

What Would it Cost to End Extreme Poverty?

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

December 9, 2025

Abstract

We study poverty minimization via direct transfers, framing this as a statistical learning problem while retaining the information constraints faced by real-world programs. Using nationally representative household consumption surveys from 23 countries that together account for 50% of the world’s poor, we estimate that reducing the poverty rate to 1% (from a baseline of 12% at the time of last survey) would cost \$170B nominal per year. This is 5.5 times the corresponding reduction in the aggregate poverty gap, but only 19% of the cost of universal basic income. Extrapolated globally, the results correspond to a cost of (approximately) ending extreme poverty of roughly 0.3% 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 BREAD/NBER Program Meeting, Brown University Bravo/SNSF Workshop, Cornell University Young Researchers Workshop, Duke University, UChicago AI in Social Science Conference, UChicago Machine Learning in Economics Summer Conference, Harvard Data Science Institute, University of Maryland Workshop on AI and Analytics for Social Good, INFORMS, Stanford Econometrics, Stanford Statistics, and Yale University. We thank Apollinaire Abi, Alia Aghajanian, Vincent Armentano, Anik Ashraf, Oscar Barriga Cabanillas, Lina Cardona-Sosa, S. Chandrasekhar, Oeindrilla Dube, Elizabeth Foster, Tanzed Hossain, Daniel Mahler, Juan Manuel Monroy Barragan, Jose Mpesayabo, Abhiroop Mukhopadhyay, Fabrice Mukendi Mutombo, Iary Rakotondradany, Khaled Saad, David Stifel, and Jonathan Weigel for their generous assistance accessing and interpreting data; Puneet Bhaskar, Qiqi Chen, Nolan Chu, Michael Cleary, Karl Gorski, Neela Kaushik, Henry Lopez Rodriguez, Matthew Mesfin, Francesca Vescia, and Victor Zhenyi Wang for help with data preparation; and Aditi Acharya, Erick Rosas Lopez, 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 living in extreme poverty—defined as living on less than \$2.15 in 2017 PPP dollars per day¹—has declined dramatically, from an estimated 41% in 1981 to 8% in 2024.² This progress has historically taken place in conjunction with economic growth (Pritchett, 2024a). Yet it may not continue: analyzing current growth trends, Roser (2025) concludes that “progress against extreme poverty will come to a halt.” And in either case the fact that hundreds of millions of people live below even very low consumption thresholds today raises an immediate policy question: could extreme poverty be eliminated more rapidly through redistribution?

Recent proposals explicitly appeal to this possibility (e.g., the wealth transfers envisioned by Kharas and McArthur, 2023). And back-of-the-envelope calculations raise hope that the sums involved might be manageable. If policymakers could identify every household below the poverty line and transfer exactly the income needed to reach it, the total annual cost would equal the global poverty gap, estimated as of 2023 to be just 0.08 percent of global GDP.³

However, such transfers are not feasible. Poverty gap estimates rely on household surveys that typically cover only 0.001 to 0.01 percent of a country’s population. Thus, for almost all households in low- and middle-income countries, there are no direct estimates of their standard of living. Collecting that information through regular surveys would be prohibitively expensive (Kilic et al., 2017) and, if used to determine transfer eligibility, would create incentives for misreporting hard-to-verify quantities like consumption (Martinelli and Parker, 2009; Banerjee et al., 2020). These cost and incentive considerations help explain why many existing programs instead determine eligibility using proxies for living standards that are cheaper to measure and verify (Hanna and Olken, 2018). For example, in the common Proxy Means Test (PMT) approach, administrators use a small amount of survey data to model consumption as a function of proxies (such as housing materials) and then collect those proxies in a census in order to determine eligibility in the full population (Grosh and Baker, 1995).

Learning to Target Income 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 recent work that

¹Throughout the paper we refer to extreme poverty as having a material standard of living below some given threshold, a definition which may differ from local communities’ concepts of poverty (Alatas et al., 2012). We focus primarily on the traditional threshold of \$2.15 in 2017 PPP dollars, but also report results for the newer threshold of \$3.00 in 2021 PPP dollars.

²World Bank Poverty & Inequality Platform, <https://pip.worldbank.org/poverty-calculator>, accessed 7 November 2025. See also Armentano et al. (2025) for descriptive facts about poverty decline.

³Authors’ calculations using countries’ poverty gap indices from the 2023 World Poverty Clock, population in 2023, 2017 market exchange rate, and PPP conversion factor for all countries in the World Bank PIP, and estimate of 2023 global GDP of 107T from <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>.

proposes to learn decision rules via loss minimization (Athey and Wager, 2021; Bertsimas and Kallus, 2020; Kitagawa and Tetenov, 2018). The policymaker learns a transfer policy that maps a household’s observable characteristics to a nonnegative transfer amount, minimizing expected loss subject to a budget constraint (or, equivalently, minimizing cost subject to a constraint on loss). We consider both unrestricted transfer policies, which allow each household to receive any nonnegative transfer, and binary policies in which each household receives either zero or some common positive amount. The domain of the loss function is the household’s post-transfer standard of living.⁴

Choosing the right loss function is a central concern. While the headcount poverty rate is the most easily understood and most widely used metric in public discourse about poverty,⁵ it is well known that, with full information, minimizing the poverty rate directly can generate inequitable allocations (Atkinson, 1987; Bourguignon and Fields, 1990; Sen, 1976). Intuitively, the most efficient way to reduce the rate would be to allocate transfers to households just below the poverty line, rather than to the poorest households. We show that this undesirable property generalizes to the incomplete-information case, and that minimizing the poverty gap instead is equitable in the sense that the resulting policies direct greater support to poorer households. Gap-minimization also delivers formal guarantees on poverty rate reduction: it is equivalent to minimizing the worst-case (across population subgroups) conditional poverty rate, and equivalent to rate minimization in some settings where rate minimization is itself equitable. Our main discussion therefore emphasizes the performance of gap-minimizing policies, while reporting rate-minimizing results for comparison.

The Cost of Targeted Transfers We first characterize solutions to the gap- and rate-minimization problems analytically in order to (a) reduce the dimensionality of space of possible transfer policies that must be searched, and (b) identify specific functionals of the population distribution that must be learned. Gap minimization requires learning conditional quantiles of the distribution of living standards, letting us use modern approaches based on neural networks and large predictor sets, while rate minimization requires learning conditional densities, which require a bespoke semi-parametric approach and smaller predictor sets. We use these results to construct computationally feasible algorithms for implementing both approaches.

We then quantify the aggregate cost of reducing poverty in a sample of 23 countries, including essentially all of those which (i) have an extreme poverty rate above 10% or account for more than 1% of the world’s extreme poor, and for which (ii) a recent, high-quality, nationally representative living standards survey that

⁴To 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. We abstract from dynamics, but consider 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.

⁵For 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).

includes a consumption aggregate and is large enough to cover at least 1,000 poor households is publicly available. The first condition aims to focus the analysis on the countries that matter most, while the second aims to ensure the reliability of our living standards measures and to ensure that we observe enough poor households for effective statistical learning.⁶ Each could be relaxed in the future. Collectively this sample accounts for 50% of the world’s extreme poor.⁷

We learn policies that target transfers to households based on observable characteristics that have been used by, or are analogous to others that have been used by, existing real-world PMT exercises. We further omit those that we deem too difficult to verify in practice (e.g., income) or likely to be politically controversial (e.g., ethnicity). That said, one should think of the exercise as characterizing what is technically feasible; if policies needed to satisfy additional criteria in order to be feasible politically, this might affect their costs.

We estimate that reducing the \$2.15 2017 PPP poverty rate to 1% in all countries in our sample, using a gap-minimizing (and hence weakly equitable) policy, would cost \$170B nominal per year. This is substantially lower than the cost of blunter, uniform policies: it is 19% of the cost of providing a Universal Basic Income at the level of the poverty line and 52% of the cost of achieving the same poverty rate using supplemental income at amounts that are universal within countries but vary across them (“Universal Supplemental Income”). These comparisons highlight the relative benefit of targeting with observable proxies of consumption over not targeting at all. On the other hand, the cost of the feasible gap-minimizing policy is 5.5 times the achieved reduction in the poverty gap, or equivalently 5.5 times the cost of achieving the same gap reduction if all households’ living standards were perfectly observable (the “oracle” scenario). This comparison quantifies the relative cost of targeting with proxies versus perfect information on consumption. Both ratios vary substantially across countries: in countries with relatively low poverty rates the benefits of targeting relative to universal transfers are very high, for instance, while in the some of the poorest countries (e.g. Malawi) universal transfers cost little more than the optimal targeted policy.

Real-world transfer policies are often simpler than the optimal gap-minimizing one. If we restrict attention to binary policies, the cost of achieving poverty rates of 1% increases by 43%, to \$243B nominal per year. This implies that, for a given poverty-reduction goal, the fiscal cost of using a simpler “eligibility rule” structure is meaningful. Reducing the poverty rate implied by the newer \$3.00 2021 PPP poverty line to 1%, meanwhile, costs 1.6 times as much as doing so under the older \$2.15 2017 PPP line, reflecting the fact that the revision

⁶Given its importance for global poverty we also include India, where we constructed the consumption aggregate by hand. The resulting full list of countries covered as of this draft is Benin, Burkina Faso, Bangladesh, Côte d’Ivoire, Colombia, Ethiopia, Ghana, Guinea-Bissau, Indonesia, India, Kenya, Madagascar, Mali, Malawi, Niger, Nigeria, Rwanda, Senegal, Togo, Tanzania, Uganda, Yemen, and South Africa. See Section 3 for further discussion and Appendix D for details of the data sources.

⁷The 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 corresponding survey year.

implied a meaningful real shift in the definition of poverty (Foster et al., 2025) and correspondingly more ambitious goals, as Pritchett (2024b) and others have called for. 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 weakly *more* expensive, for a given reduction in the rate, than policies learned to minimize the gap. The driving factor behinds this somewhat surprising result is that empirical rate minimization must solve a more challenging statistical problem than empirical gap minimization, requiring that we estimate the full conditional density functions as opposed to only their quantiles. 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. For our estimates to be valid, financing for the policies we learn would need to be incremental, not diverted from existing programming that separately affects living standards. A meaningful share of that incremental financing could plausibly come, in many countries, from domestic revenue. But most countries would likely also 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 recent evidence on these margins. Recent work on the general equilibrium effects of cash transfers, for example, has documented substantial expansionary effects (Egger et al., 2022; Gerard et al., 2024), which we necessarily abstract from here. We discuss what scale of wealth transfers, as proposed by Kharas and McArthur (2023), might be required to achieve income gains equivalent to the income transfers we study here. We also discuss the potential administrative costs of implementing policies like those we learn; while non-trivial, these would likely be orders of magnitude smaller than the costs of the transfers themselves.

Finally, we ask what the results suggest about the cost of (approximately) ending poverty globally, using the results from the 23 countries where we have data to predict the cost of ending poverty in countries where we do not. This extrapolation implies that it would cost \$318B per year to reduce the global poverty rate to 1%.⁸ This is 0.3% of global GDP, and 0.49% of OECD GDP.

Related Work Our work sits at the intersection of two long-standing traditions of thought about economic development. The first begins with an ambitious development goal and then seeks to calculate the resources required to reach it. 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 social objectives, the UN

⁸Note that achieving a 1% global rate requires reducing country-specific poverty rates to 1.4%, since some countries have essentially no extremely poor people.

Millennium Project estimated that its proposed Millennium Development Goals (MDGs) could be achieved if rich countries increased foreign aid to 0.54% of their GNI in 2015, from 0.25% in 2003 (Sachs, 2005). Other contemporary analyses focused specifically on the cost of achieving the MDG for extreme poverty, and are thus more closely related to ours. These studies calculated the growth in GDP per capita required to achieve a given poverty rate (assuming that the relative distribution of income stayed fixed) and then inferred the aid needed to generate that growth, using either a “two-gap” model (e.g., Devarajan et al., 2002) or 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 empirically rejected.⁹ This was necessarily, in the words of Devarajan et al. (2002), “a highly speculative exercise.” Our approach, 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. A natural point of departure is the body of older theoretical work on poverty measurement that assumed full knowledge of the income or consumption distribution (e.g., Foster et al., 1984; Greer and Thorbecke, 1986; Sen, 1976). A logical next question would then be how to extend these ideas to environments where policymakers observe only limited information about households. Along these lines, Kanbur (1987) and Ravallion and Chao (1989) discussed 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 observable covariates via the plug-in principle, i.e., by minimizing a given poverty metric in-sample. In practice, however, most subsequent work has focused on the more straightforward problem of classifying households as poor or non-poor to determine their eligibility for a fixed transfer. Common methods for doing so include proxy means testing (Alatas et al., 2012; Brown et al., 2018; Grosh and Baker, 1995; Noriega-Campero et al., 2020) including recent approaches that use non-traditional data (Aiken et al., 2022; Smythe and Blumenstock, 2022); community-based targeting (Alatas et al., 2012; Coady et al., 2004); and geographic targeting (Baker and Grosh, 1994; Bigman and Fofack, 2000). Our approach differs in that we seek to learn a policy that directly minimizes a poverty criterion subject to resource constraints, rather than predicting a binary eligibility outcome. With modern computational tools and high-dimensional statistical learning methods, revisiting Glewwe’s (1992) formulation is now feasible.

⁹For 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. 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.”

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

The methods we develop also contribute to a broader 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). The problems we study are “prediction-policy problems” in the sense of Kleinberg et al. (2015), and do not require estimates of causal effects, as in recent work on learning treatment assignment policies (e.g., 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 one methodological parallel is that it can be useful to reformulate budget-constrained treatment assignment 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 fixed-size cash transfers in the presence of treatment effect heterogeneity.¹¹ Future work might combine methods for learning heterogeneous effects with approaches like ours to learn dynamic transfer policies that efficiently minimize poverty over time, taking advantage of the fact that some recipients have predictably higher investment returns (i.e., are more likely to “graduate”) in response to a large initial transfer.

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. A country has a population distribution F over these variables; while the results in this section will involve this population distribution, they anticipate the idea that we will be able to draw representative samples from it for learning.

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. $t : \mathcal{X} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. This must satisfy the budget constraint $\mathbb{E}_F [t(X; B)] \leq B$.

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

¹¹For an overview of empirical evidence on the effects of cash transfers, see Bastagli et al. (2019) and Crosta et al. (2024).

¹²There is also a distinct question whether the *rule* that determines eligibility for transfers might reduce Y_i by disincentivizing

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 $\mu(x) := \mathbb{E}_F[Y_i | X_i = x]$. The conditional mean consumption μ is then thresholded to determine the allocation of transfers to units the population, i.e.

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

for some transfer size $\bar{t} \in \mathbb{R}_+$ selected by the policymaker (but not necessarily learned from data) and eligibility threshold $\eta(\bar{t}, B)$ chosen so that $\mathbb{P}_F[\mu(X) \leq \eta(\bar{t}, B)] = B/\bar{t}$.

Note that the PMT policy (1) is a binary policy: each unit can receive zero or some common amount \bar{t} .

The PMT has a connection to *statistical* loss minimization in the sense that $\mu(\cdot)$ is the function that minimizes $\mathbb{E}_F[(Y_i - \hat{\mu}(X_i))^2]$, the mean-squared error between Y_i and a predictor $\hat{\mu}(X_i)$. But this does not imply that the transfer policy it yields minimizes any particular post-transfer loss. More generally, there is no obvious formal link between this procedure and minimization of any *welfare* loss. 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 $L : \mathbb{R} \rightarrow \mathbb{R}$, an optimal unrestricted transfer policy is one which minimizes the expected post-transfer loss $\mathbb{E}_F[L(Y_i + t(X_i))]$ subject to a budget constraint B , i.e. which solves

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

Given data $(X_i, Y_i) \sim F$ i.i.d. for units $i = 1, 2, \dots, n$ 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 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.

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 $L : \mathbb{R} \rightarrow \mathbb{R}$, an optimal binary transfer policy is one which minimizes the expected loss $\mathbb{E}_F [L(Y_i + t(X_i))]$ subject to a budget constraint B and takes on only one non-zero value, i.e. which solves

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

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 Testing 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 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

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

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

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

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. The latter property generalizes the requirement that transfers be monotone decreasing in (pre-transfer) consumption.¹⁴ To build intuition for the generalization it may be useful to consider gap minimization in the full-information case; here, the weakly equitable policy amounts to a floor on post-transfer consumption.

¹³In 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

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

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

¹⁴To see this, consider (6) with $B > 0$ and $B' = 0$; then

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

Clearly, rate minimization is not per se weakly equitable in the perfect information regime, i.e. when $X_i = Y_i$. 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 1 (Imperfect Proxies). The conditional distribution $F_{Y|X=x}$ is supported on an interval $\mathcal{Y} \subseteq \mathbb{R}_+$ that contains $[0, c]$, has continuously differentiable density $f_{Y|X=x}$, 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. *Consider Definition 2 with some loss function L that is decreasing, integrable on \mathbb{R}_+ , and bounded on \mathbb{R}_+ . If L is nonconvex, then there exists a population distribution F , satisfying Assumption 1, under which the optimal policy family $t_L(x; B)$ is not weakly equitable.*

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. *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 1, 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 holds. Consider gap minimization as specified in Example 2. The optimal policy that solves (2) also solves*

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

¹⁵Lemma 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.

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

Assumption 2 (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 2 holds. 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.

3 Data Sources and Uses

3.1 Data sources and preparation

Appendix D.1 describes in full detail the criteria and exceptions by which we select countries for inclusion in the analysis. In brief, we begin with those with poverty rates above 10% or that account for more than 1% of the world’s poor, to focus attention on the countries that matter most for global poverty. Among these we limit attention to countries with data publicly available from a recent, high-quality, nationally representative LSMS-style survey, including a consumption aggregate, and that we expect (given the sample size and published poverty rate estimates) to contain at least 1,000 poor households. These conditions are meant to ensure that we can reliably learn accurate national-level policy costs. We make a few further adjustments to the resulting list: including Bangladesh and India because of their size, and omitting Afghanistan and Somalia because we were not able to obtain reliable PPP conversion rates for the relevant years. In total this yields a sample of 23 countries that vary meaningfully in their headcount poverty rates, from 3% (in Colombia) to 70% (in Malawi), and that collectively account for 50% of the world’s extreme poor.¹⁶

Table E.2 lists the specific household surveys, in the indicated countries and years, for which we learn transfer policies. These are in most cases the surveys from which recent estimates of the headcount poverty rate in the World Bank’s Poverty and Inequality Platform (PIP, pip.worldbank.org) were calculated, and are part of the World Bank’s flagship Living Standards Measurement Study (LSMS) program to collect high-quality, standardized household surveys on poverty and welfare (World Bank, 2014). In a few cases, the data underlying PIP estimates are not publicly available, and we therefore use an alternative or earlier source. All surveys are designed to be nationally representative; that said, one caveat is that they are thought

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

to under-represent refugees, internally displaced persons, and stateless populations, particularly when these groups are concentrated in institutional settings such as camps (EGRIS, 2024).¹⁷

Several years have passed since most surveys were conducted, and our estimates should thus be interpreted as the cost of reducing poverty at the time of the most recent survey. More generally, because the estimates draw on information about living standards and predictors collected at the same time, one should think of them as representing the costs of ending poverty if poverty measurement surveys were conducted regularly. One could also use data collected in one year to learn policies administered in future years, but because poverty tends to fluctuate substantially over time (Merfeld and Morduch, 2024; Armentano et al., 2025) this would reduce their accuracy, and one would want to learn policies that explicitly account for this, as for example Aiken et al. (2025b) do for proxy means tests.¹⁸

All surveys collected detailed information on household expenditure and self-production, and in most cases their public data releases include a consumption aggregate constructed from this information. The exception is India’s 2022-23 Household Consumption Expenditure Survey for which we construct our own; see Appendix D.2.1. We compute household per capita consumption in 2017 PPP USD by converting the 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.¹⁹ The estimated headcount poverty rates this procedure yields are generally close to the World Bank estimates for the corresponding country and year, and match exactly in most cases where both estimates derive from the same survey. See D.2 for details on outcome construction.

Selecting covariates to use as predictors of living standards 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 actually implemented in practice in a variety of low- and middle-income countries.²⁰ We take the union of the predictors used in these PMTs, which yields

¹⁷As an upper bound on the magnitude of this potential omission, the cost of providing a UBI of \$2.15 per day in 2017 PPP to *all* refugees and IDPs identified by the UNHCR as of 2023 is 0.04% of global GDP, increase the total cost of the policy we derive below from 0.3% to 0.34%. This calculation excludes refugees residing in Cuba, Hong Kong SAR China, Turkmenistan, Liechtenstein, Sint Maarten (Dutch part), Cayman Islands, and Monaco for which the market exchange rate in 2017 and PPP conversion factor are unavailable from the World Bank. Authors’ calculations based on <https://www.unhcr.org/refugee-statistics/download>, accessed 7 December 2025.

¹⁸See also Beuermann et al. (2025).

¹⁹The 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.

²⁰These 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).

a list spanning six broad categories: household demographics, human capital, household assets, livelihood activities, geographic indicators, and community 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. We also omit indicators for the ethnicity, tribe, religion, or caste of the household or household head; while these have been used in past PMTs, they raise questions of ethical and political suitability in the given context. We prefer to err on the conservative side and omit them.²¹ Appendix D.3 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.²²

3.2 Empirical methods

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

²¹We thank Lant Pritchett for suggesting this.

²²Our datasets are available [here](#).

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 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 $d = 40$ and $d = 411$ predictors (last column of Table E.2). That said, we will also find (below) that we can use substantially smaller predictor sets without meaningful performance degradation. 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;²³ 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 by district 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,


²³Our data use plan is available [here](#).

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.

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, and verify that our budget constraint is enforced on the actual set of test set units we consider deploying the transfers on. This paradigm, which is often referred to as the transductive setting (Vapnik, 1998), mirrors typical practice in real-world transfer program design. Finally, we use the held-out outcomes Y_i from the test set to evaluate the transfer schedules. This is the relevant paradigm when the policymaker expects to 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 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, we will see below that the estimated benefits of targeting in Malawi (as measured by the ratio of the cost of optimized transfers to the cost of universal ones; see Figure 5) is the *lowest* of any country in our sample, the opposite of what one might expect had our experimentation led to over-fitting. 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 have been modifying the analysis weights to reflect household size as well as survey sampling weights (which we had initially neglected to specify), and removing indicators related to ethnicity, race, caste, religion, or other such demographic groups (in response to audience feedback).

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



This figure reports the costs, in billions of nominal 2023 USD annually, of the indicated optimal transfer policies learned for Togo ($n = 6171$, $d = 162$). 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.

4 Empirical Results


To build intuition we first present results for a single country, Togo, in Section 4.1. The choice of Togo is largely arbitrary, but it is a representative country in the sense that our results place it near the middle of the sample in terms of the “difficulty” of targeting, as quantified below in Figure 4. We then apply our methods to the full sample and present the aggregate policy cost of achieving uniform poverty goals in Section 4.2.

4.1 Togo Case Study

Figure 1 presents results based on Togo’s Harmonized Survey on Households Living Standards 2018-2019, a survey of $n = 6171$ households that contains $d = 162$ plausibly verifiable covariates. The figure 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 (either the 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 the Togo case study, which in general will also hold in the other countries we analyze. First, as one would expect, marginal costs increase as we approach zero poverty according to either metric. Literally ending extreme poverty requires a policy effectively amounting to UBI in order to reach every last poor household, but one can come close at far lower costs. Second, unrestricted policies are less costly than binary ones (i.e., the solid lines fall below 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

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 resulting population-weighted post-transfer poverty rate, and the right-hand panel plots these against the resulting population-weighted post-transfer poverty gap index. See Section 4.2 for series definitions.

are fairly modest. And third, while gap minimization is as or more effective at reducing the gap as rate minimization, the converse does not hold: rate minimization is generally no more effective at minimizing the rate than gap minimization. Practically, this implies that one can reduce the poverty rate using the weakly equitable policies obtained from gap minimization at no penalty in terms of cost-effectiveness.

4.2 Multi-Country Analysis

We now present results aggregated across the full sample of 23 countries, with the underlying country-specific results available in Appendix E. We aggregate by calculating the cost of achieving a given poverty rate in *every* country in our sample, and summing these. These are thus conservative estimates relative to the cost of achieving the same poverty rate on *average*, since this approach imposes 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 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 population-weighted headcount poverty rate (left-hand panel) or poverty gap index (right-hand panel). 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 “USI 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, which we refer to as Universal Supplemental Income. 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),²⁴ but it is not a policy that could be feasibly implemented since program administrators do not directly observe the exact amount needed by every household in the population.

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 \$170B per year, which is 5.5 times the oracle cost of (\$31B), but 19% of the cost of a \$2.15 UBI, and 52% of the cost of a country-specific Universal Supplemental Income. 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 59% the cost of proxy means testing, highlighting the performance benefits of the loss minimization approach. Lastly, the binary gap-minimizing policy costs 43% 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. The right-hand panel shows, on the other hand, that a policy

²⁴We 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.

optimized to reduce the gap using unrestricted transfers actually perform somewhat *better* at reducing the rate than empirical rate minimization. In this sense there is little reason to prefer rate minimization in terms of performance.

The driving factor behind this somewhat surprising result is that empirical rate minimization must solve a more challenging statistical problem than empirical gap minimization, requiring that we estimate the full conditional density functions as opposed to only their quantiles. A first, immediate consequence of this challenge is that empirical rate minimization has worse sample complexity, i.e., its errors decay more slowly with sample size. A second, more subtle consequence is that conditional density estimation requires use of a bespoke algorithm (unlike conditional quantile estimation which can be implemented using off-the-shelf neural networks)—and unsurprisingly we thus find our algorithm to be less robust at handling complex and/or high dimensional data. We seek to mitigate this issue by implementing data-driven feature selection for our rate minimization algorithm;²⁵ realistically, however, it is unlikely this procedure is able to make use of as many features as gap minimization.²⁶


The full country-specific details underlying these results are presented in Appendix E in Figure E.2. Generally speaking, the comparisons above for the sample of countries as a whole generally hold within each country as well²⁷. Reducing the poverty rate to 1% with an optimal policy costs between 2.4 and 13.6 times

²⁵No such feature selection is required for the quantile regression algorithm underlying our gap-minimization procedure, as off-the-shelf neural network implementations are already highly robust to receiving high-dimensional inputs.

²⁶That said, we find that in practice gap minimization using the restricted set of features use for rate minimization performs almost identically (and, in particular, continues to reduce the rate more cost-effectively) indicating that algorithmic performance per se is the main driver (Figure E.1).


²⁷For certain countries and targeting methods, we observe that the cost of learned policies can be non-monotone decreasing in the post-transfer poverty rate and gap. This behavior arises because our nuisance parameters and metrics are estimated with error in finite samples. Binary gap and rate minimization are especially prone to exhibiting non-monotonocities because the optimal solutions for these problems can rely on different nuisance parameters even for similar budgets. In contrast, unrestricted gap minimization is less fragile in this respect because similar budgets will yield solutions that rely on similar

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.

Figure 5: Cost Ratio of Universal Supplemental Income 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.

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 a Universal Supplemental Income 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.

nuisance parameters with correlated errors.

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.

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, 2024b; 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 \$268B per year in nominal 2023 dollars. This is 1.6 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.

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

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.


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. Figure 7 plots the amount of money spent in each country, under the learned policy which reduces the national poverty rate to 1%, as a proportion of that country’s GDP (top panel) and total government revenue (bottom panel). In the average country in our sample, implementing such a scheme would cost 11% 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, it appears plausible that many countries could contribute substantially to the cost of policies like these, but implausible that many could finance them entirely on their own. This 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

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).

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 5% of GDP in 2023 and would require 11% 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.²⁸ Any such knock-on benefits need not accrue proportionally, however—in fact, Haushofer et al. (2025) 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. 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. (2025) 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

²⁸While 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.

from year to year (or even month to month; see [Merfeld and Morduch, 2024](#)). A related issue is that, as noted above, policies learned from data at a given point in time may not perform as well in subsequent years if the relationship between consumption and PMT covariates changes over time ([Aiken et al., 2025a](#)). A natural extension of our framework would incorporate techniques for addressing such ‘covariate shift’ in the statistical learning framework of Section 2 (cf. [Quinonero-Candela et al., 2008](#); [Koh et al., 2021](#)).

Finally, the gap-minimizing policies we estimate quantify the total cost of transfers required to eliminate poverty, but abstract from the administrative costs of implementing such transfers at scale. In practice, real-world programs will incur costs to identify eligible households, enroll them, and disburse payments. Evidence from costing exercises for national social registries indicates that one-off, census-style registration campaigns generally cost on the order of US\$1–3 per household registered ([Lindert et al., eds, 2020](#)).²⁹ A second component of implementation cost arises from payment delivery: governments typically compensate payment providers either through small flat per-transaction fees (fractions of a US cent to a few US cents) or through ad valorem fees of roughly 1–3% of the amount transferred, as in Pakistan’s BISP, India’s Direct Benefit Transfer programs, and the Philippines’ 4Ps ([Lindert et al., eds, 2020](#)).³⁰ Relative to the total value of cash transfer policies we evaluate, these fixed and recurring costs thus appear likely to be modest. In countries without foundational identification or payment infrastructure, larger one-time investments may be required to establish those systems.

5.1 Costs at global scale

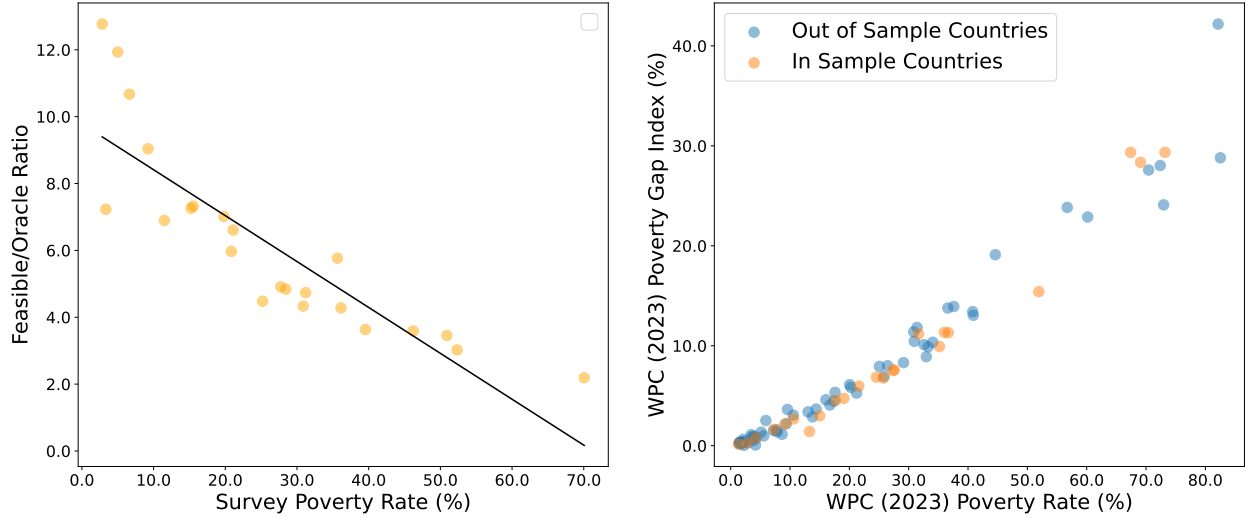
While the analysis above estimates costs for a sample of countries, as of the dates they were surveyed, it is natural to wonder what it implies for the cost of ending poverty at a global scale today. We consider in particular the cost of reducing the global poverty rate in 2023 to 1%. To do so it is sufficient to obtain a national poverty rate of 1.4% in all countries that currently have higher rates, as some other countries already have rates lower than 1% which bring down the (population-weighted) average.

To estimate the costs of doing so, we first estimate the cost of (equitably) achieving a 1.4% rate in each country in our sample. We then estimate, within that sample, the relationship between (i) the ratio of the

²⁹For example, Malawi’s Unified Beneficiary Registry reports projected Phase 2 registration costs of approximately US\$1.27 million for 785,000 households, or US\$1.74 per household ([Lindert et al., 2018](#)). For Turkey’s Integrated Social Assistance System, system development costs of roughly US\$13.1 million covering an estimated 10 million households imply a cost of about US\$1.30 per household ([Leite et al., 2017](#)). For Brazil’s *Cadastro Único*, reported registration and re-certification costs range between US\$2.03 and US\$2.06 per household ([Mostafa and Sátyro, 2014](#)). For Pakistan’s National Socioeconomic Registry, per-household registration costs are estimated between US\$2.30 and US\$2.70 ([Lindert et al., eds, 2020](#)).

³⁰Transaction fee structures vary across programs but typically fall within a modest range. In Pakistan’s BISP, the program now pays a 3% transaction fee ([Lindert et al., eds, 2020](#), pp. 239–240). In India’s Direct Benefit Transfer system for MGNREGS payments, the Government pays a flat Re 0.50 per transaction (split across the sponsor bank, destination bank, and the National Payments Corporation of India), with an additional 1% service fee permitted for payment service providers ([Lindert et al., eds, 2020](#), p. 240). In the Philippines’ 4Ps program, the government negotiated a simplified service agreement that kept bank fees “under 1%” of the payment value ([Lindert et al., eds, 2020](#), p. 239).

Figure 8: In-sample fit and common support



The left-hand panel plots the in-sample relationship between the feasible-to-oracle policy cost ratio and the poverty rate, calculated from the various surveys we use. The right-hand panel plots the relationship between the poverty gap index and the poverty rate, using estimates as of 2023 from the World Poverty Clock, differentiating between in-sample and out-of-sample countries.

feasible policy cost to the poverty gap—which one can think of as a measure of how much one must “overpay” due to imperfect information—and (ii) the poverty rate. Intuitively, we expect the ratio (i) to be higher in countries with lower poverty rates, since targeting there is a harder problem. The left-hand panel of Figure 8 shows that this is the case, and that we obtain a tight and approximately linear fit, with an R^2 of 0.74.³¹ We then use this estimated relationship, combined with estimates of poverty gaps and rates as of 2023 for all countries from the World Poverty Clock, to project the cost achieving a poverty rate of at most 1.4% in each country. This exercise necessarily involves extrapolation across space and across time; that said, the supports of the distributions involved overlap well (see the right-hand panel of Figure 8).

The net result is an estimated cost of reducing the global poverty rate to 1% equal to 0.3% of global GDP.³² This is large in relation to some relevant benchmarks: OECD donors gave foreign aid equal to 0.21% of global GDP in 2023,³³ and aid levels have since fallen sharply, particularly in the United States. Whether resuming or even increasing historical levels of foreign aid could be a winning political stance is unclear, though Fabre (2025) has recently estimated that there is majority stated support in eleven high-income countries for even higher levels of global redistribution. In terms of sheer fiscal feasibility, however, there is

³¹Note that we truncate fitted values so that (conservatively) they never fall below 1. We have also re-estimated results using a quadratic fit, in which case the total estimated cost changes from 0.3% to 0.34% of global GDP. Hong Kong SAR China, Kosovo, Macao SAR China, Puerto Rico, Turkmenistan, Taiwan, and U.S. Virgin Islands are excluded from this calculation because the corresponding conversion factors are not available from the World Bank.

³²Or 0.49% of OECD GDP. Including China, the world’s second-largest economy and now a major international aid donor, lowers this figure to 0.38%.

³³In 2023, OECD foreign aid totalled \$223B (<https://www.oecd.org/en/events/2025/01/official-development-assistance-oda-2023-final-figures.html>, accessed 6 December 2025) and global GDP totalled 107T (<https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>, accessed 6 December 2025).

no question that wealthy countries could finance most or all of policies such as this one. And in a broader sense it implies that ending poverty is no more costly than other, arguably less critical global priorities: for instance, the world spends seven times as much—2.2% of global GDP—on alcoholic beverages.³⁴

³⁴Estimated 2024 global alcoholic beverages market size: \$2,413B as per <https://www.fortunebusinessinsights.com/alcoholic-beverages-market-107439>, accessed 1 September 2025. Estimated 2024 GDP: \$110,000B as per <https://www.imf.org/external/datamapper/NGDPD@WEO/OEMDC/ADVEC/WEOWORLD>, accessed 1 September 2025.

References

- Adler, Gustavo, Noémie Lisack, and Rui C. Mano**, “Unveiling the effects of foreign exchange intervention: A panel approach,” *Emerging Markets Review*, 2019, 40, 100620.
- Aiken, Emily, Joshua E Blumenstock, and Tim Ohlenberg**, “Moving Targets: Can Machine Learning Help Proxy Means Tests Better Account for Poverty Dynamics?,” *Working Paper*, 2025.
- 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.
- , **Tim Ohlenburg, and Joshua Blumenstock**, “Moving Targets: Can Machine Learning Help Proxy Means Tests Better Account for Poverty Dynamics,” Technical Report, UC Berkeley January 2025.
- Alatas, Vivi, Abhijit Banerjee, 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.
- Baker, Judy L and Margaret E Grosh**, “Poverty reduction through geographic targeting: How well does it work?,” *World development*, 1994, 22 (7), 983–995.
- 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.
- Beuermann, Diether W., Bridget Hoffmann, Marco Stampini, David L. Vargas, and Diego Vera-Cossio**, “Shooting a moving target: Evaluating targeting tools for social programs when income fluctuates,” *Journal of Development Economics*, 2025, *172*, 103395.
- Bhattacharya, Debopam and Pascaline Dupas**, “Inferring welfare maximizing treatment assignment under budget constraints,” *Journal of Econometrics*, 2012, *167* (1), 168–196.
- Bigman, David and Hippolyte Fofack**, “Geographical Targeting for Poverty Alleviation: An Introduction to the Special Issue,” *The World Bank Economic Review*, 01 2000, *14* (1), 129–145.
- 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 Rüschepöhler, and Christopher R. Udry**, “Unconditional Cash Transfers: A Bayesian Meta-Analysis of Randomized Evaluations in Low and Middle Income Countries,” August 2024.
- 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.
- EGRIS**, “Capturing Priority SDG Indicators in Refugee, Internal Displacement and Statelessness Contexts,” Technical Report 3, Expert Group on Refugee, IDP and Statelessness Statistics October 2024.
- Fabre, Adrien**, “Public Acceptance of International Redistribution in High-Income Countries,” Technical Report, CIRED November 2025.
- Fernandez, Luisa and Gracia Hadiwidjaja**, “Do Household Socioeconomic Status 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.
- 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.

Koh, Pang Wei, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao et al., “Wilds: A benchmark of in-the-wild distribution shifts,” in “International Conference on Machine Learning” PMLR 2021, pp. 5637–5664.

Leite, Phillippe, Tina George, Changqing Sun, Theresa Jones, and Kathy Lindert, “Social Registries for Social Assistance and Beyond: A Guidance Note & Assessment Tool,” Social Protection and Jobs Discussion Paper 1704, World Bank, Washington, DC 2017.

Lindert, Kathy, Colin Andrews, Chipo Msowoya, Boban Varghese Paul, Elijah Chirwa, and Anita Mittal, “Rapid Social Registry Assessment: Malawi’s Unified Beneficiary Registry (UBR),” Social Protection and Jobs Discussion Paper 1803, World Bank, Washington, DC 2018.

—, **Tina George Karippacheril, Inés Rodríguez Caillava, and Kenichi Nishikawa Chávez**, eds, *Sourcebook on the Foundations of Social Protection Delivery Systems*, Washington, DC: World Bank, 2020.

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.
- Martinelli, César and Susan Wendy Parker**, “Deception and Misreporting in a Social Program,” *Journal of the European Economic Association*, 2009, 7 (4), 886–908. Publisher: Oxford University Press.
- Merfeld, Joshua and Jonathan Morduch**, “Poverty at Higher Frequency,” Technical Report, NYU August 2024.
- Mostafa, Joana and Natália G. D. Sátyro**, “Cadastro Único: A Registry Supported by a National Public Bank,” Working Paper 126, International Policy Centre for Inclusive Growth, Brasilia, Brazil 2014.
- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny**, “Industrialization and the Big Push,” *Journal of Political Economy*, 1989, 97 (5), 1003–1026.
- National Sample Survey Office**, “Survey on Household Consumption Expenditure: 2022-23,” Technical Report No. 591, Government of India 2024.
- 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**, “The Case for Economic Growth as the Path to Better Human Wellbeing,” Technical Report, Harvard University October 2024.
- , “‘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.
- Quinonero-Candela, Joaquin, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence**, *Dataset shift in machine learning*, Mit Press, 2008.
- Ravallion, Martin and Calvin Chao**, “Targeted policies for poverty alleviation under imperfect information: algorithms and applications,” *Journal of Policy Modeling*, 1989, 11 (2), 213–224.
- Roser, Max**, “The end of progress against extreme poverty?,” *Our World in Data*, 2025. <https://archive.ourworldindata.org/20251125-173858/end-progress-extreme-poverty.html>.
- 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.
- Vapnik, Vladimir N.**, *Statistical Learning Theory*, New York: Wiley, 1998.
- 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.
- , “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) satisfies*

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

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

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

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 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

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

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

problem

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

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 $\widehat{\mathbb{E}}_F [t_{\lambda_j}(X_i)] = \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \hat{t}_{\lambda_j}(X_i)$ 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 3. 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

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

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 3 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 optimization 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 3 holds. If $L(z) := \mathbb{I}(z < c)$, (2) can be re-written as

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

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 4. 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 $z : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^K$, 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 $p : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^K$, where $p_k(x; \alpha) = \mathbb{P}_F [z_k(X; \alpha) + Y < c \mid X = x]$. 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 $\pi : \mathcal{X} \rightarrow [0, 1]^K$. Under Assumption 4, 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:

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

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 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 $[\alpha_1, \alpha_2, \dots, \alpha_J]$ 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.

<p>Input : Conditional distributions $\{f_{Y X=x}(\cdot)\}_{x \in \mathcal{X}}$, Covariate distribution $f_X(\cdot)$, Budget B, Threshold c, Grid size m</p> <p>Output: Transfer policy t^*</p> <pre> 1 costs $\leftarrow \emptyset$ 2 policies $\leftarrow \emptyset$ 3 Define α_{\max} using $\{f_{Y X=x}(\cdot)\}_{x \in \mathcal{X}}$. 4 for $\alpha \in [\frac{\alpha_{\max} + \delta}{m}, \frac{2\alpha_{\max}}{m}, \dots, \frac{(m-1) \cdot \alpha_{\max}}{m}, \alpha_{\max}]$ do 5 for $x \in \mathcal{X}$ do 6 $\mathcal{T}_\alpha(x) \leftarrow \text{ComputeAlphaValidTransfers}(f_{Y X=x}(\cdot), \alpha, c)$ 7 $\mathcal{C}_\alpha(x), \rho(x; \alpha) \leftarrow \text{ComputeLowerConvexHull}(f_{Y X=x}(\cdot), \mathcal{T}_\alpha(x))$ 8 end 9 <i>Apply Algorithm 2</i> 10 $t_\alpha, \text{cost} \leftarrow \text{SolveFractionalMCKnapsack}(f_X(\cdot), \{(\mathcal{C}_\alpha(x), \rho(x; \alpha))\}_{x \in \mathcal{X}}, B)$ 11 costs.append(cost) 12 policies.append(t_α) 13 end 14 $t^* \leftarrow$ Minimum cost policy in policies (based on costs) 15 return t^* </pre>
--

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 $k = 1, \dots, K$

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

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

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

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

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 4, 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 $B > \mathbb{E}_F[z_{k_{|\mathcal{C}_\alpha(x)|}}(X)]$, 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 $|\mathcal{X}| \cdot K \log |\mathcal{X}| \cdot K$.

<p>Input : Convex hulls and incremental cost-loss ratios $\{(\mathcal{C}(x), \rho(x))\}_{x \in \mathcal{X}}$, Covariate distribution $f_X(\cdot)$, Budget B</p> <p>Output: Transfer policy t, Poverty rate loss</p> <pre> 1 Initialize priority queue. 2 cost $\leftarrow 0$, loss $\leftarrow 0$ 3 queue $\leftarrow \emptyset$ 4 for $x \in \mathcal{X}$ do 5 Assign unit to first point on convex hull. 6 $k_1(x) \leftarrow$ First point on convex hull $\mathcal{C}(x)$ 7 cost $+= f_X(x) \cdot z_{k_1(x)}(x)$, loss $+= f_X(x) \cdot p_{k_1(x)}(x)$ 8 Enqueue next point on hull. 9 $k_2(x) \leftarrow$ Second point on convex hull $\mathcal{C}(x)$ 10 queue.add($(x, k_2(x))$ with priority $\rho_{k_1(x)}$) 11 end 12 while cost $< B$ and queue.size() > 0 do 13 $(x, k_j(x)) \leftarrow$ queue.pop() 14 Subtract current transfer to unit x from cost and loss. 15 cost $-= f_X(x) \cdot z_{k_{j-1}(x)}(x)$, loss $-= f_X(x) \cdot p_{k_{j-1}(x)}(x)$ 16 Allocate transfer $k_j(x)$ to unit x, record new cost and loss. 17 $t(x) \leftarrow z_{k_j(x)}(x)$ 18 cost $+= f_X(x) \cdot z_{k_j(x)}(x)$, loss $+= f_X(x) \cdot p_{k_j(x)}(x)$ 19 if cost $> \epsilon$ then 20 Perform fractional adjustment for unit x— updating t, cost, loss. 21 break 22 end 23 if there remain points on convex hull for unit x then 24 $k_{j+1}(x) \leftarrow$ Next point on $\mathcal{C}(x)$ 25 queue.add($(x, k_{j+1}(x))$ with priority $k_j(x)$) 26 end 27 end 28 return t, loss </pre>

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}

$$h(y | x) = h_0(y) \cdot \exp(s(y)^T \Theta x - \psi(\Theta x)). \quad (14)$$

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 $h(y | x)$ 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

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

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

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

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 testing as in Definition 1, 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

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

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

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

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 $t_{\bar{t}}(\cdot; B) := \bar{t} \cdot \pi_{\bar{t}}(\cdot; B)$.

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

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

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 $\rho_{\bar{t}, B}^* \in \mathbb{R}$, $a_{\bar{t}, B} \in [0, 1]$, 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, $\mathbb{P}_F [\rho_T(X) = \rho_{T, B}^*] = 0$ 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 $[\bar{t}_1, \bar{t}_2, \dots, \bar{t}_J]$. For each $j = 1, 2, \dots, J$, 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.

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

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

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

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

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

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

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

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}_{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 F is a population distribution over (X_i, Y_i) , $X_i = Y_i$ (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. $B \geq \mathbb{E}_F[(c - Y_i)_+]$.

B Proofs

B.1 Additional Lemmas

We define new notation for the following lemmas. Define

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

Lemma 9. *Suppose Assumption 1 holds. 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 holds. 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

We consider two cases. We first focus on the case where L is nonconvex in a region where it is continuously differentiable. Then, we will consider the case that L is nonconvex in a region where it is not continuously differentiable.

We define the truncated Gaussian kernel, which will be used in our proof.

$$dK_h(y; m) = \frac{\phi((y-m)/h)}{h \cdot (1 - \Phi(-m/h))} \quad y \geq 0 \tag{21}$$

In addition, as $h \rightarrow 0$, $dK_h(y; m)$ converges weakly to δ_m , a point mass at m . We also note that if $m_1 < m_2$, then $K_h(\cdot; m_1) \preceq_{\text{SD}} K_h(\cdot; m_2)$.

B.2.1 Continuously Differentiable Case

In the case where L is nonconvex in a region where it is continuously differentiable, we show that it is possible to construct a population distribution F_0 with the following properties

1. F_0 is a mixture distribution over two types, i.e. $F_X = \frac{1}{2} \cdot \mathbb{I}(X = x_1) + \frac{1}{2} \cdot \mathbb{I}(X = x_2)$.
2. The conditional distribution $F_{0,Y|X=x_i}$ is a point mass, i.e. $F_{0,Y|X=x_i} = \mathbb{I}(y = y_i)$. In addition, $y_1 < y_2$, so $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$.
3. F_0 has a strict unique optimal transfer policy under F_0 and budget $B > 0$ called $t_0^*(\cdot; B)$. The optimal transfer policy $t_0^*(\cdot; B)$ is not weakly equitable under F_0 . In particular, $t_0^*(x_1; B) - t_0^*(x_2; B) < 0$.

The construction is provided at the end of this subsection.

We will use F_0 to construct a population distribution that satisfies Assumption 1 under which the optimal policy that solves (2) is not weakly equitable. Given that F_0 has conditional distribution $F_{0,Y|X=x_i} = \mathbb{I}(y =$

y_i), we define F_h be a distribution over (X, Y) where $F_{h,X} = F_{0,X}$ and $F_{h,Y|X} = K_h(y; y_i)$, where we recall that K_h is the truncated Gaussian kernel defined in (21).

We consider solving (2) under population distribution F_h , i.e.

$$\min_{t: \mathcal{X} \rightarrow \mathbb{R}_+} \{ \mathbb{E}_{F_h} [L(t(X) + Y)] : \mathbb{E}_{F_h} [t(X)] \leq B, \quad t(x) \geq 0 \quad \forall x \in \mathcal{X} \}. \quad (22)$$

We note that for any $h > 0$, F_h satisfies Assumption 1, so it remains to show that there exists some $h > 0$ under which the solution to (22) is not weakly equitable.

We show that in a neighborhood of $h = 0$, the above optimization problem has a unique solution t_h^* that is continuous in h . To see this, we first establish that $\mathbb{E}_{F_h} [L(t + Y) \mid X = x_i]$ is continuous in h in a neighborhood about $h = 0$ for all $t \in [0, 2B]$. This holds because we must have that

$$\int_0^\infty L(t + y) dF_h(y; y_i) \rightarrow \int_0^\infty L(t + y) \delta_{y_i} dy = L(t + y_i)$$

as $h \rightarrow 0$ for any $t \in \mathbb{R}_+$ because L is continuous, bounded. Since $\mathbb{E}_{F_h} [L(t + Y) \mid X = x_i]$ is continuous in h in a neighborhood about $h = 0$ and (22) has a unique solution when $h = 0$, then (22) must have a unique solution in a neighborhood about $h = 0$, and this solution must also be continuous in h . As a result, the following function is well-defined and continuous in h on this neighborhood:

$$g(h) := t_h^*(x_1; B) - t_h^*(x_2; B). \quad (23)$$

This function measures the amount of inequity of the optimal policy that solves (22). If $g(h)$ is negative, then the optimal policy is not be weakly equitable. We note that $g(0) < 0$ by construction of F_0 . If for all h in the neighborhood, we have that $g(h) < 0$, then we certainly have that our desired claim holds: there exists some $h > 0$ such that the optimal policy under F_h is not weakly equitable. On the other hand, if there exists some $h > 0$ in the neighborhood such that $g(h) \geq 0$, then, by continuity, there must exist some h' in the neighborhood so that $g(0) \leq g(h') < 0$ by the Intermediate Value Theorem. Thus, there is some $h' > 0$ so that the solution to (22) yields an optimal policy that is not weakly equitable.

Construction of F_0 Since L is non-convex on a continuously differentiable region, there exists some open interval $I \subset \mathbb{R}$ where L' is strictly decreasing. Pick $y_1 \in I$. Since L' is continuous on this interval, for any $\epsilon > 0$, there exists $\delta_1(\epsilon) > 0$, so that $|L'(y_1) - L'(z)| < \epsilon$ for $|y_1 - z| < \delta_1(\epsilon)$. We choose ϵ_1 sufficiently small so that the corresponding $B_{\delta_1(\epsilon_1)}(y_1) \subset I$. Next, we pick $y_2 \in I$ such that $L'(y_2) < L'(y_1) - \epsilon_1$, 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 $|L'(y_2) - L'(z)| < \epsilon$ for $|y_2 - z| < \delta_2(\epsilon)$. We can choose ϵ_2 sufficiently small so that $L'(y_2) + \epsilon_2 < L'(y_1) - \epsilon_1$. Let $\epsilon = \min_{i \in \{1, 2\}} \epsilon_i$. It is straightforward to see that

$$L'(y_2) + \epsilon < L'(y_1) - \epsilon. \quad (24)$$

Let $\Delta = \min(\delta_1(\epsilon), \delta_2(\epsilon))$. For $z_i \in B_\Delta(y_i)$, we have that

$$L'(z_2) < L'(y_2) + \epsilon < L'(y_1) - \epsilon < L'(z_1).$$

This implies that for $t \in [0, \Delta]$, we have that

$$L'(y_1 + t) \geq L'(y_1) - \epsilon \quad (25)$$

and

$$L'(y_2 + t) \leq L'(y_2) + \epsilon. \quad (26)$$

Combining (25), (26) and (24), implies that for all $t \in [0, \Delta]$

$$L'(y_2 + t) < L'(y_1 + t).$$

Set $F_{0,Y|X=x_i} = \mathbb{I}(y = y_i)$ for $i \in \{1, 2\}$. For this choice of population distribution F_0 and $B = \Delta/2$, 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 $F_{0,Y|X=x_1} \preceq_{\text{SD}} F_{0,Y|X=x_2}$ but $t^*(x_2) > t^*(x_1)$ and is thus, not weakly equitable.

This follows from the fact that L is concave on the feasible set because L' is strictly decreasing. The minimizer of a concave function over a closed, convex set 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

$$\mathbb{E}_{F_0} [L(\tilde{t}(X) + Y)] > \mathbb{E}_{F_0} [L(t^*(X) + Y)].$$

Since we have that

$$L'(y_1 + t_1) > L'(y_2 + t_2)$$

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

$$(L(y_1 + \Delta) - L(y_1)) \cdot \Delta > (L(y_2 + \Delta) - L(y_2)) \cdot \Delta.$$

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

$$\frac{1}{2} \cdot (L(y_1 + \Delta) + L(y_2)) > \frac{1}{2} \cdot (L(y_1) + L(y_2 + \Delta)).$$

This implies that

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

B.2.2 Non-Continuously Differentiable Case

Now, we consider the case where L is not continuously differentiable in the region where it is non-convex. In this case, we will be able to show that it is possible to construct a smoothed loss function \tilde{L} that preserves the non-convexity of L but is also continuously differentiable by using the truncated Gaussian kernel $K_b(\cdot; 0)$ for smoothing. For this smoothed loss function \tilde{L} , we can construct a distribution F_h that satisfies Assumption 1 using the above proof technique under which the optimal policy that solves (2) with population distribution F_h and loss function \tilde{L} . Using the smoothing kernel $K_b(\cdot; 0)$ and the distribution F_h , we can show that there is a distribution F under which the optimal policy that solves (2) with the original loss function L is not weakly equitable.

Since L is decreasing, it is continuous almost everywhere. Since L is also nonconvex and is continuous almost everywhere, there must exist t_0, t_1, t_2 that are continuity points of L such that $t_0 = \lambda t_1 + (1 - \lambda)t_2$ for some $\lambda \in (0, 1)$ and

$$L(t_0) - \lambda L(t_1) - (1 - \lambda)L(t_2) = \epsilon > 0.$$

If there exists A_0, A_1, A_2 such that $|L(t_1) - A_1| < \frac{\epsilon}{4}$, $|L(t_2) - A_2| < \frac{\epsilon}{4}$, and $|L(t_0) - A_0| < \frac{\epsilon}{4}$. Then,

$$A_0 - \lambda A_1 - (1 - \lambda)A_2 \geq \frac{\epsilon}{2}. \quad (27)$$

We define the following loss function.

$$\tilde{L}_b(t) := \int_0^\infty L(t + z) dK_b(z; 0),$$

where $K_b(z; 0)$ is the truncated Gaussian kernel. Since L is integrable on \mathbb{R}_+ , we have that $\tilde{L}(t)$ is continuously differentiable on \mathbb{R}_+ for any $b > 0$. By weak convergence, $b \rightarrow 0$, $\tilde{L}_b(t) \rightarrow L(t)$ for continuity points t of L . This implies that there exists b sufficiently small so that $|\tilde{L}_b(t_1) - L(t_1)| < \frac{\epsilon}{4}$, $|\tilde{L}_b(t_2) - L(t_2)| < \frac{\epsilon}{4}$, and lastly $|\tilde{L}_b(t_0) - L(t_0)| < \frac{\epsilon}{4}$. This implies that \tilde{L}_b is non-convex by (27), i.e.

$$\tilde{L}_b(\lambda t_1 + (1 - \lambda)t_2) - \lambda \tilde{L}_b(t_1) - (1 - \lambda)\tilde{L}_b(t_2) \geq \frac{\epsilon}{2}.$$

Thus, \tilde{L}_b is nonconvex in a region where it is continuously differentiable.

Since \tilde{L}_b is nonconvex in a region where it is continuously differentiable, we can apply the proof in Section B.2.1 to show that there is a distribution F_h under which the optimal policy under \tilde{L} is not weakly equitable. This distribution F_h has covariate distribution $F_{h,X}(x) = \sum_{i \in \{1,2\}} \frac{1}{2} \mathbb{I}(x = x_i)$ and $F_{h,Y|X=x_i} = K_h(\cdot; y_i)$. Define a population distribution F over (X, Y) where $F_X = \sum_{i \in \{1,2\}} \frac{1}{2} \mathbb{I}(x = x_i)$ and $F_{Y|X=x_i}$ is the distribution over $Z_1 + Z_{2,i}$, where $Z_1 \sim K_b(\cdot; 0)$ and $Z_{2,i} \sim K_h(\cdot; y_i)$. We note that $F_{Y|X=x_1} \preceq_{\text{SD}} F_{Y|X=x_2}$ in this case. We observe that

$$\begin{aligned} \mathbb{E}_F [L(t + Y) \mid X = x_i] &= \mathbb{E} [L(t + Z_1 + Z_{2,i}) \mid X = x_i] \\ &= \int_0^\infty \int_0^\infty L(t + z_1 + z_{2,i}) dK_b(z_1; 0) \cdot dF_{h,Y|X}(z_{2,i}) \\ &= \int_0^\infty \tilde{L}(t + z_{2,i}) dF_{h,Y|X}(z_{2,i}) \\ &= \mathbb{E}_{F_{h,Y|X}} [\tilde{L}(t + Y) \mid X = x_i]. \end{aligned}$$

Thus, the optimal policy t^* that solves (2) with loss function \tilde{L} and population distribution F_h is equal to the optimal policy that solves (2) with loss function L and population F described above. Since t^* is not weakly equitable under F_h , t^* is also not weakly equitable under F .

B.3 Proof of Theorem 2

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

$$t_\lambda(x) = ((G'_x)^{-1}(\lambda))_+, \quad (28)$$

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

$$\mathcal{L}(t, \lambda) = \mathbb{E}_F[L(Y + t(X))] + \lambda(\mathbb{E}_F[t(X)] - B) = -\mathbb{E}_F[G_X(t(X))] + \lambda(\mathbb{E}_F[t(X)] - B). \quad (29)$$

Define the dual function as

$$h(\lambda) = \min_{t: \mathcal{X} \rightarrow \mathbb{R}_+} \{\mathcal{L}(t, \lambda) : t(x) \geq 0 \quad \forall x \in \mathcal{X}\}. \quad (30)$$

Then the dual problem is given by

$$\max_{\lambda \geq 0} h(\lambda). \quad (31)$$

To compute $h(\lambda)$, we solve the optimization problem on the right side of (30)

$$\min_{t: \mathcal{X} \rightarrow \mathbb{R}_+} \{-\mathbb{E}_F[G_X(t(X))] + \lambda(\mathbb{E}_F[t(X)] - B) : t(x) \geq 0 \quad \forall x \in \mathcal{X}\}. \quad (32)$$

Importantly, the constraints of the above program are separable, so the policy that solves (32) 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 (32) is equivalent to the solution to

$$\min_{t: \mathcal{X} \rightarrow \mathbb{R}_+} \{-\mathbb{E}_F[G_X(t(X))] + \lambda(\mathbb{E}_F[t(X)] - B) : 0 \leq t(x) < c \quad \forall x \in \mathcal{X}\} \quad (33)$$

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

$$G'_x(t(x)) = \lambda. \quad (34)$$

The optimal solution to this program has form (28).

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

$$h(\lambda) = -\mathbb{E}_F[G_X(t_\lambda(X))] + \lambda(\mathbb{E}_F[t_\lambda(X)] - B). \quad (35)$$

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

$$\begin{aligned}
\frac{dh}{d\lambda} &= -\mathbb{E}_F \left[G'_X(t_\lambda(X)) \cdot \frac{d}{d\lambda} t_\lambda(X) \right] + \lambda \mathbb{E}_F \left[\frac{d}{d\lambda} t_\lambda(X) \right] + (\mathbb{E}_F[t_\lambda(X)] - B) \\
&= -\mathbb{E}_F \left[G'_X((G'_X)^{-1}(\lambda)) \cdot \left(\frac{1}{G''_X((G'_X)^{-1}(\lambda))} \cdot \mathbb{I}(\lambda < G'_X(0)) \right) \right] \\
&\quad + \lambda \cdot \mathbb{E}_F \left[\frac{1}{G''_X((G'_X)^{-1}(\lambda))} \cdot \mathbb{I}(\lambda < G'_X(0)) \right] + (\mathbb{E}_F[t_\lambda(X)] - B) \\
&= \mathbb{E}_F[t_\lambda(X)] - B.
\end{aligned} \tag{36}$$

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

$$\Lambda = \{\lambda \geq 0 \mid \exists x \in \mathcal{X} \text{ such that } \lambda < G'_x(0)\}. \tag{37}$$

We can show that if $B > 0$, then any optimal λ^* that solves (31) must lie in Λ , where Λ is defined in (37). 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 (31) has a unique solution λ^* that lies in Λ and the optimal policy that solves (2) is unique and given by (28) 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 (31) 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 (28) 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 (28) 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 (28) 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 $\mathbb{E}_F[L'(Y + t) \mid X = x] \leq \mathbb{E}_F[L'(Y + t) \mid X = x']$. 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 (28) with $\lambda(B)$ being the unique choice of λ that satisfies $\mathbb{E}_F[(G'_X)^{-1}(\lambda)] = B$.

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

$$\min_{t: \mathcal{X} \rightarrow \mathbb{R}_+} \{ \mathbb{E}_F[L(t(X) + t_{\lambda(B')}(X) + Y)] : \mathbb{E}_F[t(X)] \leq B - B', t(x) \geq 0 \quad \forall x \in \mathcal{X} \}. \quad (38)$$

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 (38) is given by

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

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

$$\tilde{t}_\lambda(x) = ((G'_x)^{-1}(\lambda) - t_{\lambda(B')}(x))_+. \quad (39)$$

We show that λ that satisfies $\mathbb{E}_F[\tilde{t}_\lambda(X)] = B - B'$ 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

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

Thus, we not only observe that the optimal policy that solves (38) is given by (39) 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

$$\min_{t: \mathcal{X} \rightarrow \mathbb{R}_+} \mathbb{E}_F [t(X)] \text{ subject to } \mathbb{P}_F [t(X) + Y < c \mid X = x] \leq \lambda \quad \forall x \in \mathcal{X} \quad (40)$$

is also given by (9). This is straightforward to see because the constraints of (40) are separable.

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

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

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 $t_{\text{gap}}(x) = (c - F_{Y|X=x}^{-1}(\lambda))_+$ for some $\lambda(B) \in \mathbb{R}$ such that $\mathbb{E}_F [(c - F_{Y|X=X}^{-1}(\lambda(B)))_+] = B$.

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

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

This derivation also shows that

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

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

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

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

$$\frac{d}{dt} \mathbb{E}_F [L(t + Y) \mid X = x] + \lambda = 0. \quad (41)$$

We note that

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

So, we can rewrite (41) 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 1 and 3, (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.

$$L(t, \lambda, \nu) = \mathbb{P}_F[t(X) + Y < c] + \lambda(\mathbb{E}_F[t(X)] - B) - \sum_{x \in \mathcal{X}} \nu_x t(x)$$

$$\nabla L(t, \lambda, \nu) = f_X(x) \cdot (-f_{Y|X=x}(c - t(x)) + \lambda) - \nu_x = 0.$$

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

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

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

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

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 $\operatorname{argmin}_{\alpha: \alpha \geq 0} \mathbb{E}_F[t_\alpha^*(X)]$.

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 $\pi(x) : \mathcal{X} \rightarrow \{0, 1\}^K$ 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 $\mathbb{E}_F[\langle \pi(X), p(X; \alpha) \rangle]$ and the feasibility constraint that $\mathbb{P}_F[t(X) + Y < c] \leq \epsilon$. (11) can be written as $\mathbb{E}_F[\langle \pi(X), z(X; \alpha) \rangle] \leq B$. 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 $g : \mathcal{T} \times \mathcal{Y} \rightarrow \mathbb{R}$ satisfies the following conditions.

1. $g(t, y)$ be a Lebesgue integrable function of y for each $t \in \mathcal{T}$.
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 $\theta : \mathcal{Y} \rightarrow \mathbb{R}_+$ 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} . 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

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

Since L is monotone and differentiable, for 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 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

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

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

$$G'_x(t) = - \int_{\mathcal{Y}} L'(t+y) f_{Y|X=x}(y) dy. \quad (42)$$

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 (42). We apply integration by parts to $G'_x(t)$. First, suppose that $\mathcal{Y} \subseteq [\underline{y}, \bar{y}]$.

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

We aim to compute

$$G''_x(t) = L'(\underline{y}+t) f_{Y|X=x}(\underline{y}) - L'(\bar{y}+t) f_{Y|X=x}(\bar{y}) + \frac{d}{dt} \int_{\mathcal{Y}} L(y+t) \cdot f'_{Y|X=x}(y) dy. \quad (43)$$

We note that the first two terms are differentiable on \mathbb{R} .

Similarly, if $\mathcal{Y} \subseteq [\underline{y}, \infty)$, then

$$G''_x(t) = -L'(\underline{y}+t) f_{Y|X=x}(\underline{y}) + \frac{d}{dt} \int_{\mathcal{Y}} L(y+t) \cdot f'_{Y|X=x}(y) dy. \quad (44)$$

We apply Leibniz Rule to the integral that appears on the right of (44) and (43). 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 $C > -\infty$.

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

The last line follows from Assumption 1. 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 $y \in \mathcal{Y}$. Finally, we show that there is an integrable function $\theta : \mathcal{T} \rightarrow \mathbb{R}$ 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

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

The last inequality follows from Assumption 1. Thus, we can apply Leibniz rule to show that

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

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

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

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

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

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

$$G'_x(t) = - \int_{\mathcal{Y}} L'(t+y) f_{Y|X=x}(y) dy \tag{45}$$

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 $f : \mathbb{R} \rightarrow \mathbb{R}$ 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

$$L(y' + \lambda t_1 + (1 - \lambda)t_2) < \lambda L(y' + t_1) + (1 - \lambda)L(y' + t_2). \quad (46)$$

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 (46) holds when y' is replaced by $y' + \epsilon$ and $y' - \epsilon$. Let $I = [y' - \epsilon, y' + \epsilon]$. We note that under Assumption 1, $f_{Y|X}(y) > 0$ on I .

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

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 (45).

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

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 Country Selection

This section details our criteria for including countries in the analysis, and any exceptions or caveats thereto. Generally speaking, we aim to include countries that account for as large a share of the world’s poor as possible subject to various feasibility constraints. The specific criteria we impose are that countries must:

1. Have an extreme poverty rate greater than or equal to 10%, or account for more than 1% of the world’s extreme poor, based on World Poverty Clock estimates as of 2023. This filter aims to focus our search effort on the countries that most influence aggregate costs. Globally, 96.2% of poor households live in countries that meet this criterion.
 - We make an exception at this stage to include Côte d’Ivoire and Senegal. They are covered by the EHCVM survey series, which also covers six other West African countries already in our sample, and the surveys are harmonized in such a way that including them is unusually easy.
2. Have a recent, high-quality, nationally-representative LSMS survey or equivalent. This is a prerequisite for our analysis. Whenever possible we use the same survey as that on which estimates in the World Bank’s Poverty and Inequality Platform are based. 95.6% of poor households live in countries that additionally meet this criterion.
3. Are estimated to have at least 1,000 extremely poor households covered by this survey, given the survey sample size and published poverty rate. This filter aims to ensure that the survey contains enough poor households for reliable statistical learning; relaxing it would require countries to conduct larger surveys (or surveys more focused in the regions containing the most poor households). 90.6% of poor households live in countries that additionally meet this criterion.
 - We make an exception at this stage to include Bangladesh (which we anticipated would contain slightly fewer than 1,000 households) because of its size.
4. Have made the relevant survey data, including a measure of aggregate household consumption, publicly available, or have provided it upon request. 63.1% of poor households live in countries that additionally meet this criterion.
 - We make an exception at this stage to include India (using data from the HCES 2022–23), although this requires us to reconstruct a consumption aggregate ourselves, because of its size.
5. Have usable PPP conversion factors available. Specifically, we remove Afghanistan and Somalia on the basis of advice from World Bank staff that reliable PPP conversion factors for the years in which the relevant surveys were conducted are not available. This leaves a sample containing 57.3% of poor households.

Among the remaining countries, analysis was still in progress for 10 (Brazil, Liberia, Mexico, Namibia, The Gambia, Timor-Leste, Pakistan, Sierra Leone, South Sudan, Zimbabwe) as of this draft. This leaves a final sample of 23 countries which collectively account for 50% of the world’s extreme poor.

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 convert consumption aggregates from all surveys from local currency units (LCU) to 2017 PPP USD using Consumer Price Index and PPP conversion factors. In some cases we have access to precise conversion factors used by the Poverty and Inequality Platform (PIP) team to make international poverty rate estimates. In others, including when we are using a survey not used by PIP, we use factors drawn from the World Bank’s development indicators (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 cases where it a survey spans two years and no further information about currency base year is available, we average conversion factors across the two years. Table D.1 describes properties of the consumption aggregates we use and indicates deflator usage.³⁵

Table E.2 displays the resulting poverty headcount rate estimates (column titled “Survey” under column group “Poverty Rates”) alongside poverty rate estimates for the same country and year from PIP. Where the PIP estimate was derived from the same underlying survey as our estimate, we include it in the column labeled “World Bank (same source);” where it was derived from a different survey (including by interpolating between estimates from different surveys in different years) we include it in the column labeled “World Bank (different source).” Generally speaking, whenever we use the same source, we are able to match the corresponding PIP estimate exactly (sometimes after using non-public deflators or, in Nigeria, omitting a region where survey enumeration was disrupted). In four cases we do not: Bangladesh (where our estimate is 1.4 percentage points higher than the PIP rate), South Africa (2.5 percentage points lower), India (1.0 percentage points higher), and Rwanda (3.5 percentage points higher). When we use a different source the average absolute difference between the estimates is 2.4 percentage points and the average difference is 0.1 percentage points.

D.2.1 Manual consumption aggregation of India’s HCES 2022-23 survey

For India’s HCES 2022–23 survey, no pre-existing aggregate is available as of this writing, so we construct one, following the methodology described in the official report (National Sample Survey Office, 2024, Chapter 2). This involves first constructing four sub-components as follows:

- **Reported food and non-food expenditure:** We convert household expenditures elicited with varying recall periods (one week for food items, one month for non-food items, and one year for education, health, rent, etc) to daily values.
- **Durables:** We include expenditure on durables during the last year in the consumption aggregate, in

³⁵We thank Daniel Mahler and Elizabeth Foster in particular for their generous guidance and assistance obtaining and interpreting the deflators and conversion rates required for a number of these operations.

alignment with the official methodology.³⁶

- **Housing costs:** We include only actual rent paid for housing and garage, and do not impute rent equivalence for non-rental housing, for consistency with the official methodology.
- **Transfers:** We include goods and services received in kind or as perquisites in household consumption when they were received in exchange for services, except for meals received from other households. Public transfers, such as food or durable goods from welfare schemes, are valued at state-by-sector (rural/urban) modal prices.

Aggregating these components into a household per capita total involves some nuance, as the underlying survey collected information about them via three separate questionnaires administered in consecutive months, and recorded household size sometimes varied across those visits. Following the official methodology, we divide the expenditure reported in each questionnaire by the size of the household at the time it was administered, and then sum those per-capita components. We then convert the resulting figure to 2017 USD PPP using the same conversion approach we use across surveys, described above.

The resulting aggregate tracks the official values reported in [National Sample Survey Office \(2024\)](#) closely. That report provides two separate state-by-sector (rural/urban) “monthly per capita consumption expenditure” aggregates: one that excludes the value of free goods received from welfare schemes (other than the Public Distribution System), and one that includes them valued at local prices. We convert our household daily per-capita aggregates into monthly equivalents, aggregate to state-by-sector weighted averages, and compare against both official series. Our estimates align almost one-to-one with the official figures which exclude the goods listed above: 16 of 18 states match exactly for both sectors, and for the remaining two, the only differences are in the rural sector (1.5% in Bihar and 0.06% in Andhra Pradesh). When comparing our aggregate to the official figures which include the goods listed above, mean differences are typically less than 0.5% of the official figure, with the largest gap being 1.6% (Bihar, rural). In most states the discrepancy is under 0.1%, and 5 of 18 states match exactly for both sectors. We think these minor differences most likely reflect the granularity of the local prices used in valuation.

³⁶The data do not include enough information to calculate standard flow measures of use value.

Table D.1: Consumption Aggregate Details

Country	Category of Consumption		Deflator Usage
	Grants and Transfers	Durable Goods ^a Non-Rental Housing ^b	
Bangladesh	Unspecified	Unspecified	None used
Benin	Included	User cost	Temporal, from WB ^d
Burkina Faso	Included	User cost	Temporal, from WB
Colombia	Included	Acquisition approach	None used
Côte d'Ivoire	Included	User cost	Temporal, from WB
Ethiopia ^e	Included	Acquisition approach	Spatial, from dataset
Ghana	Included	User cost	None used
Guinea-Bissau	Included	User cost	Temporal, from WB
India	Included	Acquisition approach	None used
Indonesia	Unspecified	Unspecified	None used
Kenya	Food included; non-food excluded	Unspecified	None used
Madagascar	Food included; non-food excluded	User cost	None used
Malawi	Food included; non-food excluded	User cost ^g	None used
Mali	Included	User cost	Temporal, from WB
Niger	Included	User cost	Temporal, from WB
Nigeria ⁱ	Food included; non-food unspecified	Acquisition approach ^j	Temporal and spatial, from dataset
Rwanda	Included	User cost	Temporal, from dataset
Senegal	Included	User cost	Temporal, from WB
South Africa	Included	Acquisition approach	Temporal, from dataset
Tanzania	Food included; non-food excluded	User cost	Temporal and spatial, from dataset
Togo	Included	User cost	Temporal, from WB
Uganda	Included	Acquisition approach	Temporal, from dataset
Yemen	Included	User cost	Temporal, from WB

^aThis column uses terminology introduced by [Amendola and Vecchi \(2022\)](#).

^bIncludes any housing not rented on the market: Owned housing, family- or employer-provided housing, etc.

^cA function predicting rental value based on observable housing characteristics and amenities; fit using self-reported rents of respondents living in rental housing.

^dHere "from WB" indicates deflator data obtained from the World Bank which differs from, or is not included in, the corresponding survey dataset.

^eOther irregularity: Responses from Tigray Region are not included in survey microdata.

^fAll (rental or non-rental) housing consumption is excluded, per WB methodology.

^gAssumes linear depreciation and disregards interest rate.

^hOutliers imputed using hedonic function.

ⁱOmitting responses from state of Borno per WB methodology.

^jLong-term ("lumpy") durables excluded

^kOutliers imputed by means unspecified.

^lDisaggregated by housing type and province.

D.3 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. If the survey does not pre-define age-based categories then we will ourselves count the number of children 17 and under, and the number of older adults 65 and older. 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.

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.³⁷ 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

³⁷We take the maximum rather than including the attainment of all members individually in order to produce a per-household predictor.

- (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.
- (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 financial 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.4 Secondary Data

This section compiles, for convenient reference, the full list of secondary data sources (i.e., those other than LSMS-style survey data) referenced elsewhere in the text.

- 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 IMF³⁸ ([IMF, 2025](#)): Taiwan (2017, 2021); Venezuela (2021); Yemen (2017, 2021).

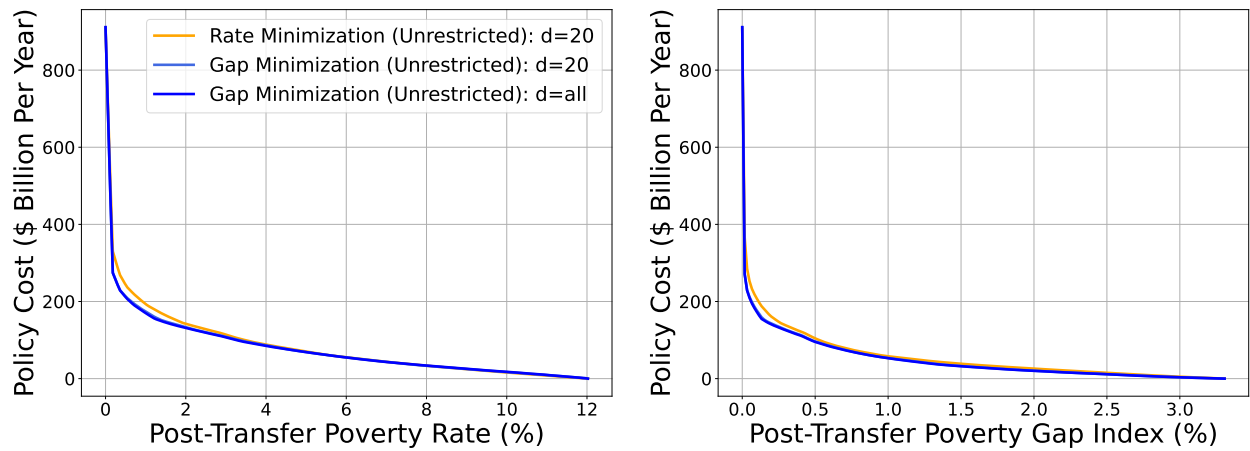
Both datasets were missing Venezuela’s 2017 conversion factor at the time of access. We estimate this value by assuming that the nominal currency exchange rate and the conversion factor from LCU to USD PPP (international dollars) have the same ratio as in the nearest year with complete data. By assuming 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](#)).

³⁸Indicator name: Rate, Domestic currency per international dollar in PPP terms, ICP benchmarks 2017-2021

E Additional Exhibits

Figure E.1: Comparison of Rate Minimization, Gap Minimization with Restricted Feature Set



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. We compare the performance of rate minimization with the restricted feature set ($d = 20$), gap minimization with the restricted feature set ($d = 20$), and gap minimization with the full feature set. We find that the policy costs of gap minimization with the full and restricted feature sets coincide.

Table E.2: Survey characteristics

Country	Survey Name	Survey Year	n	d	Poverty Rate		
					World Bank (same source)	World Bank (other source)	Survey
Bangladesh	HIES	2022	14268	197	3.6	-	5.0
Benin	EHCVM	2018	8012	323	20.8	-	20.8
Burkina Faso	EHCVM	2018	7010	379	31.2	-	31.2
Colombia	ENPH	2016	87201	110	-	4.9	2.8
Côte d'Ivoire	EHCVM	2018	12992	366	11.5	-	11.5
Ethiopia	Socio-Economic Panel Survey	2021	4959	148	-	32.0	35.6
Ghana	Living Standards Survey 7	2016	14009	327	25.2	-	25.2
Guinea-Bissau	EHCVM	2018	5351	408	21.7	-	21.1
India	Household Consumption Expenditure Survey	2022	261746	41	2.3	-	3.3
Indonesia	SUSENAS	2018	295155	99	4.4	-	6.6
Kenya	Continuous Household Survey	2021	16963	40	36.1	-	36.1
Madagascar	Permanent Household Survey	2021	16540	171	52.3	-	52.3
Malawi	Fifth Integrated Household Survey	2019	11434	152	70.1	-	70.1
Mali	EHCVM	2018	6602	407	15.2	-	15.2
Niger	EHCVM	2018	6024	215	50.9	-	50.9
Nigeria	Living Standards Survey	2018	21580	129	30.9	-	30.9
Rwanda	EICV7	2023	15054	156	24.2	-	27.7
Senegal	EHCVM	2018	7156	411	9.2	-	9.2
South Africa	Income and Expenditure Survey	2010	25328	71	18.0	-	15.5
Tanzania	National Panel Survey, Wave 5	2020	4709	166	-	44.9	46.2
Togo	HSHLS	2018	6171	162	28.4	-	28.4
Uganda	National Panel Survey	2019	3074	101	-	42.1	39.5
Yemen	Household Budget Survey	2014	9376	357	19.8	-	19.8

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. The “Poverty Rate” columns report the headcount poverty rates we estimate (“Survey”) and those reported by the World Bank, with the latter differentiated into two cases: those in which the World Bank estimate is based on the same survey as ours, and those in which it is based on or extrapolated from other sources ([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,” EICV7 is an abbreviation for “Integrated Household Living Conditions Survey”, HIES is an abbreviation for Household Income and Expenditure Survey, SUSENAS corresponds to the National Socio-Economic Survey.

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

Country	Poverty Rate	Share of World's Poor
Bangladesh	4.2	1.2
Benin	10.6	0.2
Burkina Faso	27.4	1.0
Colombia	13.3	1.1
Côte d'Ivoire	7.3	0.3
Ethiopia	15.0	3.0
Ghana	17.7	1.0
Guinea-Bissau	25.7	0.1
India	1.3	3.1
Indonesia	2.8	1.3
Kenya	27.6	2.4
Madagascar	67.4	3.3
Malawi	73.2	2.7
Mali	19.1	0.7
Niger	51.9	2.3
Nigeria	31.7	11.6
Rwanda	35.2	0.8
Senegal	9.1	0.3
South Africa	21.6	2.2
Tanzania	36.7	4.1
Togo	24.6	0.4
Uganda	36.0	3.1
Yemen	69.1	3.8

This table reports the poverty rate and the share of the world's poor in 2023 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
Bangladesh	2022	4.85	479.07	42.72	0.01	0.11
Benin	2018	1.57	17.31	2.35	0.09	0.67
Burkina Faso	2018	2.84	19.28	3.81	0.15	0.74
Colombia	2016	1.59	358.93	99.49	0.00	0.02
Côte d'Ivoire	2018	1.87	71.01	10.41	0.03	0.18
Ethiopia	2021	25.72	125.11	13.80	0.21	1.86
Ghana	2016	5.10	71.28	9.37	0.07	0.54
Guinea-Bissau	2018	0.22	1.89	0.28	0.12	0.80
India	2022	17.24	3483.84	701.39	0.00	0.02
Indonesia	2018	12.70	1264.73	188.71	0.01	0.07
Kenya	2021	9.66	123.36	20.75	0.08	0.47
Madagascar	2021	4.94	16.14	1.79	0.31	2.76
Malawi	2019	4.27	13.17	1.94	0.32	2.19
Mali	2018	1.95	24.72	3.85	0.08	0.51
Niger	2018	4.72	15.58	2.83	0.30	1.67
Nigeria	2018	31.18	511.75	43.52	0.06	0.72
Rwanda	2023	1.72	14.33	3.15	0.12	0.55
Senegal	2018	1.11	28.05	5.30	0.04	0.21
South Africa	2010	8.97	583.21	138.96	0.02	0.06
Tanzania	2020	13.67	77.78	11.58	0.18	1.18
Togo	2018	1.40	8.53	1.55	0.16	0.90
Uganda	2019	6.28	42.14	5.67	0.15	1.11
Yemen	2014	6.14	55.64	13.16	0.11	0.47


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.2: Policy Cost vs. Post-Transfer Poverty Measures by Country


(a) Bangladesh




(b) Benin




(c) Burkina-Faso



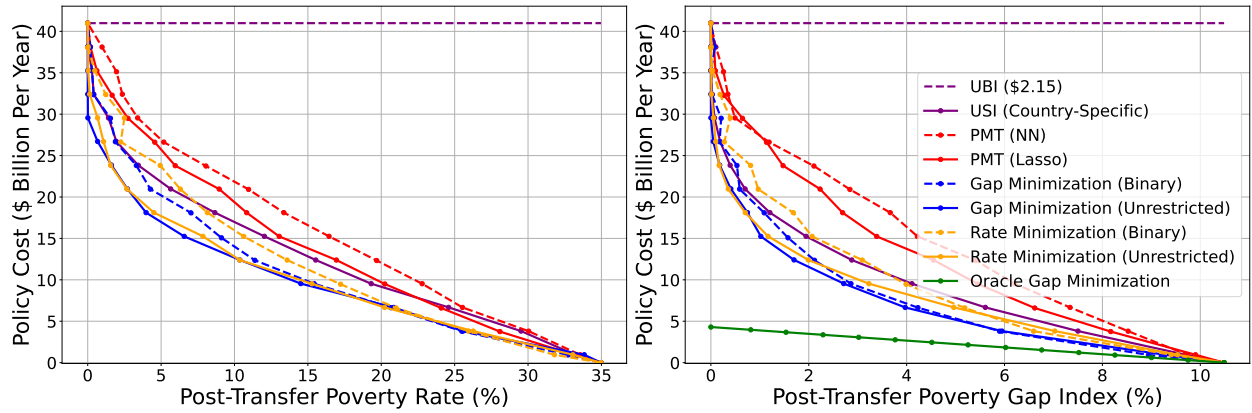
(d) Colombia




(e) Côte d'Ivoire




(f) Ethiopia




(g) Ghana




(h) Guinea-Bissau




(i) India




(j) Indonesia




(k) Kenya




(l) Madagascar




(m) Malawi




(n) Mali




(o) Niger




(p) Nigeria




(q) Rwanda




(r) Senegal




(s) South Africa




(t) Tanzania




(u) Togo



(v) Uganda



(w) Yemen



This figure reports country-by-country results for gap minimization (binary/unrestricted), rate minimization (binary/unrestricted), PMT (with neural network / with lasso), UBI \$2.15, and Universal Supplemental Income.