Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries

Rema Hanna and Benjamin A. Olken

Of the 17 Sustainable Development Goals articulated by the United Nations, number one is the elimination of extreme poverty by 2030. In recent years, the largest reductions in extreme poverty worldwide have resulted from substantial economic growth in many emerging economies, particularly China and India (Page and Pande, in this journal issue). While future growth should continue to reduce poverty, it will not solve the problem by itself—both because a substantial amount of poverty remains even in rapidly growing countries, and because sustained economic growth over decades of the type seen in China is the exception, not the norm (Jones and Olken 2008).

Given these realities, there is a potentially important role for national-level transfer programs that assist poor families in developing countries. Such programs are often run by developing country governments. For example, China's rural minimum living standard guarantee (Dibao) program reaches nearly 75 million individuals (Golan, Sicular, and Umapathi 2017), while Mexico's conditional cash transfer program reaches 32 million individuals (World Bank 2014). As countries become wealthier, a greater share of GDP usually goes to social transfer and insurance programs (Chetty and Looney 2006), suggesting that these types of safety net programs are likely to expand in many developing countries.

■ *Rema Hanna is the Jeffrey Cheah Professor of South-East Asia Studies, Harvard Kennedy School, Cambridge, Massachusetts. Benjamin A. Olken is Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Their email addresses are rema_hanna@hks.harvard.edu and bolken@mit.edu.*

Many countries have implemented transfer programs that seek to target beneficiaries: that is, to identify who is poor and then to restrict transfers to those individuals. In developed countries, because one rarely observes true income-earning ability, targeting is usually based on income (for example, Mirrlees 1971). But in developing countries, governments do not observe income for the vast majority of the population who work in the informal sector, which typically includes most of the poor. Imperfect targeting using various proxy measures for income leads to both *inclusion errors* (giving the transfer to those who are not poor) and *exclusion errors* (failing to deliver the transfer to poor individuals who slip through the cracks in the targeting protocol).

Given these challenges, some have begun to advocate for "universal basic income" programs, which dispense with trying to identify the poor and instead provide transfers to everyone. A universal basic income program is comparatively straightforward to implement—each individual receives a fixed transfer, regardless of income—so the main challenge is to ensure that each person receives the transfer only once. Although universal basic income programs distribute the same value of transfer to everyone, including the very rich, if they are financed through proportional or progressive taxation, they can still result in a substantial redistribution to the poor.

We begin by considering the universal basic income as part of the solution to an optimal income-taxation problem, focusing on the case of developing countries, where there is limited income data and inclusion in the formal tax system is low. We examine how the targeting of transfer programs is conducted in these settings, and provide empirical evidence on the tradeoffs involved between universal basic income and targeted transfer schemes using data from Indonesia and Peru—two countries that run nationwide transfer programs that are targeted to the poor. We conclude by linking our findings back to the broader policy debate on what tools should be preferred for redistribution, as well as the practical challenges of administering them in developing countries.

Conceptual Connections and Underpinnings

Universal Basic Income and the Tax Schedule

A universal basic income is usually conceived of as a fixed transfer given to everyone, regardless of income level. Such programs are fairly rare in practice, with prominent examples coming from places with substantial natural resource revenues. For example, Alaska's Permanent Fund transfers a fixed amount—usually between $1,000 and $2,000 per person, per year—to every citizen of the state. Iran implemented a similar program starting in 2011 (Salehi-Isfahani and Mostafavi-Dehzooei 2017). Common arguments that are made for universal basic income programs include their ease of implementation and low administrative costs, because the government does not have to verify income. Another common claim is that such

programs do not distort labor supply, because the payments do not decline if you work more, but as we discuss below, this conclusion is not as straightforward as it may at first appear.

For most developing economies, a substantial universal basic income would need to be financed via domestic taxation, because official development assistance is a small fraction of government budgets. For example, for countries ranked as "upper middle-income" by the World Bank, such as Peru, Lebanon, and the Dominican Republic, overall tax revenue was approximately 157 times the amount of net official development assistance received in 2010. This pattern also holds for "lower middle-income" countries, such as Indonesia, India, and Morocco, where overall tax revenue in 2010 was approximately 14 times the amount of development assistance received. It is only with "low-income" countries, such as Afghanistan, Ethiopia, and Mozambique, that official development assistance exceeded tax revenue—it was approximately 1.2 times the amount of tax revenue in 2010.[1]

Given that a universal basic income would need to be financed through taxation, one can recast universal basic transfers as a particular feature of a tax schedule, as described by Saez (2002) and illustrated by Figure 1. We assume that a universal basic income program is fully financed by a progressive income tax schedule. Pre-tax income y for an individual appears on the horizontal axis, and after-tax income $\alpha(y)$ for that individual is on the vertical axis. The tax paid by a household is therefore given by $y - \alpha(y)$). The universal basic income—the transfer given to everyone, regardless of their income—is thus equal $\alpha(0)$ to on the vertical axis. This perspective suggests that any tax-and-transfer system that has the feature that $\alpha(0) > 0$ can be thought of as including a universal basic income.

Figure 1 plots a progressive tax schedule with and without a fully financed universal basic income. The universal basic income acts as an intercept shifter, raising post-tax income at $y = 0$. However, it will not raise everyone's after-tax income by that same amount, because to finance the universal basic income out of income taxes, marginal tax rates have to increase for *someone*. This illustrates three points about the universal basic income: 1) it can make use of income data collected through the tax system; 2) although a universal basic income acts as an intercept shifter, raising post-tax income at $y = 0$, it will not raise everyone's after-tax income by the same amount; and 3) as emphasized by Saez (2002), it is *not* the case that a universal basic income has no labor supply distortions, since there is still a tax on labor income, and in particular adding the universal basic income to the tax schedule requires adjusting the tax schedule elsewhere.

---

[1] Data on taxes and official development assistance come from the World Development Indicators. Data on social spending come from the International Labor Organization Social Security Inquiry, http://www.ilo.org/dyn/ilossi/ssimain.home?p_lang=en. For each comparison above, we include only countries with non-missing data for both variables (for example, we sum and compare tax revenue and official development assistance only among countries with non-missing data for both). Country classifications are from the World Bank, where low-income countries have 2015 gross national per-capita income of $1,025 or less; lower middle-income, $1,026–4,035; upper middle-income, $4,036–12,475; and high-income, $12,476 or more.

Figure 1

Example of Progressive Post-Tax Income Schedules With and Without a Universal Basic Income (UBI)
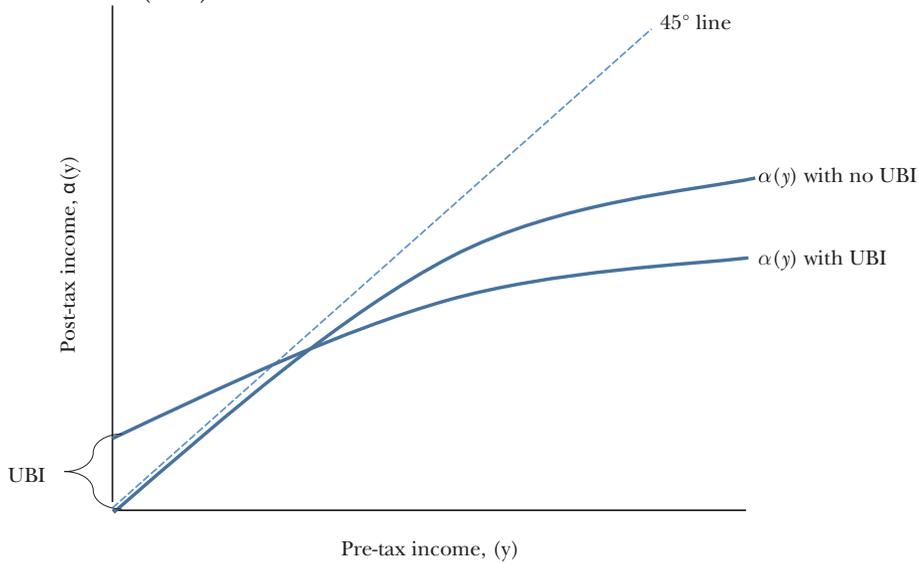

From this perspective, the tradeoffs involved in adding a universal basic income to an income tax schedule, therefore, depend on how the *rest* of the income tax schedule is adjusted to satisfy the government budget constraint—both in terms of overall redistribution and potential distortionary effects from the increased marginal tax rates that introducing a universal basic income will necessitate. Saez (2002) discusses these conditions in the context of a Mirrlees (1971) optimal income tax model in detail. He also discusses the conditions under which an optimal income tax schedule is likely to feature a universal basic income—essentially, when *intensive* labor supply elasticities (the effects of changes in wages on hours of existing workers) are larger than *extensive* labor supply elasticities (the effect of changes in wages on hours of workers entering or leaving the labor force).

Having a program that assures a transfer for those with zero or very low income could also be important for reasons outside the basic Mirrlees (1971)/Saez (2002) model. For one, the program could provide insurance for individuals to take risks: for example, choosing to plant a crop that could be high-return but has a nontrivial probability of zero income might be a more attractive option with a universal basic income in place. The same ideas could be applied to other types of risky decisions, such as human capital investment (Stantcheva 2014).

What about Poor Countries, Where Tax Systems Are Less Developed?

In a developed country, where the government observes income for most people and can redistribute through the tax system, this framework is broadly

Figure 2

Example of Post-Tax Income Schedules With and Without a Universal Basic Income (UBI), with a Tax-Exempt Region
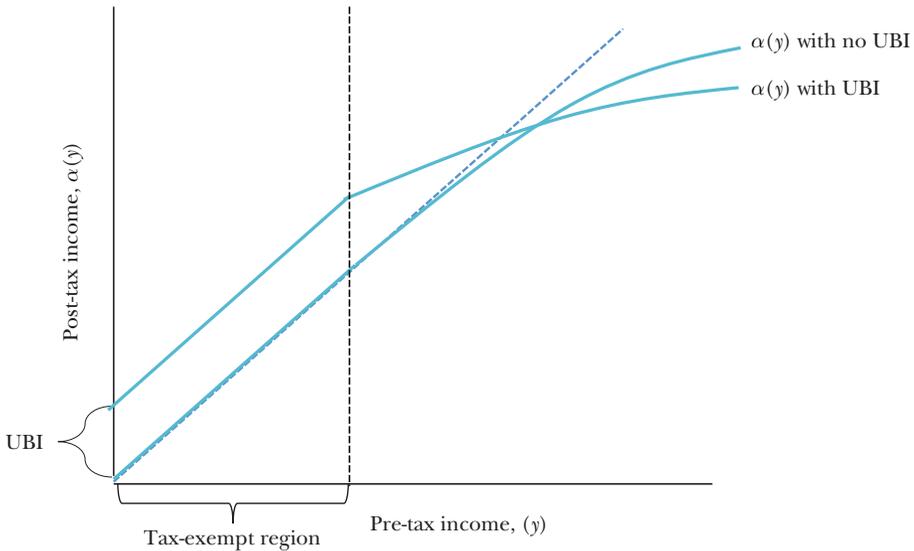
applicable. However, if we are interested in the poor in developing countries, it is more complicated.

In less-developed countries, substantial activity occurs in the "informal sector" that includes casual labor, undocumented firms, and small farms. Thus, most individuals are outside the tax net. Jensen (2016) documents that the share of employment for which people pay *any* income taxes (and therefore for which the government observes income) rises substantially with real per-capita national GDP. Conversely, in most poor countries, the government does not observe any information about income for most people, and in particular, for the poor. In Indonesia and Peru—the two countries that we explore in detail below—Jensen reports that 87.5 percent and 79 percent of the countries' employed populations, respectively, have incomes below the tax exclusion thresholds. Jensen argues that these thresholds emerge from the work environment—when most people work in large firms, the government can easily observe incomes (Kleven, Kreiner, and Saez 2016). But when there is a large informal sector, or many people work on their own or in tiny firms, as is true in most developing countries (Hsieh and Olken 2014), this is not possible.

Why does this matter? Figure 2 illustrates an example tax schedule with and without a universal basic income when most individuals—in developing countries, often upwards of 80 percent of people—fall into a tax-exempt region. In the tax-exempt region, everyone receives the same net transfer regardless of whether they are in the 5th or 70th percentile of income. However, to finance a universal basic income of a size similar to the case in Figure 1, one would need to increase marginal tax rates substantially more for those relatively few people inside the tax net. To the

extent that these individuals are particularly productive, or that, for a corporate income tax, the firms paying it are more productive, such a tax increase may have disproportionally larger efficiency consequences.

In this way, a universal basic income will function somewhat differently in a country with a large number of people outside the tax net. In particular, 1) the *net* transfer could be the same for the majority of the population (all those below the tax-exempt cutoff), which may not be the optimal resource allocation if we want to focus on the poorest of the poor; and 2) relatively few households will be able to contribute to the financing, which may restrict the overall level of funds available for the universal basic income. In countries with large numbers of people outside the tax net, any attempt to target resources toward the poorest of the poor requires using an alternative data collection approach beyond the income tax system.

The discussion thus far has assumed that a universal basic income is financed entirely through income taxes (individual or corporate), which in developing countries are paid primarily by the very top earners. These taxes combined with import duties (which are also likely to fall on the rich), make up about 42 percent of tax revenue for low- and middle-income countries (according to the most recent World Bank *World Development Indicators*), with the rest of tax revenue in these countries coming from consumption taxes, which are typically proportional to consumption. A universal basic income financed through a combination of an income and consumption tax will look closer to a combination of Figures 1 and 2. In this case, one can tax back some of the transfer on households below the income tax exemption through the consumption tax, but the fact that the tax is generally proportional to consumption limits the flexibility in doing so.

How to Target the Poor in Developing Countries

How can developing country governments target the poor when income is not observable? When the government observes a proxy for income and targets households based on that, one has what is termed a "proxy-means test." The implementation of a proxy-means test is typically based on large, periodic quasi-censuses of the population, focusing on those most likely to be poor. In these censuses, government enumerators go door-to-door, often visiting millions of households. Some countries, such as India, do target government assistance based on self-reported and unverified income, but this is the exception, not the rule, because people can easily lie if there is no way of verifying it. Instead, government census enumerators typically ask about assets, all of which are easy to observe directly. Examples might include ownership of items such as televisions and refrigerators, the type of material used in one's roof, floor, and walls, the number of rooms in one's house, and so on.

The government uses these assets to predict incomes (or per-capita consumption, which can be easier to measure in survey data). Thus, eligibility for benefits is based on *predicted*, rather than actual, income. This method is quite common, and is used in both large countries such as Indonesia, Pakistan, Nigeria, Mexico, and the Philippines, as well as a number of smaller countries, ranging from Burkina Faso to Ecuador to Jamaica (Fiszbein and Schady 2009).

There are a variety of methods to predict income (or consumption), but they all share basic features: The government takes another dataset that was collected in a low-stakes context (for example, one collected not for targeting, but rather just for research purposes), and therefore for which households have no strong incentive to lie. In this dataset, the government observes the same asset variables as in the proxy-means census and also observes a measure of poverty, such as a household's monthly income or per-capita expenditure. The government then estimates a regression with the measure of poverty as the dependent variable and the assets as explanatory variables. The proxy-means score is the predicted income or expenditure, which the government can calculate for any household using the coefficients from that regression.

The government then can set a threshold for eligibility and distribute benefits to all households with predicted incomes below the threshold.[2] If the government sets the threshold level for eligibility in such a way that it covers the entire population, then a "targeted" program effectively becomes a universal basic income. Because predictions of income, of course, happen with error, targeted systems based on a proxy-means test will feature both inclusion and exclusion errors.

Comparing Targeted Programs to a Universal Basic Income: Examples from Indonesia and Peru

Depending on the policy context, either a universal basic income or a targeted program may be optimal. Besley (1990) provides an early discussion of tradeoffs between the two approaches. Here, we illustrate the tradeoffs using data from Indonesia and Peru, two countries that have targeted, nationwide transfer programs. The simulations described in this section (particularly Figures 4 and 5) are broadly similar to other analyses of targeting performance, such as Ravallion (2009), Alatas, Banerjee, Hanna, Olken, and Tobias (2012), and particularly Klasen and Lange (2016).

Indonesia has a number of targeted transfer programs, ranging from conditional cash transfers (Program Keluarga Harapan or PKH) to scholarships for poor students (Bantuan Siswa Miskin or BSM) to subsidized health insurance for the poor. In addition, the government has periodically conducted nationwide, unconditional cash transfers (previously called Bantuan Langsung Tunai or BLT, and more recently called Bantuan Langsung Sementara Masyarakat, or BLSM). The

---

[2] Once the eligibility threshold for a program is set, most developing countries do not vary the level of assistance among the eligible group. Part of the reason is that the imprecision of estimating income based on assets is well understood. Also, it is easier for reasons of politics and transparency just to explain that a given program provides all who are eligible with a certain benefit. However, countries can link benefits levels to predicted income levels by having programs with different eligibility cutoffs. For example, in Indonesia, the conditional cash transfer program called PKH targets roughly the bottom 7–10 percent of households (the "very poor"), but those who are in roughly the bottom 40 percent of the national income distribution (the "near-poor") can qualify for subsidized health insurance even if they cannot access PKH.

government has conducted nationwide targeting censuses approximately every three years since 2005, and then uses proxy-means testing.

Several different types of transfer programs also exist in Peru, ranging from nutritional subsidies (Vaso de Leche) to subsidized health insurance (Seguro Integral de Salud, or SIS) to conditional cash transfers for poor families (Programa Nacional de Apoyo Directo a los Más Pobres, or Juntos). Eligibility for these programs is determined by a proxy-means targeting system, the Sistema de Focalización de Hogares (SISFOH), which assigns households a numerical index to identify eligible households and then allows each individual program to select its final list of beneficiaries. Recertification occurs every three years via a new targeting survey.

For simplicity, we choose one targeted program per country to explore (although the basic results would continue to hold if we chose a basket of programs). In Indonesia, we focus on the Bantuan Langsung Tunai (BLT) program, a temporary and periodic unconditional cash transfer to poor households. The program was implemented first in 2005–2006, and then again in 2008–2009 and in 2014, to help offset shocks in fuel prices. Beneficiaries receive the equivalent of approximately $100 over the course of one year (World Bank 2012). For Peru, we focus on the Juntos program, which is a conditional cash transfer to mothers designed to subsidize child health and education. Beneficiary households receive a monthly transfer of 100 soles (approximately $30). Both programs target roughly one-third of the population.

Predicting Income with Proxy Measures

To simulate targeting in the programs, we obtained household-level data from the Indonesian National Socioeconomic Survey (SUSENAS) and the Peruvian National Household Survey (ENAHO), both for the years 2010–2011. Both datasets are used for targeting in the respective countries. Our sample contains 263,705 households in Indonesia and 46,305 households in Peru. Both surveys contain the complete set of asset variables used in targeting, as well a measure of *actual* per-capita consumption for the household, so we can use these datasets to examine targeting accuracy. For each country, we randomly divide the observations into equally sized "training" and "test" sets. In the training set, we regress monthly household per-capita consumption on the actual indicator variables used in each country's respective proxy-means test formula: 82 variables in Indonesia; 72 variables in Peru (regressions provided in the online Appendix). We then predict monthly per-capita consumption for each household in the test sets using the coefficients from the training regressions, and use this "predicted per-capita consumption" for targeting purposes in the simulations below.

The predictors used are not perfect. The typical fit we found of these regressions (the R^2) is between 0.53 and 0.66—so while the regressions have a good amount of explanatory power, they also lead to prediction errors: *exclusion errors* of excluding some who should have been eligible (households with true per-capita consumption below the poverty line) and *inclusion errors* of including some who should not have been eligible (households with true per-capita consumption above it).

We apply the formula to the half of our datasets that was not used for estimation (the "test" datasets). We can then use the actual per-capita consumption measure within the survey to assess whether the formulas accurately predict a household's place in the income distribution. Figure 3 shows the results, plotting actual per-capita consumption against predicted per-capita consumption. The figure shows four regions: correctly included (bottom-left region), correctly excluded (top-right region), inclusion error (top-left region), and exclusion error (bottom-right region).

For the exposition here, we are simplifying the actual process by which each country determines benefit eligibility. In modeling these two programs, we intend to target households below 1.5 times the poverty line in Indonesia (about 33 percent of our sample), and households below the poverty line in Peru (about 28 percent of our sample). However, in reality, the monetary poverty line varies by year, urban/rural zoning, and province (in Indonesia) or geographic region (in Peru). In Peru, instead of predicting income directly, the national household targeting system (SISFOH) actually produces a numerical index of household poverty using proxy means-testing variables; it then also incorporates certain other household circumstances (such as consumption of water and electricity) when determining final eligibility. In Indonesia, the proxy means-testing models are estimated separately by region of the country to allow the model to account for regional variation in how consumption of goods predicts income across different areas. In both countries, prediction using proxy-means tests is used to determine eligibility for various transfer programs, but each individual program may edit the list of ultimate beneficiaries.

The Inclusion versus Exclusion Error Tradeoff

Our first step is to explore the tradeoffs in the errors of inclusion and exclusion. One way to think about the government's problem is that by setting different cutoffs for program eligibility c, the government can choose the balance between the inclusion and exclusion errors that it makes. Figure 3 plots the results with one example c, but one can see how varying (that is, shifting the vertical line in Figure 3 to the left or to the right) would change the balance of inclusion and exclusion error.[3] Say that the government aims to assist those who are actually poor. Intuitively, not giving out the program to anyone (and setting $c = 0$) means no transfers, and hence very high exclusion error, because you are excluding everyone below the poverty line, but it also means no inclusion error, since no higher-income people who should not be receiving assistance are getting it. In contrast, a universal basic income (setting $c = \infty$) implies no exclusion error since all of the poor will get it, but very high inclusion error, because all the higher-income people are getting assistance, too. For values in between, varying the cutoff value c allows us to trace out the tradeoffs between inclusion error and exclusion error available to the government.

---

[3]Note that these definitions of inclusion and exclusion error are with respect to a household's true poverty status (above or below poverty line based on true per-capita household consumption; that is, holding the horizontal line in Figure 3 fixed), not with respect to the program design (that is, the eligibility choice c, shown by the vertical line).

Figure 3
Predicted versus Actual per-capita Consumption for Households in Test Set Data

A: Indonesia
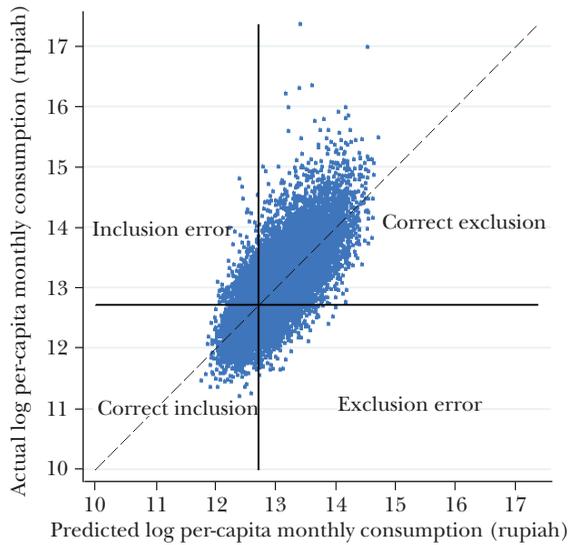
B: Peru
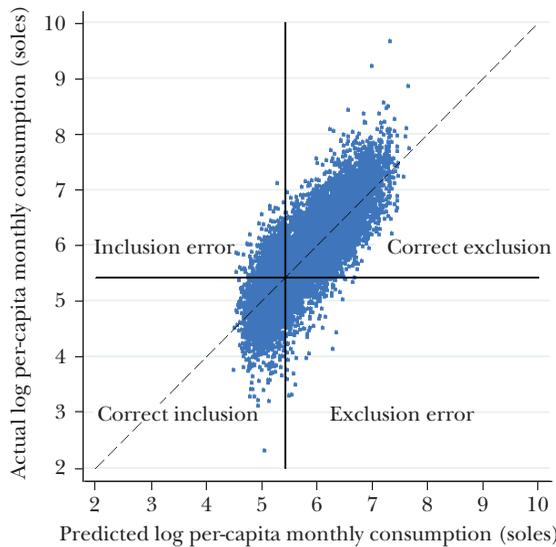
Note: The figures above plot actual log per-capita monthly consumption against predicted log per-capita monthly consumption for households in the test set data for both Indonesia and Peru (see text for details). The solid vertical and horizontal lines represent eligibility cutoffs for the programs we model here. (In Indonesia, we set the cutoff approximately at the 33rd percentile for consumption for the sample; and in Peru, we set the cutoff at the 28th percentile.) The dashed line is a 45° line. For legibility, the points plotted above represent a random sample of 10 percent of our full data from Indonesia and 50 percent of our full data from Peru.

These tradeoffs can be seen in Figure 4A, where we plot data from the proxy-means test regressions that we estimated for Indonesia and Peru. We define a target poverty level y^* for each country, with the idea that the program is trying to target all individuals with household per-capita consumption below y^*.[4] We then define exclusion errors as failing to give benefits to those with household per-capita incomes below y^*, and inclusion errors as giving benefits to those with household per-capita incomes above y^*. These curves, also known as ROC curves,[5] plot inclusion error on the *x*-axis and 1 – exclusion error on the *y*-axis, and trace out the tradeoff between exclusion error and inclusion error obtained by varying the cutoff value *c*, holding the target poverty level for the program (y^*) fixed.

Figure 4A confirms the fundamental tradeoff between inclusion and exclusion error. In both Indonesia and Peru, to reach 80 percent of the intended beneficiaries (exclusion error of 20 percent), the government would need to tolerate an inclusion error of between 22 and 31 percent. Further reductions in exclusion error come at the cost of much higher inclusion error.

While more sophisticated prediction methods (like more flexible prediction equations, machine learning methods, and so on) can improve on these predictions (McBride and Nichols 2016), the fundamental tradeoff between inclusion and exclusion error that we document here remains. Also, these graphs ignore the very real problems of incomplete take-up by those who are deemed eligible, which could affect both the targeted programs and universal transfers. We will revisit this topic below.

Narrowly versus Widely Targeted Programs with a Fixed Budget

If the government only cares about making sure the poor have access to transfers of some type, a universal basic income—which eliminates exclusion error entirely—may seem attractive. However, as we move to the right in Figure 4A, by raising the cutoff level *c*, more and more people become eligible for the transfer. If the total budget for the program is fixed, then the transfer given per-person falls as we move to the right on this curve. Therefore, moving to the right on the curve not only changes the tradeoff between exclusion error and inclusion error, but holding the total program budget fixed, also changes the amount given.

Figure 4B shows this other aspect of the tradeoff, graphing the level of cash transfers disbursed per beneficiary under Indonesia's BLT program and Peru's Juntos program under different targeting schemes, as defined by different levels of the cutoff income *c*. The transfer level to eligible households steeply falls as you allow more people into the program. Again, the universal basic income can

---

[4]For Indonesia, we plot the ROC curve with y^* equal to the 1.5× the government's official poverty line, and for Peru, we plot the curve with y^* equal to the government's official poverty line. In the Appendix, we compare curves setting y^* equal to 0.75× the government's official poverty line, which we call the "extreme poor"; 1.0× the government's official poverty line; and 1.5× the government's official poverty line, or the "near poor." The tradeoff is similar in all three cases.

[5]ROC stands for "receiver operating characteristic." It refers to a situation in which there is binary classifier—in this case, the proxy-means score combined with different cutoff levels *c*—and shows the tradeoff between true positives and false positives as one varies the cutoff level *c*.
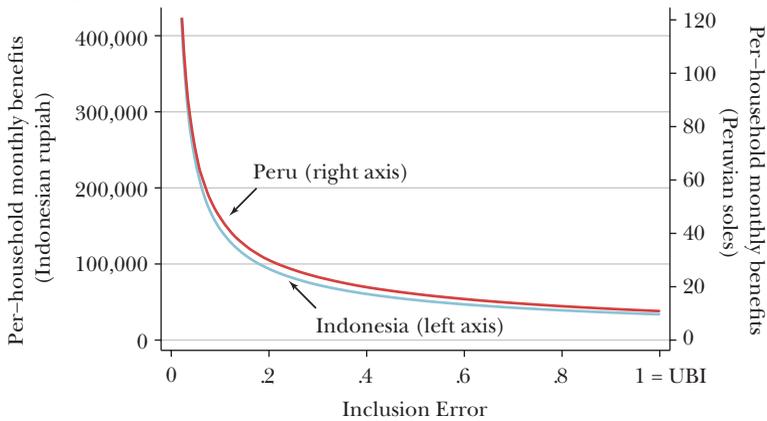
Figure 4
Tradeoffs between Inclusion Error and Exclusion Error by Varying Eligibility Cutoff
(exclusion error is the portion of intended beneficiaries excluded; inclusion error is the portion not intended to be beneficiaries who are included)

A: ROC Curves for Program Targeting (by Household)
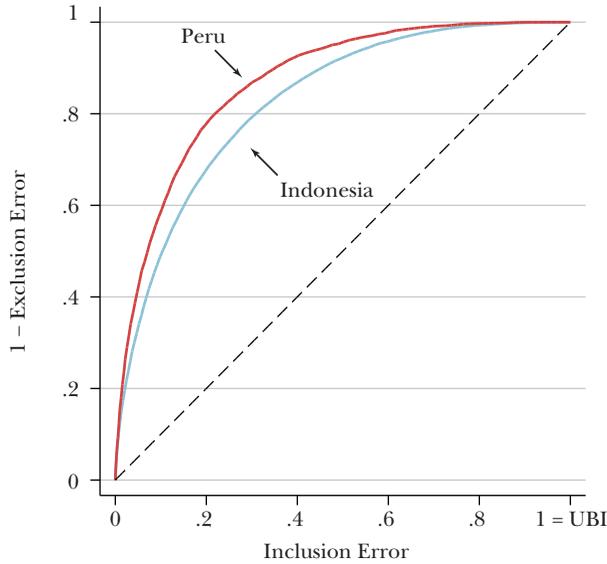
B: Benefits per Household versus Inclusion Error for Different Eligibility Cutoffs

Notes: ROC stands for "receiver operating characteristic." It refers to a situation in which there is binary classifier and shows the tradeoff between true positives and false positives as one varies the cutoff level *c*. Figure 4A plots the ROC curves for both Peruvian and Indonesian test set data, using a target household per-capita consumption value equal to the government poverty line in Peru and 1.5 times the government poverty line in Indonesia. Figure 4B traces the per-household benefit amount as inclusion error increases, assuming a fixed transfer budget of approximately $1.83 billion per year in Indonesia and $274 million per year in Peru (modeled after the sizes of BLT and Juntos, respectively). We include cost-savings from administrative targeting at the universal basic income (UBI) point, where inclusion error equals 1. See text for details.

be thought of as the rightmost extreme point, with zero exclusion error and the smallest per-capita transfer for a given budget.

One argument for a universal basic income is that you can save on the administrative costs of targeting, because you do not need to do the periodic targeting census. To capture this, in Figure 4B we add back in the costs of targeting for the UBI, so the benefits per-capita ticks up very, very slightly at the universal basic income point ($c = \infty$). However, the administrative cost of targeting is very low relative to the benefits given out: Indonesia spends roughly $42 million every three years for updating its Unified DataBase, with annual operating costs of about $1.1 million (for further information, see Bah, Nazara, and Satriawan 2015). Peru spends roughly $10.8 million every three years with annual operating costs of about $1.1 million (Ministerio de Economía y Finanzas 2008). Per year, this translates to an additional 0.8 and 1.7 percent of the overall transfer budget in Indonesia and Peru, respectively. Thus, administrative costs are not a large driver of the benefits level, especially in comparison to the level of inclusion error one chooses, which is why the discontinuity at the universal basic income point is visually imperceptible in Figure 4B (although you will be able to see it in Figure 5).

Welfare Comparisons

For a given total budget, which point on the graph is "best" depends on various social welfare weights. For example, it is necessary to put a weight on inclusion errors, while remembering that just because someone is above the cutoff line does not imply that the social value of them receiving a transfer is zero—indeed, generically such a transfer still has a positive value. It is necessary to put a weight on exclusion errors: how costly is it to the government for some deserving people not to obtain the transfer, merely because of errors in the targeting formula? It is necessary to put a weight on the implied reduction in per-capita transfer to the very poorest that is required to accommodate the increased number of beneficiaries. It is also necessary to take a stand on how much more the poor value a given dollar than the rich (that is, the difference in marginal utilities for poor and rich), which is given by the curvature on the utility function. For a given amount of targeting, one can then determine the per-capita transfer.

Deciding how to parse the tradeoffs between inclusion error, exclusion error, and per-capita benefits requires specifying a social welfare function that allows the government to evaluate the social benefits from these different decision rules while holding the overall size of the transfer constant. Here, we will evaluate the total social welfare using what is commonly known as a constant relative risk-aversion CRRA-utility function:

$$U = \frac{\Sigma (y_i + b_i)^{1-\rho}}{1-\rho}$$

where y_i is household i's pre-tax per-capita income, b_i is the per-capita benefits assigned to household i, and ρ is a coefficient of relative risk-aversion, where higher values of ρ put higher weights on transfers received by the very poor. (Note that

Figure 5
Social Welfare versus Inclusion Error
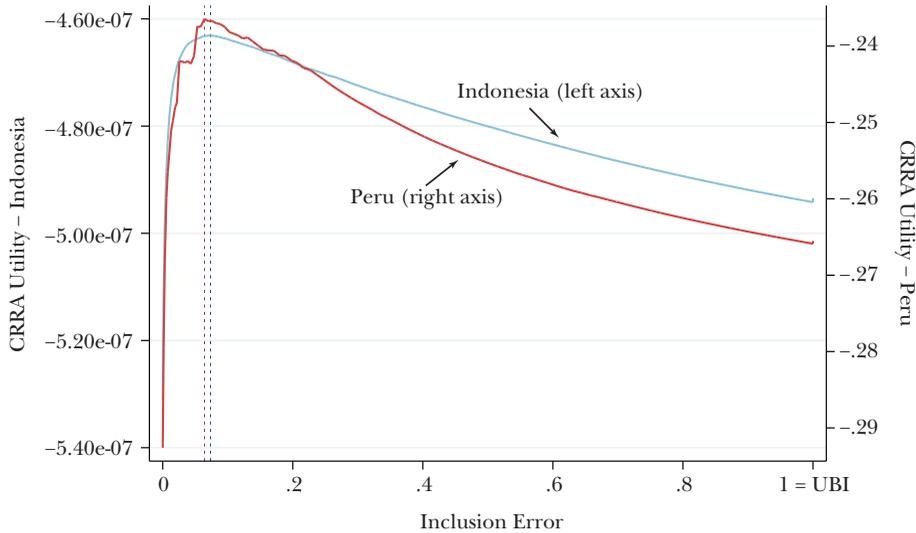
Notes: Figure 5 uses the same pooled household data as in Figure 4. For each cutoff value c, we calculate the per-capita benefit amount for included households, and then calculate constant relative risk-aversion utility (CRRA) using the formula presented in the text with $\rho = 3$. I. Utility is negative given the CRRA assumptions, with higher (that is, less negative) values indicating higher utility. If a household is not included in the program at a given value of c, we simply set $b_i = 0$, but their utility is still included. Dashed lines indicate the point of maximum social welfare in each country. We include cost-savings from administrative targeting at the universal basic income (UBI) point, where inclusion error equals 1.

utility is negative given the CRRA assumptions with less negative values indicating higher utility.) By comparing the sum of *U* for different configurations of benefits, one can see which scheme produces the highest net social welfare judged by this metric.

Of course, conclusions will depend on the particular value of ρ one chooses. Here, we calculate the social welfare from programs with different levels of targeting in Figure 5 for each country, using $\rho = 3$. (As a robustness check, we also calculated social welfare with $\rho = 1$ and $\rho = 5$, and the conclusions are qualitatively similar.) Specifically, in Figure 5, we plot social welfare evaluated against the amount of inclusion error for a fixed budget, for both countries. As before, social welfare will increase discontinuously for the universal basic income because no targeting costs are incurred in this program.

The key finding from this graph is that narrowly targeted programs—those focused on distributing large benefits-per-capita to the poorest of the poor—appear to achieve much higher utility levels than less narrowly targeted programs, including but not limited to a universal basic income. In Indonesia, the socially optimal program calculated in this way targets about 19 percent of the population, with inclusion error of 7.4 percent and exclusion error of 58.2 percent;

for Peru, the socially optimal program targets approximately 18 percent of the population, with inclusion error of 6.4 percent and exclusion error of 52.4 percent.

Even programs that appear somewhat badly targeted could still dominate a universal basic income. To understand where the programs in Indonesia and Peru fall in terms of social welfare, we can see where the current level of inclusion error roughly falls. Bah, Bazzi, Sumarto, and Tobias (2018) calculate that inclusion error in the 2008 Indonesian BLT program was roughly 34 percent, which is higher than the social optimum, but still substantially more narrowly targeted than a universal basic income. Peru's Juntos is at roughly 6.4 percent inclusion error, according to Robles, Rubio, and Stampini (2015), suggesting that the rate of inclusion error is close to the social optimum as calculated here.

In short, this analysis illustrates that targeting could be an effective mechanism in improving overall social welfare, even with high levels of targeting error. However, note that this discussion focuses only on tradeoffs involved in choosing whether to target transfers based on one survey—that is, how to identify the poor in the static situation of a given point in time. In reality, income status varies over time. Thus, the size of the relative errors will also depend on how frequently the government collects asset data from households and how much churn in and out of poverty occurs over time. Because targeting surveys are infrequent, actual targeting errors may be higher than what we report here, and so one may want to account for this factor in determining where to place the cutoff.

Other Challenges with Targeted Transfers

In choosing between a broader approach like a universal basic income and a more narrowly targeted approach, several other important considerations arise in addition to inclusion and exclusion errors. We consider four issues: breadth of political support, horizontal equity, transparency, and labor market distortions.

First, we have focused on the case where the government has a fixed budget available for the transfer program. However, political support for redistribution may be much higher if everyone gets a "piece of the pie" (Gelbach and Pritchett 2002). Thus, as inclusion error increases, the overall amount of funding available for transfers may increase. As the elasticity between the total budget and the number of beneficiaries increases, universal transfers become more attractive (Klasen and Lange 2016).

Second, the principle of horizontal equity holds that each person who has the same relevant conditions should be treated the same. Transfer programs using proxy-means tests will have errors in the prediction algorithm—and so they do not treat all people with the same utility of receiving the benefits identically. For programs where the income cutoff for providing benefits is near-zero, horizontal equity violations are (mechanically) close to zero, because almost no households receive the program. Violations of horizontal equity are also mechanically zero for universal basic income: by definition, if all households receive the transfer, then all households are treated the same.

Figure 6
Horizontal Equity versus Inclusion Error in the Programs from Peru and Indonesia
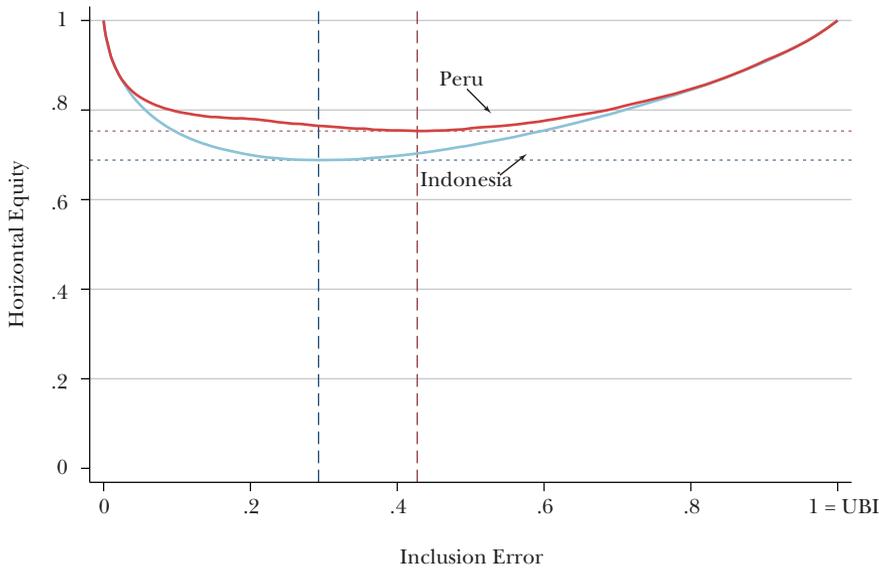
Note: Figure 6 traces horizontal equity as a function of inclusion error. At each cutoff *c*, we calculate, for each household, the percentage of households within ±5 income percentiles (based on actual income) that received the same benefit status—included or excluded—based on the results of proxy-means test prediction. In other words, for households that were included in the program at a given *c*, we calculate the percentage of similar households that were also included; for households that were excluded, we calculate the percentage of similar households that were also excluded. At each cutoff point, we average these percentages over all households to compute the horizontal equity variable graphed above. Dashed lines indicate the location and value of minimum horizontal equity.

　　To calculate a measure of horizontal equity and see how it varies with the cutoff value *c*, we do the following for both Peru and Indonesia. For each cutoff value *c*, we determine which households would receive the program and which would not. For each household income level, we then compute the percentage of households 5 percentiles above or below who have the same outcome (that is, for households receiving the program, the percentage of households 5 percentiles above and below who also would receive the program; for households not receiving the program, the percentage of households 5 percentiles above and below who would not receive the program). For each cutoff (and hence allocation rule) *c*, we average these two percentages over all households to get a measure of horizontal equity. We then repeat this exercise for each possible value of *c*.

　　Figure 6 shows the resulting measure of horizontal equity and how it varies as we vary the cutoff *c*, and hence the inclusion error, of the program. We find that for cutoff levels in an intermediate range, in which inclusion errors were roughly 30–40 percent, it is common to find that horizontal equity is violated about 20 percent of the time in Peru and 30 percent of the time in Indonesia. In addition to the

unfairness they represent, violations of horizontal equity can also lead to significant political problems.

A third important consideration involves the logistics of implementing targeting schemes in a low-capacity (and potentially corruption-prone) environment. A proxy-means test has an inherent lack of transparency, because eligibility is determined based on a weighted sum of many different variables, which would be hard for ordinary citizens to understand even if the weights were public information. Moreover, the proxy-means test formula (essentially the weight that the coefficients in the regression would give to each asset) are kept secret because if they are known, households (perhaps in cooperation with better-informed agents) may strategically misreport or hide assets to make sure they fall under the cutoff. Using data from Colombia, Camacho and Conover (2011) show that over time as the formula became known, there was substantial bunching of reporting right below the eligibility cutoff. Given this concern, governments not only keep the formula confidential, but they also tweak the coefficients every few years.

Thus, it is very difficult for citizens to verify that a proxy-means test scheme is being implemented properly. If a local official says to a citizen, "Sorry, you're not on the list of eligible beneficiaries," villagers have little recourse, because their true eligibility is hard to verify. On the other hand, a universal basic income is inherently transparent, in the sense that *everyone* is entitled to a transfer.

The evidence suggests that greater transparency of the lists of eligible beneficiaries may be important to effective implementation of a proxy-means test-based scheme. In Banerjee, Hanna, Kyle, Olken, and Sumarto (2018), we and our coauthors worked with the Indonesian government on a randomized experiment in which heads of all villages received a list of who was eligible to receive subsidized rice, but in a treatment group of villages, the central government also mailed out an "identification card" directly to citizens deemed eligible. Mailing identification cards to the beneficiaries resulted in treatment villages receiving 26 percent more subsidized rice compared to control villages—and reduced "leakage" (rice that, as far as we can tell, was not distributed to anyone) by between 33 and 58 percent.

Finally, targeting may introduce distortions. These distortions can, in theory, reduce the quantity of work, either because income effects decrease the incentive to work, or because phasing out the transfer over some income range decreases the after-tax-and-transfer effective marginal wage in that range. To study the income effects, we and our coauthors re-examined seven different randomized trials of government-run cash transfer programs throughout the developing world (Banerjee, Hanna, Kriendler, and Olken 2017), which provided between 4 and 20 percent of household consumption to beneficiaries for many years. We found no systematic evidence that any of these programs reduced labor supply.

Regarding the second channel (effects that come through changing the after-tax-and-transfer effective marginal wage), a vast literature in developed economies has studied the alteration in work incentives that arises either from the changes in the income tax rate formula needed to raise the revenue to pay for the transfer, or from the effective marginal tax rate imposed by the phase-out of the transfers.

In the US context, for example, the earned-income tax credit provides a negative marginal income tax rate at the very bottom of the tax schedule (thus subsidizing work), and was designed with these concerns in mind, given that there is evidence in the United States that labor supply elasticities are higher on the extensive margin (working at all versus staying out of the labor force) than on the intensive margin (how much you work) (for example, Eissa and Leibman 1996; Meyer and Rosenbaum 2001; Saez 2002). Relatedly, a major goal of the 1996 US welfare reform was to redesign the welfare system to reduce very high marginal tax rates that came from rapid phase-outs of benefits over certain income tax ranges. As discussed above, a universal basic income program does not eliminate this issue entirely, although the relevant distortions to think about for a universal basic income are the increases in marginal income tax rates required to finance the universal basic income.

These issues are related to, but somewhat different from, what happens with a proxy-means test in a developing economy. In this setting, the labor supply distortion will involve the implied tax rates from the phase-out of the program for those who are eligible, but will also be smoothed out by the error rate in the formula for those who have equivalent income (or consumption) but are not designated as eligible by the formula.[6] Perhaps ironically, a more accurate proxy-means test formula, with fewer errors, may actually induce *more* labor supply distortions, since fewer errors in the proxy-means formula will imply that some households face a steeper effective tax rate over some range. More noise in the proxy-means test formula, making the proxy-means formula less predictive, tends to smooth out the phase-out region, and will thus tend to reduce the implicit tax rate that a population faces from a targeted benefit.
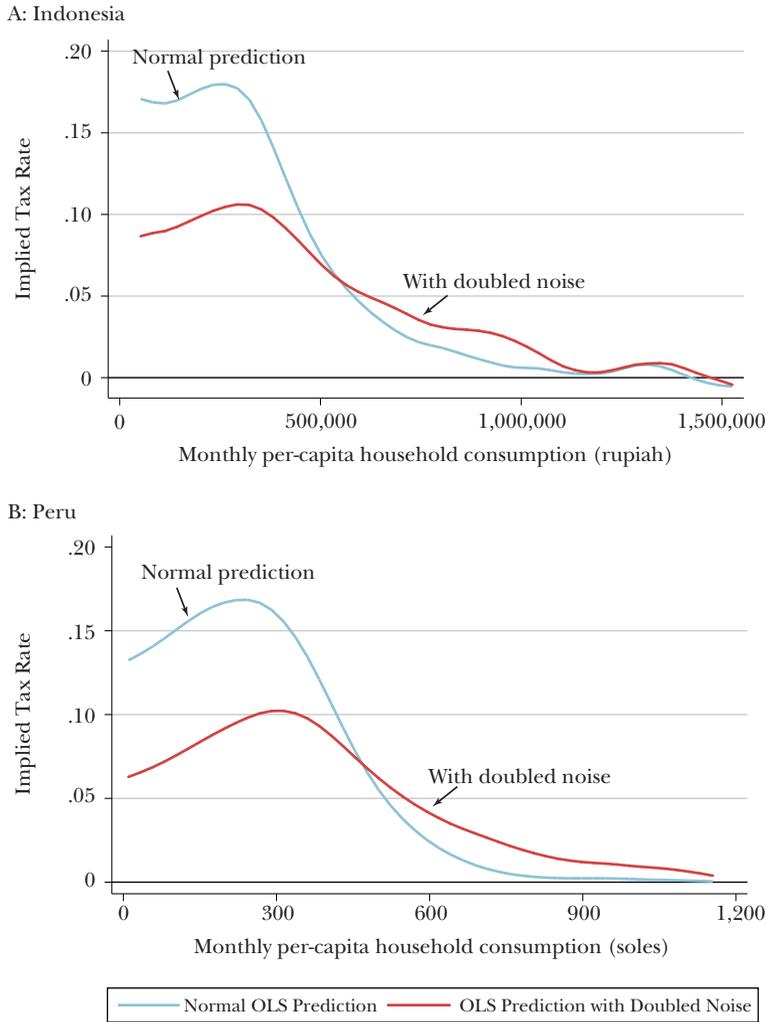
For example, Figure 7 illustrates the implied marginal tax rates from actual proxy-means test formulas and from a counterfactual formula where we double the amount of noise, making the proxy-means test less predictive than it is in practice. As the figures show, doubling the amount of noise in the proxy-means test formula reduces the implied marginal tax rate for the poor, but it increases the implied marginal tax rate somewhat for the middle-class, who face some probability of receiving the benefits (due to inclusion error) that is declining in their income. If households are aware that increasing their income or consumption may reduce their chances of being eligible for benefits, they may reduce their labor supply accordingly, though documenting to what extent this happens empirically remains an important direction for future research.

The proxy-means test puts particularly heavy weight on certain assets. While the precise proxy-means test formula may not be known, households do observe which variables the government asks about—that is, they know that the poverty census asks about televisions, refrigerators, and so on. The potential for linkage from ownership

---

[6] To see this, recall that for each income level y, one can think of the proxy-means test as producing a probability of receiving the benefit $p(y)$. The fact that the proxy-means test has predictive power implies that p will be decreasing with income ($p'(y) < 0$), but the fact that there are inclusion and exclusion errors means that $p(y)$ is continuous. Abstracting away from the rest of the tax system, one can then write a household's expected after-tax income as $\alpha(y) = y + p(y)b$. One can therefore think of the "expected tax rate" imposed by the proxy-means test system as $-p'(y)b$.

Figure 7
Implied Tax Rate versus Household Income, with Noisy Proxy Means Test Formulas

A: Indonesia

B: Peru

Note: Figure 7 illustrates the implied marginal tax rates from actual proxy-means test formulas and from a counterfactual formula where we double the amount of noise, making the proxy-means test less predictive than it is in practice. The implied marginal tax rates are plotted against actual per-capita household consumption for the Indonesian and Peruvian cases described in the text. We designated households as benefit recipients if *predicted* income was less than 1.5 times the poverty line (in Indonesia) or less than 1 times the poverty line (in Peru). Households above the 95th income percentile were dropped from both country samples. We then performed a local polynomial regression of benefit receipt status on *actual* income, using a bandwidth of 50,000 rupiah (Indonesia) and 75 soles (Peru). We took the first derivative of this function and multiplied it by (minus) the benefit amount b in order to calculate the implied tax rate as a function of household income $(-p'(y)b)$. We then applied the same process after adding a random, normally distributed noise term (with the same standard deviation as our original predicted incomes) to these predicted income values. OLS is ordinary least squares.

of such assets to reduced eligibility for benefits puts a particularly high implicit tax rate on these particular assets, which could lead households to reduce their consumption of these assets. For example, from 1696 to 1851, Britain imposed a tax on windows, which were easily observable to the taxman—but this led to the construction of buildings with very few windows, low light, and poor ventilation (Glantz 2008; Oates and Schwab 2015). It is not clear that households understand the formula in a proxy-means test enough to distort their purchases of assets, and we are currently studying this question in cooperation with the government of Indonesia.

Alternative Methods of Targeting

Our discussion has assumed that the government produces an eligibility list for targeted transfers through a proxy-means test based on household-level data. However, we now consider two other methods of targeting: community-based targeting and a system that imposes costs on beneficiaries and lets people assist in targeting themselves. We also consider how conditional cash transfer programs, popular throughout the developing world, can be thought of in relation to the targeting challenge.

Community-Based Targeting

In a community-based targeting process, a fixed number of program slots are allocated to a given community, such as a hamlet (a part of a village). The hamlet, through a participatory meeting of some type, decides who among them is most in need of the transfer.

One might think that such a system would be unpopular: after all, who wants to discuss their income in front of their neighbors and peers? (Not Americans, as discussed in Card, Mas, Moretti, and Saez 2012!) In addition, one might fear that such a system is prone to elite capture: for example, a powerful headman might manipulate the meeting to put his nephew on the list of beneficiaries. However, when we and our coauthors tried such an approach experimentally in Indonesia and compared it experimentally, we found that citizens vastly preferred the community based-approach to a data-based proxy-means test. In addition, the community approach did a much better job of identifying those households who self-assessed themselves to be poor, and was only slightly worse at identifying the poor based on per-capita consumption levels. Furthermore, there was no detectable elite capture (Alatas, Banerjee, Hanna, Olken, and Tobias 2012).

A community targeting approach avoids some key pitfalls of the proxy-means test related to transparency and horizontal equity. The process for allocating beneficiaries under the community approach was completely public and transparent—unlike the proxy-means test that relies on secret formulas to produce a list of eligible and ineligible households that villagers do not really understand. It also addresses the perception of a horizontal equity problem—from the perspective of the villagers, if it is broadly agreed that person *A* is poorer than person *B*, they can ensure that person *A* receives the transfer.

Community targeting also has the potential to be less distortionary. In Alatas et al. (2012), we showed that communities appear to target based on earning potential, rather than actual earnings. For example, widows—who have lower earning potential than households headed by men—appear to be more likely to be ranked as poor, even conditional on their actual income level. The same is true for education: those with high education, and hence high earning power, are ranked as wealthier than their actual income alone would imply. To the extent that communities target based on *earning potential* rather than actual income, such a system has the potential to be less distortionary not only than a proxy-means test, but also than a regular income tax.

Community targeting has important limitations—in particular, it does not readily allow for comparisons *across* locations, and if implemented at scale, could lead households to take actions to hide their income from their neighbors. But its potential advantages may mean that it is worth wider consideration as part of the targeting toolkit.

Differential Costs of Take-Up and Self-Targeting

A substantial fraction of households who are eligible for a government benefit program do not enroll in the program. Moffitt (2003) and Currie (2006a) summarize an extensive US-based literature studying the reasons for incomplete take-up, including the stigma of receiving benefits, incomplete information, and both the psychic and real costs of applying for programs.

A universal basic income may reduce the take-up problem, because the government can in theory send the checks to everyone automatically. This is, of course, easier said than done. Currie (2006b) shows that take-up is incomplete even in programs with universal eligibility in the United States, and a developing country setting raises additional issues. For example, the government would need to have a list of every individual without duplications. Most developing countries do not have such lists, although India has been working over the past decade to issue a unique, unduplicated and biometrically authenticated identity number through the Aadhaar program to each of its more than 1 billion citizens. Even with such a list, a country would need a way of automatically delivering the transfer. In most developing country contexts, most people do not have an address, let alone a bank account, so the idea of "mailing a check to the address on file" would not work. This suggests that even if enrollment is in theory automatic, in practice some take-up decisions and enrollment actions will need to occur.

However, having some type of costly take-up can also provide a benefit by helping in the targeting of a program. If the cost/benefit of taking up a program is higher for the rich than for the poor, the poor may be more likely to take up the program. This idea—that having an "ordeal" associated with receiving a program can improve the ability to screen among rich and poor—was developed by Nichols and Zeckhauser (1982), and further studied in the context of take-up decisions by Kleven and Kopczuk (2011). A key result is that one may want to offer a transfer program to all takers, but with the catch that signing up for the program entails a cost that is differentially costly for the rich.

Workfare programs, such as those offered by the US Works Progress Administration in the 1930s and through the National Rural Employment Guarantee Act (NREGA) in India today, operate on this principle. In the NREGA program, for example, any Indian in a rural area is entitled to 100 days of work at the official minimum wage, and currently about 180 million people per year are employed through the program, making it one of the largest (if not the largest) anti-poverty programs in the world today. This is a universal entitlement—there is no poverty screening, and there is a guarantee to everyone who wants the work. On the other hand, people with better job options tend to prefer doing something other than performing manual labor for minimum wage in the hot sun. Whether these types of programs—universal in nature, but achieving screening through so-called ordeal mechanisms—are better or worse than a universal basic income program depends on whether the substantial utility costs imposed on beneficiaries (like manual labor in the hot sun in order to receive benefits) are outweighed by the cost savings from improved targeting.

However, it may be possible to obtain many of the benefits from self-selection with much smaller utility costs to recipients if one marries self-selection to a proxy-means test. In Alatas et al. (2016), we and our coauthors study what happens when, instead of having the government come to everyone's house and conduct the poverty census automatically, households instead have to come and apply for the program, at which point they are screened using the proxy-means test.[7] Some fraction of households who pass the in-person eligibility test have their eligibility verified via a home visit, but this allows the government to skip the home visits for those not eligible and to check the eligible at home probabilistically, avoiding substantial costs. The proxy-means test is the same—what differs is the take-up step. Compared to manual labor in the sun, the cost of applying is relatively small—about three hours total in our context—and pales in comparison to the potential benefits, which average about $150 per year for six years in the context we studied. Yet, despite the relatively small costs, we show that self-selection yielded substantially improved screening; the beneficiaries selected by the application-based method were about 20 percent poorer than those selected through automatic enrollment.

The key seems to be that the application process substantially reduces the inclusion errors in the proxy-means test process while holding the official eligibility cutoff c fixed. A relatively well-off household forecasts correctly that the probability they pass the proxy-means test is small, and hence even the relatively small application cost can be enough to discourage them from applying. Since there are so many more relatively well-off households than very poor households, inducing these households to self-select out can actually lead to a substantial improvement in

---

[7]Both the automatic door-to-door approach and the application-based approach are used in other contexts. For example, Colombia uses an automatic enrollment door-to-door proxy-means test; Costa Rica and Chile use application-based proxy-means tests; and both Mexico and Brazil use different approaches in different contexts.

program targeting. This logic suggests that the benefits from adding self-selection to a proxy-means test system may be largest when the nontarget population is large relative to the target population. Moreover, an additional benefit is that some of the poorest households may live at the very margins of society and therefore be off the government's radar screen, and may be missed by enumerators trying to do the universal door-to-door screen. When beneficiaries can come to apply, such households can make themselves known.

However, application processes need not necessarily always improve targeting. For example, a complicated application form may dissuade those who are less literate or comfortable with bureaucracy from filling it out, leading to worse targeting (Gupta 2017). The degree to which these approaches can improve social welfare depend on whether the selection they induce enables the government to improve the efficiency with which it can deliver assistance to those who need it most, and also on the extent to which the efficiency gains achieved thereby are sufficiently large (and hence can translate into larger benefits) to offset the costs they impose on beneficiaries.

Conditional Transfers

A number of transfer programs in developing countries have explicit conditions that beneficiaries must meet in order to receive assistance. For example, one of the first of these programs, Mexico's *Prospera* program (originally called *Progresa*), gives cash transfers to poor households (screened via a proxy-means test) who also meet basic maternal and child health and education conditions, including regular pre- and postnatal care, regular growth monitoring, immunizations, school enrollment, and school attendance. Conditional cash transfer programs have spread throughout the world, and now are present in more than 63 countries, reaching tens of millions of households annually (Bastagli et al. 2016).

A number of randomized trials of these programs, including the original launch of *Progresa*, showed that these programs led to substantial improvements on the conditioned indicators, both initially (for example, see Behrman and Todd 1999; Gertler 2004) and over the medium term (Behrman, Parker, and Todd 2011; Barham, Macours, and Maluccio 2017; Kugler and Rojas 2018; Cahyadi, Hanna, Olken, Prima, Satriawan, and Syamsulhakim 2018). These conditions also may help make the programs more politically palatable, since voters in many countries may prefer that individuals do something in return for receiving aid.

While much attention has been given to the *incentive* effects of the health and education conditions, they can also have *targeting* effects. In particular, if some very poor households are unable to meet the conditions, and if the conditions are actually enforced (which they are not always), then some very poor households may be excluded from the programs. Baird, McIntosh, and Özler (2011) studied this issue, randomly comparing unconditional and conditional transfer programs in Malawi. While they found that the conditions increased school enrollments, they noted that the unconditional program actually had a larger effect on reducing teen pregnancy. The reason was that the unconditional program (unlike the conditional program)

still provided cash transfers to girls who dropped out of school anyway and were at the highest risk of pregnancy, and the cash transfers reduced the likelihood of early pregnancy for these girls.

While in principle conditional cash transfers could be either universal (anyone who meets the education and health conditions receives the transfer) or targeted using a proxy-means test, almost all are targeted, most often using proxy-means tests (Fiszbein and Schady 2009). One reason for this is that, in general, richer households are more likely to meet the health and education conditions than the poor. Thus, in the absence of explicit additional restrictions limiting these programs to the poor, the fact that the conditions are more likely to be fulfilled by wealthier households would actually make a universal version of these programs regressive rather than progressive.

Summing Up

Government-led anti-poverty programs have a crucial role to play in helping to eliminate extreme poverty worldwide. However, important questions remain about the form of these transfers, and in particular, whether they should be universal or targeted more narrowly to the poor. While some government programs are universal, such as publicly funded primary schools that are free to the families using them, most cash transfer programs today involve some type of targeting mechanism.

In this paper, we explore how to think about this tradeoff in developing countries, where incomes are unobserved for very large portions of the population. In this setting, universal transfers cannot be easily taxed back as one moves up the income distribution. As a result, universal transfer programs will give out the same *net* transfer quite high up the income distribution—and therefore the choice between a universal basic income program and transfers targeted through other means implies a very substantial tradeoff between eliminating exclusion error and giving much smaller transfers on a per-beneficiary basis.

Our evidence from Indonesia and Peru shows that existing targeting methods in developing countries, while imperfect, appear to deliver substantial improvements in welfare compared to universal programs, because they can transfer much more on a per-beneficiary basis to the poor as compared with universal programs. The primary downside of these programs is horizontal equity—because targeting is imperfect, there will be a substantial number of poor households who slip through cracks and are excluded. Nevertheless, for many developing countries, our simulations suggest the welfare gains from targeting may be substantial.

References

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi. 2016. "Self-Targeting: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 124(2): 371–427.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias. 2012. "Targeting the Poor: Evidence from a Field Experiment in Indonesia." *American Economic Review* 102(4): 1206–40.

Bah, Adama, Samuel Bazzi, Sudarno Sumarto, and Julia Tobias. 2018. "Finding the Poor vs. Measuring Their Poverty: Exploring the Drivers of Targeting Effectiveness in Indonesia." World Bank Policy Research Working Paper 8342.

Bah, Adama, Suahasil Nazara, and Elan Satriawan. 2015. "Indonesia's Single Registry for Social Protection Programmes." Research Brief 49, International Policy Centre for Inclusive Growth.

Baird, Sarah, Craig McIntosh, and Berk Özler. 2011. "Cash or Condition? Evidence from a Cash Transfer Experiment." *Quarterly Journal of Economics* 126(4): 1709–53.

Banerjee, Abhijit, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto. 2018. "Tangible Information and Citizen Empowerment: Identification Cards and Food Subsidy Programs in Indonesia." *Journal of Political Economy* 126(2): 451–91.

Banerjee, Abhijit, Rema Hanna, Gabriel E. Kreindler, and Benjamin A. Olken. 2017. "Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs Worldwide." *World Bank Research Observer* 32(2): 155–85.

Barham, Tania, Karen Macours, and John Maluccio. 2017. "Are Conditional Cash Transfers Fulfilling Their Promise? Schooling, Learning, and Earnings after 10 Years." CEPR Discussion Paper DP11937.

Bastagli, Francesca, Jessica Hagen-Zanker, Luke Harman, Valentina Barca, Georgina Sturge, and Tanja Schmidt. 2016. *Cash Transfers: What Does the Evidence Say? A Rigorous Review of Programme Impact and the Role of Design and Implementation Features.* With Luca Pellerano. Overseas Development Institute Report, July.

Behrman, Jere R., Susan W. Parker, and Petra E. Todd. 2011. "Do Conditional Cash Transfers for Schooling Generate Lasting Benefits?: A Five-Year Follow-up of PROGRESA/Oportunidades." *Journal of Human Resources* 46(1): 93–122.

Behrman, Jere R., and Petra E. Todd. 1999. "Randomness in the Experimental Samples of PROGRESA (Education, Health, and Nutrition Program)." IFPRI Research Report.

Besley, Timothy. 1990. "Means Testing versus Universal Provision in Poverty Alleviation Programmes." *Economica* 57(225): 119–29.

Cahyadi, Nur, Rema Hanna, Benjamin A. Olken, Rizal Adi Prima, Elan Satriawan, and Ekki Syamsulhakim. 2018. "Cumulative Impacts of Conditional Cash Transfer Programs: Experimental Evidence from Indonesia." NBER Working paper 24670.

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

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review* 102(6): 2981–3003.

Chetty, Raj, and Adam Looney. 2006. "Consumption Smoothing and the Welfare Consequences of Social Insurance in Developing Economies." *Journal of Public Economics* 90(12): 2351–56.

Currie, Janet M. 2006a. *The Invisible Safety Net: Protecting the Nation's Poor Children and Families.* Princeton University Press.

Currie Janet. 2006b. "The Take-up of Social Benefits." Chap. 3 in *Public Policy and the Income Distribution*, edited by Alan Auerbach, David Card, and John M. Quigley. New York: Russell Sage Foundation.

Eissa, Nada, and Jeffrey B. Liebman. 1996. "Labor Supply Response to the Earned Income Tax Credit." *Quarterly Journal of Economics* 111(2): 605–637.

Fiszbein, Ariel, and Norbert Schady. 2009. "Conditional Cash Transfers: Reducing Present and Future Poverty." With Francisco H. G. Ferreira, Margaret Grosh, Niall Keleher, Pedro Olinto, and Emmanuel Skoufias. World Bank Policy Research Report. World Bank.

Gelbach, Jonah B., and Lant Pritchett. 2002. "Is More for the Poor Less for the Poor? The Politics of Means-Tested Targeting." *Topics in Economic Analysis & Policy* 2(1).

Gertler, Paul. 2004. "Do Conditional Cash Transfers Improve Child Health? Evidence from PROGRESA's Control Randomized Experiment." *American Economic Review* 94(2): 336–41.

Glantz, Andrew E. 2008. "A Tax on Light and Air: Impact of the Window Duty on Tax Administration and Architecture, 1696–1851." *Penn History Review* 15(2): Article 3.

Golan, Jennifer, Terry Sicular, and Nithin Umapathi. 2017. "Unconditional Cash Transfers in China: Who Benefits from the Rural Minimum Living Standard Guarantee (Dibao) Program?" *World Development* 93: 316–36.

Gupta, Sarika. 2017. "Perils of the Paperwork: The Impact of Information and Application Assistance on Welfare Program Take-Up in India." Available at: https://scholar.harvard.edu/sarikagupta/publications/mediation-deficit-when-information-isn%E2%80%99t-enough-take-welfare-programs.

Hsieh, Chang-Tai, and Benjamin A. Olken. 2014. "The Missing 'Missing Middle.'" *Journal of Economic Perspectives* 28(3): 89–108.

International Labor Organization. No date. "ILO Social Security Inquiry." http://www.ilo.org/dyn/ilossi/ssimain.home?p_lang=en.

Jensen, Anders. 2016. "Employment Structure and the Rise of the Modern Tax System." Available at: https://www.semanticscholar.org/paper/Employment-Structure-and-the-Rise-of-the-Modern-Tax-Jensen/e7adac3d8699faca906b8f834be6b88cf870fc30.

Jones, Benjamin F., and Benjamin A. Olken. 2008. "The Anatomy of Start-Stop Growth." *Review of Economics and Statistics* 90(3): 582–87.

Klasen, Stephan, and Simon Lange. 2016. "How Narrowly Should Anti-poverty Programs Be Targeted? Simulation Evidence from Bolivia and Indonesia." Courant Research Center: Poverty, Equity and Growth—Discussion Papers, no. 213.

Kleven, Henrik Jacobsen, and Wojciech Kopczuk. 2011. "Transfer Program Complexity and the Take-Up of Social Benefits." *American Economic Journal: Economic Policy* 3(1): 54–90.

Kleven, Henrik Jacobsen, Claus Thrustrup Kreiner, and Emmanuel Saez. 2016. "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries." *Economica* 83(330): 219–46.

Kugler, Adriana D., and Ingrid Rojas. 2018. "Do CCTs Improve Employment and Earnings in the Very Long-Term? Evidence from Mexico." NBER Working Paper 24248.

McBride, Linden, and Austin Nichols. 2016. "Retooling Poverty Targeting Using Out-of-Sample Validation and Machine Learning." *World Bank Economic Review.* https://doi.org/10.1093/wber/lhw056.

Meyer, Bruce D., and Dan T. Rosenbaum. 2001. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." *Quarterly Journal of Economics* 116(3): 1063–1114.

Ministerio de Economía y Finanzas. 2008. *Marco Macroeconómico Multianual 2009–2011.* Lima, Peru: MFF.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38(2): 175–208.

Moffitt, Robert A. (ed.) 2003. *Means-Tested Transfer Programs in the United States.* University of Chicago Press.

Nichols, Albert L., and Richard J. Zeckhauser. 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review* 72(2): 372–77.

Oates, Wallace E., and Robert M. Schwab. 2015. "The Window Tax: A Case Study in Excess Burden." *Journal of Economic Perspectives* 29(1): 163–80.

Ravallion, Martin. 2009. "How Relevant is Targeting to the Success of an Antipoverty Program?" *World Bank Research Observer* 24(2): 205–231.

Robles, Marcos, Marcela G. Rubio, and Marco Stampini. 2015. "Have Cash Transfers Succeeded in Reaching the Poor in Latin America and the Caribbean?" Inter-American Development Bank Policy Brief no. IDB-PB-246.

Saez, Emmanuel. 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *Quarterly Journal of Economics* 117(3): 1039–1073.

Salehi-Isfahani, Djavad, and Mohammad H. Mostafavi-Dehzooei. 2017. "Cash Transfers and Labor Supply: Evidence from a Large-Scale Program in Iran." Economic Research Forum Working Paper 1090.

Stantcheva, Stefanie. 2014. "Optimal Income Taxation with Adverse Selection in the Labour Market." *Review of Economic Studies* 81(3): 1296–1329.

World Bank. 2012. "BLT Temporary Unconditional Cash Transfer." Social Assistance Program and Public Expenditure Review 2.

World Bank. 2014. *The State of Social Safety Nets 2014.* World Bank.

World Bank. Various years. "World Development Indicators."