. When solving the problem in the population case, the time-complexity of the algorithm is log-linear in $|\mathcal{X}| \cdot K$. While this time-complexity may seem prohibitively large when the covariate space is continuous, in practice, this algorithm is only applied to units in the test set, so the time-complexity is at most $n_{\text{test}} \cdot K$.

The time complexity of this algorithm is $|\mathcal{X}| \cdot K \log |\mathcal{X}| \cdot K$. The worst-case run time arises when $B > \mathbb{E}_F\left[z_{k_{|\mathcal{C}_\alpha(X)|}}(X)\right]$, so all points on the convex hull for all units will be added and removed from the

priority queue. Since there are at most $|\mathcal{X}| \cdot K$ points added and removed this yields a complexity of $|\mathcal{X}| \cdot K \log |\mathcal{X}| \cdot K$.

Input : Convex hulls and incremental cost-loss ratios $\{(\mathcal{C}(x), \rho(x))\}_{x \in \mathcal{X}}$,
Covariate distribution $f_X(\cdot)$,
Budget B
Output: Transfer policy t, Poverty rate loss

Initialize priority queue.
cost $\leftarrow 0$, loss $\leftarrow 0$
queue $\leftarrow \emptyset$
for $x \in \mathcal{X}$ **do**
  Assign unit to first point on convex hull.
  $k_1(x) \leftarrow$ First point on convex hull $\mathcal{C}(x)$
  cost += $f_X(x) \cdot z_{k_1(x)}(x)$, loss += $f_X(x) \cdot p_{k_1(x)}(x)$
  Enqueue next point on hull.
  $k_2(x) \leftarrow$ Second point on convex hull $\mathcal{C}(x)$
  queue.add($(x, k_2(x))$ with priority $\rho_{k_1}(x)$))
end
while cost $< B$ *and* queue.*size()* > 0 **do**
  $(x, k_j(x)) \leftarrow$ queue.pop()
  Subtract current transfer to unit x from cost *and* loss.
  cost-= $f_X(x) \cdot z_{k_{j-1}(x)}(x)$, loss-= $f_X(x) \cdot p_{k_{j-1}(x)}(x)$
  Allocate transfer $k_j(x)$ to unit x, record new cost *and* loss.
  $t(x) \leftarrow z_{k_j(x)}(x)$
  cost += $f_X(x) \cdot z_{k_j(x)}(x)$, loss += $f_X(x) \cdot p_{k_j(x)}(x)$
  if cost $> \epsilon$ **then**
    Perform fractional adjustment for unit x– updating t, cost, loss.
    break
  end
  if *there remain points on convex hull for unit x* **then**
    $k_{j+1}(x) \leftarrow$ Next point on $\mathcal{C}(x)$
    queue.add($(x, k_{j+1}(x))$ with priority $k_j(x)$)
  end
end
return t, loss

Algorithm 2: Solve fractional multi-choice knapsack.

A.2.2 Finite-Data Regime

Thus far, we have given a characterization of the optimal deterministic policy that minimizes the poverty rate and a practical algorithm for poverty rate minimization. However, in realistic settings, we only have access to data sampled from F. In this section, we assume the true data distribution F is unknown but we have access to a training set $\mathcal{D}_{\text{train}} = \{(X_i, Y_i)\}_{i=1}^{n_{\text{train}}}$ where $(X_i, Y_i) \sim F$ i.i.d., and we aim to optimize allocations for unlabeled samples $\mathcal{D}_{\text{test}} = \{X_i\}_{i=1}^{n_{\text{test}}}$ where $X_i \sim F$ i.i.d.

Algorithm 1 relies on the conditional density function $f_{Y|X}$ to compute the optimal transfer policy. In finite samples, we consider the plug-in estimator for the optimal transfer policy, which is obtained by running Algorithm 1 with an estimator of the conditional density $\hat{f}_{Y|X}$ and the empirical covariate distribution $\hat{f}_X(x) = \frac{1}{n_{\text{test}}} \cdot \mathbb{I}(X_i \in \mathcal{D}_{\text{test}})$ instead of the true conditional and covariate distributions.

To estimate $\hat{f}_{Y|X}$, we apply an extension of Lindsey's method (Efron and Tibshirani, 1996) to the training

data $\mathcal{D}_{\text{train}}$. Lindsey's method is a popular technique for marginal density estimation, and our approach is a straightforward extension of this method for conditional density estimation. To model the conditional densities, we consider the following exponential family of densities on $\mathcal{Y}$

$$h(y \mid x) = h_0(y) \cdot \exp(s(y)^T \Theta x - \psi(\Theta x)). \tag{14}$$

Here, $h_0(y)$ represents a carrier density, $s(y)$ is a J-dimensional vector of sufficient statistics, Θ is a $J \times d$ matrix of parameters, and ψ is the log partition function, the normalizing function that ensures $h(y \mid x)$ integrates to 1 over $\mathcal{Y}$. Many choices of carrier density and sufficient statistics are possible. Following Lindsey's method, we set the carrier density to be a nonparametric estimate of f_Y, the marginal distribution of Y. In addition, we consider a basis of B-Spline functions to form the sufficient statistics.

We compute the conditional density estimator as follows.

1. We estimate f_Y, the marginal density of Y, by computing a kernel density estimate using $\{Y_i\}_{i=1}^{n_{\text{train}}}$. We denote this estimate as $\hat{f}_Y$.

2. Setting $\hat{f}_Y$ as the carrier density in (14), we estimate Θ via maximum likelihood. To estimate the log-likelihood of (14), we first approximate the log-partition function by discretizing $\mathcal{Y}$ into M bins with midpoints $[y_1, \dots y_M]$ and corresponding widths $[\Delta_1, \Delta_2, \dots \Delta_M]$. We note that the log partition function can be approximated by

$$\psi(z) = \log(\int \hat{f}_Y(y) \cdot \exp(s(y)^T z)) \approx \log(\sum_{j=1}^{M} \hat{f}_Y(y_j) \cdot \exp(s(y_j)^T z) \cdot \Delta_j).$$

   Thus, the log-likelihood of (14) can be written as

$$\log h(y_i \mid x_i; \Theta) = s(y_i)^T \Theta x_i - \sum_{j=1}^{M} \hat{f}_Y(y_j) \cdot \exp(s(y_j)^T \Theta x_i) \cdot \Delta_j + c, \tag{15}$$

   where c is a constant factor that does not depend on Θ. Notably, the log-likelihood is concave in Θ, so standard optimization tools can be applied to obtain an estimate of Θ. In our work, we optimize Θ via stochastic gradient descent.

Like Lindsey's method, this is a hybrid approach between a parametric and nonparametric method because it models the relationship between high-dimensional covariates and the density value with a generalized linear model but allows the density function to take on a flexible form by fitting the carrier density nonparametrically.

A.3 Binary Policies

In this section, we solve (2) with the additional constraint that the policy must be binary-valued. In contrast to proxy-means test thresholding as given in Definition 1, here we minimize a particular poverty measure and optimize not only who receives the transfer but also the optimal transfer size.

Learning the optimal binary transfer policy that minimizes a particular loss function L among all possible

binary transfer policy can be framed as the following optimization problem

$$\begin{aligned} &\text{minimize } \mathbb{E}_F\left[L(t(X)+Y))\right] \\ &\text{subject to } \mathbb{E}_F\left[t(X)\right] \leq B \\ &\qquad\qquad t(x) \in \{0, \bar{t}\} \quad \forall x \in \mathcal{X} \\ &\qquad\qquad \bar{t} \in [0, c]. \end{aligned} \tag{16}$$

However, the above optimization problem is non-convex because it involves jointly optimizing over the maximum transfer value $\bar{t}$ and the transfer policy t.

We can rewrite (16) as follows

$$\min_{\bar{t}\in[0,c]} \left\{ \min_{\pi:\mathcal{X}\to\{0,1\}} \mathbb{E}_F\left[L(\bar{t}+Y)\cdot\pi(X) + L(Y)\cdot(1-\pi(X))\right] \text{ subject to } \mathbb{E}_F\left[\pi(X)\right] \leq B/\bar{t} \right\}. \tag{17}$$

Solving the inner minimization of (17) allows us to compute the optimal binary transfer policy among binary transfer policies where the maximum transfer value is a fixed value $\bar{t}$. The inner minimization of (17) is convex and can be viewed as a capacity-constrained classification problem. Let $\pi_{\bar{t}}(\cdot; B)$ be the solution to the inner minimization of (17). The corresponding binary transfer policy can be defined as $t_{\bar{t}}(\cdot; B) := \bar{t}\cdot\pi_{\bar{t}}(x; B)$.

The function $\pi_{\bar{t}}(\cdot; B)$ identifies the units who will benefit most (in terms of conditional loss reduction) from the $\bar{t}$-valued transfer. We can define the conditional loss reduction function for a $\bar{t}$-valued transfer

$$\rho_{\bar{t}}(x) := \mathbb{E}_F\left[L(\bar{t}+Y) - L(Y) \mid X = x\right]. \tag{18}$$

For poverty measures, L is decreasing so $\rho_{\bar{t}}(x) \leq 0$ for any $\bar{t} \in \mathbb{R}_+$.

Following from Theorem 1 from Sun et al. (2021), there are constants $\rho^*_{\bar{t},B} \in \mathbb{R}, a_{T,B} \in [0,1]$, such that the optimal solution to the inner minimization of (17) has the form

$$\pi_{\bar{t}}(x; B) = \begin{cases} 1 & \rho_{\bar{t}}(x) < \rho^*_{\bar{t},B} \\ a_{\bar{t},B} & \rho_{\bar{t}}(x) = \rho^*_{\bar{t},B} \\ 0 & \rho_{\bar{t}}(x) > \rho^*_{\bar{t},B}, \end{cases} \tag{19}$$

where either $\rho^*_{\bar{t},B} = a_{\bar{t},B} = 0$ (i.e., we have sufficient budget to treat all units) or $\rho^*_{\bar{t},B} < 0$ and the pair $(\rho^*_{\bar{t},B}, a_{\bar{t},B})$ is the unique pair for which the policy cost is B. If X_i has continuous density, $\mathbb{P}_F\left[\rho_T(X) = \rho^*_{T,B}\right] = 0$ and the policy $\pi_{\bar{t}}(\cdot; B)$ is both deterministic and the unique optimal solution.

We operationalize this result to provide an algorithm for optimal binary policies. Our overall algorithm first discretizes the interval $[0, c]$ into a grid of possible maximum transfer amounts $[\bar{t}_1, \bar{t}_2, \ldots \bar{t}_J]$. For each $j = 1, 2, \ldots J$, we compute $\pi_{\bar{t}_j}(\cdot; B)$ by ranking units by $\rho_{\bar{t}_j}(\cdot)$ in increasing order and allocating transfers $\bar{t}_j$ sequentially until the budget B is exhausted. We set the optimal binary policy t^*_{binary} under loss function L to be the policy that yields the lowest expected loss over the grid, i.e.

$$t^*_{\text{binary}}(x; B) = \bar{t}_{j^*} \cdot \pi_{\bar{t}_{j^*}}(x; B), \text{ where } j^* = \operatorname*{argmin}_{j\in[J]} \mathbb{E}_F\left[L(t_{\bar{t}_j}(X_i; B) + Y_i)\right].$$

We note that this procedure will yield the optimal binary-valued policy under any decreasing poverty measure L and does not require the loss function L to be convex. As a result, we will use this approach to learn optimal binary-valued policies under both the poverty rate and the poverty gap index.

A.3.1 Finite-Data Regime

Again, we consider the setting where we only have access to data sampled from F. In this section, we assume the true data distribution F is unknown but we have access to a training set $\mathcal{D}_{\text{train}} = \{(X_i, Y_i)\}_{i=1}^{n_{\text{train}}}$ where $(X_i, Y_i) \sim F$ i.i.d., and we aim to optimize allocations for unlabeled samples $\mathcal{D}_{\text{test}} = \{X_i\}_{i=1}^{n_{\text{test}}}$ where $X_i \sim F$ i.i.d. We use the following procedure to estimate the optimal binary policy for a budget B.

First, we discretize the interval $[0, c]$ into a grid of possible maximum transfer amounts $\bar{t}_1, \bar{t}_2, \dots \bar{t}_J$. Second, we estimate nuisance parameters $\rho_{\bar{t}_j}(\cdot)$ for $j = 1, 2, \dots J$. Using the training set $\mathcal{D}_{\text{train}} = \{(X_i, Y_i)\}_{i=1}^{n_{\text{train}}}$, we define pseudo-labels

$$Z_i = L(\bar{t} + Y_i) - L(Y_i)$$

for $i = 1, 2, \dots, n_{\text{train}}$. Estimate $\hat{\rho}_{\bar{t}_j}(\cdot)$ via deep learning using $\mathcal{D}_{\text{train}}$ and pseudolabels. Observe that

$$\rho_{\bar{t}} \in \operatorname*{argmin}_{\rho} \mathbb{E}_F \left[(Z_i - \rho(X_i))^2 \right].$$

Our implementation represents $\rho_{\bar{t}}$ with a neural network and solves the following empirical risk minimization problem

$$\hat{\rho}_{\bar{t}_j}(\cdot) \in \operatorname*{argmin}_{\rho} \widehat{\mathbb{E}}_F \left[(Z_i - \rho(X_i))^2 \right].$$

Now, we must learn the optimal transfer size $\bar{t}_j$. For each $j = 1, 2, \dots J$, create a ranked list of units by sorting $\hat{\rho}_{\bar{t}_j}(X_i)$ in increasing order for X_i in the training set. We can estimate the threshold $\rho^*_{\bar{t}_j, B}$ by allocating transfers to the $B/\bar{t}_j$-fraction of units at the top of the ranked list for each $j = 1, 2, \dots J$. After that, we can form $\hat{t}_{\bar{t}_j}(\cdot; B)$ using the estimated conditional loss reduction function $\hat{\rho}_{\bar{t}_j}(\cdot)$ and estimated threshold $\hat{\rho}^*_{\bar{t}_j, B}$. Then, we can select

$$\hat{j}^* \in \operatorname*{argmin}_{j \in [J]} \widehat{\mathbb{E}}_F \left[L(\hat{t}_{\bar{t}_j}(X_i; B) + Y_i) \right].$$

We assign transfers to units in the test set $\mathcal{D}_{\text{test}}$ by ranking units in the test set by $\hat{\rho}_{j^*}$ in increasing order and allocating $\bar{t}_{\hat{j}^*}$-valued transfers to units starting at the top of the ranked list until we hit the budget constraint B, which ensures that the budget constraint is exactly satisfied on the test set.

A.4 Weakly Equitable Oracle Policy

We highlight the behavior of the weakly equitable oracle policy. Suppose that Assumption 1 holds, $X_i = Y_i$ (i.e. perfect information), and F_Y has continuous support on a compact set $\mathcal{Y} \subset \mathbb{R}_+$. Let $t(x; B)$ be a policy family that satisfies $\mathbb{E}_F [t(Y_i; B)] \leq B$ for all $B \in \mathbb{R}_+$. If t is weakly equitable with respect to F then

$$t(y; B) = (\lambda(B) - y)_+,$$

where $\lambda(B)$ is monotone increasing in B and $\mathbb{E}_F [(\lambda(B) - Y_i)_+] \leq B$. The requirement that t is weakly equitable ensures that the incremental transfer is allocated to the worst-off group. In other words, the budget is first spent on allocating a transfer to the poorest unit in terms of pre-transfer consumption to raise them to the pre-transfer consumption of the second-poorest. After equalizing the consumption of the poorest and second-poorest, the remaining budget is spent on transfers to the poorest and second-poorest units in terms of pre-transfer consumption to raise them to the level of the third-poorest, and so on.

We realize that if t allocates transfers based on perfect information on consumption and is weakly equitable in the sense of Definition 4, then t can only obtain a reduction in the post-transfer poverty rate by lifting all

units with $Y_i < c$ above the poverty line. This is because the weakly equitable oracle policy first equalizes consumption among units who lie below the poverty line, so a reduction in the poverty rate will not be observed until we have enough budget to close the poverty gap, i.e. $B \geq \mathbb{E}_F\left[(c - Y_i)_+\right]$.

B Proofs

B.1 Additional Lemmas

We define new notation for the following lemmas. Define

$$\begin{aligned} G_x(t) &= -\mathbb{E}_F\left[L(Y+t) \mid X = x\right] \\ G'_x(t) &= -\frac{d}{dt}\mathbb{E}_F\left[L(Y+t) \mid X = x\right]. \end{aligned} \tag{20}$$

Lemma 9. *Suppose Assumption 1, 2 hold. Let L be decreasing, strictly convex, differentiable, bounded below by $C > -\infty$. Then, $G_x(t)$ is differentiable on $\mathbb{R}_+$, strictly concave, and twice differentiable on $\mathbb{R}_+$.*

Lemma 10. *Suppose Assumption 1, 2 hold. Let L denotes a FGT index with $\alpha \geq 1$. Then, $G_x(t)$ is differentiable on $\mathbb{R}_+$, strictly concave on $[0, c]$, and twice differentiable on $[0, c]$*

B.2 Proof of Theorem 1

Since L is piecewise continuously differentiable, we say that there exists a countable collection of intervals $\{I_i\}$ that cover $\mathbb{R}_+$, and on each interval $\{I_i\}$, L is continuously differentiable. We can say that a piecewise continuously differentiable function L is nonconvex if L violates convexity on the interior of an interval I on which it is continuously differentiable or if L has a jump discontinuity or a corner between two intervals I_1, I_2 and is convex on I_1 and I_2.

B.2.1 Convexity Violation on Continuously Differentiable Region

We consider the first case, where L has a convexity violation on a region where L is continuously differentiable. Our goal is to construct a population distribution F with the following properties:

1. F satisfies Assumption 2
2. F is a mixture distribution over two types. $F_X = \frac{1}{2} \cdot \mathbb{I}(X = x_1) + \frac{1}{2} \cdot \mathbb{I}(X = x_2)$, $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$
3. Define $\ell_i(t) := \mathbb{E}_{F_{Y|X=x_i}}\left[L(t+Y) \mid X = x_i\right]$ and $\ell'_i(t)$. We require that F yields $\ell'_1(t_1) > \ell'_2(t_2)$ for every $t_1, t_2 \in [0, \Delta]$ for some $\Delta > 0$.

Then, we will show that the optimal policy under F is not monotone decreasing in consumption.

Since L is non-convex and continuously differentiable region there exists some open interval $I \subset \mathbb{R}$ where L' is strictly decreasing. Pick $\mu_1 \in I$. Since L' is continuous, for any $\epsilon > 0$, there exists $\delta_1(\epsilon) > 0$, so that $|L'(\mu_1) - L'(z)| < \epsilon$ for $|\mu_1 - z| < \delta_1(\epsilon)$. We choose ϵ_1 sufficiently small so that the corresponding $B_{\delta_1(\epsilon_1)}(\mu_1) \subset I$. Next, we pick $\mu_2 \in I$ such that $L'(\mu_2) < L'(\mu_1) - \epsilon_1$, which is possible because L' is strictly decreasing on I. Again, by the continuity of L', for any ϵ, there exists $\delta_2(\epsilon) > 0$, so that $|L'(\mu_2) - L'(z)| < \epsilon$ for $|\mu_2 - z| < \delta_2(\epsilon)$. We can choose ϵ_2 sufficiently small so that $L'(\mu_2) + \epsilon_2 < L'(\mu_1) - \epsilon_1$.

Let $\epsilon = \min_{i\in\{1,2\}} \epsilon_i$. It is straightforward to see that

$$L'(\mu_2) + \epsilon < L'(\mu_1) - \epsilon. \tag{21}$$

Let $\delta = \min(\delta_1(\epsilon), \delta_2(\epsilon))$. For $z_i \in B_\delta(\mu_i)$, we have that $|L'(z_i) - L'(\mu_i)| < \epsilon$, so

$$L'(z_2) < L'(\mu_2) + \epsilon < L'(\mu_1) - \epsilon < L'(z_1).$$

This implies that for $y_1 \in B_{\frac{\delta}{2}}(\mu_1), t \in [0, \frac{\delta}{2}]$, we have that

$$L'(y_1 + t) \geq L'(\mu_1) - \epsilon \tag{22}$$

and for $y_2 \in B_{\frac{\delta}{2}}(\mu_2), t \in [0, \frac{\delta}{2}]$, we have that

$$L'(y_2 + t) \leq L'(\mu_2) + \epsilon. \tag{23}$$

Define

$$L'_{\min} := \inf_{y\in\mathcal{Y}, t\in[0,\frac{\delta}{2}]} L'(t+y)$$
$$L'_{\max} := \sup_{y\in\mathcal{Y}, t\in[0,\frac{\delta}{2}]} L'(t+y).$$

Let

$$q := \frac{L'(\mu_1) - L'(\mu_2) - 2\epsilon}{L'_{\max} - L'_{\min} + L'(\mu_1) - L'(\mu_2) - 2\epsilon}. \tag{24}$$

Note that by (21), $L'(\mu_1) - L'(\mu_2) - 2\epsilon > 0$, so $0 < q < 1$.

We can define F as follows. Let F be a mixture distribution over two types x_1, x_2, where $F_X = \frac{1}{2}\cdot\mathbb{I}(X = x_1) + \frac{1}{2}\cdot\mathbb{I}(X = x_2)$. Let $F_{Y|X=x_i}$ are truncated Gaussians on $\mathcal{Y}$ with mean μ_i and scale parameter $\sigma > 0$ chosen so that the distributions are sufficiently concentrated about their means. In particular, we set σ so that $\mathbb{P}_F[Y \notin B_\delta(\mu_i) \mid X = x_i] < q$. This population distribution F must satisfy criteria (a), (b) above. It remains to show that this population distribution satisfies criterion (c).

Given a conditional distribution $F_{Y|X=x_1}$ and $t_1 \in [0, \frac{\delta}{2}]$, we have that

$$\begin{aligned}
\ell'_1(t_1) &= \mathbb{E}_{F_{Y|X=x_1}}[L'(t_1 + Y) \mid X = x_1] \\
&= \int_{B_{\frac{\delta}{2}}(\mu_1)} L'(t_1+y) f_{Y|X=x_1}(y) + \int_{\mathcal{Y}\setminus B_{\frac{\delta}{2}}(\mu_1)} L'(t_1+y)\cdot f_{Y|X=x_1}(y) \\
&\geq (L'(\mu_1) - \epsilon)\cdot \mathbb{P}_{F_{Y|X=x_1}}\left[Y \in B_{\frac{\delta}{2}}(\mu_1) \mid X = x_1\right] + L'_{\min}\cdot \mathbb{P}_{F_{Y|X=x_1}}\left[Y \notin B_{\frac{\delta}{2}}(\mu_1)\right] \\
&> (L'(\mu_1) - \epsilon)\cdot(1-q) + L'_{\min}\cdot q.
\end{aligned}$$

In addition, given a conditional distribution $F_{Y|X=x_1}$ and $t_2 \in [0, \frac{\delta}{2}]$, we have that

$$\begin{aligned}
\ell_2'(t_2) &= \mathbb{E}_{F_{Y|X=x_1}}\left[L'(t_2+Y) \mid X = x_2\right] \\
&= \int_{B_{\frac{\delta}{2}}(\mu_2)} L'(t_2+y) f_{Y|X=x_2}(y) + \int_{\mathcal{Y}\setminus B_{\frac{\delta}{2}}(\mu_2)} L'(t_2+y)\cdot f_{Y|X=x_2}(y) \\
&\leq (L'(\mu_2)+\epsilon)\cdot \mathbb{P}_{F_{Y|X=x_2}}\left[Y \in B_{\frac{\delta}{2}}(\mu_2) \mid X = x_2\right] + L'_{\max}\cdot \mathbb{P}_{F_{Y|X=x_1}}\left[Y \notin B_{\frac{\delta}{2}}(\mu_2)\right] \\
&< (L'(\mu_2)+\epsilon)\cdot(1-q) + L'_{\max}\cdot q.
\end{aligned}$$

Plugging in q as defined in (24) yields that the lower bound on $\ell_1'(t_1)$ equals the upper bound on $\ell_2'(t_2)$ for $t_1, t_2 \in [0, \Delta]$. As a result, for $t_1, t_2 \in [0, \Delta]$, where $\Delta = \frac{\delta}{2}$ we have provided a population distribution F has the property that $\ell_1'(t_1) > \ell_2'(t_2)$ as desired.

For this choice of population distribution F and $B = \Delta$, the optimal policy that solves (2) is $t^*(x) = 0\cdot\mathbb{I}(X = x_1) + \Delta\cdot\mathbb{I}(X = x_2)$. Note that this policy is non-monotone increasing in consumption because $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$ but $t^*(x_2) > t^*(x_1)$.

To see this, we note that ℓ is concave on the feasible set. The minimizer of a concave function subject must occur on an extremal point of the feasible set. The two extremal points are the policy t^* and the policy $\tilde{t}(x) = \Delta\cdot\mathbb{I}(X = x_1) + 0\cdot\mathbb{I}(X = x_2)$. It suffices to show that

$$\mathbb{E}_F\left[\tilde{t}(X)+Y)\right] > \mathbb{E}_F\left[t^*(X)+Y)\right].$$

Since we have that

$$\ell_1'(t_1) > \ell_2'(t_2)$$

for $t_1, t_2 \in [0, \Delta]$. By the Mean Value Theorem, we must have that

$$(\ell_1(\Delta) - \ell_1(0))\cdot\Delta > (\ell_2(\Delta) - \ell_2(0))\cdot\Delta.$$

Multiplying both sides by $\frac{1}{2}$ and rearranging, we have that

$$\frac{1}{2}\cdot(\ell_1(\Delta) + \ell_2(0)) > \frac{1}{2}\cdot(\ell_1(0) + \ell_2(\Delta)).$$

This implies that

$$\mathbb{E}_F\left[L(\tilde{t}(X)+Y)\right] > \mathbb{E}_F\left[L(t^*(X)+Y)\right].$$

B.2.2 Convexity Violation due to Corner

As in the previous case, our goal is to construct a population distribution F with the following properties:

1. F satisfies Assumption 2
2. F is a mixture distribution over two types. $F_X = \frac{1}{2}\cdot\mathbb{I}(X = x_1) + \frac{1}{2}\cdot\mathbb{I}(X = x_2)$, $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$
3. Define $\ell_i(t) := \mathbb{E}_{F_{Y|X=x_i}}\left[L(t+Y) \mid X = x_i\right]$ and $\ell_i'(t)$. We require that F yields $\ell_1'(t_1) > \ell_2'(t_2)$ for every $t_1, t_2 \in [0, \Delta]$ for some $\Delta > 0$.

Then, we will show that the optimal policy under F is not monotone decreasing in consumption.

Suppose L is nonconvex about a corner z_0 between two intervals $I_1 = [a, z_0)$ and $I_2 = [z_0, b)$ on which it is convex and continuously differentiable. Let $L'_-(z)$ denote the left derivative of L at z and $L'_+(z)$ denote the right derivative at z. A point z_0 is a corner if $L'_-(z_0) \neq L'_+(z_0)$. We observe that L is nonconvex about a corner z_0 if $L'_-(z_0) > L'_+(z_0)$.

Let $J = L'_-(z_0) - L'_+(z_0)$. Let $\epsilon = J/3$. Since L' is continuously differentiable on $[a, z_0)$ and on (z_0, b), we have that there exists $\delta_1 > 0$ so that for $z \in \mathcal{N}_1$, where $\mathcal{N}_1 = \{z \mid z_0 - z < \delta_1\}$, we have $|L'(z) - L'_-(z_0)| < \epsilon$ and for $z \in \mathcal{N}_2$, where $\mathcal{N}_2 : \{z \mid z - z_0 < \delta_2\}$, we have $|L'(z) - L'_+(z_0)| < \epsilon$. As a result, we have that for any $z_1 \in \mathcal{N}_1, z_2 \in \mathcal{N}_2$, $L'(z_1) > L'(z_2)$. We can pick $\Delta = \min(\delta_1, \delta_2)/8$ and points $y_1 \in \mathcal{N}_1$ and $y_2 \in \mathcal{N}_2$ so that $[y_i - \Delta, y_i + \Delta] \subset \mathcal{N}_i$ for $i = 1, 2$.

We can set $F_{Y|X=x_i} = \text{Uniform}(y_i - \Delta, y_i)$. Clearly, we have that $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$. We find that $\ell'_1(t_1) > \ell'_2(t_2)$ for any $t_1, t_2 \in [0, \Delta]$, as a result, the optimal transfer policy will allocate higher transfers to x_2 and is thus not weakly equitable.

B.2.3 Convexity Violation due to Jump Discontinuity

Our goal is to construct a population distribution F with the following properties:

1. F satisfies Assumption 2
2. F is a mixture distribution over two types. $F_X = \frac{1}{2} \cdot \mathbb{I}(X = x_1) + \frac{1}{2} \cdot \mathbb{I}(X = x_2)$, $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$
3. Define $\ell_i(t) := \mathbb{E}_{F_{Y|X=x_i}}[L(t + Y) \mid X = x_i]$. We require that F yields $\ell_2(0) - \ell_2(t) > J$ for any $t \in [\Delta/2, \Delta]$ and $\ell_1(0) - \ell_1(t) < J$ for all $t \in [0, \Delta]$.

Then, we will show that the optimal policy under F is not monotone decreasing in consumption.

Suppose L has a jump discontinuity at z_0 between $I_1 = [a, z_0)$ and $I_2 = [z_0, b)$, where L is continuously differentiable on I_1, I_2. Let $L_-(z)$ denote the left limit of L at z and let $L_+(z)$ denote the right limit of L at z. Let $J := L_-(z_0) - L_+(z_0)$.

We can pick $y_1 < z_0$ so that $L(y_1) - L(z_0) < J/2$. By continuity, we can also pick $\delta > 0$ so that for $z - y_1 < \min(\delta, \frac{z_0 - y_1}{4})$, we have that $|L(z) - L(y_1)| < J/4$. Thus, for all $z \in [y_1, y_1 + \delta]$, we have that $L(y_1) - L(z_0) < \frac{3J}{4}$. Then we can set $\Delta = \frac{z_0 - y_1 - \delta}{2}$. Finally, we can pick $y_2^-, y_2^+ \in (z_0 - \Delta, z_0 - \frac{\Delta}{2})$ such that $y_2^- < y_2^+$. Let $F_{Y|X=x_2} = \text{Uniform}[y_2^-, y_2^+]$ and let $F_{Y|X=x_1} = \text{Uniform}[y_1, y_1 + \delta]$. Then, we have that $\ell_1(0) - \ell_1(t) < J$ for $t \in [0, \Delta]$ and $\ell_2(0) - \ell_2(t) > J$ for all $t \in [\Delta/2, \Delta]$. Thus, the optimal policy is not weakly equitable.

B.3 Proof of Theorem 2

We show that if $B > 0$, then (2) has a unique solution with form given by

$$t_\lambda(x) = ((G'_x)^{-1}(\lambda))_+, \tag{25}$$

where $\lambda \geq 0$ is the unique choice that satisfies $\mathbb{E}_F[t_\lambda(X)] = B$.

We can dualize the budget constraint of (2) as follows. Let

$$\mathcal{L}(t, \lambda) = \mathbb{E}_F[L(Y + t(X))] + \lambda(\mathbb{E}_F[t(X)] - B) = -\mathbb{E}_F[G_X(t(X))] + \lambda(\mathbb{E}_F[t(X)] - B). \tag{26}$$

Define the dual function as

$$h(\lambda) = \min_{t:\mathcal{X}\to\mathbb{R}_+} \{\mathcal{L}(t,\lambda) : t(x) \geq 0 \quad \forall x \in \mathcal{X}\}. \tag{27}$$

Then the dual problem is given by

$$\max_{\lambda\geq 0} h(\lambda). \tag{28}$$

To compute $h(\lambda)$, we solve the optimization problem on the right side of (27)

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \{-\mathbb{E}_F\left[G_X(t(X))\right] + \lambda(\mathbb{E}_F\left[t(X)\right] - B) : t(x) \geq 0 \quad \forall x \in \mathcal{X}\}. \tag{29}$$

Importantly, the constraints of the above program are separable, so the policy that solves (29) satisfies the KKT conditions for each $x \in \mathcal{X}$. Suppose that the first set of conditions on L hold, i.e. L is decreasing, strictly convex, differentiable, bounded below by a constant $C > -\infty$. In this case, we can apply Lemma 9 to see that $G_x(t)$ is differentiable and is strictly concave for every $x \in \mathcal{X}$. This implies that the optimal solution to this program is unique. To characterize the optimal solution, we note that G_x' is strictly decreasing so is invertible.

Suppose that the second set of conditions on L hold, i.e. L is an FGT index with $\alpha \geq 1$. In this case, we can apply Lemma 10 to see that $G_x(t)$ is differentiable on $(0,\infty)$ and strictly concave on $(0,c)$ for every $x \in \mathcal{X}$. Since L is an FGT index, the solution to (29) is equivalent to the solution to

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \{-\mathbb{E}_F\left[G_X(t(X))\right] + \lambda(\mathbb{E}_F\left[t(X)\right] - B) : 0 \leq t(x) < c \quad \forall x \in \mathcal{X}\} \tag{30}$$

because $L(z) = 0$ for $z \geq c$. We observe that $G_x(t)$ is differentiable and is strictly concave on this region for every $x \in \mathcal{X}$. As a result, the optimal solution to the program is unique. We also note that G_x' is strictly decreasing on $[0,c]$ so is invertible on this region.

In both cases, the unique solution will satisfy the first-order stationarity condition of this program or lie on the boundary, i.e. $t_\lambda(x) = 0$. The first-order stationarity condition is

$$G_x'(t(x)) = \lambda. \tag{31}$$

The optimal solution to this program has form (25).

We note that $h(\lambda)$ is then given by

$$h(\lambda) = -\mathbb{E}_F\left[G_X(t_\lambda(X))\right] + \lambda(\mathbb{E}_F\left[t_\lambda(X)\right] - B). \tag{32}$$

Furthermore, we argue that a unique choice of λ determines the solution to the program. To see this, note that the derivative of h is

$$\begin{aligned}
\frac{dh}{d\lambda} &= -\mathbb{E}_F\left[G_X'(t_\lambda(X)) \cdot \frac{d}{d\lambda} t_\lambda(X)\right] + \lambda \mathbb{E}_F\left[\frac{d}{d\lambda} t_\lambda(X)\right] + (\mathbb{E}_F\left[t_\lambda(X)\right] - B) \\
&= -\mathbb{E}_F\left[G_X'((G_X')^{-1}(\lambda)) \cdot \left(\frac{1}{G_X''((G_X')^{-1}(\lambda))} \cdot \mathbb{I}(\lambda < G_X'(0))\right)\right] \\
&\quad + \lambda \cdot \mathbb{E}_F\left[\frac{1}{G_X''((G_X')^{-1}(\lambda))} \cdot \mathbb{I}(\lambda < G_X'(0))\right] + (\mathbb{E}_F\left[t_\lambda(X)\right] - B) \\
&= \mathbb{E}_F\left[t_\lambda(X)\right] - B.
\end{aligned} \tag{33}$$

The second derivative of the dual objective is given by

$$\frac{d^2h}{d\lambda^2} = \mathbb{E}_F\left[\frac{1}{G''_X(t_\lambda(X))} \cdot \mathbb{I}(\lambda < G'_X(0))\right].$$

Under the first set of conditions on L, i.e. that L is decreasing, strictly convex, differentiable, and bounded below by a constant $C > -\infty$, we can again apply Lemma 9 to see that G_x is twice-differentiable and strictly concave, so $G''_x(t_\lambda(x)) > 0$. This implies that h is strictly concave on

$$\Lambda = \{\lambda \geq 0 \mid \exists x \in \mathcal{X} \text{ such that } \lambda < G'_x(0)\}. \tag{34}$$

We can show that if $B > 0$, then any optimal λ^* that solves (28) must lie in Λ, where Λ is defined in (34). Suppose for the sake of contradiction that $\lambda^* \notin \Lambda$. Since λ^* is an optimal solution to the dual then it must satisfy the first-order stationarity condition, so we must have that $\mathbb{E}_F[t_{\lambda^*}(X)] = B$. We note that if $\lambda^* \notin \Lambda$, then $\lambda^* > G'_x(0)$ for all $x \in \mathcal{X}$. This implies that $t_{\lambda^*}(x) = 0$. However, this contradicts the first-order stationarity condition. Thus, if $B > 0$, the optimal λ^* lies in Λ. We note that h is continuous and Λ is compact, so h has at least one maximizer on Λ. In addition, h is strictly concave on Λ and Λ is convex, so h must have at most one maximizer on Λ. Thus, if $B > 0$, then (28) has a unique solution λ^* that lies in Λ and the optimal policy that solves (2) is unique and given by (25) with this choice of λ^*.

Similarly, under the second set of conditions on L, we can apply Lemma 10 to see that G_x is twice-differentiable on $[0, c]$ and strictly concave on $[0, c]$. When L is an FGT index for $\alpha \geq 1$, we must have that $t_\lambda(x) \in [0, c]$, so $G''_x(t_\lambda(x)) > 0$. As above, we can similarly show that for $B > 0$, the optimal λ^* that solves (28) must lie in Λ and h is strictly concave on Λ, so λ^* must be unique.

Now, we verify the two properties of weak equity.

B.3.1 Transfer is monotone increasing in budget.

We can characterize how the optimal solution $\lambda(B)$ that determines the optimal solution (25) varies with B. By Lemma 9, G_x is strictly concave, so G'_x must be a decreasing function. This implies that $(G'_x)^{-1}$ is also a decreasing function. So, $t_\lambda(x)$ as defined in (25) is decreasing in λ. This means that as B increases, the choice of $\lambda(B)$ for which $\mathbb{E}_F[t_\lambda(X)] = B$ decreases. Thus, $t_{\lambda(B)}(x)$ is monotone increasing in B.

B.3.2 Incremental transfer is monotone decreasing in post-transfer consumption.

First, we show that the simpler property in (7) holds. Second, we show that the optimal policy that is obtained by solving (2) with budget B is the same as the cumulative optimal policy that is obtained by summing the optimal policy $t(x; B')$ that solves (2) with a small budget B' and the policy $\Delta t(x; B - B')$ that solves (2) with the remaining budget $B - B'$ after administering the transfers $t(x; B')$. In other words,

$$t(x; B) = t(x; B') + \Delta t(x; B - B').$$

The first result implies that

$$F_{Y+t(X;B')|X=x} \preceq_{\text{SD}} F_{Y+t(X;B')|X=x'} \implies \Delta t(x; B - B') \geq \Delta t(x'; B - B').$$

The second result implies that

$$\Delta t(x; B - B') = t(x; B) - t(x; B')$$

for all $x \in \mathcal{X}$. Taken together, these two results imply that (6).

Transfer is monotone decreasing in consumption. We show that $t_\lambda(x) \geq t_\lambda(x')$ for any $\lambda \geq 0$. Since the optimal policy has the form in (25) for some choice of λ, then the optimal policy must be monotone decreasing in $Y \mid X = x$.

We note that L' exists and is monotone increasing because L is decreasing and convex. As a result, if $Y \mid X = x \preceq_{\text{SD}} Y \mid X = x'$, then $\mathbb{E}_F\left[L'(Y+t) \mid X = x\right] \leq \mathbb{E}_F\left[L'(Y+t) \mid X = x'\right]$. This implies that $G'_x(t) \geq G'_{x'}(t)$ for any $t \in \mathbb{R}_+$. Since G'_x is decreasing in t, then $(G'_x)^{-1}$ is also decreasing. Since $G'_x(t) \geq G'_{x'}(t)$ for all t and both are decreasing in t, then we also have that $(G'_x)^{-1}(\lambda) \geq (G'_{x'})^{-1}(\lambda)$. Thus, we have that $t_\lambda(x) \geq t_\lambda(x')$. Thus, the optimal solution to (2) is monotone decreasing in consumption.

Cumulative policy is equal to sum of incremental policies. It is straightforward to see that the optimal policy under the former is given by $t_{\lambda(B)}$, where t_λ is given by (25) with $\lambda(B)$ being the unique choice of λ that satisfies $\mathbb{E}_F\left[((G'_X)^{-1}(\lambda))_+\right] = B$.

The optimal policy under budget B' is similarly given by $t_{\lambda(B')}$, where $\lambda(B)$ being the unique choice of λ that satisfies $\mathbb{E}_F\left[((G'_X)^{-1}(\lambda))_+\right] = B'$. We now consider solving

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \{\mathbb{E}_F\left[L(t(X) + t_{\lambda(B')}(X) + Y)\right] : \mathbb{E}_F\left[t(X)\right] \leq B - B',\, t(x) \geq 0 \quad \forall x \in \mathcal{X}\}. \tag{35}$$

Let

$$\begin{aligned}
\tilde{G}_x(t) &= G_x(t + t_{\lambda(B')}(x)),\\
\tilde{G}'_x(t) &= G'_x(t + t_{\lambda(B')}(x))\\
\tilde{t}_\lambda(x) &= ((\tilde{G}'_x)^{-1}(\lambda))_+.
\end{aligned}$$

The optimal policy that solves (35) is given by

$$\tilde{t}_\lambda(x) = ((\tilde{G}'_x)^{-1}(\lambda))_+$$

for some unique choice of λ that satisfies $\mathbb{E}_F\left[\tilde{t}_\lambda(X)\right] = B - B'$. We observe that

$$\tilde{t}_\lambda(x) = ((G'_x)^{-1}(\lambda) - t_{\lambda(B')}(x))_+. \tag{36}$$

We show that λ that satisfies $\mathbb{E}_F\left[\tilde{t}_\lambda(X)\right] = B - B'$ must be equal to $\lambda(B)$. We note $B > B'$, so $(G'_x)^{-1}(\lambda(B')) < (G'_x)^{-1}(\lambda(B))$. We consider three cases

1. $(G'_x)^{-1}(\lambda(B)) > (G'_x)^{-1}(\lambda(B')) > 0$. In this case,

$$\begin{aligned}
t_{\lambda(B')}(x) + \tilde{t}_{\lambda(B)}(x) &= t_{\lambda(B')}(x) + (G'_x)^{-1}(\lambda(B)) - t_{\lambda(B')}(x))\\
&= (G'_x)^{-1}(\lambda(B))_+.
\end{aligned}$$

2. $(G'_x)^{-1}(\lambda(B)) > 0 > (G'_x)^{-1}(\lambda(B'))$. In this case,

$$\begin{aligned}
t_{\lambda(B')}(x) + \tilde{t}_{\lambda(B)}(x) &= 0 + (G'_x)^{-1}(\lambda(B))\\
&= ((G'_x)^{-1}(\lambda(B)))_+.
\end{aligned}$$

3. $0 > (G'_x)^{-1}(\lambda(B)) > (G'_x)^{-1}(\lambda(B'))$. In this case,

$$\begin{aligned} t_{\lambda(B')}(x) + \tilde{t}_{\lambda(B)}(x) &= 0 + 0 \\ &= ((G'_x)^{-1}(\lambda(B)))_+. \end{aligned}$$

Rearranging, we have that $\tilde{t}_{\lambda(B)}(x) = ((G'_x)^{-1}(\lambda(B)))_+ - t_{\lambda(B')}(x)$. To conclude, we note that

$$\begin{aligned} \mathbb{E}_F\left[\tilde{t}_{\lambda(B)}(X)\right] &= \mathbb{E}_F\left[((G'_x)^{-1}(\lambda(B)))_+ - t_{\lambda(B')}(X)\right] \\ &= \mathbb{E}_F\left[t_{\lambda(B)}(X)\right] - \mathbb{E}_F\left[t_{\lambda(B')}(X)\right] \\ &= B - B'. \end{aligned}$$

Thus, we not only observe that the optimal policy that solves (35) is given by (36) with $\lambda = \lambda(B)$, but we see that $t_{\lambda(B')}(x) + \tilde{t}_{\lambda(B)}(x) = t_{\lambda(B)}(x)$. So, we have shown that our desired claim holds.

B.4 Proof of Lemma 3

We note that any optimal gap-targeting policy must have the form given in Lemma 5. At the same time, we note that the optimal policy that solves

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \mathbb{E}_F\left[t(X)\right] \text{ subject to } \mathbb{P}_F\left[t(X) + Y < c \mid X = x\right] \leq \lambda \quad \forall x \in \mathcal{X} \tag{37}$$

is also given by (9). This is straightforward to see because the constraints of (37) are separable.

Thus, we can view the optimal gap targeting policy as optimizing

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \max_{x\in\mathcal{X}} \mathbb{P}_F\left[t(X) + Y < c \mid X = x\right] \text{ subject to } \mathbb{E}_F\left[t(X)\right] \leq B.$$

B.5 Proof of Lemma 4

Let m_x denote the location of $F_{Y|X=x}$. By Lemma 5, the optimal gap minimizing policy is given by $t_{\text{gap}}(x) = (c - F^{-1}_{Y|X=x}(\lambda))_+$ for some $\lambda(B) \in \mathbb{R}$ such that $\mathbb{E}_F\left[(c - F^{-1}_{Y|X=X}(\lambda(B)))_+\right] = B$.

We note that when $F_{Y|X=x}$ has monotone increasing density on $[0, c]$, the post-transfer poverty rate $\mathbb{E}_F\left[F_{Y|X=X}(c - t(X))\right]$ is convex. In this regime, the optimal rate-minimizing policy is given by $t_{\text{rate}}(x; p) = (c - f^{-1}_{Y|X=x}(p))_+$ for some p such that $\mathbb{E}_F\left[t_{\text{rate}}(X; p)\right] = B$. We show that when $p(B) = g(G^{-1}(\lambda(B)))$, $\mathbb{E}_F\left[t_{\text{rate}}(X; p(B))\right] = B$.

$$\begin{aligned} \mathbb{E}_F\left[(c - f^{-1}_{Y|X=X}(p(B)))_+\right] &= \mathbb{E}_F\left[(c - f^{-1}_{Y|X=X}(g(G^{-1}(\lambda(B)))))_+\right] \\ &= \mathbb{E}_F\left[(c - f^{-1}_{Y|X=X}(f_{Y|X=X}(G^{-1}(\lambda(B)) + m_X)))_+\right] \\ &= \mathbb{E}_F\left[(c - (G^{-1}(\lambda(B)) + m_X))_+\right] \\ &= \mathbb{E}_F\left[(c - F^{-1}_{Y|X=X}(\lambda(B)))_+\right] \\ &= B. \end{aligned}$$

This derivation also shows that

$$t_{\text{rate}}(x;\lambda(B)) = (c - f^{-1}_{Y|X=x}(\lambda(B)))_+ = (c - F^{-1}_{Y|X=x}(p(B))_+ = t_{\text{gap}}(x;p(B)),$$

so the optimal rate-minimizing policy is the optimal gap-minimizing policy.

B.6 Proof of Lemma 5

We compute the optimal policy when $L(z) = (c-z)_+$. Since L is convex, the optimal policy solves

$$\min_{t:\mathcal{X}\to\mathbb{R}_+} \{\mathbb{E}_F\left[L(t(X)+Y)\right] + \lambda(\mathbb{E}_F\left[t(X)\right] - B)\}$$

for some $\lambda \geq 0$. The first-order stationarity condition for a given x is

$$\frac{d}{dt}\mathbb{E}_F\left[L(t+Y) \mid X = x\right] + \lambda = 0. \tag{38}$$

We note that

$$\begin{aligned}
\frac{d}{dt}\mathbb{E}_F\left[L(t+Y) \mid X = x\right] &= \frac{d}{dt}\mathbb{E}_F\left[(c-t-Y)\mathbb{I}(Y < c-t)\right] \\
&= \frac{d}{dt}\int_0^{c-t}(c-t)f_{Y|X=x}(y)dy - \frac{d}{dt}\int_0^{c-t} y f_{Y|X=x}(y)dy \\
&= \frac{d}{dt}((c-t)\cdot F_{Y|X=x}(c-t)) - (c-t)\cdot f_{Y|X=x}(c-t) \\
&= -(c-t)\cdot f_{Y|X=x}(c-t) - F_{Y|X=x}(c-t) + (c-t)\cdot f_{Y|X=x}(c-t) \\
&= -F_{Y|X=x}(c-t).
\end{aligned}$$

So, we can rewrite (38) as

$$F_{Y|X=x}(c-t) = \lambda.$$

We note that the policy must also be feasible, so $t \geq 0$. So, the optimal policy has the form in (9).

B.7 Proof of Theorem 6

Under Assumptions 2 and 4, (2) with $L(z) = \mathbb{I}(z < c)$ has a continuous objective. In addition, the feasible set can be restricted to a compact space because any optimal policy that solves (2) will satisfy $0 \leq t(x) < c$. Thus, there exists a minimizer of (2) with $L(z) = \mathbb{I}(z < c)$.

We note that any minimizer of (2) must satisfy the KKT conditions, which we derive below. We define the Lagrangian and its derivative below.

$$\begin{aligned}
L(t,\lambda,\nu) &= \mathbb{P}_F\left[t(X) + Y < c\right] + \lambda(\mathbb{E}_F\left[t(X)\right] - B) - \sum_{x\in\mathcal{X}} \nu_x t(x) \\
\nabla L(t,\lambda,\nu) &= f_X(x)\cdot(-f_{Y|X=x}(c-t(x)) + \lambda) - \nu_x = 0.
\end{aligned}$$

Let $\lambda \geq 0$ and $\nu_x \geq 0$ for all $x \in \mathcal{X}$. The KKT conditions of this optimization problem are as follows

$$\begin{aligned} f_{Y|X=x}(c - t(x)) - \lambda &= -\frac{\nu_x}{f_X(x)} \quad \text{(Stationarity)} \\ \nu_x \cdot t(x) &= 0 \quad \text{(Complementary Slackness)} \\ \lambda(\mathbb{E}_F\left[t(X)\right] - B) &= 0 \quad \text{(Complementary Slackness)} \\ t(x) &\geq 0 \quad \text{(Primal Feasibility)} \\ \mathbb{E}_F\left[t(X)\right] &\leq B \quad \text{(Primal Feasibility)}. \end{aligned}$$

The optimal solution must satisfy the KKT conditions for some $\lambda \geq 0$. By complementary slackness, we must have that this solution satisfies $\mathbb{E}_F\left[t(X)\right] - B = 0$ (the inequality is tight) or $\lambda = 0$.

We consider a policy that satisfies the KKT conditions for $\lambda = 0$. When $\lambda = 0$, we must have that

$$\begin{aligned} f_{Y|X=x}(c - t(x)) &= -\frac{\nu_x}{f_X(x)} \quad \text{(Stationarity)} \\ \nu_x \cdot t(x) &= 0 \quad \text{(Complementary Slackness)} \\ t(x) &\geq 0 \quad \text{(Primal Feasibility)} \\ \mathbb{E}_F\left[t(X)\right] &\leq B \quad \text{(Primal Feasibility)}. \end{aligned}$$

The only policy that can satisfy these conditions simultaneously must have $t(x) \in \{0, c\}$. This policy must be α-valid for $\alpha = 0$.

We can also characterize policies that satisfy the KKT conditions for $\lambda > 0$:

1. If t satisfies the KKT conditions and $t(x) > 0$, then $\nu_x = 0$ to ensure complementary slackness is satisfied. Then we must have that $f_{Y|X=x}(c - t(x)) = \lambda$. Also, we have that $t(x) \geq 0$. So, $t(x) \in \mathcal{T}_\alpha(x)$ for $\lambda = \alpha$.

2. If t satisfies the KKT conditions and $t(x) = 0$, then we must have that $f_X(x) \cdot (f_{Y|X=x}(c) - \lambda) = \nu_x$. So, $t(x) \in \mathcal{T}_\alpha(x)$ for $\lambda = \alpha$.

As a result, a policy t that satisfies the KKT conditions for a given value λ will be α-valid for some choice of $\alpha \geq 0$.

B.8 Proof of Corollary 7

By Theorem 6, an optimal deterministic policy must be α-valid for some $\alpha \geq 0$. So, for each value α, we can compute the policy t^*_α. Then, the optimal policy can be obtained by computing $\operatorname{argmin}_{\alpha:\alpha\geq 0} \mathbb{E}_F\left[t^*_\alpha(X)\right]$.

B.9 Proof of Corollary 8

Recall that if a policy t is deterministic and α-valid then $t(x) \in \mathcal{T}_\alpha(x)$ for every $x \in \mathcal{X}$, and we can write that $t(x) = \langle \pi(x), z(x;\alpha) \rangle$ where $\pi(x) : \mathcal{X} \to \{0,1\}^K$ and $\sum_{k\in[K]} \pi_k(x) = 1$. Since α-valid policies have a highly structured form, the constraint of (11) can be written as $\mathbb{E}_F\left[\langle \pi(X), p(X;\alpha) \rangle\right]$ and the feasibility constraint that $\mathbb{P}_F\left[t(X) + Y < c\right] \leq \epsilon$. (11) can be written as $\mathbb{E}_F\left[\langle \pi(X), z(X;\alpha) \rangle\right] \leq B$. This yields (12).

C Technical Proofs

C.1 Proof of Lemma 9

C.1.1 Once Differentiable

We apply Leibniz rule to show that $G_x(t)$ is differentiable and find an explicit form for the derivative.

Leibniz Rule (Klenke, 2013): Let $\mathcal{T}$ be an open subset of $\mathbb{R}$ and let $\mathcal{Y}$ be a measure space. Suppose that $g : \mathcal{T} \times \mathcal{Y} \to \mathbb{R}$ satisfies the following conditions.

1. $g(t, y)$ be a Lebesgue integrable function of y for each $t \in \mathcal{X}$.
2. For almost all $y \in \mathcal{Y}$, the partial derivative $\frac{\partial}{\partial t} g(t, y)$ exists for all $t \in \mathcal{T}$.
3. There is a Lebesgue integrable function $\theta : \mathcal{Y} \to \mathbb{R}_+$ for $|\frac{\partial g}{\partial t}(t, \cdot)| \leq \theta(y)$ almost everywhere for all $t \in \mathcal{T}$.

Then, for all $t \in \mathcal{T}$,

$$\frac{d}{dt} \int_{\mathcal{Y}} g(t, y) dy = \int_{\mathcal{Y}} \frac{\partial}{\partial t} g(t, y) dy.$$

Let $\mathcal{T} = (0, \infty)$. Recall that $F_{Y|X=x}$ has positive density on $\mathcal{Y}_x$. We can set $g(t, y) := L(y+t) \cdot f_{Y|X=x}(y)$. Let $\delta \leq \inf \mathcal{T}$. We observe that $|L(y + t)| \leq \max(|L(\delta)|, |C|)$ because L is bounded below by $C > -\infty$ and $L(y+t) \geq L(\delta)$ for $y \in \mathbb{R}_+, t \in \mathcal{T}$ since L is decreasing. We use this fact to show that g is Lebesgue integrable as follows

$$\begin{aligned}
\int_{\mathcal{Y}_x} |g(t, y)| dy &= \int_{\mathcal{Y}_x} |L(y + t) \cdot f_{Y|X=x}(y)| dy \\
&= \int_{\mathcal{Y}_x} |L(y + t)| \cdot f_{Y|X=x}(y) dy \\
&\leq \max(|L(\delta)|, |C|) \int_{\mathcal{Y}_x} f_{Y|X=x}(y) dy \\
&< \infty.
\end{aligned}$$

Since L is monotone and differentiable, for all $y \in \mathcal{Y}_x$, $\frac{\partial g}{\partial t}(t, y) = L'(y + t) \cdot f_{Y|X=x}(y)$ exists for all $t \in \mathcal{T}$.

Since L is convex, differentiable, and decreasing, $L'(z) \leq 0$ and L' is an increasing function. Thus, there exists $\delta \leq \inf \mathcal{T}$ so $|L'(y + t)| \leq |L'(\delta)| < \infty$ for all $t \in \mathcal{T}$. Let $\theta(y) := |L'(\delta)| \cdot f_{Y|X=x}(y)$. We have that $\theta(y)$ is Lebesgue integrable because

$$\int_{\mathcal{Y}_x} \theta(y) \cdot f_{Y|X=x}(y) dy = |L'(\delta)| \int_{\mathcal{Y}_x} f_{Y|X=x}(y) dy = |L'(\delta)| < \infty.$$

Thus, $G_x(t)$ satisfies the conditions of Leibniz rule, and

$$G'_x(t) = -\int_{\mathcal{Y}_x} L'(t + y) f_{Y|X=x}(y) dy. \tag{39}$$

C.1.2 Strict Concavity

Since L is strictly convex, it follows that $G_x(t)$ is strictly concave.

C.1.3 Twice Differentiable

We recall the form of $G'_x(t)$ given in (39). We apply integration by parts to $G'_x(t)$. Suppose $\mathcal{Y}_x := [\underline{y}, \overline{y}]$.

$$\begin{aligned}
G'_x(t) &= -\int_{\mathcal{Y}} L'(y+t) f_{Y|X=x}(y) dy \\
&= -\left(L(y+t) f_{Y|X=x}(y) \Big|_{\underline{y}}^{\overline{y}} - \int_{\mathcal{Y}} L(y+t) \cdot f'_{Y|X=x}(y) dy \right) \\
&= -L(\overline{y}+t) f_{Y|X=x}(\overline{y}) + L(\underline{y}+t) f_{Y|X=x}(\underline{y}) + \int_{\mathcal{Y}} L(y+t) \cdot f'_{Y|X=x}(y) dy.
\end{aligned}$$

We aim to compute

$$G''_x(t) = L'(\underline{y}+t) f_{Y|X=x}(\underline{y}) - L'(\overline{y}+t) f_{Y|X=x}(\overline{y}) + \frac{d}{dt} \int_{\mathcal{Y}_x} L(y+t) \cdot f'_{Y|X=x}(y) dy.$$

We note that the first two terms are differentiable on $\mathbb{R}$. We apply Leibniz Rule to the integral in the above expression. We first show that $h(t,y) = L(y+t) \cdot f'_{Y|X=x}(y)$ is Lebesgue integrable. Recall that L is bounded on below by a constant $C > -\infty$.

$$\begin{aligned}
\int_{\mathcal{Y}_x} |h(t,y)| dy &= \int_{\mathcal{Y}_x} |L(y+t) \cdot f'_{Y|X=x}(y)| dy \\
&= \int_{\mathcal{Y}_x} \left| L(y+t) \cdot \frac{f'_{Y|X=x}(y)}{f_{Y|X=x}(y)} \right| \cdot f_{Y|X=x}(y) dy \\
&\leq \max(C, L(0)) \cdot \int_{\mathcal{Y}} \left| \frac{f'_{Y|X=x}(y)}{f_{Y|X=x}(y)} \right| \cdot f_{Y|X=x}(y) dy \\
&= \max(C, L(0)) \cdot \mathbb{E}_F \left[\left| \frac{d}{dY} \log f_{Y|X=x}(Y) \right| \mid X = x \right] \\
&< \infty.
\end{aligned}$$

The last line follows from Assumption 2. Next, we note that by the Monotone Differentiation Theorem, L is differentiable almost everywhere. As a result, $\frac{\partial}{\partial t} h(t,y) = L'(t+y) \cdot f'_{Y|X=x}(y)$ exists for all $t \in \mathcal{T}$ for almost all $y \in \mathcal{Y}_x$. Finally, we show that there is an integrable function $\theta : \mathcal{T} \to \mathcal{Y}_x$ for $|\frac{\partial h}{\partial t}(t, \cdot)| \leq \theta$. Since L is convex and decreasing, $L'(z) \leq 0$ almost everywhere and L' is an increasing function. Thus, there exists $\delta < 0$ so that $|L'(y+t)| \leq |L'(\delta)| < \infty$ almost everywhere for all $t \in \mathcal{T}$. Let $\theta(y) = |L'(\delta)| \cdot f'_{Y|X=x}(y)$. We have that $\theta(y)$ is Lebesgue integrable because

$$\begin{aligned}
\int_{\mathcal{Y}_x} |\theta(y)| \cdot f_{Y|X=x}(y) dy &\leq |L'(\delta)| \cdot \int_{\mathcal{Y}_x} \left| \frac{f'_{Y|X=x}(y)}{f_{Y|X=x}(y)} \right| \cdot f_{Y|X=x}(y) dy \\
&= |L'(\delta)| \cdot \mathbb{E}_F \left[\left| \frac{d}{dY} \log f_{Y|X=x}(Y) \right| \mid X = x \right] \\
&< \infty.
\end{aligned}$$

The last inequality follows from Assumption 2. Thus, we can apply Leibniz rule to show that

$$G''_x(t) = L'(t) \cdot f_{Y|X=x}(y) + \int_{\mathcal{Y}_x} L'(y+t) \cdot f'_{Y|X=x}(y) dy.$$

C.2 Proof of Lemma 10

C.2.1 Once Differentiable

Similar to the proof of Lemma 9, we can apply Leibniz Rule (Klenke, 2013) to show that $G_x(t)$ is once differentiable on $\mathcal{T} = (0, \infty)$. Recall that $F_{Y|X=x}$ has positive density on $\mathcal{Y}$. We can set $g(t, y) := L(y + t) \cdot f_{Y|X=x}(y)$. Since L is an FGT index, we observe that $0 \leq L(y + t) \leq L(0)$ for $y \in \mathbb{R}_+, t \in \mathcal{T}$ since L is decreasing and bounded below by 0. We use this fact to show that g is Lebesgue integrable as follows

$$\begin{aligned}
\int_{\mathcal{Y}_x} |g(t, y)| dy &= \int_{\mathcal{Y}_x} |L(y + t) \cdot f_{Y|X=x}(y)| dy \\
&= \int_{\mathcal{Y}_x} |L(y + t)| \cdot f_{Y|X=x}(y) dy \\
&\leq L(0) \cdot \int_{\mathcal{Y}_x} f_{Y|X=x}(y) dy \\
&< \infty.
\end{aligned}$$

Thus, $g(t, y)$ is Lebesgue integrable.

Since L is an FGT index, it is differentiable except at c. So, for almost all $y \in \mathcal{Y}$, $\frac{\partial g}{\partial t}(t, y) = L'(y + t) \cdot f_{Y|X=x}(y)$ exists for all $t \in \mathcal{T}$.

Since L is an FGT index for $\alpha \geq 1$, it is convex, differentiable except at c, and decreasing. Thus, $L'(z) \leq 0$ almost everywhere and L' is an increasing function. Thus, $|L'(y + t)| \leq |L'(0)| < \infty$. Let $\theta(y) := |L'(\delta)| \cdot f_{Y|X=x}(y)$. We have that $\theta(y)$ is Lebesgue integrable because

$$\int_{\mathcal{Y}_x} \theta(y) \cdot f_{Y|X=x}(y) dy = |L'(0)| \int_{\mathcal{Y}} f_{Y|X=x}(y) dy = |L'(0)| < \infty.$$

Thus, $G_x(t)$ satisfies the conditions of Leibniz rule, and

$$G'_x(t) = -\int_{\mathcal{Y}_x} L'(t + y) f_{Y|X=x}(y) dy \tag{40}$$

holds.

C.2.2 Strict Concavity

We observe that the FGT indices for $\alpha \geq 1$ are convex. In addition, they are strictly convex at c in the sense defined below.

Definition 8 (Strict convexity at a point)**.** A convex function $f : \mathbb{R} \to \mathbb{R}$ is strictly convex at a point c if for all $z_1, z_2 \in \mathbb{R}$ such that $z_1 < c < z_2$ and $\lambda \in (0, 1)$,

$$f(\lambda z_1 + (1 - \lambda) z_2) < \lambda f(z_1) + (1 - \lambda) f(z_2).$$

Let $t_1, t_2 \in [0, c]$ such that $t_1 < t_2$. We define $y' = c - \lambda t_1 - (1 - \lambda) t_2$. We note that $y' + t_1 < c < y' + t_2$. Applying Definition 8, we have that

$$L(y' + \lambda t_1 + (1 - \lambda) t_2) < \lambda L(y' + t_1) + (1 - \lambda) L(y' + t_2). \tag{41}$$

We note that this result also holds in a small neighborhood about y'. Let $\epsilon = \frac{1}{2} \cdot \min(y' + t_2 - c, c - y' - t_1)$. In this case, $y' + \epsilon + t_1 < c < y' + \epsilon + t_2$ and $y' - \epsilon + t_1 < c < y' - \epsilon + t_2$, so (41) holds when y' is replaced by $y' + \epsilon$ and $y' - \epsilon$. Let $I = [y' - \epsilon, y' + \epsilon]$. We note that under Assumption 2, $f_{Y|X}(y) > 0$ on I.

$$
\begin{aligned}
&G_x(\lambda t_1 + (1-\lambda)t_2) \\
&= -\int_{\mathcal{Y}_x} L(y + \lambda t_1 + (1-\lambda)t_2) f_{Y|X=x}(y) dy \\
&= -\int_I L(y + \lambda t_1 + (1-\lambda)t_2) f_{Y|X=x}(y) dy - \int_{\mathcal{Y}\setminus I} L(y + \lambda t_1 + (1-\lambda)t_2) f_{Y|X=x}(y) dy \\
&> -\int_I [\lambda L(y + t_1) + (1-\lambda) L(y + t_2)] f_{Y|X=x}(y) dy - \int_{\mathcal{Y}\setminus I} L(y + \lambda t_1 + (1-\lambda)t_2) f_{Y|X=x}(y) dy \\
&\geq -\int_I [\lambda L(y + t_1) + (1-\lambda) L(y + t_2)] f_{Y|X=x}(y) dy - \int_{\mathcal{Y}_x\setminus I} [\lambda L(y + t_1) + (1-\lambda) L(y + t_2)] f_{Y|X=x}(y) dy \\
&= \lambda G_x(t_1) + (1-\lambda) G_x(t_2).
\end{aligned}
$$

Thus, G_x is strictly concave on $[0, c]$.

C.2.3 Twice Differentiable

We note that the FGT indices for $\alpha \geq 1$ can be written as $L(z) = \ell(z) \cdot \mathbb{I}(z < c)$ for a twice differentiable function ℓ. We recall the form of $G'_x(t)$ given in (40).

$$
\begin{aligned}
G'_x(t) &:= -\int_{\mathcal{Y}_x} L'(t + y) f_{Y|X=x}(y) dy \\
&= -\int_{\mathcal{Y}_x} \ell'(t + y) \cdot \mathbb{I}(t + y < c) f_{Y|X=x}(y) dy \\
&= -\int_{\underline{y}}^{c-t} \ell'(t + y) f_{Y|X=x}(y) dy.
\end{aligned}
$$

We can use Leibniz integral rule to compute $G''_x(t)$ as follows.

$$
G''_x(t) = \ell'(c) f_{Y|X=x}(c - t) - \int_{\underline{y}}^{c-t} \ell''(t + y) f_{Y|X=x}(y) dy.
$$

D Data

D.1 Predictor Selection

This section describes our rubric for selecting covariates from household surveys for inclusion as predictors when learning transfer policies. Our goal is to select only characteristics that are plausibly verifiable. To that end, we build a rubric from an initial list of covariates actually used in PMTs implemented in low- and middle-income countries, as described in Section 3. In particular, we define categories of covariates, generalizing where appropriate from individual covariates that have been used. For example, if previous PMTs used the material out of which the house's walls are made, then we take that as justification for also using the material out of which other parts of the house, such as the roof, are made. The resulting categories (labelled (a), (b), (c), etc. below) are exhaustive of the categories of variables within each top-level group (e.g., household demographics) that we include.

We also exercise judgment in excluding some variables that have been used in real-world PMTs but that we deem too difficult to plausibly verify at large scale. Examples include: Age of dwelling; indicators of food security, such as whether household members had skipped meals in the past week; use of fertilizer; use of an internet connection; income and private transfers (like foreign remittance); and ownership of a mobile phone or usage of someone else's mobile phone.

We use the resulting rubric as a guide when selecting predictors to use from each survey. The rubric, edited for clarity, is as follows:

1. Household demographics
    (a) Counts of household members by type, including the overall number of household members, the number of male and female members, the numbers in various age-based subcategories such as children, adults, elderly adults / senior citizens, and the number of specially-abled members. Examples of prior usage for the number of elderly adults: 60 and over for Below Poverty Line (BPL) classification in India (Planning Commission, 2011) or 65 and over (Alatas et al., 2012; Hanna and Olken, 2018).
    (b) Household head: Age, gender, and marital status. Examples of prior usage: Alatas et al. (2012) include married/unmarried; Fernandez and Hadiwidjaja (2018) include single/married/divorced; Kidd and Wylde (2011) include widow status.
    (c) Ethnicity, tribe, and caste of the household/household head.
2. Human capital
    (a) Educational attainment of the household head. Depending on the source this may take the form of years completed, highest academic level completed, or simply a literacy indicator. Examples of prior usage: Alatas et al. (2012), Planning Commission (2011, 2012).
    (b) Maximum educational attainment of all adult members. Again, this may be measured in years, in level completed, etc.[28] Example of prior usage: Kidd and Wylde (2011)
    (c) Maximum educational attainment of all female adult members. Example of prior usage: Kidd and Wylde (2011).
    (d) Number of children currently enrolled in school.
3. Household asset: presence and number owned
    (a) We do *not* include ownership-structure details beyond the fact of asset ownership.
    (b) Dwelling characteristics such as per capita number of rooms (Kidd and Wylde, 2011) or floor space (Alatas et al., 2012), material used to construct floors, walls, or roof, etc., as well as an indicator for homeownership itself. Examples of prior usage: Alatas et al. (2012) use binary indicators for home-ownership and roof materials. Hanna and Olken (2018) consider more granular categories.
    (c) Presence and physical characteristics of amenities such as: Type of latrine, water source, lighting source, drainage system, waste collection, access to electricity and gas, type of cooking fuel, cable connection. Examples of prior usage: Alatas et al. (2012) use a simple binary indicator for availability of clean water within the house. Hanna and Olken (2018) use more granular categories of water source.

[28] We take the maximum rather than including the attainment of all members individually in order to produce a per-household predictor.

(d) Ownership of consumer durables such as appliances (e.g. radio, television, refrigerator, generator, cooker, heater, fan, air conditioner); transportation (e.g. car, bicycle, motorbike); furniture (e.g. sofa, bed etc); and devices (e.g. computers).

(e) Productive agricultural assets including land, livestock, irrigation facilities, and farm machinery. We exclude cultivation details, such as crop types and amounts, as they may not be verifiable. We include the amount of land owned. Examples of prior usage: Planning Commission (2011, 2012) include farm machinery.

(f) Productive non-agricultural assets, e.g a sewing machine.

(g) We generally do *not* include finantial assets, as we expect the kinds of financial assets held by poor households to be hard to verify, but we include any that they hold as part of a government scheme for which the government might plausibly hold records. Example of prior usage: Planning Commission (2011, 2012) include usage of the Kisan credit card scheme in India.

(h) We *exclude* crop stores, as they may not be easy to verify.

4. Livelihood activities

(a) Primary sector of employment of the household head (e.g., agriculture, manufacturing, services).

(b) Primary occupation of the household head (e.g., self-employed, salaried employee, casual laborer, etc.).

(c) We do *not* include the occupations of other members of the household.

(d) Ownership of enterprises.

(e) Receipt of other public transfers, including amount, if this could plausibly be verified by merging in other administrative records. We omit transfers when it is unclear whether they are administratively documented, such as child support or transfers from non-government institutions. Example of prior usage: (Camacho and Conover, 2011) describes a poverty census which includes social security information.

5. Geographic indicators

(a) Urban/rural status, including any further available classifications such as peri-urban.

(b) Administrative geographic information: Which administrative division a household is in. Note that some care is required in the treatment of these indicators, as surveys may include region identifiers at a finer granularity than that at which they are representative. To obtain results which reflect expected performance across households anywhere in each country, we include geographic identifiers only at granularities at which the survey is representative. For example, if a survey were to sample households from all districts but only a subset of subdistricts, we would include district identifiers but not subdistrict identifiers as predictors.

(c) Distance to important locations such as district centers, markets, or public facilities such as post offices. Examples of prior usage: Alatas et al. (2012) use distance to district centers and markets. Hanna and Olken (2018) use distance to a post office.

(d) Environmental conditions such as rainfall history. Del Ninno and Mills (2015) provides a few case studies that illustrate feasibility.

6. Community characteristics

   (a) Presence of publicly or privately provided services such as healthcare. Examples of prior usage: Alatas et al. (2012) uses an indicator for presence of a doctor. Kidd and Wylde (2011) uses an indicator for presence of a midwife.

   (b) Presence of publicly or privately provided infrastructure such as paved roads, banking facilities, or regional government offices. Examples of prior work: Alatas et al. (2012); Kidd and Wylde (2011)use an indicator for presence of banking facilities. Kidd and Wylde (2011) use an indicator for presence of regional government offices.

   (c) Population characteristics such as headcount or population density. Example of prior work: Kidd and Wylde (2011).

D.2 Outcome Construction

We calculate consumption per capita by dividing total household consumption by household size without an adult equivalence scale adjustment, following the World Bank Poverty and Inequality Platform's methodology (World Bank, 2025a).

Wherever possible we use the consumption aggregates provided along with the disaggregated data for this purpose. Accounting for transfers, housing, and other durable goods are among the thornier issues in measuring consumption (Amendola and Vecchi, 2022), and the sources vary somewhat in the ways they do this (see Table D.1). That said, large majorities include the value of both cash and in-kind transfers; include an estimate of the value of housing services consumed; and include an estimate of the value of services from non-housing durable goods. We apply temporal and spatial price deflators when they are provided as part of official data. In cases where one or both is not present, we do not attempt our own temporal or spatial correction.

We convert consumption aggregates from all surveys from local currency units (LCU) to 2017 PPP USD using Consumer Price Index and PPP conversion factors provided by the World Bank (World Bank, 2023). This requires identifying the base year of the LCUs in which the consumption aggregate is defined. The base year is clear from survey context and/or documentation in most cases; in others we infer it by examining data. For example, the eight surveys that make up the Enquête Harmonisée sur le Conditions de Vie des Ménages (EHCVM) series, run in eight West African countries (Benin, Burkina Faso, Côte d'Ivoire, Guinea-Bissau, Mali, Niger, Senegal, and Togo), were all conducted between 2018 and 2019. Their documentation does not specify one or the other year as the currency base year, but the data include a temporal deflator; these take on values very near 1 (within roughly 3%) and have mean values extremely close to 1. Most of the surveys have values near 1 around the middle of the time series (early 2019). A majority of samples are from 2019 in all surveys, and the deflator having a mean of 1 indicates that the base currency is an average of currency units across surveyed samples—also suggesting 2019, the year during which the average sample is drawn. Given these observations, we use a currency unit of 2019 West African francs for all EHCVM surveys.

Table D.1: Consumption Aggregate Details

Country	Category of Consumption		
	Grants and Transfers	**Durable Goods**[a]	**Non-Rental Housing**[b]
Benin	Included	User cost	Hedonic function[c]
Burkina Faso	Included	User cost	Hedonic function
Côte d'Ivoire	Included	User cost	Hedonic function
Ethiopia	Included	Acquisition approach	Unspecified
Ghana	Not included	Acquisition approach	Excluded
Guinea-Bissau	Included	User cost	Hedonic function
Kenya[d]	Food gifts included; non-food gifts excluded	Unspecified	Hedonic function
Malawi	Food gifts included; non-food gifts excluded	User cost[e]	Self-reported rental value[f]
Mali	Included	User cost	Hedonic function
Niger	Included	User cost	Hedonic function
Nigeria	Food gifts included; non-food gifts unspecified	Acquisition approach[g]	Hedonic function
Senegal	Included	User cost	Hedonic function
South Africa	Included	Acquisition approach	Average self-reported rental yield by housing type and province
South Sudan[h]	Unspecified	User cost	All housing consumption is excluded
Tanzania	Food gifts included; non-food gifts excluded	User cost	Hedonic function
Togo	Included	User cost	Hedonic function
Uganda	Included	Acquisition approach	Market rental prices (details unspecified)

[a]This column uses terminology introduced by Amendola and Vecchi (2022).
[b]Includes any housing not rented on the market: Owned housing, family- or employer-provided housing, etc.
[c]A function predicting rental value based on observable housing characteristics and amenities; fit using self-reported rents of respondents living in rental housing.
[d]Additional irregularity: All (rental or non-rental) housing consumption is excluded in rural areas.
[e]Assumes linear depreciation and disregards interest rate
[f]Outliers imputed using hedonic function
[g]Long-term ("lumpy") durables exluded
[h]Additional irregularity: All (rental or non-rental) housing consumption is excluded.

D.3 Non-Survey Data

In addition to the household survey data that forms the primary basis for our empirical results, we rely on some additional data. Specific usages are reported alongside corresponding results and exhibits, and we compile the full list here for reference.

- Country-level poverty data from the World Bank Poverty and Inequality Platform (World Bank, 2025b):
  - Poverty headcount rate and poverty gap index for each country for which we use a survey, from the year of that survey, interpolated if necessary.
  - Poverty headcount rate and poverty gap index for all countries in the world where data are available, from the most recent year for which each rate/gap index is available.
- Country-level poverty headcount data from the World Poverty Clock.
- Country-level population estimates from the World Bank Development Indicators (World Bank, 2023).
- Currency exchange rates between local currency units (LCU) and US dollars. These rates are drawn from the World Bank's Development Indicators (World Bank, 2023). The World Bank's figures are in turn drawn from the IMF's International Financial Statistics database, whose indicators are hosted at IMF (2025). We use IMF data directly for exchange rates for Taiwan and Tuvalu, both of which were missing from the World Bank's exchange rate dataset at the time of access.

  Some exchange rate values were missing from both sources at the time of access: Guinea's 2021 exchange rate, Myanmar's 2021 exchange rate, Venezuela's 2021 exchange rate, and Zimbabwe's 2017 exchange rate. We estimate these missing values by assuming that the nominal currency exchange rate and the conversion factor from LCU to USD PPP (international dollars) have the same ratio as in the nearest year with complete data. By assuming that ratio, we are able to obtain an estimate for the nominal exchange rate based on the conversion factor from LCU to USD PPP, which in all cases is present for the given year.
- Conversion factors between LCU and USD PPP (international dollars). These factors are drawn from the World Bank's Development Indicators (World Bank, 2023). We fill some gaps in the World Bank dataset using data from the IMF[29] (IMF, 2025): Taiwan (2017, 2021); Venezuela (2021); Yemen (2017, 2021).

  Both datasets were missing Venezuela's 2017 conversion factor at the time of access. We estimate this value by assuming that the nominal currency exchange rate and the conversion factor from LCU to USD PPP (international dollars) have the same ratio as in the nearest year with complete data. By assuming that ratio, we are able to obtain an estimate for the conversion factor based on the nominal exchange rate for that year, which is present in IMF data.
- Country-level Consumer Price Index data (World Bank, 2023), to adjust for inflation.
- Net official development assistance and official aid received, by country (World Bank, 2023).
- GDP by country and region (World Bank, 2023).
- Revenue as percentage of GDP by country (IMF, 2025) and by region (World Bank, 2023).

[29] Indicator name: Rate, Domestic currency per international dollar in PPP terms, ICP benchmarks 2017-2021

E Additional Exhibits

Table E.2: Survey characteristics

Country	Survey Name	Survey Year	n	d	Survey Poverty Rate	WB Poverty Rate
Benin	EHCVM	2018	8012	345	16.21	20.83
Burkina-Faso	EHCVM	2018	7010	394	24.71	31.20
Côte d'Ivoire	EHCVM	2018	12992	377	8.52	11.50
Ethiopia	Socio-Economic Panel Survey	2021	4959	148	42.89	31.96
Ghana	Socioeconomic Panel Survey	2009	4953	302	19.80	32.88
Guinea-Bissau	EHCVM	2018	5410	424	18.35	21.66
Kenya	Continuous Household Survey	2021	16885	43	36.38	36.15
Malawi	Fifth Integrated Household Survey	2019	11434	140	72.60	70.06
Mali	EHCVM	2018	6602	424	8.82	15.22
Niger	EHCVM	2018	6024	212	45.38	50.91
Nigeria	Living Standards Survey	2018	22110	132	29.86	30.86
Senegal	EHCVM	2018	7156	422	7.54	9.22
South Africa	Income and Expenditure Survey	2010	25403	66	16.06	18.04
South Sudan	High Frequency Survey	2015	3505	100	52.78	62.64
Tanzania	National Panel Survey, Wave 5	2020	4709	164	46.21	44.95
Togo	HSHLS	2018	6171	164	24.63	28.43
Uganda	National Panel Survey	2019	3074	90	40.70	42.12

This table describes the countries and surveys used in the empirical analysis. The column n is the number of households in the survey data. The column d is the number of covariates that we use from the survey for learning transfer policies. "Survey Poverty Rate" is the estimated share of people living in the country who are below the extreme poverty line in the survey year, based on our estimates from the survey data. "WB Poverty Rate" is the estimated poverty rate in the survey year from the World Bank World Bank's Poverty and Inequality Platform, accessed 14 July 2025. We note that EHCVM is an abbreviation for "Enquête Harmonisée sur le Conditions de Vie des Ménages (EHCVM) 2018-2019", HSHLS is an abbreviation for "Harmonized Survey on Households Living Standards."

Table E.3: Poverty Rate and Share of World's Poor for In-Sample Countries

Country	Poverty Rate	Share of World's Poor
Benin	14.00	0.30
Burkina-Faso	33.00	1.24
Côte d'Ivoire	7.00	0.32
Ethiopia	14.00	2.82
Ghana	21.00	1.13
Guinea-Bissau	24.00	0.08
Kenya	21.00	2.00
Malawi	73.00	2.74
Mali	15.00	0.54
Niger	53.00	2.36
Nigeria	31.00	11.38
Senegal	7.00	0.21
South Africa	21.00	2.09
South Sudan	79.00	2.02
Tanzania	35.00	3.97
Togo	32.00	0.47
Uganda	35.00	3.05

This table reports the poverty rate and the share of the world's poor by country for each country in our sample. "Poverty Rate" is poverty rate (%) of each country based on World Poverty Clock (data source, accessed 25 August 2025) and "Share of World's Poor" is the fraction of the world's poor living that lives in each country in our sample. The "Share of World's Poor" is computed from World Poverty Clock country-level poverty rates and populations (data source, accessed 25 August 2025).

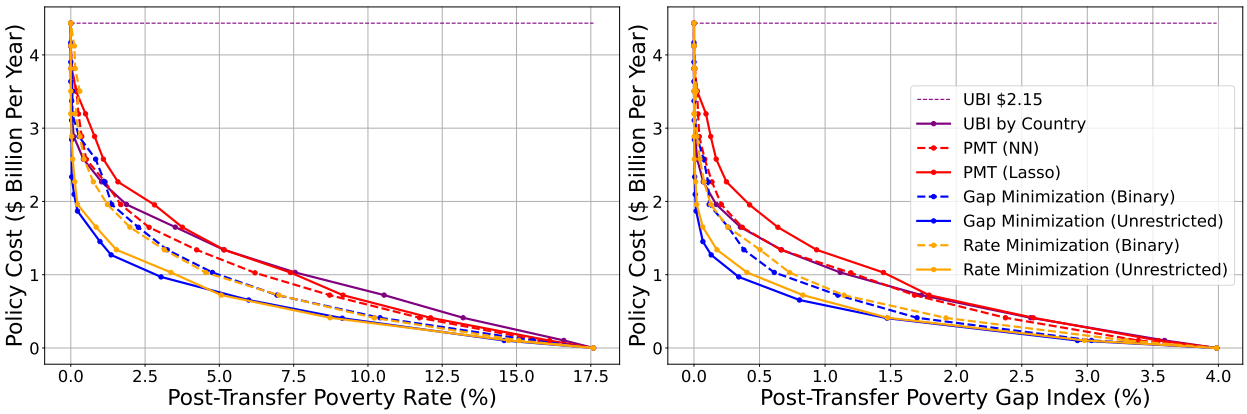
Table E.4: Comparison of Policy Costs for In-Sample Countries to Country GDP and Government Revenue.

Country	Reference Year	Policy Cost	GDP	Gov't Revenue	Policy Cost / GDP	Policy Cost / Gov't Revenue
Benin	2018	1.29	17.06	2.32	0.08	0.56
Burkina-Faso	2018	2.30	18.25	3.61	0.13	0.64
Côte d'Ivoire	2018	1.49	71.08	10.42	0.02	0.14
Ethiopia	2021	24.14	129.12	14.24	0.19	1.69
Ghana	2009	3.97	41.90	4.83	0.09	0.82
Guinea-Bissau	2018	0.21	1.70	0.25	0.13	0.85
Kenya	2021	8.56	116.23	19.55	0.07	0.44
Malawi	2019	4.64	13.94	2.06	0.33	2.26
Mali	2018	1.58	19.67	3.06	0.08	0.52
Niger	2018	4.42	14.82	2.69	0.30	1.64
Nigeria	2018	26.74	640.68	54.49	0.04	0.49
Senegal	2018	1.00	27.72	5.24	0.04	0.19
South Africa	2010	8.62	400.65	95.46	0.02	0.09
South Sudan	2015	2.46	15.42	2.64	0.16	0.93
Tanzania	2020	13.14	79.15	11.78	0.17	1.12
Togo	2018	1.46	8.40	1.53	0.17	0.96
Uganda	2019	6.51	50.91	6.85	0.13	0.95

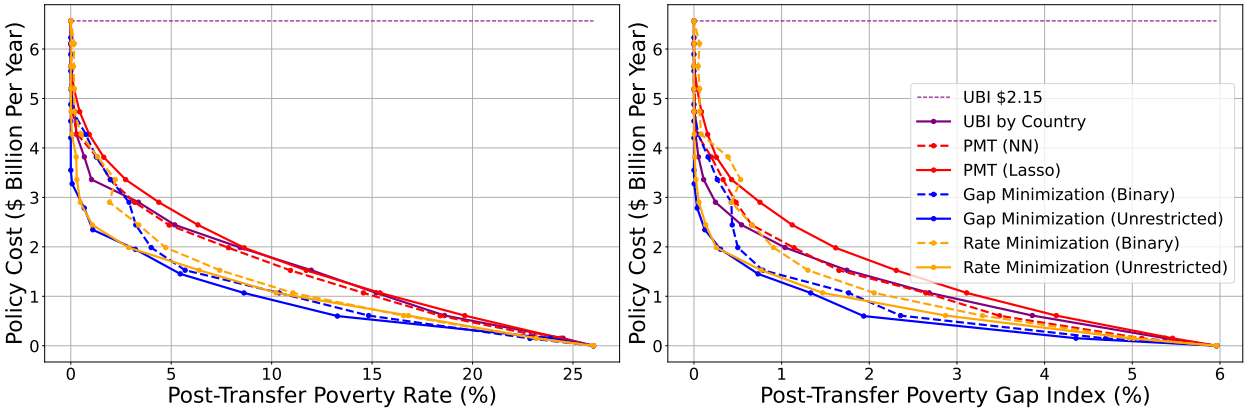
This table reports the cost of the gap-minimizing policy that reduces the poverty rate to 1% in each country, and compares it to national GDP and government revenue. "Reference Year" corresponds to the year that the Policy Cost, GDP, and government revenue are reported and is taken to be the survey year. "GDP" corresponds to the country GDP in the survey year in units of billions of 2023 nominal USD. "Policy Cost" is the cost of the policy, in billions of 2023 nominal USD, to attain a 1% poverty rate in the country. "Gov't Revenue" corresponds to country government revenue in the survey year in units of billions of 2023 nominal USD. We obtain the survey year GDP from the World Bank (data source, accessed 13 August 25). We obtain the survey year Gov't Revenue by multiplying survey year government revenue percentages from the IMF(data source, accessed 14 July 25) by the country GDP in the preceding column."Policy Cost / GDP" is the ratio of the policy cost and GDP. "Policy Cost / Gov't Revenue" is the ratio of policy cost and government revenue.

Figure E.1: Policy Cost vs. Post-Transfer Poverty Measures by Country

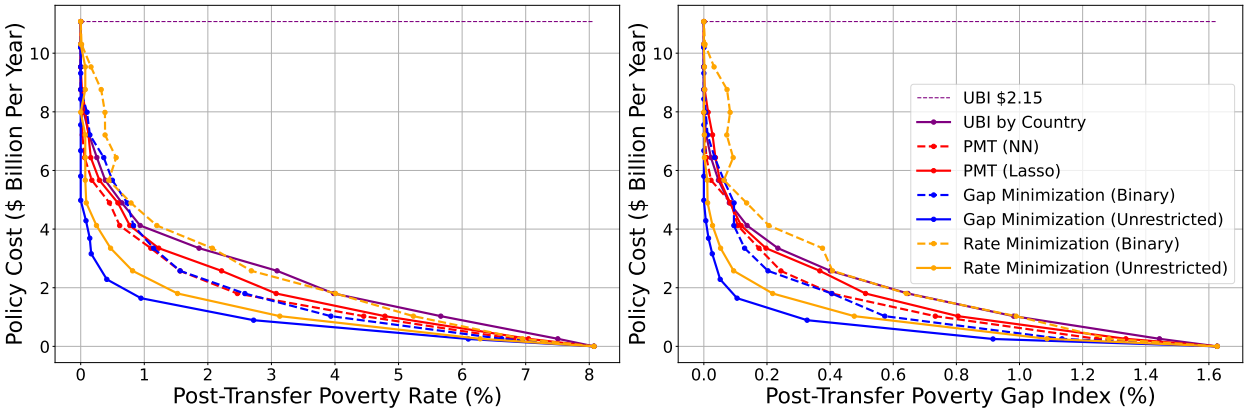
(a) Benin



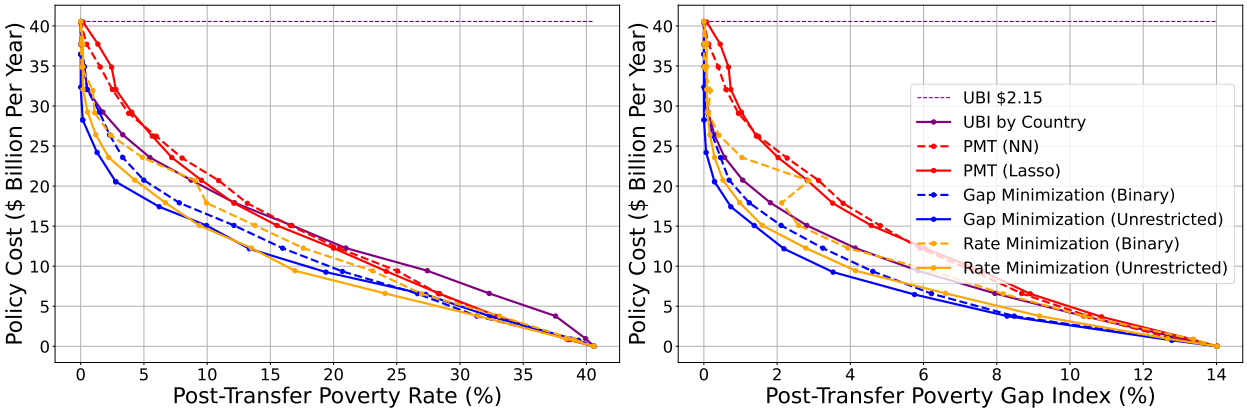
(b) Burkina-Faso



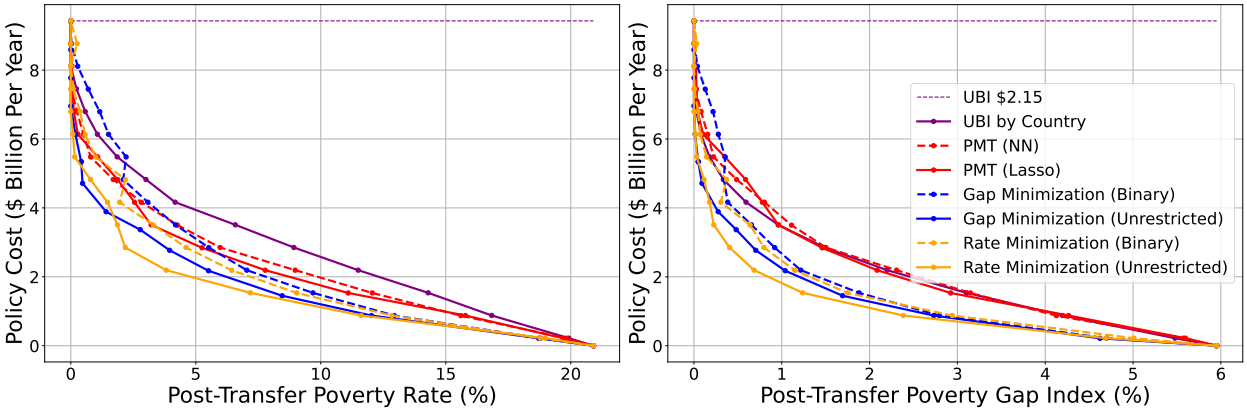
(c) Côte d'Ivoire



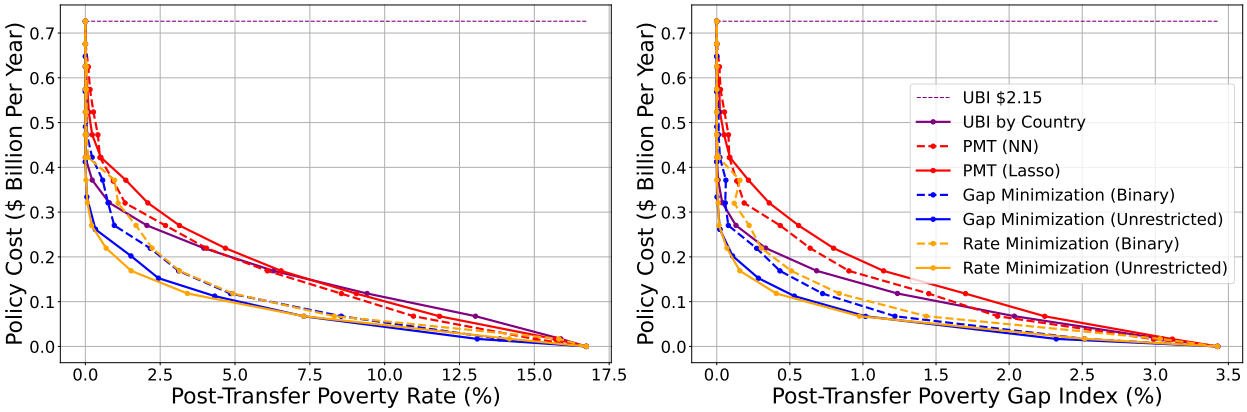
(d) Ethiopia



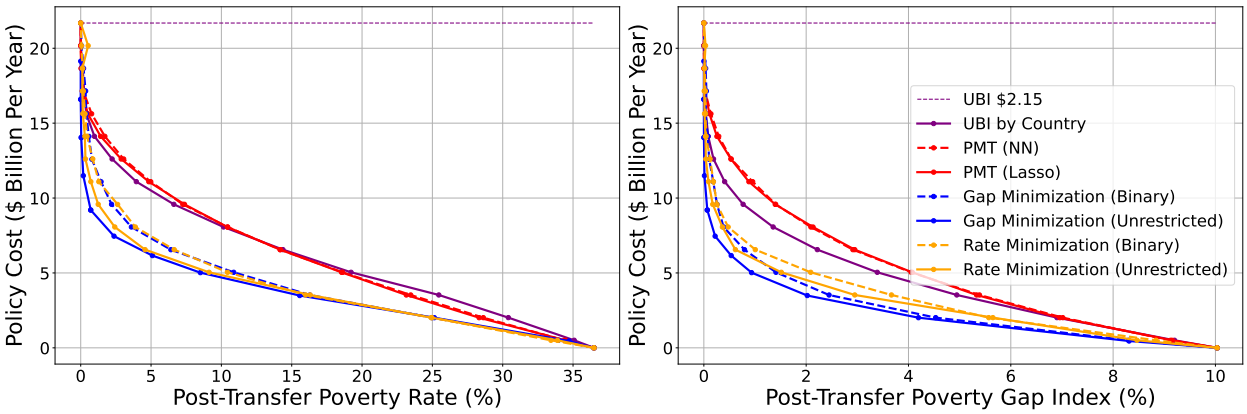
(e) Ghana



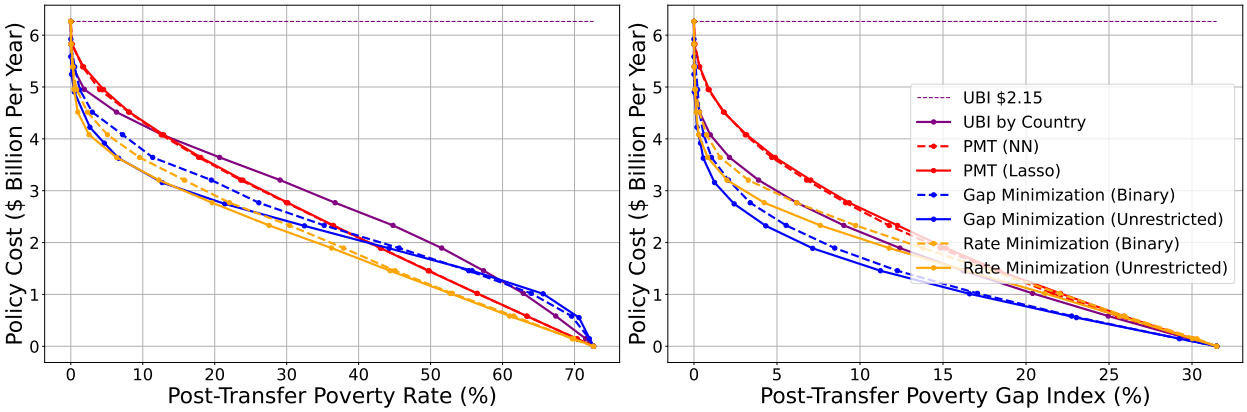
(f) Guinea-Bissau



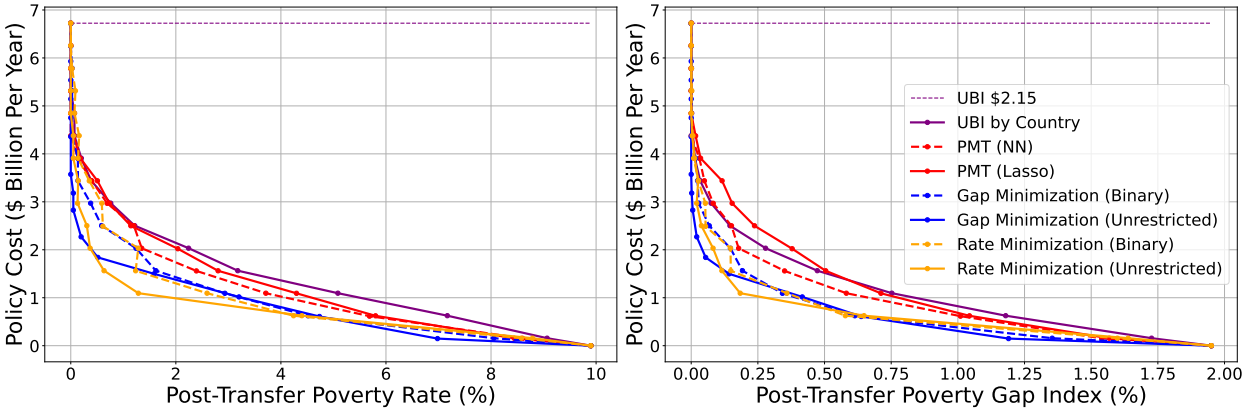
(g) Kenya



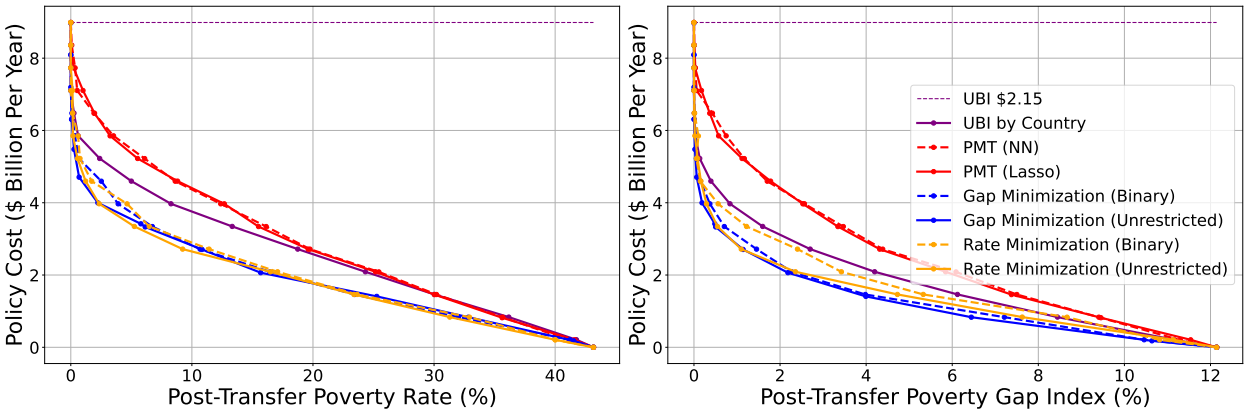
(h) Malawi



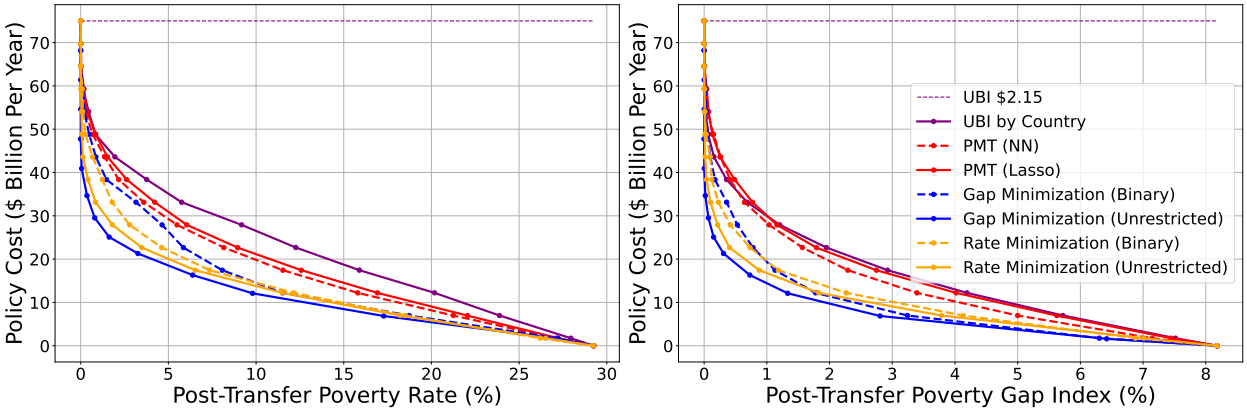
(i) Mali



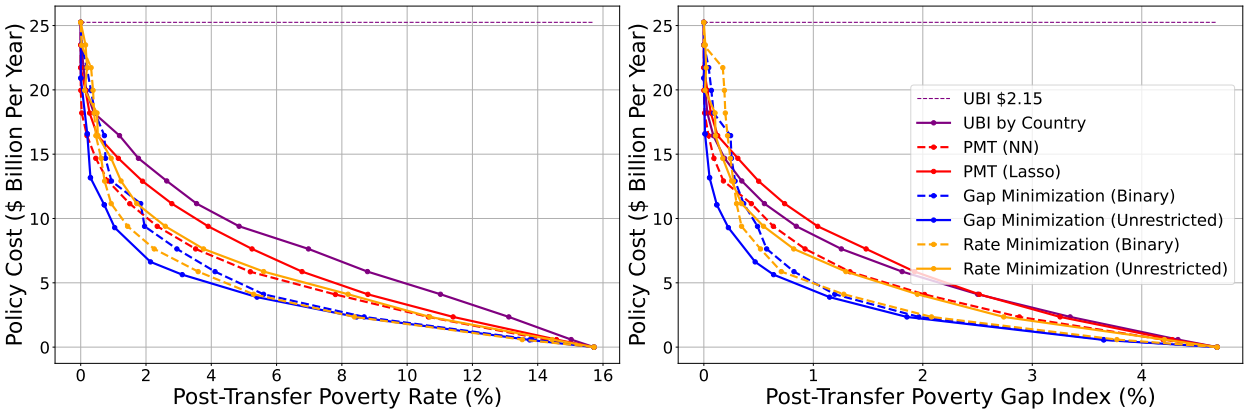
(j) Niger



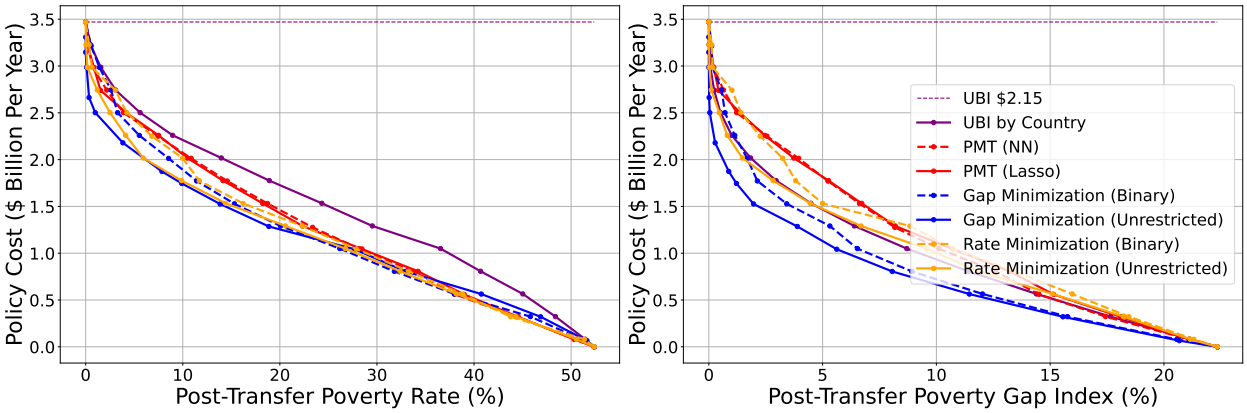
(k) Nigeria



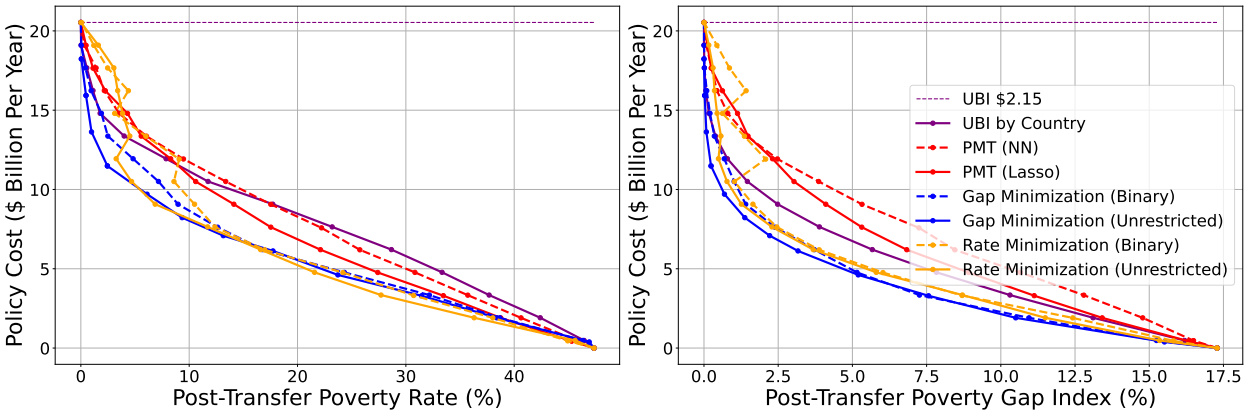
(l) South Africa



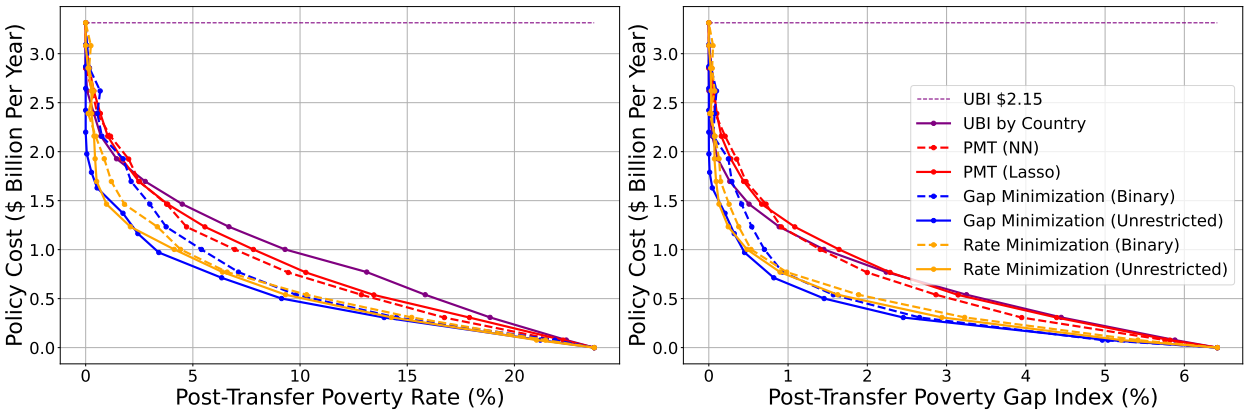
(m) South Sudan



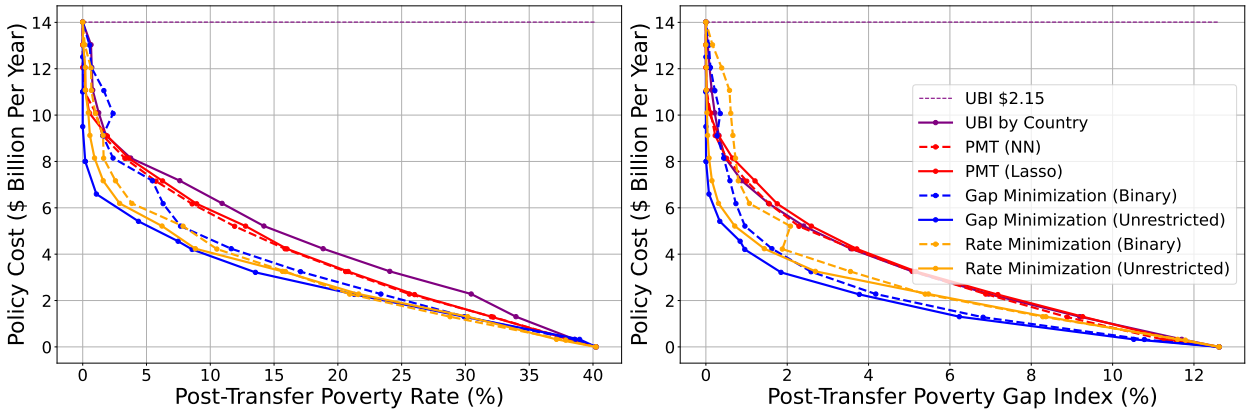
(n) Tanzania



(o) Togo



(p) Uganda



This figure reports country-by-country results for gap minimization (binary/unrestricted), rate minimization (binary/unrestricted), PMT-thresholding (with neural network / with lasso), UBI $2.15, and UBI at a country-specific amount.