NBER WORKING PAPER SERIES

SOCIAL PROTECTION IN THE DEVELOPING WORLD

Abhijit Banerjee
Rema Hanna
Benjamin A. Olken
Diana Sverdlin Lisker

Working Paper 32382
http://www.nber.org/papers/w32382

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024

We thank Amy Finkelstein, Nathan Hendren, the JEL editors Steven Durlauf and David Romer, and four anonymous referees for helpful comments. We thank Cristine von Dessauer, Tilman Graff, Emily Romano, Beatriz Velho and especially Kim Lan Mallon for outstanding research assistance. Financial support from the Wellspring Philanthropic Fund is gratefully acknowledged. This research was generously supported by the National Institute on Aging, Grant Number T32-AG000186 and is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. The views expressed herein are those of the authors and do not necessarily reflect the views of the individuals or organizations acknowledged here or of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w32382

Social Protection in the Developing World
Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Diana Sverdlin Lisker
NBER Working Paper No. 32382
April 2024
JEL No. I38,O15

ABSTRACT

Social protection programs have become increasingly widespread in low- and middle-income countries, with their own distinct characteristics to match the environments in which they are operating. This paper reviews the growing literature on the design and impact of these programs. We review how to identify potential beneficiaries given the large informal sector, the design and implementation of redistribution and income support programs, and the challenges and potential of social insurance. We use our frameworks as a guide for consolidating and organizing the existing literature, and also to highlight areas and questions for future research.

Abhijit Banerjee
Department of Economics, E52-540
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
banerjee@mit.edu

Benjamin A. Olken
Department of Economics, E52-542
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
bolken@mit.edu

Rema Hanna
Kennedy School of Government
Harvard University
79 JFK Street
Cambridge, MA 02138
and NBER
Rema_Hanna@hks.harvard.edu

Diana Sverdlin Lisker
The World Bank
1818 H St, NW
Washington, DC 20433
dsverdlinlisker@worldbank.org

A data appendix is available at http://www.nber.org/data-appendix/w32382

1 Introduction

Social assistance programs have been rapidly evolving in the developing world. Today, safety net programs cover about 2.5 billion people. More than 120 low- and middle-income countries run cash transfer programs for poor families, and more than 70 of them run social pension programs (World Bank, 2018). These programs have increasingly taken the place of broad-based subsidies of basic commodities, promising increased efficiency and an enhanced ability to redistribute income to the poor. The COVID-19 pandemic further expanded these policies, suggesting their increased role in policy in the future (Gentilini et al., 2020).

Designing these programs for a developing country context, however, entails challenges that differ from those faced in more developed economies. For example, one of the biggest questions in the public economics literature is how to target social assistance programs. But, unlike in high-income countries, information on income is hard to observe, partly because informal work and self-employment is common, resulting in many households being excluded from the income tax system (Jensen, 2022). This means that the standard frameworks used for thinking about targeting benefits—e.g., income eligibility thresholds, endogenous labor responses to effective marginal tax rates (e.g., Mirrlees, 1971; Saez, 2002)—do not readily apply in a developing country setting. This has led to the rise of other targeting approaches in developing countries, which in turn, has necessitated research designed specifically to understand these contexts (e.g., Alatas et al. (2012); Hanna and Karlan (2017)).

Similarly, program design may also need to be radically different in a developing country context. For example, one successful development intervention is to give households a productive livestock asset—chickens, a goat, a cow—and support them in taking care of it (e.g., Banerjee et al., 2015, 2016; Banerjee, Duflo and Sharma, 2021; Bandiera et al., 2017; Bedoya et al., 2019). This kind of program is based on the theory of poverty traps—that is, the idea that with non-convexities in production and savings constraints, a large lumpy transfer may push households into a higher steady-state income level. While poverty traps could occur in more developed settings as well, the particular nature of the interventions to relieve these shocks may differ: an intervention that provides productive livestock may make sense in low-income countries given that 63 percent of the population is engaged in agriculture (ILOSTAT, 2019), but would make less sense in a developed nation where the vast majority of the poor and near-poor are not in agriculture.

This paper aims to both outline theoretical frameworks for understanding the social protection challenges in developing countries, as well as provide a review of the literature to date. We focus both on redistributive programs, which seek to assist the poor, and social insurance programs, which aim to insure against shocks, while acknowledging that many programs could incorporate both features and that the two goals are often interrelated.[1]

The paper is organized as follows. Sections 2 and 3 focus on programs with primarily redistributive goals, i.e., programs that aim to assist the poor. In Section 2, we begin by outlining what we call the *'ex-post'* problem; that is, after types have been realized, how should the

[1] While active labor market programs, such as minimum wage policies or job retention programs, often fall under social protection programs in the policy space, we do not cover them in this review.

social planner think about improving welfare through redistribution? In particular, Section 2 asks the question of who should receive transfers and the degree to which transfers should be dependent on estimated or actual incomes for programs with re-distributive aims. We briefly discuss the standard Mirrlees-type model as a benchmark, but then mainly focus on what happens in contexts where the usual proxies for type in developed countries (such as income) are not easily observable. We briefly describe a theoretical framework for thinking though targeting issues in this context, and then review the empirical evidence on a wide range of approaches for determining who should benefit from these types of programs.

Once you decide who needs help, how should you provide such help? How do you know the help is "working?" In Section 3, we discuss both how to evaluate the impact of redistributive programs, and what we know about the impacts of the various types of programs. Specifically, we discuss how both the design and evaluation of programs may be influenced by frictions between the social planner's utility function and individuals' utility functions, when there is a cost of take-up, and when non-convexities mean that the differential timing of transfers may have different implications. We then review the existing empirical literature on some of the additional choices that one can make in designing these programs in light of this framework. We close out this section with a discussion of governance concerns, which affect how programs are in the end delivered, particularly in the context of low state capacity.

In Section 4, we then turn to an *'ex-ante'* framework—that is, before types are realized—that allows us to consider the value of social insurance. Again, we start from the standard frameworks and discuss the reasons private insurance markets may be incomplete, and then discuss how the problems and solutions may differ in countries with large informal sectors. We then review the evidence on the set of common programs that insure against expenditure loss and income loss. Section 5 concludes.

Importantly, in reviewing the literature, we seek to draw attention to important, but yet unanswered questions and new frontiers. To take just a few examples: the large informal sector in developing countries demands a different and perhaps more sophisticated approach to workplace accident and unemployment insurance schemes; how governments should modify their social assistance programs as income grows; and how the rise of digital technology and mobile payments will change options for both targeting and implementation. We elaborate on these and other topics throughout the review.

2 Targeting for redistribution once income has been realized

2.1 Targeting when income is observed: the classic public finance approach and its applicability to developing countries

We begin by thinking about social protection from the *'ex-post'* perspective; that is, conditional on the realizations of incomes, how to transfer resources from those with low marginal utility of consumption to those with high marginal utility. In this section, we start with what is often the key question in such models – how to identify who has high marginal utility and

would benefit from additional resources – and discuss how this problem differs in low- and middle-income countries from the way it has been conceptualized for high-income countries.

The benchmark public finance model assumes there is unobserved heterogeneity among taxpayers and the social planner would like to tax those with high ability to give transfers to those with lower ability, considering the incentives for high-ability earners to claim lower incomes in order to pay fewer taxes and get more transfers. This is complicated by the fact that the planner does not observe earning ability, and therefore bases decisions on realized earnings. This introduces a reason for high earners to claim lower earnings potential and undersupply effort (Mirrlees, 1971).[2]

As highlighted by Saez (2002), this approach identifies the elasticity of labor supply with respect to taxes and with respect to income as the key parameters in the design of welfare programs. There is a large literature on this point for developed countries (e.g. Saez, Slemrod and Giertz, 2012; Meghir and Phillips, 2010; Currie and Gahvari, 2008) motivated partly by the fact that many US welfare programs (especially prior to the 1996 welfare reform bill) implicitly put a high tax rate on earnings through a steep phasing-out of benefits.

However, the literature suggests that labor supply distortions might not be as relevant for developing countries. Two recent experimental studies of the wage elasticity, both from Malawi, report low values of short-run labor supply elasticity (Goldberg, 2016; McKenzie, 2017). Moreover, Banerjee et al. (2017) re-analyze data from seven different experimental trials of large-scale, conditional cash transfer programs where earnings were explicitly not part of the conditionality and find no evidence of income effects on work. Baird, McKenzie and Özler's (2018) systematic literature review came to a similar conclusion. The fact that both income effects and the wage elasticity are small suggests that it is less important to focus on the endogeneity of earnings in designing transfer schemes in developing countries compared to in high-income countries.

Governments in the developing world face a different challenge: they cannot observe earnings for informal workers and for those with relatively low earnings, which collectively comprise a large share of the population (Glewwe, 1992; Jensen, 2022). For workers in the formal sector – which encompasses most of the workforce in high-income countries – third-party reporting means that governments can obtain reliable information on people's incomes. But since most workers in low- and middle-income countries are informal, or below income tax reporting limits, this approach does not work. Many governments of these countries therefore worry that if program eligibility was based solely on income or employment status, the government would need to rely on self-reports, and many non-poor people would simply report a low income to claim benefits. The key issue in many low- and middle-income countries is therefore not how income-based eligibility rules distort actual labor supply, but rather the difficulty of observing income at all.

For this reason, instead of targeting transfers based on income or employment, as in high-income countries, many developing countries use a combination of three options, at least for those outside the formal sector: data-based approaches using easily observed characteristics

[2]We focus here on income as the key metric, following the public finance tradition of Mirrlees (1971). However, the government may have broader metrics of utility, including physical and mental health.

of households (e.g., location and asset ownership), self-selection, and community information. In the rest of this section, we discuss the theory and practice of these methods.

2.2 A model of targeting when income is imperfectly observed

What to do when governments cannot observe incomes? In a classic paper, Akerlof (1978) discusses how 'tags'—i.e., fixed individual characteristics—can be used to improve the targeting performance of self-selection schemes. The idea of using 'tags' forms the theoretical basis of what is known as proxy-means testing, in which assets and other observable characteristics are used to predict people's incomes for targeting. This approach forms the basis for targeting in many developing countries, as we discuss in more detail in Section 2.3 below.

These proxy-means tests allow the government to target the poor without income data. However, one challenge is mistargeting: targeting in such a system is based on the statistical model relating the tags to incomes. This model, like all statistical models, is imperfect, and the data on the tags is often noisy, leading to both inclusion errors(including those who are not poor) and exclusion errors (excluding those who are poor). This is exacerbated by the fact that the data for targeting is only updated every few years, due to concerns about collection costs (in contrast, most rich countries use frequently updated administrative data).

Governments therefore face a choice between using these noisy proxies, using potentially biased self-reports, or doing universal transfers and not trying to target at all (or some combination of these approaches).

In this section, we present a simple model to think through these tradeoffs. Our model accounts for the two features of developing countries: first, incomes and employment are hard to measure precisely since much work is in the informal sector (Jensen, 2022). Second, as discussed above, labor supply elasticities tend to be low. Unlike in traditional models such as Saez (2002), where the tradeoffs are primarily driven by the labor supply response, we will assume that labor supply is fixed and unaffected by redistribution. Instead, the key driver in our model is the government's ability to learn people's true incomes in a low-information environment, potentially with noise, when people can potentially misreport their incomes.[3]

2.2.1 The baseline model

The problem is one of how to best redistribute to a population of mostly low-income citizens who are outside the formal sector.

We assume that the government can identify the population that is broadly the target of this redistribution scheme in that it can exclude those with incomes above the poverty line y^*.[4] Let $h(y)$ represent the density of the true underlying income distribution of those who might be eligible for the program, so $\int_0^{y^*} h(y)\,dy = 1$.

[3]For related work exploring optimal targeting schemes, see also Kleven and Kopczuk (2011).

[4]This helps clarify that we are not designing the entire tax and transfer system here; we focus on the part that matters for the low income population. We also could consider a threshold $\hat{y} > y^*$, or even set $\hat{y} = \infty$ and consider the entire population, but that is not the object of interest here.

We presume B is funded by a combination of direct taxes on the formal sector and indirect taxes potentially on the entire population. If the indirect taxes were linear in the amount spent, nothing would change, so think of the incomes in our model as income net of indirect taxes. We also presume that the budget B is not so large that it could totally eliminate poverty with perfectly targeted transfers to everyone, i.e., $B < \int_0^{y^*} h(y)(y^* - y^p)\,dy$.

A key question that we will examine is what data the government should use as the basis for the targeted transfers, that is, what it should use as the 'predicted income' y^p. Define the 'predicted income' to be a linear combination of a person's self-reported income $\tilde{y}$ and the government's estimate of his or her income, y^e, where the superscript e refers to the government's income estimate from an 'audit':

$$y^p = \alpha\tilde{y} + (1-\alpha)\,y^e \tag{1}$$

The 'audit' here encompasses whatever procedure the government uses to observe an estimate of income, such as a proxy-means test (described in detail in Section 2.3 below) where the government proxies for income are based on assets it can observe. The estimate from the audit data is noisy, so the estimated income $y^e = y + \epsilon$, where ϵ has expectation zero and a fixed finite support. We study the choice of α, i.e., how much weight to put on self-reported income vs. the noisy audit-based estimate of income.

If people misreport their income, there is a cost $F = \frac{a}{2}(y^e - \tilde{y})^2$, based on the difference between audited income estimate y^e and self-reported income. The cost may be a probability of getting caught multiplied by some kind of a penalty. Here we assume that the penalty does not contribute to the program budget, so it could be a psychic cost or a bribe to a government official.[5]

We follow Saez (2002) and Saez and Stantcheva (2016) to introduce the government's preference for redistributing to the neediest by assuming that the government's maximand puts weight $g(y)$ on the after-transfer utility of someone with true income y, with $g'(y) < 0$. We adopt the normalization $\int_0^{y^*} h(y)\,g(y)\,dy = 1$.[6] In addition, we allow for the government to care about equity *among those who have the same true income* (that is, we allow for 'horizontal equity' considerations). We introduce this by assuming that the government's utility function for those who have a post-transfer income of z is $u(z)$, where u is a (potentially strictly) concave function. This formulation of the government's maximand nests the standard utilitarian social planner if the amounts involved in the redistribution are relatively

[5] We find similar results in the opposite extreme case where the penalty is financial, such as a fine, and the penalties help relax the government's budget constraint, as in Chetty (2009). This would also be the case if, for example, the fine was a bribe, as in the example in the text, but the salaries of the bureaucrats are adjusted downwards to take account of their income from bribes in the spirit of Becker and Stigler (1974), so in effect a bribe is just a payment to the government; see Appendix.

[6] As described by Saez and Stantcheva (2016), for considering small tax changes, the approach of using social welfare weights $g(y)$ nests the idea of using a utilitarian social welfare function, since one can set $g(y)$ equal to the marginal utility of income at y. However, they argue for this approach as it is more general. It also allows us to differentiate between a preference for redistribution (i.e., people with lower $g(y)$) and a preference for horizontal equity, which turns out to be important, as we describe below.

small but is much more general, allowing for a range of political economy considerations in the government's maximand.

One important way that we differ from traditional analyses (such as Feldstein (1999)) is that, in our baseline model, we assume that the government's maximand does not take into account the beneficiary's disutility from misrepresenting their income (i.e., does not include $F = \frac{a}{2}(y^e - \tilde{y})^2$).[7] This simplifies our analysis considerably and perhaps fits the context better as well. For example, if the penalty takes the form of a bribe to the government official who discovered the discrepancy, as it often may, it is not unreasonable to assume that from the point of view of the government it nets out to zero.[8]

The formal analysis of the model is in the Online Appendix. The first result applies to the baseline model where we assume that everyone has the same cost of misreporting their income and that the government cares about horizontal equity (u is strictly concave).

Result 1 *In the baseline model with homogeneous preferences over the penalty for misreporting, the optimal redistributive scheme always has the maximum feasible slope with respect to earnings given the budget constraint and has no universal component* ($T = 0$)*. The optimal scheme relies entirely on self-reports* ($\alpha = 1$)*. Moreover, this scheme delivers the first-best level of targeting.*

This stark result has a simple intuition. As long as all potential beneficiaries have the same cost of misreporting incomes, the amount by which they distort is predictable, and the optimal targeting mechanism can take that fully into account, and hence can undo any misreporting that exists. The government can thus avoid using the noisy audit estimates because it can perfectly infer people's real incomes from their (uniformly) misreported self-reports. Given that the government can perfectly back out incomes, it is optimal to use targeted transfers rather than make sub-optimal transfers to richer households.

This stark result also implies that *heterogeneity* in the cost of misreporting, i.e., heterogeneity in a, is key to understanding why governments may want to use audit-based estimates, why there are inclusion and exclusion errors in the targeting process, and why there may be reasons to limit the extent to which transfers are income-dependent, because heterogeneity in a means that the government can no longer perfectly back out true incomes from distorted self-reports. We explore this next.

2.2.2 Introducing heterogeneity in the cost of misreporting

Suppose now that a takes two values, a_1 and a_2, with $a_1 < a_2$. Let μ denote the fraction of the population that are a_1 types. The cost of misreporting a is unobserved by the government. We assume that both types have the same income distribution $h(y)$. Heterogeneity in the cost of prevarication introduces an important additional consideration: transfer schemes that rely on self-reports, in effect, redistribute from those with high values of a (those who do

[7] This is related to the discussion in Chetty (2009) on socials costs of taxation versus transfers.

[8] Even if the cost is primarily psychic, it seems unreasonable to assume that the government is very concerned about its population being forced to lie about its income. Alternatively, we explore what happens when misreporting costs enter the planner's welfare function in Appendix A.5.

not misreport incomes very much) to those with low values of a (those who do). That is, the government can no longer fully unravel misreports, which helps those who are willing to misreport more.

This model therefore features a tradeoff. On the one hand, as the government relies more on the unbiased but noisy audit-based estimates and less on self-reports (i.e., as we lower α), the redistribution from a_2 types to a_1 types falls. On the other hand, as before, the more the government relies on audit-based estimates, the greater the social welfare loss induced by the noise (i.e., the ϵ's) in the audit process. The net effect of lowering α depends on the relative sizes of these effects. In the limit, holding other parameters fixed, as the audit process becomes better and better ($Var(\epsilon) \to 0$), the government will rely entirely on the audit-based estimates (i.e., $\alpha = 0$); conversely, as the two types become more similar, the government will rely entirely on self-reports (i.e., $\alpha = 1$).

When both forces are present—i.e., when $Var(\epsilon) > 0$ but the two types are sufficiently different—it can be optimal for the government to use the audit-based estimates (i.e., it can be the case that $\alpha < 1$, unlike before), but it will always be optimal to rely at least a bit on self-reports, so $0 < \alpha < 1$. This is because the welfare loss from redistribution from a_2 types to a_1 types depends on α^2 and therefore, for α small enough, raising α has a second order negative welfare effect through this channel but a first order positive welfare effect because it reduces the impact of the noise (ϵ). Hence, in this model, the government will not completely ignore self-reports, though it can clearly now be optimal to combine self-reports with audit-based estimates.[9]

Turning next to t, as before, a higher t still means more transfers go to the poor. However, unlike in the previous model, there are now two forces pushing in the opposite direction. First, the social losses from redistributing from a_2 to a_1 types are greater as t is higher. Second, since we no longer have $\alpha = 1$, there are also losses due to targeting noise ϵ, and these are also stronger when t is higher. Hence, relative to the case where there is no heterogeneity, the optimal value of t may be lower.[10] That is, there is less redistribution when there is more heterogeneity in the ability to misreport income.

Summing up, we have the following result:

Result 2 *If there is heterogeneity in the cost of misrepresenting incomes, it can be optimal for the government to use noisy audit-based estimates as part of the targeting process. The weight placed on audit-based estimates relative to self-reports is increasing in heterogeneity in misreporting costs and decreasing in noise in the audit process. Moreover, it can now be optimal for the benefits to have a universal component, and not be entirely income-based.*

A numerical example showing how the optimal choices of α and t change as we change the extent of heterogeneity and $Var(\epsilon)$ can be found in Online Appendix A.4.

[9] In practice, if there is a fixed cost of collecting the self-report data, and if the optimal α is very small absent these fixed costs, the government may not bother paying the fixed costs and may choose to ignore self-reports altogether.

[10] Of course, we may still be at the maximum value of t, but it is now possible that the optimal solution will feature some lump-sum component T and less than the maximum level of t.

We have so far assumed that those who have a lower cost of misreporting have the exact same income distribution as those with a higher cost. In Online Appendix A.3, we also examine the case where heterogeneity in misreporting costs is correlated with income. Doing so yields the additional insight that, in this case, we can drop the assumption that the government cares about horizontal equity (i.e., that u is strictly concave) and still get the result that the government relies at least partially on audit data to deal with the possibility that those with higher income have a lower misreporting cost. Indeed, we show that with linear u, if there are two types, with the cost of misreporting negatively correlated with income, it will be optimal to completely ignore self-reports and set $\alpha = 0$.

2.2.3 Summary

Our model, though stylized, helps highlight the tradeoffs in targeting transfers in the presence of noisy data. We highlight that the key challenge is heterogeneity in misreporting costs, which prevents the government from 'backing out' people's true incomes from their self-reports. Given this, even the optimal transfer scheme will typically involve both inclusion and exclusion errors. The fact that the government responds to the imperfections in self-reported data by using noisy audit-based estimates of income generates additional targeting errors.

We also see why it may be optimal to rely at least partially on self-reports even when they are biased, but also why if the beneficiaries are heterogeneous in terms of the cost of misreporting, self-reported data may be down-weighted or even ignored in favor of the data the government collects, even if that data is noisy. The model also suggests reasons why even a redistribution-minded government may opt for some universal component in social transfers, because of a combination of the concern for equity (concave u), heterogeneity in misreporting costs, and noise in the audit process ($Var(\epsilon) > 0$).

2.3 Targeting on Observable Characteristics in Practice

2.3.1 Geographic targeting

The simplest version of characteristic-based targeting is **geographic targeting**, i.e., targeting programs to poorer regions. Its main attraction is its simplicity and the administrative ease of implementing it, since it does not require differentiating between individuals within a region. The 'characteristic' used for targeting is just the region where someone lives.

The main question involved in geographic targeting is how to choose the geographies in a data-driven way (Baker and Grosh, 1994). If one has representative survey data at a particular administrative level, one can use that dataset to estimate average incomes at that administrative level, and target based on that. To use finer geographic resolution, one can combine survey sample data with census data to predict the poverty status of smaller geographical areas (Elbers, Lanjouw and Lanjouw, 2003; Elbers et al., 2007), noting the statistical caveats raised by Tarozzi and Deaton (2009), or by using newer remote-sensed satellite or administration datasets that help predict smaller regions (e.g., Jean et al., 2016; Blumenstock, Cadamuro and On, 2015). These geographic-based approaches typically

produce less accurate estimates than individual targeting (since they miss all the poor people not in chosen regions), but are much less data intensive and are administratively easier to implement.

2.3.2 Proxy-means tests

Characteristic-based targeting at the individual/household level is typically referred to as a **proxy-means test** (PMT). The implementation of a PMT is based on two datasets. One starts with a training dataset that includes measures of what the government is actually trying to target (e.g., per-capita household consumption, income, etc.), as well as the characteristics—usually demographics and assets—that will be used for targeting. The idea is to select a set of observable characteristics ($X^{'}$) which, unlike income, the government can actually collect. This training dataset is usually a nationally representative household survey that is collected for research or general statistical purposes; as such, households typically have no reason to lie about their actual consumption or income. Using these data, one then estimates a model of the form $y_i = X^{'}\beta + \epsilon_i$.

To identify beneficiaries, one then needs a much larger dataset of the X_i on the entire relevant population, i.e., a census of the entire country or a 'social register' that contains information from anyone who may plausibly be eligible. This step may also include some self-selection. For example, in some cases, the government will collect the characteristics X_i only for those who apply, and hence only applicants may be deemed eligible. In other cases, the government attempts to collect the X_i for almost everyone by having census enumerators go door-to-door throughout the country. Using this much larger dataset, one calculates a predicted poverty score, $\hat{y}_i$, using the characteristics X_i and the estimates of $\hat{\beta}$, for each individual i that is then used for targeting.

This type of proxy-means testing was first used in the 1980s in Chile for the targeting of its Ficha CAS program, and is now quite common (Coady, Grosh and Hoddinott, 2004*a*). Well-known examples of proxy-means tests include the SISBEN in Colombia, SISFOH in Peru, Listahanan in the Philippines, the NSER in Pakistan and Takaful and Karama in Egypt.

The way proxy-means tests are implemented connects to the challenges highlighted by the model in section 2.2.1. First, the predicted poverty scores, which correspond to the audit-based estimates in the model, are imperfect predictors of poverty. To illustrate this, Hanna and Olken (2018) simulate this type of targeting for two countries—Indonesia and Peru—using household survey data and a set of assets and household characteristics that are typically used in proxy-means tests. They show that the prediction equation has substantial predictive power—the R^2 is between 0.53 and 0.66—but it nonetheless has errors of both inclusion and exclusion. Other papers show similar findings (see, e.g., Brown, Ravallion and van de Walle (2018)).

There is also no reason to think that the relationship between underlying income and the types of assets used in the PMT algorithm is the same for everyone. The model used to predict PMT scores makes no allowance for differences in preferences: someone might end up in a bigger house than what would normally correspond with their income because, for example, they feel obligated to provide temporary housing to their extended family. As in

the model, such heterogeneity in the link between true income and PMT-predicted income is a major challenge in targeting algorithms.

Recent advances have sought to reduce noise in PMTs in two ways. First, one can replace the OLS prediction equation $y_i = X'\beta + \epsilon_i$ with more sophisticated machine-learning prediction algorithms. However, most papers find little to no improvements using these techniques relative to the common methods currently used (e.g., Mcbride and Nichols, 2018; Baez, Kshirsagar and Skoufias, 2019; Areias and Wai-Poi, 2022).

Second, one can use new types of administrative data rather than collecting a door-to-door census, which allows more frequent, and hence up-to-date, data. For example, Abelson, Varshney and Sun (2014) use remote sensing to differentiate between thatched and metal roofs in Kenya's GiveDirectly program. Using data from Togo, Aiken et al. (2022) find that using mobile phone metadata for targeting has substantive predictive power, though it is not as accurate as a more traditional PMT. The fact that one can use administrative data such as these, which is continually and automatically updated, suggests that targeting using these data may be more responsive to shocks.[11] In fact, some countries (e.g., Pakistan and Togo) used this approach to target assistance during the 2020 COVID-19 crisis (Gentilini et al., 2020). Understanding whether this method was effective is an important area of ongoing work.

2.3.3 Endogenous characteristics?

A further challenge is that many characteristics used in this type of targeting are, themselves, choice variables. This introduces an additional potential source of inefficiency, since households could distort their consumption to remain eligible. For example, after England imposed a "window tax" in 1696—a tax levied on the number of windows a house had, which proxied for wealth—people built new houses with very few windows, leading to many dark houses (Oates and Schwab, 2015). Even geography is not an immutable characteristic: households could choose where to live in response to geographic targeting incentives. In practice, are these distortions substantial in modern contexts?

To investigate these issues with individual-level targeting in a modern proxy-means test, Banerjee et al. (2020*b*) conduct a nationwide randomized experiment, in cooperation with the Indonesian Government's Central Bureau of Statistics, which administers the census used for actual PMT targeting. To test for endogenous distortions, the statistics bureau randomly added questions on flat-screen televisions and/or the number of cell phone SIM cards owned to the targeting census in some randomly-selected provinces, but not in others. While self-reports of television ownership fell in treated provinces six months later, the effect quickly died out, and more importantly, there were no changes in actual television purchases or active SIM cards in treated areas.

In the framework of the model in Section 2.2, the fact that people do not seem to change their consumption decisions in response to the PMT questionnaire may correspond to the case where the cost of distorting choices (what we call a in the model) is very large. But this is

[11] Others (e.g., Baez, Kshirsagar and Skoufias, 2019; del Ninno and Mills, 2015) suggest augmenting PMT targeting system with remote sensed and/or administrative data to help capture shocks.

not the only interpretation. First, recall that in the model, a referred to the cost of changing one's reported income, not one's real income. Along these lines, one interpretation is that reported consumption can change even if actual consumption does not (e.g., people can hide the television when the enumerator comes around). The Banerjee et al. (2020*b*) result had hints of this: reported ownership of televisions declined in government data six months later, even though actual television ownership was unchanged. Similarly, Martinelli and Parker (2009) examine self-reported versus enumerator-verified PMT data in Mexico, and find more under-reporting of assets in the self-reported data compared to the enumerator-verified data when the monetary stakes are higher.

Beyond this, there is a second important consideration. In the model, we assume that the household knows exactly what to do to claim the benefits. In practice, the fact that there are many different variables, interacting in complex ways in the PMT formula, means that households have limited control over the outcome and therefore low returns to manipulation. In fact, governments typically keep the formulas secret for this reason: studying the proxy-means test in Colombia, Camacho and Conover (2011) find that local politicians were only able to manipulate the data after the formula became public in order to sign up more of their constituents for federal government transfers.[12]

Finally, in contrast to PMTs, there is some evidence that geographic targeting can lead to real distortions. In particular, place-based policies can induce migration to areas with generous transfer programs and/or prevent out-migration from these areas. For example, Imbert and Papp (2020*a*) show that India's public works program, the Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA)—which was geographically targeted to rural areas, and is discussed more in depth in 2.4—reduced seasonal migration to urban areas.

2.4 Self-selection mechanisms

Another targeting approach is **self-selection**. The idea is to pair benefits with an action, where the required action is relatively less costly for the targeted population (the poor) than for the rest (the non-poor) (Nichols and Zeckhauser, 1982). If so, then as in the model, the selection rule allows the government to make an inference about people's income (perhaps noisily) from their self-chosen actions. The required action can include buying low-quality food, work requirements, and administrative ordeals (e.g., standing in line, filling out lengthy paperwork, or going through an interview process). However, this approach imposes costs on the poor—e.g., lower quality food than they would prefer or time wasted standing in pointless lines. A key question is whether the targeting benefits, which would allow governments to deliver more aid to the poor, outweigh these additional costs, so that the poor individuals themselves can nevertheless be better off even once these costs are taken into consideration.

Workfare. One of the most common ordeal mechanisms is workfare, i.e., requiring beneficiaries to work in exchange for payments. These programs have a long history, dating back at least to the 19$^{\text{th}}$ century in England, where transfers were granted through residence in a

[12]Recent theoretical advances suggest ways to improve targeting by anticipating problems and taking the potential responses of agents into account (Björkegren, Blumenstock and Knight, 2021; Ball, 2023).

workhouse (Besley and Coate, 1992), and to the United States' Civilian Conservation Corps and Works Progress Administration in the 1930s (Aizer et al., 2020).

Workfare programs are a classic example of an ordeal mechanism because work takes time. If the wages for the workfare job are low, or the tasks are unpleasant or difficult, this will generate self-selection: those who can get a more attractive job will select out (Ravallion, 1991; Besley and Coate, 1992). However, if the wage is set too high, the program runs the risk of crowding out more productive private sector work. The net efficiency consequences also depend on whether the work being done is productive. Using workfare labor to build roads may have little efficiency loss; using workfare labor to dig ditches no one needs would also generate self-selection, but at much larger social efficiency costs.

Perhaps the largest workfare program in the world is India's MGNREGA, which offers 100 days of paid employment per year to anyone in rural areas who is willing to do casual manual labor. Perhaps unsurprisingly, given the hard work the program entails, Dutta et al. (2014) find substantial self-selection in program take-up.

Given that workfare can crowd out other work, estimating the efficiency costs are challenging. Murgai, Ravallion and van de Walle (2016) investigate this in the Indian state of Bihar by asking workfare participants their best estimate of what their earnings would have been in the absence of workfare. They find that those who joined MGNREGA gave up income equivalent to about 30-35 percent of the workfare income received. Similarly, Bertrand et al. (2021) experimentally examine a youth workfare program in urban Cote d'Ivoire. They find that while the program led to earnings gains, earning gains are only 53 percent of the transfer due to the crowding-out of private employment. Both programs set the workfare wage at the formal minimum wage, which—unlike in many developed country settings—may be *above* the prevailing wage, given the large informal sector. This suggests that workfare programs may have more effective targeting properties in places with a lower formal minimum wage and with less informal employment. We further discuss workfare programs in Section 3.2.4.

Self-selection with small costs. One challenge with ordeal mechanisms such as workfare is that they impose very large costs on beneficiaries—e.g., hard labor all day under the hot sun. Is it possible to get the gains from self-selection without imposing such large costs?

One possible selection mechanism is to ask people to apply for a program, rather than to be automatically enrolled. Upon application, programs then often add additional screening mechanisms, such as a PMT. This can affect selection in two ways. First, to the extent that prospective beneficiaries understand the screening mechanism, those who know they will not be eligible will not bother to apply. This saves the government the cost of verifying these applicants. Second, beneficiaries will use their own estimate of their likelihood of passing the screening when deciding whether to apply. This estimate will be a mix of the information used in the screening mechanisms, *plus* their private information about their true income level. Thus, if people have imperfect information or beliefs about the screening mechanism, asking them to decide to apply may induce them to reveal some of this private information to the government.

Alatas et al. (2016*b*) develop a model that captures this idea, and then experimentally test a self-targeting mechanism (application with verification done by PMT) against automatic

Figure 1: Probability of receiving benefits for those in automatic enrollment vs. self-selection villages, based on Alatas et al. (2016*b*)

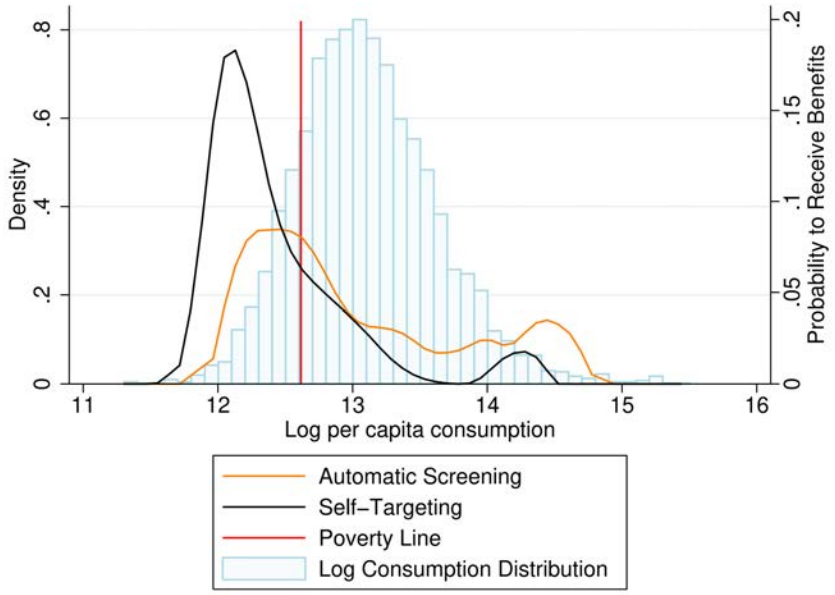

*Notes: The graph shows the probabilities of receiving benefits in both the automatic enrollment and self-targeting treatments described in Alatas et al. (2016*b*) (right axis), along with a histogram of log per-capita consumption (x-axis). The red vertical line represents the poverty line.*

enrollment based on a PMT done by census enumerators within the context of the expansion of the government's real conditional cash transfer program. They found that self-targeting led to dramatically poorer beneficiaries, driven in large part by those who are wealthier on dimensions that would be missed by the PMT being less likely to apply. Surprisingly, this approach *also* reduces exclusion error compared to the PMT – while the government makes its best efforts in the automatic enrollment system, some of the very poor, who live on the margins of society, can be missed; in the self-targeting group, these individuals apply (see Figure 1). This suggests that self-selection at the application stage can reveal important information, potentially decreasing both inclusion and exclusion errors.

2.5 Community-information based targeting

An alternative approach seeks to leverage local information about people's poverty status by allowing communities to participate in choosing beneficiaries (Alderman, 2002; Galasso and Ravallion, 2005). But, if governments try to use this information for **community-based targeting**, is the outcome better than a proxy-means test?

Alatas et al. (2012) experimentally compare proxy-means test-based targeting to community-based targeting. In their setting, community meetings were held in each sub-village, and households were ranked based on their relative poverty levels. Those ranked as among the poorest households were given transfers.

The results suggest that the optimal approach may depend on the government's objective function, and in particular, on whether per-capita consumption or a different, perhaps more locally-based, notion of poverty is the 'correct' measure of neediness. On the one hand, the authors find that the proxy-means test did better at identifying households based on per-capita consumption—the proxy-means test had about a 10 percent lower error rate in terms of identifying households below PPP $2/day, although the households whose identification switched were sufficiently close to the poverty line that both approaches would perform similarly in terms of most social welfare functions. On the other hand, community targeting did better in terms of identifying households who self-identify as poor. As a result, citizens in villages randomized to community targeting reported higher satisfaction with both the targeting process and outcomes than in those with PMT-based targeting.[13]

These results suggest that community-based targeting can work. However, there are often concerns about elite capture in community-based targeting. That is, local leaders might disproportionately choose themselves, family members, or others who are not necessarily the most in need. There is therefore a potential theoretical tradeoff between the superior local information and greater capture of these programs by local elites (e.g., Bardhan and Mookherjee, 2000).

Empirically, however, elite capture seems to be small, at least in comparison to other sources of targeting failures. In a field experiment in Indonesia, Alatas et al. (2019) find that increasing elite decision-making over targeting does not substantially change outcomes. More generally, while they do find some evidence of elite preference in some programs, it is small in magnitude, so that the gains one would achieve by reducing elite capture are entirely dwarfed by the gains that one could make by improving the quality of PMT data.[14] Similarly, Basurto, Dupas and Robinson (2020) study the targeting of subsidies for both agricultural inputs and food in Malawi, which is done by local chiefs, and compare it to a hypothetical targeting done through a PMT. While they find that nepotism exists, the welfare consequences are small, again because those who receive preferential treatment are also relatively poor; moreover, there are potential gains as chiefs use their local knowledge to target.

While these studies suggest that community-based targeting has potential, especially when assessed against local objective functions (such as self-assessed welfare), or to fill in the gaps in the PMT (e.g., detecting shocks, finding productive households), it does not strictly dominate PMTs, and it may not be appropriate in all settings. Indeed, several studies find results that confirm that PMTs tend to do at least a somewhat better job at identifying households with lower objective per-capita consumption, and there is some heterogeneity in the extent to which community-based targeting does better on other metrics, such as self-assessed welfare, which in itself may capture broader metrics of well-being relative to consumption.

Premand and Schnitzer (2021), for example, ran an experiment in Niger randomizing villages

[13]The gain in legitimacy from community targeting may be important. For example, Cameron and Shah (2014) show that mistargeting in Indonesia's cash transfer programs increased crime.

[14]The welfare losses from elite capture are essentially the product of a) how much richer elites are (9 percent), b) how much more likely elites are to obtain benefits (19 percent in the worst program considered), and c) elites' proportion in the general population. Multiplying yields a welfare cost less than 1 percent.

about to receive a national cash transfer to be targeted using a PMT, a community-based approach, or a formula that aims for food-insecure households. They find that households selected by the PMT are 8 percentage points more likely to be poor than those selected by the community. Like Alatas et al. (2012), they also find that community-based targeting does a better job matching self-assessed welfare status and the perceptions of others (though the difference is not large) and identifying households with recent shocks.

Several other studies (Stoeffler, Mills and del Ninno (2016), Beaman et al. (2021), Dupas, Fafchamps and Houeix (2022)) also find that community-based targeting does worse than proxy-means tests when the desired outcome is per-capita consumption, though these studies do not typically assess their comparative effectiveness vis-à-vis self-assessed welfare. Other studies suggest that the efficacy of community-based targeting compared to a PMT depends on the variables included in the PMT (e.g., Sabates-Wheeler, Hurrell and Devereux (2015)).

These papers suggest directions for future research. First, are there ways to make community selection more effective? For example, Alatas et al. (2012) find that households that were randomized to be ranked first are ranked more accurately than those ranked later, suggesting people get tired. This suggests that the micro-structure of the algorithm used to elicit the community ranking can affect the ranking's accuracy, and understanding this better can help improve the information elicitation process. Second, does the efficacy of community targeting differ systematically? For example, Alatas et al. (2016*a*) find that the community-based targeting treatment in Alatas et al. (2012) was more effective at identifying households within areas with more diffusive inter-personal network structures. More generally, being able to predict in which types of locations community-based targeting is likely to be effective is an important direction. A third question is whether community-based targeting can be compatible with cross-community targeting. That is, one can allocate a fixed number of slots to a community (for example, based on geographic targeting, a past census, or another PMT approach), and ask them to choose *which* households should receive those slots. But if a community is asked to identify who is poor on an absolute level, it has little incentive not to list many people as poor in order to maximize government transfers. Designing incentive-compatible procedures for community-based targeting that reveal the level of poverty, not just who is the poorest, seems like an important mechanism design problem for future research.

2.6 Incomplete take-up, and what we can learn from it

The flip side of self-targeting—through which take-up decisions can help improve targeting —is that many people do not take up the program, even when they are entitled to it. How should we think about this? Does this reflect a rational cost-benefit calculus, information, or other optimization frictions, behavioral constraints (e.g., procrastination), stigma, or something else? Is this optimal?

2.6.1 How prevalent are take-up challenges?

The fact that not everyone takes up programs is not, per se, surprising. Self-selection costs vary across people. For example, finding appropriate childcare in order to spend an afternoon

signing up for a program might be more difficult for some. Moreover, not everyone who we think is poor is actually poor—as described above, the data is hardly perfect.

However, the magnitude of incomplete take-up in social protection programs suggests that the problem goes beyond optimal self-targeting. For example, (Bhattacharya et al., 2015) report that two-thirds of women eligible for a non-contributory social pension for life are not enrolled in the program. The generosity of the pension (the equivalent of 70% of median per-capita consumption across urban areas) makes it unlikely that the low take-up could be rationalized by a cost-benefit calculation.

2.6.2 Application costs and information barriers

The fact that so many of the poor miss out on benefits suggests that some non-take-up might be inefficient. Finkelstein and Notowidigdo (2019) argue that in the US, misperception of the benefits and costs might explain part of the problem. They find that benefit information nearly doubles take-up (from a low base); adding sign-up assistance triples take-up.[15]

There are a number of related experiments in the developing world that examine how much of the take-up challenge can be attributed to the (perceived) hassles of applying to programs by providing enhanced take-up assistance to targeted households. These studies typically find that while take-up assistance can meaningfully increase program take-up compared to not having assistance, it does not close the majority of the take-up gap.

For example, Carneiro, Galasso and Ginja (2019) study a program known as Chile Solidario that targeted extremely poor households with extensive home visits by a social worker over a two-year period, including guidance on how to apply for government programs, as well as a financial stipend. The program itself had imperfect take-up; only about 20 percent of those eligible for Chile Solidario enrolled. For those who availed themselves of Chile Solidario assistance, take-up of a separate government family child allowance program (SUF) had increased by 17 percentage points, from a base of 53 percent, 4 years later. The intensive take-up assistance of Chile Solidario mattered, but it only closed about one-third of the take-up gap.

In a second example, Gupta (2017) looks at widows and divorcees in Delhi, India, and investigates why they fail to enroll in the pension programs. Applying is a substantial hassle: widows must fill out forms, provide supporting documentation (e.g., proof of marital status, unique ID card, proof of 5-year residency), get a local politician's signature, and submit all this to the local district government. Intensive assistance also helped in this case, but it did not substantially close the take-up gap: 22 percent of eligible women receiving intensive mediation successfully obtained the pension, compared to 15 percent of control households. While this is a 47 percent increase, it still left behind 80 percent of eligible women.

Banerjee et al. (2021) examine related questions in Indonesia in the context of national health

[15] (Finkelstein and Notowidigdo, 2019) also note that improving take-up on the margin does not automatically improve welfare, if households are fully optimizing and rationally not signing up for programs given the costs of doing so. However, this result depends on the fact that households are fully optimizing, which is often not the case in many contexts.

insurance. Universal health insurance was launched in 2014 and by 2015, the contributory portion of the program had enrolled less than 20 percent of the intended population despite being officially mandatory. Banerjee et al. (2021) find that while application assistance increased *attempted* enrollments substantially—by 16 percentage points, compared to just 1.8 percent in the control group—the vast majority of those who attempted to enroll failed. As a result, providing assistance only increased enrollment by 2.4 percentage points.

A common theme among these papers is that the take-up problem is not easily solved by helping households apply. Instead, the fact that there is any paperwork may itself be an important barrier. This suggests that programs may want to radically simplify application requirements, such as by pre-filling forms using administrative data. However, this only works if the underlying state systems are accurate. Many of the problems Indonesian households encountered in Banerjee et al. (2021) came about because the government's underlying family records were inaccurate, a problem that could only be fixed by a costly trip to a government office. Improving the underlying administrative capacity to allow for simpler enrollment procedures may be challenging, but important.

2.6.3 Stigma

A common refrain in the developed world is that low take-up reflects low demand for government services due to the stigma associated with receiving benefits (Moffitt, 1983). While there is comparatively little research on stigma in developing countries, what exists suggests less concern about stigma. For example, Osman and Speer (2024) conduct experiments in Egypt to recruit young people to labor market assistance programs, varying the recruitment messages to distinguish the effects of different types of stigma. Several messages designed to reduce 'professional' or 'social' stigma actually seem to have primed respondents to be more alert to these concerns. However, an explicit 'welfare stigma' framing—i.e., that the cost of the job training program was subsidized "to help those in financial hardship"—had no effect whatsoever on program take-up. Nevertheless, understanding the role of stigma in more detail—and in particular whether or not program participation is public information—is an important direction for future research.

2.7 Targeting on treatment effects, not levels

The discussion thus far has focused primarily on providing programs to the poor. Implicit in this view is that anti-poverty programs are more effective at increasing private or social utility when the household is poorer. In our model, this was the result of the assumptions that $g'(y) < 0$ and/or that $u(y)$ is concave.

Targeting is more complicated, however, when there is heterogeneity in program treatment effects. Imagine that a program's only effect is to generate income, but that the effect on income is heterogeneous by individuals, i.e., we write dy_i. Now, optimal targeting would want to find those individuals for whom $g(y_i)u^{'}(y_i)\,dy_i$ is highest, which in this particular formulation is a combination of the individual's specific treatment effect and income. Even for programs that do not target income, the point is that one is no longer interested in targeting the program just based on an observable characteristic (i.e., poverty score), but

also on the program's treatment effect, dy_i.

This is a harder problem, because generally speaking dy_i is unobserved. However, there are still several options. One option is to predict dy_i based on observable covariates, for example, from a randomized trial. Bhattacharya and Dupas (2012) consider this approach in the context of subsidies for anti-malarial bednets, where they account for the probability of bednet use. Other recent work shows how to use machine learning to estimate heterogeneity in treatment effects (e.g., Wager and Athey, 2018; Chernozhukov et al., 2018). Several papers use these techniques to show that they can identify heterogeneous treatment effects to improve the targeting of immunizations (Chernozhukov et al., 2018), workfare (Bertrand et al., 2021), and cash transfers (Haushofer et al., 2022). Caria et al. (2023) combine this approach with a dynamic experimental design to both simultaneously learn which treatments work best optimally, and which work best for which types of respondents, for job training programs for Syrian refugees in Jordan.

A second option is to encourage self-selection on the basis of treatment effects. Dupas et al. (2016), for example, investigate a program that provides free chlorine in Kenya. They compare the free distribution of chlorine versus the distribution of vouchers redeemable at nearby shops for free chlorine. Redeeming a voucher is a screening mechanism as only those who intend to use the chlorine will redeem it. Indeed, they find similar rates of chlorine in water in the two groups, yet the vouchers saved 60 percent of the cost by not distributing chlorine to those who would not use it.

These examples illustrate the potential to apply these approaches to social protection more broadly. For example, the graduation approach (discussed in Section 3.2.6 below) may be appropriate for some types of households and not others. Can one identify them based on ex-ante characteristics? Or find a way for people to self-select in? On the flip side, for social insurance products, one does not want to encourage either adverse selection or so-called 'selection on moral hazard' (Einav et al., 2013). For example, making enrollment to health insurance easier could potentially encourage the healthier to enroll (Banerjee et al., 2021).

There are two caveats worth noting. First, for this approach to work, it is important to be able to identify who could benefit the most, and that may not be feasible in all settings. For example, Carter, Tjernström and Toledo (2019) study the Rural Business Development Program in Nicaragua and show that observable characteristics do not do a good job at predicting which households would benefit the most. Second, we need to think carefully about the ethics of targeting on productivity. For example, there is a concern that targeting on work productivity for a cash transfer program could screen out people who really need help but may not be able to work, such as people who have severe disabilities. More work needs to be done to understand these issues.

2.8 Should we even target at all? Universal vs. targeted programs

All of the targeting methods discussed thus far suffer from exclusion and inclusion errors (Brown, Ravallion and van de Walle, 2018). Given these challenges, as described in Section 2.2, another option is to make programs universal, such as through a universal basic income (UBI) program or a universal in-kind program. These programs involve giving each indi-

vidual a fixed amount of money or goods—regardless of income—and are financed through proportional or progressive taxation. Conceptually, one can think of them as shifting the intercept of the tax schedule and modifying the rates (i.e., the slope of the tax schedule) in such a way that it satisfies the government's budget constraint (Saez, 2002; Hanna and Olken, 2018). These programs could literally be available to everyone, as in a universal basic income, or they could be available to everyone in a particular category (e.g., the elderly, children), but the key distinction is that they are not conditioned on income or proxies for income.

Universal transfers have several advantages. In principle, they have no exclusion error, though in practice take-up may be far from universal. As such, they also avoid horizontal inequity problems (i.e., the challenge that similar people may not receive like benefits). Universal transfers also have the potential to reduce administrative costs (by obviating the need to collect household data), though as a practical matter, these costs are usually very small compared to the benefits being given out. Finally, they may also improve the political economy of redistribution (Banerjee, Niehaus and Suri, 2019).

However, a key challenge is that for a given budget, universally-eligible programs transfer substantially fewer resources to each beneficiary than targeted programs. That is, solving the problem of eliminating exclusion error by giving a program to all can be expensive (Hoynes and Rothstein, 2019). While this is true everywhere, this tradeoff is particularly severe in developing countries because the money is not easy to tax back given the informal economy. As a result, universal transfers end up giving the same *net* transfer quite high up the income distribution in developing countries (Hanna and Olken, 2018). This makes universal programs particularly expensive relative to targeted programs.

Deciding which type of program is preferable (targeted vs. universal) ultimately depends on the social welfare function and the relative importance that the government gives to inclusion errors, exclusion errors, and per-capita benefits, as described in the theory in Section 2.2.

To examine these tradeoffs in the data, Hanna and Olken (2018) calibrate this using standard CRRA utility functions and find that for a given budget, programs targeted using standard proxy-means tests substantially outperform universal transfers because they transfer substantially more resources to the poor. But, they do entail more horizontal inequity, so the tradeoff ultimately comes down to how much one values increasing welfare of the poor on average vs. reducing exclusion errors and eliminating horizontal inequity.

One important caveat is that while targeted programs may be more effective in improving welfare on average, even in the case of imperfect targeting (Coady, Grosh and Hoddinott, 2004*b*; Hanna and Olken, 2018), this assumes that the budget is fixed. As the elasticity between the total budget and the number of beneficiaries increases, universal transfers may become more attractive politically (Klasen and Lange, 2016). Many have argued that the durability of programs like Social Security and Medicare in the United States comes from their universal eligibility, which makes them politically popular. Understanding the political response to targeted as opposed to universal programs, or programs (like Social Security) that are both universal and progressive, is an important dimension for future work.[16]

[16]More broadly, understanding the political economy of social protection programs in important for policy

2.9 Summary of targeting

Targeting in developing countries often is quite different than in developed countries, given information constraints and high levels of informality. The growing literature provides indications of the tradeoffs between methods under different contexts: i.e., what "works" really depends on the extent of development and informalization (e.g., Chile may be different than Malawi), what the government is trying to target on (e.g., systematic poverty, shocks, productivity), institutional capacity, and the kinds of frictions that impact take-up.

It is also worth noting that given that each method has different strengths and weaknesses, and many of these may be complementary, in practice, many developing countries "mix and match" methods to improve targeting efficiency and reduce administrative costs (Coady, Grosh and Hoddinott, 2004*a*). For example, many countries first geographically target programs to hard hit regions (to save on the administrative costs of going to relatively rich regions with few of the poor), and then use PMTs within these regions to find those most in need. Others augment the PMTs—which we know are imperfect—with community methods to fill in the gaps and reduce exclusion errors.

While research has been growing in this area, substantial questions exist for future work, from how to best conduct dynamic targeting; whether we can target on productivity, or target different programs towards those who would gain the most from a particular type of program; how to best use newer administrative and satellite data; how to reduce the take-up problem; how targeting affects political and budgetary outcomes; and when stigma matters most for take-up.

3 Transfers for redistribution: understanding program types and impacts

The previous section discussed *who* should receive assistance. We now turn to understanding how to think about the impacts of transfer programs for redistribution, as well as the impacts of various types of programs. In Section 3.1, we begin by clarifying some of the theoretical issues in evaluating the impact of transfers. Section 3.2 then discusses the evidence on the many types of transfer programs in low- and middle-income countries. In Section 3.3, we discuss governance challenges in making sure the aid is actually delivered, which is an overarching issue in the impacts of social protection programs, particularly in the low institutional capacity environments observed in many low- and middle-income countries.

3.1 Theory: welfare analysis of redistribution programs

3.1.1 A simple framework

We have so far assumed that a dollar of benefits is worth a dollar to the person who gets it. In the public discourse, however, there is both a lot of discussion about the form the transfer should take, and how the household makes use of the benefits. We will argue that under a

design and future research. We refer the interested readers to Appendix B for an overview.

set of quite standard conditions, the expression for the social value of a transfer is actually very simple and can clarify a number of vexing issues.

Assume that the utility function of a household is given by $u(c_1, \ldots c_n, b, l)$ where $c_1, \ldots c_n$ is a vector of consumption goods, b is publicly provided benefits and l is labor supply. The household maximizes this utility with respect to a budget constraint that is given by

$$\sum_{i=1}^{n} c_i \leq y + wl + b - t(wl + y - y^*) \tag{2}$$

where y is non-labor income, w is the wage rate, b is the monetary value associated with the government benefit, t is the tax/transfer rate and y^* is the poverty line.[17] Writing utility in this way captures a whole range of possibilities: the government benefit could be a cash transfer which is conditional on an action, in which case b enters the utility function negatively, or it could be a cash transfer combined with a pyschosocial intervention in which b would enter positively. The conditionality is also allowed to impose some time costs, which would then increase the disutility of labor (i.e., $\frac{\partial^2 u}{\partial l \partial b} < 0$).

In addition, the household may face an additional constraint on its choice which takes the form $(c_1, \ldots c_n, l) \in \Phi(b)$. This could represent a credit constraint or a constraint that comes with the program—for example, a conditional cash transfer program would require the household to make certain other choices in order to receive the benefit.

We assume that a household is fully sophisticated in its understanding of these constraints and how they are affected by b.[18] Denote the indirect utility function generated by the household's constrained utility maximization by $v(y, b, t)$. We assume no spillovers so that no one else's utility is affected by the benefits going to a particular household.

The key assumption is that the social welfare function is a weighted average of the individual indirect utilities:

$$W = \int_0^{y^*} g(y) h(y) v(y, b, t) \, dy \tag{3}$$

Under these assumptions, the social benefit from a household with income y getting a small increase in b is given by $g(y) h(y) \frac{dv}{db} db = g(y) h(y) \frac{dv}{dy} \frac{\frac{dv}{db}}{\frac{dv}{dy}} db$. The ratio $\frac{\frac{dv}{db}}{\frac{dv}{dy}}$ represents the household's willingness-to-pay for db, which we denote by $WTP(y)$. $\frac{dv}{dy}$ is the household's marginal utility of income which we denote by $MU(y)$.[19]

Denote the incremental cost to the government to providing a small additional benefit db by $\phi'(b) db$. Then the cost of providing it to everyone with income y is $h(y) \phi'(b)$. If there is no labor supply response from a change in benefits db, then the key statistic that determines the social desirability of a particular intervention under these assumptions is given as follows:

[17] Note that we use y to denote non-labor income here, whereas in Section 2.2.1 we used y to denote total income. We split out labor and non-labor income here because we want to discuss how to think about changes in labor income in response to benefit changes explicitly here.

[18] While the government can vary b by income level, i.e., b_i, we suppress this to simplify notation.

[19] Formally, this should be $WTP(y, b, t)$ and $MU(y, b, t)$, but we suppress the dependence on b and t to simplify notation.

Result 3 *If the government welfare function can be written as a function of individual utility functions with no spillovers, as in equation (3), and if individuals are optimizing when they make their work and spending choices, then denoting willingness-to-pay for a program from household with income y by $WTP(y)$, a program is socially beneficial iff*

$$\frac{g\left(y\right)MU\left(y\right)WTP\left(y\right)}{\phi'\left(b\right)} > 1 \tag{4}$$

As noted above, in developing country settings, the labor supply reduction in response to additional public benefits appears limited, so expression (4) is the key ratio from the government's perspective. Even if their labor supply does change, the utility consequences of this are second-order by the envelope theorem, so this does not change expression (4) as long as it does not affect government tax receipts. So, as long as either a) labor supply does not change or b) if it does change, this does not affect net transfers or taxes collected by the government—a reasonable assumption for most of the poor in developing countries—expression (4) is the key expression of interest.[20]

3.1.2 Implications

This simple observation—that the marginal benefits from a government intervention are given by $\frac{g(y)MU(y)WTP(y)}{\phi'(b)}$—has a number of important implications for evaluation.

Implications for evaluating cash transfers. The usual assumption with cash transfers is that the willingness-to-pay for \$1 is always \$1 (i.e., $WTP\left(y\right) = 1$), though as we will see, there may be exceptions to this rule. The cost of paying an additional \$1 is also a standard quantity in public economics. What that leaves in the government's decision problem is the term that picks up the household's poverty, $g(y)MU(y)$.

This has several key implications. First, assuming we know just how poor the household is, understanding the impacts of cash transfers on household consumption choices (e.g., whether a household buys food versus a television) is not first-order in understanding a program's welfare impacts. Second, from a welfare perspective, it does not really matter if cash transfers cause people to work less, unless doing so has implications for taxes (i.e., the transfers affect work, which in turn affects tax revenues). This is because, other than the tax revenue impact on the government, leisure is just another consumption choice.

In other words, even though there is a large literature that focuses on the impacts of cash transfers on various household outcomes—from food consumption, to health outcomes, to whether one bought a TV—from a welfare perspective, *none of this should matter* if a) the government welfare function can be written as a function of individual household utility functions and b) there are no spillovers to other households.

[20] If, however, labor supply changes, and this affects taxes or transfers, then we do need to modify expression (4) to account for this. Denote the change in labor supply as $\frac{dl}{db}$. In this case, the cost of providing the benefit in the denominator becomes $\left(\phi'\left(b\right) - tw\frac{dl}{db}\right)db$. This is then akin to the marginal value of public funds (MVPF) discussed by Hendren and Sprung-Keyser (2020), but augmented by $g\left(y\right)MU\left(y\right)$ to take distributional consequences into account.

However, understanding a program's targeting—which we discussed in Section 2—is first-order since what matters is the degree to which cash transfers are redistributed to the poor. What also matters is the program's governance and administration (discussed below), i.e., whether the poor actually got their full cash transfer or whether some got lost to corruption. But how beneficiaries choose to spend $1 of benefits—on nicer food or a new television—is to first-order not welfare relevant.

In short, the theory implies that while it may be academically interesting to understand how cash benefits affect consumption choices and work, and while these factors may be important to the political sustainability and messaging of these programs, from a pure welfare perspective, this matters less. Unless, of course, the conditions outlined above fail—which they often do, as we discuss more in detail below.

Implications for programs other than cash transfers. For programs other than cash transfers, the basic framework is the same, except that one needs to measure willingness-to-pay for the program, rather than assuming it is equal to 1. In situations where one needs to estimate the willingness-to-pay, it may be possible to do so using quasi-experimental program variation. For example, in the United States, Finkelstein, Hendren and Shepard (2019) exploit natural variation in the co-pay price of subsidized health insurance in Massachusetts to directly estimate an individual's willingness-to-pay for the government-subsidized insurance, as do other studies of health insurance demand (Thornton et al., 2010; Asuming, 2013; Banerjee et al., 2021). Alternatively, one can experimentally measure the willingness-to-pay for non-pure cash social protection programs. For example, one could offer some households the choice between the program and cash, and then vary the rate at which households trade off cash vs. the program. If one cannot estimate willingness-to-pay directly, one can try to impute it.[21]

3.1.3 Limitations to this approach: wedges, spillovers, and optimization frictions.

The willingness-to-pay approach is a benchmark with three important assumptions (stated in Result 3) for this logic to go through: a) that the government welfare function can be written as a function of individual utility functions; b) that there are no spillovers to other households; and c) that individuals are sophisticated about their own preferences when they make their work and spending choices. We discuss each of these in turn.

Wedges between individual utilities and social utility. One key assumption in the framework is that the social planner maximizes the sum of each household's utility, as given by equation 3. There are two important ways in which this assumption may be an important over-simplification, and where there may in fact be wedges between the social planner's utility and the sum of individual utilities. First, the 'household' is not necessarily a decision-maker with a well-defined utility function, and the household's choices may not reflect the utility of the individual household members. For example, husbands and wives may disagree about

[21] For example, Hendren and Sprung-Keyser (2020) give examples of how to impute willingness-to-pay for 133 programs in the United States, arguing that the willingness-to-pay for a program is equal to the change in N.P.V. income from receiving the program.

how to spend money and make other choices, and this may or may not aggregate into a well-defined utility function (Browning and Chiappori, 1998). And even if a household where the husband makes all decisions ignoring his wife's preferences has a well-defined utility function, the social planner may not choose to put that particular utility function in the social welfare function. Instead, it may base its decisions about the choice of benefits on some other preferences, that, for example, give equal weight to both spouses. In that case, it will need to look beyond the household's willingness-to-pay.

Second, the social planner may have explicit preferences over types of consumption. That is, the social welfare function may include the consumption choices c directly as an argument, above and beyond how they affect $u(c, b)$. For example, voters often seem to prefer that welfare recipients do not spend cash assistance on 'temptation goods' such as alcohol or cigarettes (Currie and Gahvari, 2008). More generally, one could imagine that voters would prefer a household receiving cash assistance to buy food rather than a TV, even if a TV generates a lot of utility. Similarly, some voters seem to prefer that those who receive assistance work. Therefore, while we adopt the view that households should be free to choose how to spend their transfers as a benchmark, we revisit this issue in which policy makers or voters may have particular preferences on beneficiaries' consumption decisions as we discuss various program types and designs below.

Spillovers. A second issue is spillovers. Households that receive benefits may spend them, creating Keynesian multipliers that can have positive spillovers onto other households (Sadoulet, de Janvry and Davis, 2001; Angelucci and Giorgi, 2009; Egger et al., 2022). In principle, the theoretical framework above can incorporate spillovers if one is careful to measure them. In the social planners' problem above, one can consider the impact of a policy change b on all households, not just on those that receive the benefits. However, in the presence of such spillovers, measuring willingness-to-pay by beneficiaries alone will be insufficient.

A related issue is whether the transfers are large enough to cause prices to rise (Cunha, Giorgi and Jayachandran, 2019; Egger et al., 2022). If so, the right conceptual approach is to add the equivalent variation EV to the direct benefits given b, since the equivalent variation captures the difference in welfare due to the price changes. Note that price changes affect everyone—beneficiaries and non-beneficiaries alike—so both need to be accounted for.

Does measured willingness-to-pay capture true household valuations? The third potential limitation of this framework is the failure of the sophistication assumption, in which case a household's choices may not reflect the true utility value of potential decisions, and hence, measured willingness-to-pay may not be an accurate measure of its true utility gain from a choice. This failure can take many forms. For example, one important constraint is information: households may not fully understand the benefit or the value that they will receive from it, and so they may make mistakes in forecasting their willingness-to-pay (Banerjee, Duflo and Hornbeck, 2018). In fact, some households may not understand a benefit's value until they actually experience it, so their willingness-to-pay before receiving the benefit may not reflect their willingness-to-pay after.

Behavioral frictions may also affect willingness-to-pay. For example, if benefits have endowment effects, it would imply that one's willingness-to-pay for a benefit would be higher once

one actually has it than before one has it (see, for example, Carney et al., 2022). Time-inconsistent preferences are another type of behavioral friction (we return to this in more detail below),

A third friction comes from credit constraints and lumpy investment opportunities, which can also affect whether measured willingness-to-pay captures the true value to the household. Consider a lumpy investment with a co-payment of $10 that would pay a return of $100 tomorrow. The value to the household of this investment (above the co-payment) should properly be $90. But a household that only has $5 on hand, and that cannot borrow, will only have a stated willingness-to-pay of $5. This is what one would find if one elicited willingness-to-pay directly.[22]

A more subtle issue arises if the poor face obstacles navigating the bureaucratic process of obtaining benefits. For example, households may be illiterate, or just intimidated by government bureaucracy (Gupta, 2017). They may highly value the benefits, but simply not be able to take them up due to these frictions.

In short, one needs to be careful to think about household decision-making when analyzing the impact of government programs. For many programs, such as cash transfers, the value of the program is simply the value of the transfer, multiplied by the marginal utility of income for those receiving it–so understanding targeting of benefits is often more important than carefully measuring which consumption choices respond. For programs other than cash, measuring the willingness-to-pay for the program, again multiplied by the marginal utility of income, again provides a useful benchmark. That said, one also needs to pay close attention to impacts on within-household allocation, the potential for spillovers, and the degree to which optimization frictions can artificially lower observed willingness-to-pay.

3.2 Empirics: evidence on design choices for transfer programs

In the sections that follow, we review the evidence on a wide range of social protection programs with this framework in mind. We start with the simplest program conceptually—unconditional cash transfers—in Section 3.2.1. We then examine different program types and features. In so doing, we focus on whether and how different programs types and designs affect welfare in the senses discussed in the previous section: affecting future income above and beyond the transfer itself (e.g., by raising wages, or allowing the household to overcome credit constraints); affecting future taxpaying decisions; changing *intra*-household allocations (e.g., between spouses, or between parents and children); changing patterns of consumption in ways a social planner may care about, even if the household does not; spillovers to other households, both pecuniary and otherwise; and affecting welfare by allowing households to overcome information constraints, credit constraints, or other frictions.

Note that we are not attempting a systematic meta-analysis of all papers; we refer interested readers to several recent reviews, such as Bastagli et al. (2016) and Ralston, Andrews and

[22]One could in principle design experiments to estimate willingness-to-pay where the credit constraint does not bind—for example, by enrolling them in a lottery with a small chance of winning enough to cover any reasonable value of their willingness-to-pay, and asking them to commit to a willingness-to-pay conditional on winning the lottery—but this is not the usual approach taken to measure willingness-to-pay.

Figure 2: Percentage of countries that enacted UCTs covering at least 1% of the population by income classification, over time

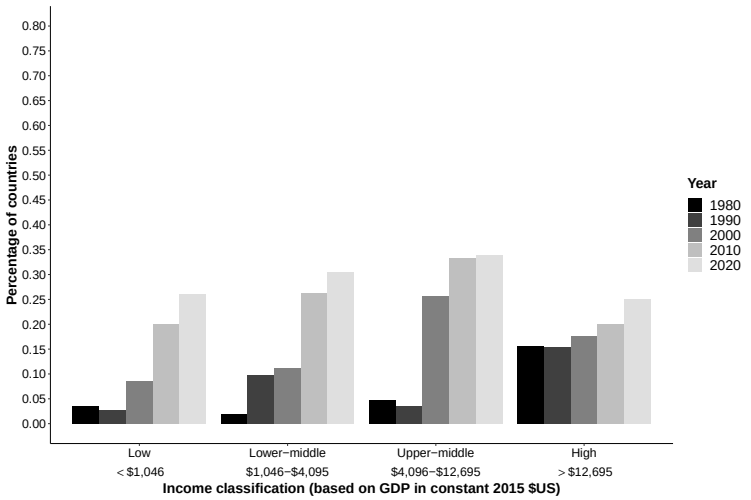

Notes: This figure graphs the percentage of countries that enacted UCTs covering at least 1% of the population, by each year considered against income classification, following the World Bank's 2021 income classification thresholds, as measured by GDP per capita (in constant 2015 $US). We define UCT as a government-implemented, large-scale program that enforces no conditions on recipients and is means-rested. *Source:* GDP per capita and population data from The World Bank (2021*a*,*b*)

Hsiao (2017). Instead, we aim to capture key themes from the literature, to relate these themes to the framework above, and to highlight where more evidence is needed.

3.2.1 Welfare effects of unconditional cash transfers

We begin by considering the simplest program: unconditional cash transfers (UCT). The benchmark program would be a transfer that, once allocated, is completely unconditional– i.e., a household's ability to receive the transfer does not depend on any action. As described above, from a welfare perspective, the basic question of an unconditional cash transfer program is who receives the benefits. This is because a household's willingness-to-pay for $1 in transfers should be $1 (subject to the caveats above), with the social value of these transfers inflated by a measure of $g(y)\,MU(y)$. To first order, the key parameter for evaluating the impact of unconditional cash transfers is therefore the targeting of the transfers, which affects $g(y)\,MU(y)$.

Figure 2 graphs the percentage of countries at different income levels with UCTs covering at least 1 percent of the population for countries at different levels of real constant per-capita GDP. The figure shows that at the beginning of the period we study (1980), such programs were mostly found in the wealthiest countries. By the end of the period, UCTs had become much more common, and almost equally likely to be found in rich or poor countries alike.

Several randomized studies directly evaluate the impacts of UCTs, such as the Haushofer and Shapiro (2016, 2018) and Egger et al. (2022) studies of the GiveDirectly program in Kenya, or consider it in comparison to other programs, such as Baird, Mcintosh and Özler (2011)'s

study in Malawi. There are also many non-randomized studies; Bastagli et al. (2016) provide a systematic review of 201 cash transfer studies (25 percent of which are UCTs).[23]

In general, these programs find substantial evidence that cash transfers improve the different aspects of the welfare of recipient households. For example, Haushofer and Shapiro (2018) document that recipients of cash transfers in Kenya report increases in self-reported psychological well-being and increases in food security, assets, and non-durable expenditures. Haushofer and Shapiro (2016) show that many of these effects persist even several years after the transfers ended. More generally, the systematic review from Bastagli et al. (2016) finds that most cash transfer programs lead to increases in both overall expenditures and food expenditures.

Effects on adult labor supply. As discussed above, changes in labor supply themselves are not welfare-relevant under the assumptions in Result 3: the labor/leisure choice can be considered one of many consumption choices, and a household choosing to work less in response to a transfer can be viewed in a social welfare sense as no different from choosing to buy new clothing or a better roof. Nevertheless, there is substantial interest in this question from a policy perspective, particularly around the political support of these programs. However, the evidence overwhelmingly suggests that cash transfers—at least how they are designed in developing country settings—do not have measurable impacts on labor supply (Banerjee et al. (2017); Handa et al. (2018); see also the discussion in Section 2.1).

Effects on consumption of 'temptation goods.' A common concern about cash transfers is that people will buy 'temptation goods,' such as alcohol and tobacco. However, the same arguments that apply to labor supply apply here: to the extent that governments' utility function is a function of individual household utilities, and households are rationally choosing to consume more of these goods, this should not matter. Nonetheless, voters appear to be less enthusiastic about funding transfer programs that would lead to increased consumption of these goods, either because of fears that these choices reflect short-term temptations rather than rational decisions (Banerjee and Mullainathan, 2010), because of intra-household issues (e.g., an addict may grab the transfer to fund his or her addiction), or simply because voters disapprove of these goods. Evans and Popova (2016) examine 19 cash transfer studies (a mix of UCTs and CCTs), however, and find no evidence that cash transfers increase expenditures on these goods.

Effects on investment and future incomes. The value of the cash transfer may understate its true welfare effects in the presence of credit constraints and lumpy investment opportunities (though if the household understands them, they may still be reflected in a properly measured willingness-to-pay). In the presence of these credit constraints and lumpy investment opportunities, one cannot use the envelope theorem to argue that future income increases are second-order. Instead, one should properly count any future income increases in addition. Given the ample evidence that the poor are credit constrained (e.g., Banerjee and Duflo, 2010), and the plausibility of fixed costs in investment (e.g., needing to buy a machine), these effects may be substantial.[24]

[23] An early and important UCT program was The South African Child Support grant (see, for example, Agüero, Carter and Woolard, 2006; Case, Hosegood and Lund, 2005).

[24] Another reason that cash transfers may increase risky investment-making is if transfers make households

Indeed, there are a number of studies that suggest that households use cash transfers to make future lumpy investments. Gertler, Martinez and Rubio-Codina (2012) study the Oportunidades cash transfer program in Mexico and find that households consume 74 percent of the transfer received, and invest the rest.[25] They then compare households who were treated 18 months earlier rather than later, and find that those treated earlier have higher incomes, even 4 years later. They interpret this increase as a return on the extra 18 months' worth of investments these early-treated households had, compared to the later-treated households.

How should this sort of investment be valued from a welfare perspective? Consider the Gertler, Martinez and Rubio-Codina (2012) case. The envelope theorem approach would be to value a transfer of \$1 at \$1. If so, the fact that the investments increase future incomes is already included in the \$1 by the envelope theorem (i.e., households were already optimizing and the solutions are continuous in the parameters), and so one should not also include the investment income, as this would be double-counting. In the presence of credit constraints and lumpy investments, however, one should instead value the initial consumption (e.g., \$0.74 out of \$1), and add to it the discounted future increases in consumption from those investments. In the Gertler, Martinez and Rubio-Codina (2012) case, they estimate that each \$1 transferred leads to a \$0.016 higher consumption per month starting 12 months after the transfer. The net welfare effect is therefore given by $0.74 + PDV(0.016 \text{ per month})$. With a 10% annual discount rate, for example, and assuming the effects started 1 year after the transfer and lasted until year 10, the net welfare effect of this transfer would be \$1.8 instead of \$1. While the precise numbers depend on the assumptions, the point is that, if one considers these effects separately from a welfare perspective, investment effects can make a big difference.

Spillovers. Cash transfers have been shown to have several important types of spillovers on other, non-recipient households, which are important to include when considering the net welfare effects of these programs. These happen through several different mechanisms.

Informal insurance: One type of spillover effect is through informal insurance arrangements in the spirit of Townsend (1994) and Ligon, Thomas and Worrall (2002). If households are implicitly insuring one another, this could mean that if one household receives a transfer, but another household does not, they may share the proceeds of the transfer to some extent. Angelucci and Giorgi (2009), for example, document these kinds of spillover effects in Mexico's Progresa program. In contrast, Evans and Kosec (2022) show that in Tanzania's CCT, while beneficiary households were more likely to have someone with a personal problem turn to them for assistance, they were no more likely to actually provide assistance to others.

Prices: A second channel through which cash transfer programs can affect the welfare of non-beneficiaries is through price changes. However, since those receiving benefits tend to spend the benefits on a wide variety of items, the demand shock for any given item is likely to

less risk-averse, which would be true if households have declining absolute risk aversion. See the related discussion in Janzen, Carter and Ikegami (2021). If there are no non-convexities in investment, however, then this effect is second-order by the envelope theorem, and so we should not additionally value investment that occurs for this reason above and beyond the value of the cash transfer.

[25] Oportunidades is a conditional cash transfer, as are several other papers that we discuss in this section. Here, we discuss the cash component of these programs; we discuss conditionality in Section 3.2.2 below.

be small. This, combined with the fact that supply of those items may be at least sometimes elastic, suggests that, in practice, price effects for pure cash transfers may be small. Indeed, several studies, such as Cunha, Giorgi and Jayachandran (2019) and Egger et al. (2022) find either zero or economically very small impacts of generalized cash transfers on prices.[26]

One important counter-example is Filmer et al. (2023), who study a cash transfer program in the Philippines in which many, but not all, households in treated areas received benefits; on average 65 percent of households were treated. As a result of this high saturation, the program raised aggregate incomes in treated villages by about 15 percent. They show that this led the prices of protein-rich perishable foods (e.g., eggs and fish) to increase by 6-8 percent, with effects concentrated in villages where the share treated was high and in remote areas where supply was less elastic. They then show that ineligible children in treated villages demonstrated increased rates of stunting, which they argue is consistent with declines in protein consumption for non-beneficiaries associated with the higher prices.

Demand effects: A third source of spillovers is through Keynesian-type demand multipliers. That is, transfers may be spent locally, boosting incomes of those who provide goods or services to beneficiaries, who then spend some of their income locally, and so on.

Egger et al. (2022) use the fact that the transfer program they study was randomized at the village level to estimate a village level 'fiscal multiplier' of this sort. They find that each dollar transferred leads to a 2.5-2.8 dollar total increase in local incomes in those villages due to this local multiplier effect. This suggests that these types of demand-side spillovers may be important in local economies, at least in the context of the large (US $1,000), one-time transfers that they study.

Note, however, that since transfers are paid for externally, this estimate does not capture the net general equilibrium 'multiplier' if one were to impose such a policy nationally. That is, if the transfers were paid for by increased taxes, the taxes *also* could have both direct and multiplier effects on economic activity. From the context of the full general-equilibrium effects of a domestically financed program, the Egger et al. (2022) results therefore may be capturing only part of the effect, as they do not include the reduction in demand in the areas where taxes would be increased to finance such a scheme. Also, note that in most models, long-run outcomes are determined by supply and not demand, so it is unclear that the experiment is providing a good estimate of what the long-run will look like.

Gerard, Naritomi and Silva (2021) study related questions in their study of the labor market effects of a large-scale expansion of a conditional cash transfer program (Bolsa Família). Unlike other programs, Bolsa Família is explicitly means-tested (i.e., those with formal sector income above a certain level are ineligible), so one may be particularly concerned about the labor market effects of such a program. However, consistent with the presence of Keynesian multipliers, they find that expansions in Bolsa Família substantially increased labor market participation, even among workers who were never part of the program. Since they do not observe consumption nor everyone's income, they do not formally estimate a Keynesian multiplier, but the results are consistent with cash transfer programs having strong local

[26] By contrast, Cunha, Giorgi and Jayachandran (2019) and Banerjee et al. (2023) do find price impacts of in-kind transfers, which are a much larger supply shock to a small set of commodities; see Section 3.2.2.

multipliers.

Other considerations. In addition to these effects, unconditional cash transfer programs may have important welfare effects within households that are not captured by this simple framework. For example, these programs may affect the education and health of children (which may or may not be fully counted in a household utility function, depending on whether parents are completely altruistic towards their children) and the relative bargaining weights and outcomes among spouses. We return to these issues in Sections 3.2.2 and 3.2.5.

3.2.2 Conditional cash transfers and intergenerational investments

Perhaps the single largest innovation in social protection programs in the developing world in the past 30 years has been the introduction of conditional cash transfer programs (CCTs). These programs provide a regular stream of cash transfers as in many UCTs, but *condition* the transfer on the household fulfilling a set of criteria, usually related to human capital investments for kids. In Appendix Table 1, we present the conditions of sixty-seven CCTs from forty-five countries. The most common conditions are school enrollment and attendance (52 CCTs), health checkups (38), complete vaccination schedule (22), and attendance of training sessions or workshops (20).

CCTs were pioneered in the 1990s, with Brazil's Bolsa Família, Mexico's Progresa, and Bangladesh's Female Secondary School Assistance Project being early examples. Since then, over 60 countries have started their own CCTs (World Bank, 2018). Figure 3 shows the percentage of countries with CCT programs (i.e., covering at least 1 percent of the population), by quartile of GDP per capita (in constant 2015 $US). CCTs went from being essentially non-existent in 1980 to existing in between 10% and 35% of countries in the bottom three GDP per capita quartiles by 2020. Unlike with UCTs shown in Figure 2, these programs are less common in the wealthiest countries.

The theory behind these programs is that the cash transfers provide assistance to the current generation, while a combination of the cash payments and the requirements for human capital incentives aim to break the inter-generational cycle of poverty. Implicit in these requirements is a statement about welfare: that the government values child human capital investments more than the household decision maker would absent the incentives (if not, then there would be no need for incentives).

Thus, from the welfare perspective outlined above, evaluating these programs requires not only measuring the immediate cash transfer value to the household, but also separately valuing the impacts from the human capital investments above the degree to which they are valued by households. To be precise, consider an action D that the planner would like the household to take. Suppose the household's utility if it takes the requisite action is $u(c) + aD$, with $0 \leq a < 1$, but the planner's utility from a household taking the action D is $u(c) + D$. Suppose further that that the size of the benefit is B, that taking the action has a consumption cost to the household F, and that the income effect of the benefit wouldn't by itself cause it to take the action. The benefit to the household of getting the benefit (in the case where it chooses the action) is $u'(y)[B - F] + aD$. However, the benefit to the planner is $u'(y)[B - F] + D$. So, what we'd want to count is the net benefit to the household –

Figure 3: Percentage of countries that enacted CCTs covering at least 1% of the population by income classification over time

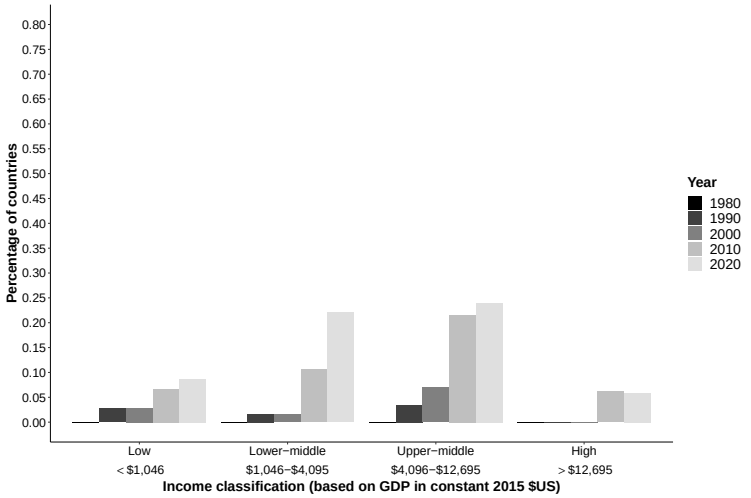

Notes: This figure graphs the percentage of countries that enacted CCTs covering at least 1% of the population, by each year considered (1980, 1990, 2000, 2020) against income classification, following the World Bank's 2021 income classification thresholds, as measured by GDP per capita (in constant 2015 $US). We define CCT as a government-implemented, large-scale program that enforces health or education-related conditions on recipients. *Source:* GDP per capita and population data from The World Bank (2021*a*,*b*)

$u'(y)[B - F] + aD$ – plus the uninternalized piece from the benefits that the planner values but the households don't – i.e., $(1 - a)D$.[27]

Mexico's Progresa program was evaluated using a phase-in design. Surveys were conducted in 320 treated municipalities as well as 186 control ones, which received the program about 18 months later (Skoufias, 2005). This phase-in design allows researchers to understand if a CCT led to human capital investments in health (Gertler, 2004), education (Schultz, 2004), and other outcomes (Skoufias, 2005). When CCTs were instituted in other countries, a number of them were experimentally evaluated during early stages.[28]

These evaluations often find impacts of CCTs on child health and schooling. In the short run, there is evidence that these programs lead to an increase in height and a reduction in stunting (Gertler, 2004; Attanasio et al., 2005; Kandpal et al., 2016); increases in elementary school enrollment (Schultz, 2004; Todd and Winters, 2011; Baird et al., 2014); and increases in cognitive outcomes for children (Macours, Schady and Vakis, 2012).[29] They also show reductions in child labor (Barrera-Osorio et al., 2011; Schady and Araujo, 2006; Benedetti, Pablo and Mcewan, 2016).

[27]We thank the editor for suggesting this formulation.

[28]There are now randomized evaluations of CCT programs in Colombia (Barrera-Osorio, Linden and Saavedra, 2019), Honduras (Benedetti, Pablo and Mcewan, 2016; Galiani and McEwan, 2013; Glewwe and Olinto, 2004; Morris et al., 2004), Indonesia (Alatas, 2011; Cahyadi et al., 2020), Nicaragua (Barham and Maluccio, 2009; Macours, Schady and Vakis, 2012; Macours and Vakis, 2014), the Philippines (Filmer et al., 2023; Kandpal et al., 2016) and Tanzania (Evans et al., 2014; Evans, Holtemeyer and Kosec, 2019).

[29]Parker and Todd (2017) provide an in-depth discussion of the many studies of the Progresa program.

Long-run impacts of conditional cash transfers. Several studies suggest that these human capital benefits may persist. Two studies, for example, examine experimental variation. Cahyadi et al. (2020) measure the impacts of a CCT program in Indonesia after six years, finding persistent human capital gains: 6 years after the start of the program, truancy and child labor fell by half, and stunting—which requires cumulative investments in health—fell by 23 percent. Similarly, Millán et al. (2020) examine the CCT in Honduras, which ran for 5 years in treatment municipalities, but never in the control areas. Eight years after the program ended, they find large increases in education—particularly secondary and university enrollment—for non-indigenous groups, but more muted effects for indigenous groups. They also find that the CCT increased international migration for men who experienced the program as older cohorts. They find little robust evidence on wages.[30]

A second empirical strategy uses the fact that some programs started earlier in some places than others. Behrman, Parker and Todd (2011) use this strategy, comparing those who received an 'extra' 18 months of Mexico's Progresa program, after 5.5 years of exposure in treatment areas, as well as comparing these areas to additional never-treated areas using a matching design. They find increases in schooling attainment for those who received the program earlier.

Barham, Macours and Maluccio (2013, 2018*a*,*b*) use a related design in Nicaragua, using a randomized design and comparing early treatment locations with late treatment locations. In Barham, Macours and Maluccio (2013), the authors focus on boys, and compare the impact of being exposed to a CCT in utero through age 2 with being exposed at ages 2-5. They find that boys exposed early in life had better cognitive outcomes, with no impacts on stunting. In Barham, Macours and Maluccio (2018*a*), they compare girls treated at ages 9-12 with those treated at ages 11-14, and examine outcomes when they were 19-22 years old. They find that girls with early exposure to the program had higher educational attainment, higher earnings, and lower fertility compared to girls with late treatment. In Barham, Macours and Maluccio (2018*b*), the authors examine similar-aged boys and find similar education and labor market effects.

More recently, Araujo and Macours (2021) use this same design to follow the experimental Progresa cohorts *about 20 years later* and find that educational attainment increases. For the children who were school-age when the program started, labor income is higher, particularly for the top of the income distribution. Those in the early treatment group were more likely to migrate, particularly to the U.S., which could in part account for the higher incomes.

A third empirical strategy is to use the general phase-in of programs, coupled with the differential ages of children when the CCT expands to their municipality. The variation in when the CCT enters generates a difference-in-difference in total exposure to the CCT, while controlling for age. Parker and Vogl (2023) use this strategy to analyze Progresa. They find that exposure to Mexico's Progresa in primary school leads to higher educational attainment,

[30]The idea that nutritional investments can have long-run effects can be seen in the long-run followup study of protein supplementation done starting in 1969 in Guatemala, which reduced stunting for 0-3 year olds by more than 50 percent. The study found that, in 2002-2004 (so more than 30 years later), wages were 46 percent higher for people who were ages 0-3 in the two villages randomized to receive protein supplements compared to comparable people from the two control villages (Hoddinott et al., 2008).

greater labor force participation, more migration, and higher earnings for women. They also find modest impacts on education and migration for men.

Does conditionality matter? Can it be a two-edged sword? These human capital improvements provide evidence that CCTs can change household behavior. But, perhaps poor households, when faced with an infusion of cash, would have increased investments in human capital due to income effects, even without the incentives. Do the conditions themselves matter? Several studies have found evidence that households respond to conditionality per se on the targeted outcomes. Studying random variation in the warnings for violating the conditions on school attendance in Brazil's Bolsa Família program, Brollo, Kaufmann and La Ferrara (2017) find that child attendance increases in response to the threat of penalties. Similarly, in Nicaragua, Macours, Schady and Vakis (2012) argue that a treatment arm that gave additional cash over the basic CCT did not lead to additional improvements in child outcomes, suggesting that the conditions drove the effects in the basic treatment.

The CCT conditions could also help those children who are more disadvantaged within families: Akresh, de Walque and Kazianga (2013) conduct a randomized trial in Burkina Faso, where villages were randomized into typical UCTs or CCTs. For boys and older children, who tend to be prioritized by parents, both programs led to increases in enrollments. However, they find that the CCT led to large increases in enrollments for girls and younger children, whom they identify as more 'marginal' for education, while the UCT had no impacts on enrollment for these groups.

Could there be downsides to conditionality? Baird, Mcintosh and Özler (2011) randomize communities in Malawi to a UCT, a CCT (based on school attendance), or a control group. The results show that conditionality matters, but also entails risks. On the one hand, the CCT conditions improved targeted indicators: while both the UCT and CCT led to a reduction in dropout rates for the adolescent girls, the reduction was more than twice as large in the CCT. Likewise, the CCT led to increases in reading comprehension, while the UCT did not. On the other hand, those in the UCT arm experienced larger impacts in terms of delaying marriage and reducing pregnancy rates, particularly among adolescent girls who dropped out of school. This suggests that the success of the conditionality could come at the cost of denying cash transfers to non-compliers who could benefit from the program. Ultimately, how one balances these tradeoffs in choosing the program design depends on the policy objective the social planner wishes to achieve.

Combined, these studies suggest that conditions may matter for child human capital outcomes, but that there may be a downside from cutting off assistance to non-compliers. This presents the intriguing possibility that a *labeled* transfer program—where households are told that there are conditions, but they are not rigorously enforced—could achieve a similar impact as conditionality without its downsides. Such an approach could also reduce monitoring costs. Benhassine, Devoto and Duflo (2015) suggest this may be the case: in Morocco, they compare a labeled cash transfer with a CCT, and find that the labeled transfer substantially reduced the drop-out rate and increased school enrollment among those who had dropped out at baseline, whereas the traditional conditional cash transfer had smaller positive effects. However, more research is needed on labeling to better understand when it works, how long it persists, and how to structure the labels.

3.2.3 In-kind vs. cash transfers vs. vouchers

Conceptual Issues. An alternative approach to delivering assistance is to provide households with goods or services directly, rather than cash. There are two ways to do so. The first is to physically transfer a good or service (e.g., hand people loaves of bread, distribute bednets, etc). A second way is to give people a voucher to purchase a specific good or service (e.g., give people a debit card to use to purchase rice, give people a coupon to pick up a bednet at a clinic). Both types of in-kind programs create frictions that can influence recipients' consumption decisions, but they potentially differ in terms of how they affect the market for the goods in question.

There are reasons why cash could be preferred to in-kind transfers; after all, households can use cash to purchase whatever will increase their utility most. Moreover, for households that would purchase more of the in-kind good than is provided by the transfer anyway, providing in-kind goods should be the same as cash, since households can just reallocate their cash elsewhere.

However, theory provides some reasons why policy makers may prefer in-kind. The first has to do with specific social preferences over types of consumption discussed in Section 3.1.2. If, for example, the social decision maker cares explicitly about certain types of 'merit goods,' such as nutrition or health care, rather than just household utility, the social decision maker may explicitly want to influence beneficiaries' consumption baskets. When households are constrained—so they are given more of the in-kind good than they would consume otherwise, and cannot easily resell it—in-kind transfers can encourage consumption of these merit goods. This strategy can also work if households are simply not optimizing or are using mental accounting (Hastings and Shapiro, 2018).

A second reason policy makers may prefer in-kind transfers has to do with how these programs interact with prices. There are two issues. First, in the case where the government provides the goods in question (i.e., provides goods directly, rather than vouchers), it can increase the overall supply of the goods in question. This depends, of course, on how goods are sourced-for example, if the government buys food in local markets and then gives it away, supply in that local market is unaffected. However, if the government buys food far away and transports it into a particular area, there may be a positive supply shock of the good in question in that area. This means that, if supply is inelastic, prices for the subsidized goods may be lower in some kinds of in-kind programs than in a cash program (Coate, Johnson and Zeckhauser, 1994; Basu, 1996). Second, governments typically set in-kind benefits in terms of quantities (i.e., 5 kg of rice per month), whereas they typically give cash in specific amounts each month or quarter (i.e., 500 Rupees per month). Setting benefits in terms of quantities rather than prices can provide implicit price insurance to beneficiaries (Gadenne et al., 2021). While the government could in principle adjust the quantity of in-kind benefits based on prices to hold value constant, or conversely adjust the value of vouchers or cash payments to keep their real value constant, these adjustments may not happen.

Third, in-kind transfers could potentially improve self-targeting if they increase participation costs more for the wealthy (Nichols and Zeckhauser, 1982). For example, if the government provides poor-quality, subsidized bread, richer households may not want it.

Finally, there are administrative differences between these programs. In direct distribution programs, the infrastructure required to deliver millions of tons of goods is different than the infrastructure required to run a voucher system using electronic debit cards (Banerjee et al., 2023) or to give out cash payments, and these administrative differences may really matter in low state capacity settings.

Experimental comparisons of in-kind and cash programs. Several studies examine these issues by experimentally allocating areas into cash, voucher, or in-kind distribution programs. In Mexico, for example, the government ran an RCT that randomized villages into receiving in-kind food transfers, approximately equivalent cash transfers, or a control group. Several studies have examined this experiment. Overall, Skoufias, Unar and González-Cossío (2008) find that both types of transfers have similar effects on food consumption, total consumption, male labor supply, and poverty outcomes; Cunha (2014) finds similar consumption effects of in-kind and cash programs.

These experiments, however, reveal two potentially important differences between cash and in-kind programs. First, Leroy et al. (2010) find (and Cunha (2014) confirms) that in-kind programs lead to increases in both total calories and micro-nutrients consumed through stickiness in consumption choices. Second, Cunha, Giorgi and Jayachandran (2019) find evidence that in-kind transfers cause price declines of about 5 percent in remote areas. They argue that, in remote locations, in-kind rather than cash increases the real value of transfers provided by 14 percent (much larger than the 5 percent price decline) because it reduces prices on all purchases of affected goods by everyone in these locations, not just those financed by the transfer program. The authors find no price changes in more developed locations.[31]

A second experiment comparing cash, in-kind food, and voucher food programs was done by Hidrobo et al. (2014) in Ecuador. They also use a randomized design to compare the impact and cost-effectiveness of the three alternatives in urban areas with well-functioning markets. They find similar effects of the three programs on both food and non-food consumption, but they find differences in food composition: relative to the in-kind program, the voucher program—which had a much wider range of products—led to higher dietary diversity. For example, voucher households consumed more vegetables, eggs, and dairy than those that received the in-kind transfer. The in-kind program was also more than twice as expensive (in terms of costs per outcome achieved) as either the cash or the voucher program.

Third, Banerjee et al. (2023) examine an experiment that compared an in-kind rice program to an electronic voucher program that allowed recipient households to purchase an approximately equivalent amount of rice and eggs (in value) from private providers. Importantly, the Banerjee et al. (2023) study is an at-scale experiment: 105 districts, with a combined 3.4 million beneficiary households, were included, with both programs implemented as usual by the government bureaucracy. They find a dramatic difference: e-voucher programs delivered concentrated assistance to targeted households, whereas in-kind aid was spread much more widely. As a result, targeted households received 45 percent more assistance in voucher dis-

[31]This is consistent with Michelson et al. (2012), who argue that rural villages have difficulty responding to cash transfers stemming from less market access and less competition among suppliers, and Jiménez-Hernández and Seira (2022) who show that direct provision by government introduces competition, and hence is effective when there is local market power in the provision of goods.

tricts than in in-kind areas. For households in the bottom 15 percent at baseline, poverty fell by 20 percent. The results suggest an important additional dimension, namely that voucher programs may be more reliably administered in settings with low state capacity.

These studies together present a nuanced understanding of the relative costs and benefits of in-kind, voucher, and food programs. Cash remains an important benchmark, since households can choose to spend it on what they prefer. To the extent that policy makers believe households are not privately optimizing their nutrition choices or otherwise prefer food programs, the evidence suggests that food vouchers can have impacts on the set of foods that households consume. In-kind programs are more costly to administer, and harder to administer faithfully, but may be useful particularly in very rural, isolated areas where supply is inelastic or non-competitive.

3.2.4 Workfare programs

In workfare programs, the government provides public employment, typically at a low wage, for those who want it. As discussed in Section 2.4 above, the idea is that work requirements can be a screening mechanism: those who have better work options than those offered by the government will opt out of the government's workfare program.

In these programs, the government is directly intervening in the labor market, which has several implications for evaluating these programs. First, in addition to differences in workers' preferences (i.e., how much they prefer a workfare job compared to their alternative private sector job), one also needs to consider the productivity of the work. If the work is productive (i.e., building public infrastructure with the same efficiency as the private market), this type of program can be efficient; if the work is unproductive (e.g., slowly building 'bridges to nowhere'), then the social efficiency loss also needs to be accounted for in any cost-benefit analysis. Second, there may be important spillovers through labor market effects on wages, or they may affect migration (see, for example, Imbert and Papp, 2020*b*,*a*; Berg et al., 2018; Azam, 2012; Bertrand et al., 2021).

In an at-scale study, Muralidharan, Niehaus and Sukhtankar (2023) examine the general equilibrium market effects of the MGNREGA program. They conduct a randomized experiment at the sub-district level that improved program administration (by providing biometric smartcards, which reduced leakage, reduced payment delays, and generally made the program function better). They find that private sector wages increased. As a result, while the reform raised the income of low-income households by 13%, 90% of the gain came from equilibrium effects on the private market. They find, perhaps surprisingly, that this also led to an increase in private sector employment, consistent with monopsony in local labor markets.

Similarly, Franklin et al. (Forthcoming) study a workfare program that was randomly phased in across neighborhoods in Addis Ababa, Ethiopia. The program provided an hourly wage that was about 64 percent higher than the private market, which led households to shift to public employment. The program led to an increase in private market wages that accounted for most of the increase in welfare of program beneficiaries. Moreover, the authors experimentally show that workfare increased amenities in treated neighborhoods, suggesting that

the investments they study were productive.

3.2.5 Programs targeting specific individuals within the household

If the household behaves like the unitary model predicts (i.e., the household solves its optimization problem as a single unit) then the decision about which household member should receive transfers would be irrelevant: no matter who receives the income, the money would be spent in the same way. However, there is reason to think that the recipient matters (see, for example, Browning and Chiappori, 1998; Robinson, 2012). Indeed, this is one of the frictions that we discuss in Section 3.1.2 that moves us away from the benchmark case. If money is spent differently in the household depending on who receives it, then the recipient matters for policy. The empirical evidence on the degree to which this matters in a real-world policy context is, however, mixed.

A number of studies explicitly randomize whether transfers are given to men or women in developing countries, finding little effect overall.[32] For example, Akresh, de Walque and Kazianga (2016) randomize whether the transfer is given to mothers or fathers in Burkina Faso. They find no differences on child health or education outcomes on average. That said, they do find that giving transfers to mothers leads to lower levels of child labor. They also find some evidence that giving money to fathers improves outcomes in poor rainfall years, and leads to more household investment in livestock, cash crops, and housing.

Likewise, Benhassine, Devoto and Duflo (2015) and Haushofer and Shapiro (2016), in their studies of cash transfers, randomize the recipient to be men or women, and both studies find few differences on average. Haushofer and Shapiro (2016) do, however, find that giving transfers to women leads to higher reported psychological well-being and greater female empowerment, though they find no differences on many other dimensions (e.g., food security, health, and education).

However, Field et al. (2021) show how female empowerment that stems from control over one's money in the form of bank accounts could have real effects. Studying MGNREGA, they experimentally show that women who were set up to receive direct deposits of their transfers into their own accounts (rather than their husbands'), as well as training on account use, worked more in both the public sector program, as well as in private sector jobs. Three years later, this treatment even shifted community norms of female employment.

Finally, another serious gender-related concern is whether cash transfers themselves, by causing discord within households over spending preferences (particularly if women receive the program), could lead to increases in intimate partner violence (IPV). Buller et al. (2018) review fourteen quantitative and eight qualitative studies and find little evidence of this. In fact, they find that most studies demonstrate evidence of decreased IPV, with only two studies showing mixed or adverse overall impacts.

[32] These studies build on Lundberg et al's (1997) quasi-experimental study in the United Kingdom. Other important non-randomized studies include Duflo (2003)'s study of pensions, discussed in Section 4.3.2.

3.2.6 Poverty traps and lumpy transfers

An important decision when designing a transfer program is the transfer size and frequency. Thinking about this question in the developing world context is important given that extreme poverty together with saving constraints may be more likely to create poverty traps (Dasgupta and Ray, 1986). This suggests a role for lumpy transfers for some types of people rather than streams of payments to spur investments.

Conceptual issues: To fix ideas, consider a very simple Solow (1956)-style model where individuals save a fixed fraction of income and invest the rest. The capital stock in period t is k_t. Each period an individual produces $f(k_t)$ and consumes $(1-s)\,f(k_t)$. Next period's capital stock is given by $k_{t+1} = (1-\delta)\,k_t + sf(k_t)$, where δ is the depreciation rate of capital. The steady state level of capital is given by setting $k_{t+1} = k_t$, which yields $f(k)/k = s/\delta$.

The key is to understand the shape of the production function. If $f(k)$ is globally weakly concave, this yields a unique steady state—i.e., there is no poverty trap. But if $f(k)$ has convex regions (i.e., has an S-shape), there is the possibility of multiple steady states. In particular, there is a tipping point k^*: if capital is just below k^*, the individual will disaccumulate capital and return to a low equilibrium, but if capital is just above k^*, the individual will accumulate capital and move to a higher equilibrium.

The key point here is that if there are these threshold tipping points k^*, one-time shocks to capital can have long-run implications for households: a one-time negative shock can lead them to a long-run poverty trap, or conversely, a one-time positive capital infusion can lead them to escape poverty and end up at a new, higher income level. From a social protection perspective, whether or not these tipping points exist—or whether, alternatively, we are in a world with a unique steady-state—matters substantially.[33]

The idea of threshold tipping points and poverty traps is very general, but the details of the transition path can differ depending on the particular model being analyzed. In particular, note that we are simplifying here by assuming that savings is a fixed fraction of income, as in Solow (1956). In a more general formulation, with lumpy investments but endogenous savings, there can still be poverty traps, but the dynamics near the critical point may differ (see, e.g., Skiba, 1978; Buera, 2009; Ikegami et al., 2018), with much more rapid transitions just above and below the tipping point than in the very simple model we have sketched here.

Empirics: Which model are we in? In an RCT in Bangladesh that provided assets to a subset of poor households, Balboni et al. (2022) document the existence of poverty traps. They sort people based on their initial wealth, consider the shock of the capital transfer, and compare assets at baseline (k_t) with assets 4 years later (k_{t+1}). They find that there is a unique steady state in control villages, but that transferring assets creates multiple steady states—so not only does the transfer temporarily increase incomes in the short run, but for some, it can tip them into a substantially higher steady state.[34]

[33] This relates to the discussion of how to value a $1 transfer. On the margin, the $1 is valued at $1, but a lumpier transfer could be valued at a different rate given how it will be spent and its potential returns.

[34] Lybbert et al. (2004), Carter and Lybbert (2012), and Banerjee et al. (2019*a*) also provide evidence of the existence of poverty traps.

More generally, a number of studies have examined big-push 'graduation' programs—called 'graduation' due the idea that they can help people 'graduate' from poverty. These types of programs were pioneered by BRAC in Bangladesh and are now present in over 40 countries (Banerjee, Duflo and Sharma, 2021). The program usually consists of a lumpy productive asset—in the form of livestock—bundled with skills training, savings, health education, coaching and cash payments for a period of time. The goal is to alleviate a household's capital and skill constraints to 'push' them out of poverty. Several studies have found positive results, including Bandiera et al. (2017) in Bangladesh; Banerjee et al. (2015) in Ethiopia, Ghana, Honduras, India, Pakistan, and Peru; and Bedoya et al. (2019) in Afghanistan. In the short run (3 to 4 year follow-up, 1 to 2 years after the program ends), there is an improvement in consumption, food security, asset holding and savings. In the medium run (7 year follow-up and 5 years after the program ends), both individual and household outcomes remain higher (Banerjee et al., 2016; Bandiera et al., 2017). In the long run (10 year follow-up), there appears to be no additional growth, but persistent effects for those who received the program (Banerjee, Duflo and Sharma, 2021). While the programs show long-run effects, they are also expensive, and so whether the returns are large enough to suggest moving households out of a poverty trap, as opposed to households just receiving regular returns from a large transfer (as in the permanent income hypothesis), depends on the discount rate assumed. In India and Bangladesh, the program costs are relatively low, and the impacts are large, so that the net effect is positive for most plausible discount rates, but this may not be true elsewhere, where the discount rate may matter more.

Two specific questions that arise from these papers are: (1) can the government implement these programs? and (2) to make the programs cheaper and more cost-effective, can you scale down some program components? Recent experimental evidence from Botea et al. (2021) in Zambia answers 'yes' to both questions. First, they show that governments can run the programs: they show increases in consumption, assets, and mental health as a result of the full government-run package. But, importantly, they show similar results between the group that received a complete graduation treatment and the group that received the asset (and a savings tool) but no training.

In Ghana, however, Banerjee et al. (2022) find that neither giving people just assets nor just the savings tool has a positive long-run impact, while the whole package does. The Zambia and Ghana papers' results are consistent if both the asset and the savings tool are important for program success, but the training is not. While the Zambia results suggest that perhaps some of the training can be stripped off these programs, reducing the cost, the Ghana results suggest that further testing may be useful in unbundling these programs.

Blattman, Fiala and Martinez (2014, 2020) study a different type of lump-sum transfer program-upon submitting a business plan, young adults form groups in Uganda and receive a one-time grant. The authors find substantial increases in investment, work and income four years later. However, in the nine year long-run follow up (Blattman, Fiala and Martinez, 2020) they see the gains dissipate, as the control group catches up with them. Although both groups converge in employment, earnings, and consumption levels, those who get the grants have a lasting impact on durable asset stocks and skilled work.

These evaluations suggest that lumpy transfers can have long-run effects. But is it the

lumpiness per se that matters, or the complementary investments that go along with it? To test this, Haushofer and Shapiro (2016, 2018) experimentally compare the effects of a lump-sum cash transfer ($404 in PPP terms) to an equivalent cash payment in 9-month installments. At 9 months, the monthly transfers increased food security, while lump-sum transfers increased assets. After 3 years, there were no differences between the lump-sum and monthly transfer groups. It is important to note that the monthly payments were only spread over 9 months and thus may understate differences when compared to a steady-state transfer program.

More generally, the evidence discussed here suggests the possibility that there are some households who are constrained, and for whom a one-time intervention could have long-lasting effects. We note two outstanding questions. First, not all households are likely to benefit from one-time interventions. Which households are 'trapped,' and how can we identify them? Second, what, exactly, keeps these households in the poverty trap? Is it a rational story given lumpy investments, or are they unable to save optimally, as suggested by, e.g., Augenblick et al. (2022)?

3.3 Governance in the delivery of assistance

A crucial challenge that developing countries face when it comes to social protection is governance, i.e., making sure that benefits get delivered to the targeted individuals. In many countries, this is a non-trivial problem. In India's MGNREGA workfare program, for example, Muralidharan, Niehaus and Sukhtankar (2016) find a 30.7 percent leakage rate in Andhra Pradesh in 2012. In Indonesia's subsidized rice program, Banerjee et al. (2018) calculate that eligible households only receive a third of the intended subsidy. Camacho and Conover (2011) find cheating on the PMT formulas, and Niehaus et al. (2013) estimate that 70 percent of ineligible households in India have below-poverty line cards, and many households report paying small bribes. Of course, these are particular examples, and researchers often study leakage where it is thought to be high, so one should not take these as unbiased estimates of overall leakage rates. And, some of this 'leakage' goes to other households near the poverty line. But, these facts do suggest that meaningful improvements to programs can be made by improving service delivery.

Corruption risk can also distort program choices, as people may prefer programs with a lower risk of corruption. For example, in a study in Indonesia, Kyle (2018) shows that the poor tend to prefer targeted social assistance programs to broad-based energy subsidies, for which the benefits accrue more to the middle class. But, in districts where local politicians are corrupt, poor citizens tend to support fuel subsidies, for which local corruption matters less.

A traditional approach to improving governance often focuses on monitoring and audits (Olken, 2007). More recent evidence suggests that reforms to payment systems, private sector involvement, improvements in bureaucracy, and information provision can also help. We next discuss the existing evidence on these approaches in the social protection context.

Digital payment systems. In many developing countries, the authentication of beneficiaries and payments systems is manual: there is a list of beneficiaries, some local official needs to

verify who receives benefits, and cash is given. This makes it hard for the central government to monitor whether targeted beneficiaries receive their full payments.

Several studies have examined the implications of digitizing this process. Muralidharan, Niehaus and Sukhtankar (2016) examine the randomized rollout of "smartcards"—biometrically authenticated payments—for receiving benefits from India's workfare program (MGNREGA) and pension programs (SSP) in Andhra Pradesh, India. The smartcards worked: they reduced payment delays, reduced leakage by 41 percent, and quite importantly increased program take-up by poor households. A second study, Banerjee et al. (2023) (described in more detail in Section 3.2.3) studied the conversion of an Indonesian food subsidy program from an in-kind program administered by local officials to digital vouchers (a debit card) redeemable at a network of bank agents. The digital vouchers led targeted households to receive substantially more benefits, which in turn bumped many poor households over the poverty line. These effects occurred even without biometric identification, suggesting that the digital card, rather than the biometrics, may have been important in this case.

A word of caution, though: Muralidharan, Niehaus and Sukhtankar (2023) show that a reform in Jharkhand, India, that began to require biometric authentication for the subsidized food program had the unfortunate side effect of dropping many poor people from the system who had not linked their biometric information to their food benefits account. Banerjee et al. (2021) similarly find that many people could not enroll in Indonesia's health insurance because of errors in the underlying identification data systems. These studies emphasized the need for universal access and robust administrative data to complement reforms.

Finally, it is worth noting that the degree to which people prefer digital systems may depend on the overall level of adoption of those systems in the country more broadly. Berkouwer et al. (2021) show that during the COVID-19 crisis, people in Kenya—where there is high mobile money adoption—preferred mobile money over electricity subsidies. However, it was the opposite in Ghana, where mobile money is less widespread: here, nearly half of recipients preferred electricity transfers. In fact, many were willing to forgo high levels of subsidy in order to receive electricity instead of mobile money.

Back-end payment systems. A related type of reform is linking delivery accounts to back-end payment systems, which are the background systems that ensure the flow of payments from government programs to people. Banerjee et al. (2020*a*) study these systems in the case of MGNREGA in Bihar, which simplified the payment process between central and local governments. The changes reduced program expenditure by 24 percent while maintaining the same payments to workers. However, it is worth noting that these programs must be implemented with care. For example, the reform studied by Banerjee et al. (2020*a*) delayed payments to beneficiaries. While these payment systems are important, such transitions must be therefore managed with care.

Information provision. If "leakages" arise from local officials not following the national government's rules, one fix is to provide information directly to beneficiaries so that they can demand their full benefits from local officials. Banerjee et al. (2018) experimentally test this hypothesis in a food transfer program in Indonesia and find that simply mailing out identification cards dramatically increases the subsidy eligible households receive. They find

additional benefits from posting information publicly, creating "common knowledge" and thus increasing accountability of local officials.

Outsourcing delivery. While social programs are typically run by the government, there is scope for the private sector to intervene. Theoretically, there can be efficiency gains from involving the private sector since the government can provide stronger incentives to contractors. Nonetheless, if there is little competition in tenders, then the private sector may provide lower quality than the government. Empirically, Banerjee et al. (2019*b*) examine what happens when they randomly allow some villages to outsource the last mile delivery for Indonesia's rice subsidy program. Privatization of delivery leads to increased efficiency with no drop in quality, though the effects are small. It is worth noting, however, that this study was about privatization to various small-scale local vendors, not large companies; larger-scale privatizations remain an open question for further study.

Improvements in bureaucratic quality. He and Wang (2017) study what happens when villages receive college graduate village officials, who are more educated than typical village officials and supposedly free from local interest groups. Using the staggered timing of assignment of these officials to local villages, they show that in villages with college graduate officials, more households are registered as poor or disabled and, most importantly, more poor households benefit from subsidies targeted to the poor.

3.4 Summing up

The basic welfare framework discussed in Section 3.1 and Result 3 is powerful in its simplicity: it says that for most programs, the value of $1 in cash transfers to a household is $1, and hence what matters most is *who* gets the support, rather than the details of the form the $1 in transfers takes. This suggests that understanding the targeting of programs, as described in Section 2, is fundamental to understanding the welfare performance of a wide variety of social protection programs. Likewise, the governance issues outlined in Section 3.3 are first-order welfare relevant since they can affect whether the $1 actually reaches the intended household.

However, there are some circumstances where this result may not hold. For example, to the extent that parents are not perfect agents for their children, or perhaps are credit-constrained, one may care about the impacts directly on children's human capital acquisition. The large literature on conditional cash transfers, for example, suggests that these programs, by explicitly conditioning transfers on children going to school and receiving regular health checkups, can improve children's human capital, with the potential to help break the cycle of poverty. These benefits may be valuable above and beyond the value of the cash to recipient households. The presence of spillovers to other households and poverty traps can also change the welfare calculus.

The discussion here also suggests important directions for additional research. First, different programs may be appropriate for different types of people. For example, conditional cash transfers may work well if households are near the margin of take-up for the incentivized behavior; for households far from the threshold, imposing conditions that they cannot comply with could make them worse off by denying them program access (Baird, Mcintosh and Özler,

2011). Figuring out how to assign the right programs to the right people is an important area for future work.

On a related theme, if households are not fully optimizing—for example, due to behavioral or cognitive frictions, perhaps made worse by poverty (Mullainathan and Shafir, 2013; Dean, Schilbach and Schofield, 2017), design elements of the program may be first order welfare-relevant even if they are not reflected in households' willingness-to-pay. Understanding how to leverage behavioral nudges to best improve program outcomes is thus important.

A third area for future research concerns poverty traps. The literature suggests the possibility of sustained, long-run changes in outcomes from 'graduation' programs. But this is clearly not for everyone—the standard 'graduation approach' model may work in a rural area for someone with some ability in animal husbandry, but this may not work for people in urban sectors, or for others who may want other types of employment. The studies reviewed suggest a 'possibility result'—poverty traps may exist and, for at least some people, a particular program can work. However, this is by no means a panacea or the right solution for everyone, and not everyone needs all the program's components. Untangling this knot and figuring out approaches that work for others is an important area for future work.

A fourth area is interactions with labor markets. Programs that provide low-wage work, such as MGNREGA in India, are one approach to linking work to social protection goals. However, one can think more broadly about how labor market interventions are part of the social protection system, such as a minimum wage, job training, job matching programs, etc. While a full review of these types of labor market policies is outside our scope, understanding these issues in developing countries is an important area for research.

Finally, the programs discussed in this section have primarily been conceptualized as responses to long-run poverty. But, to the extent that eligibility for these programs is dynamic, they can also help provide insurance against shocks. We tackle the question of how to provide insurance more systematically in the next section.

4 Mitigating risks

Poverty is not a static concept: someone may belong to the middle class one day, but a loss of income due to a bad health shock, a loss of a job, or even a single bad harvest may suddenly create real challenges. Governments worldwide often step in to provide social insurance programs to help people manage these kinds of risks. They do so in part because private insurance markets are often incomplete, so even if individuals wanted to insure against these risks on the private market, doing so would be difficult or expensive. There are often substantial aggregate shocks, especially due to climate and natural disaster-related issues—droughts, storms, floods, and earthquakes—which can pose challenges for private insurance, especially in low- and middle-income countries, which are not as well integrated into global reinsurance markets. And, for whatever reason, even if private insurance was supplied at actuarially fair prices, insurance demand can be quite limited, so people may end up leaving risks uninsured that ex-post can cause them serious challenges.

These problems, of course, are not limited to the developing world, but they end up in many

ways more severe. Life in many low- and middle-income countries is particularly risky. For example, 78 percent of poor people reside in rural areas and mostly subsist on agriculture, so they face much more income risk than employees (The World Bank, 2014). Since there is, on net, much less insurance against these shocks, households end up taking much more severe actions to smooth shocks (see, for example, Chetty and Looney (2007)).

In this section, we broaden our framework to consider social insurance policies. In Section 4.1, we begin by summarizing the basic theoretical challenges of privately providing insurance (particularly in developing country settings), discussing the basic rationale for these programs. Next, as many low-income and emerging economies are characterized by informal insurance, we discuss their limitations and the places where formal insurance could help fill the gaps. We then review the current state of knowledge on social insurance schemes in developing countries. In Section 4.2, we examine insurance against *income* shocks, e.g., unemployment insurance and parental benefits. In Section 4.3, we discuss insurance programs that help cushion *expenditure* shocks, e.g., health insurance and pensions/annuities. In Section 4.4, we discuss aggregate shocks, in particular those related to *climate and natural disasters*. These shocks affect many people at once—changing the nature of the response, and limiting the ability of traditional risk-pooling to solve the problem—and can affect both the income and expenditure sides simultaneously. Finally, in Section 4.5, we discuss how risk interacts with broader issues of social protection system design.

Note that while the literature on social insurance in developed countries is extensive, it is relatively less developed for low-income countries. There is also substantial heterogeneity across countries in the scope of government-provided social insurance, with social insurance programs more common in middle-income economies than low-income ones.[35] In light of this, we also discuss directions for future research.

4.1 The challenges of insurance provision and rationales for social insurance

4.1.1 Basic theoretical issues: moral hazard and adverse selection

In Section 2, we discussed a model where the government redistributes between households with different realized incomes. The premise of insurance markets, by contrast, is that much of this redistribution can be achieved by ex-ante contracting between households.

We start with the simplest model, where there are only two available types: with probability p_i, a person will be a high-ability type (earning higher income y_{hi}), and with probability $(1 - p_i)$, she will be a low-ability type (earning lower income y_{li}). If the probabilities are public information, then the individual can purchase insurance that pays $y_{hi} - y_{li}$ in the low state by paying an actuarially fair premium ex-ante equal to $m_i = (1 - p_i)(y_{hi} - y_{li})$. In a model with risk-averse individuals, symmetric information about types and actions, and actuarially fair insurance markets, individuals will fully smooth their consumption across states by purchasing insurance. In this world, the private market will provide many types

[35]For these reasons, social insurance programs are, for example, much more developed in countries in Latin America. See Levy and Schady (2013) for a nice review.

of insurance contracts: health insurance, unemployment insurance, crop insurance, etc.

As soon as we deviate from symmetric information, two market failures arise that undermine the private provision of insurance: adverse selection and moral hazard. Adverse selection arises because the individual has private information about their distribution of outcomes that the insurer cannot observe and thus cannot properly price; that is, either the probabilities p_i or the risk $(y_{hi} - y_{li})$ is at least partially private information. If the insurance company prices its contract to at least break even with high-risk types, then those with lower risk may not buy into the product offered. In an extreme case, the market can completely unravel such that no insurance contracts exist (Stiglitz and Weiss, 1981). Government action can help address this concern.

The second market failure is moral hazard, i.e., the idea that insuring individuals against adverse events can encourage risky behavior. For example, if individuals are insured against unemployment spells, they might put less effort into finding a new job. This can be included in the model by endogenizing p_i such that the probability of being in each state of the world is affected by an individual's behavior (Stiglitz, 1974). If there is heterogeneity in moral hazard, this can also be a rationale for government intervention (Einav et al., 2013).

4.1.2 Low formal insurance demand in developing countries

The benchmark model above would predict that, when faced with the offer of an actuarially fair insurance contract, any risk-averse individual would buy it. Indeed, private insurance markets in developed countries are large: the private insurance market accounts for about 9 percent of GDP across the OECD (OECD, 2020). By contrast, this market is much smaller in many developing countries. Insurance premium volumes represent 1.63% of GDP in low- and middle-income countries, compared to 4.63% in high-income countries (World Economic Forum, 2019). This section explores several possible explanations for low insurance *demand*, i.e., why many consumers in developing countries may not want to buy even actuarially fair insurance. These explanations include information, credit constraints, basis risk, and trust.

Information, trust, and experience. One challenge is that people may not fully understand insurance products, and therefore may be unwilling to purchase them. Buying insurance requires a large degree of trust: individuals need to pay money now in the hope that in the future, if they have a problem, the insurer will cover their loss (and not deny the claim). Sometimes a lack of trust is unwarranted, and comes from confusion over how the products work and when claims should be received. In other cases, however, profit-maximizing firms may have an incentive to limit claims, even if they are legitimate under the rules. To create more trust, many developed countries have regulation and legal processes to ensure that individuals know their rights and have mechanisms for recourse; even there, denial of insurance claims is still a common concern. These challenges may be much more severe in the developing world, where financial regulation and the legal system are much less developed. Some lack of trust may therefore be more rational on the part of citizens. For example, several papers show that lack of trust hinders insurance take-up by exploring how experience effects (either from oneself or one's network) can improve insurance take-up (Cole, Stein and Tobacman, 2014; Cai, de Janvry and Sadoulet, 2020).

It is worth noting that this problem is not limited to private insurance markets. People may not understand or trust insurance products from the government. Providing benefit information is often not sufficient if lack of trust is an issue (e.g., Dercon, Gunning and Zeitlin (2019); Banerjee et al. (2021)). In this case, it is possible that experience with insurance matters more than just information about government-provided insurance programs that require co-payments. Asuming, Kim and Sim (2021) and Banerjee et al. (2021) find that temporary subsidies can lead to longer-run insurance demand through experience.

Credit constraints and timing issues. Insurance premiums are typically paid upfront. If households are credit constrained—which many poor households in developing countries are—they may be less likely to take up insurance against income shocks if they need to pay premiums before income is realized. For example, in the context of insurance premiums for farmers in Kenya, Casaburi and Willis (2018) find that the take-up rate for pay-at-harvest insurance is 72 percent, compared to 5 percent for the standard pay-up-front contract.

Basis risk. Basis risk is another reason why demand for insurance may be low, particularly for agricultural insurance contracts (see Miranda and Farrin, 2012; Carter et al., 2017; Jensen and Barrett, 2017). Given private information concerns, instead of insuring against loss directly, agricultural insurance contracts are often written to pay out as a function of *predicted losses*. These predictions typically only account for exogenous determinants of risk (e.g., weather), and are therefore imperfect. The resulting 'basis risk'—i.e., the difference between real losses incurred and losses insured based on index values—can reduce demand for insurance products substantially.

Summing up. Experience, trust, credit constraints, and basis risk may all stifle demand for insurance products in developing countries, even if actuarially fair. There is some evidence that government-provided programs are also not immune to these issues, especially when they include co-pays. However, it is an open question whether governments can help solve issues that private insurers may not be able to fully overcome, either by providing social insurance directly or by working in conjunction with private providers to address the challenges and concerns that limit demand.

4.1.3 Insurance supply issued in developing countries

The literature on the supply challenges of insurance in low- and middle-income countries is less developed, but the problem can be characterized by two key factors. First, the limited information environments observed in developing countries could make it challenging to administer insurance products in practice. For example, in developed countries, in addition to weather-based insurance, agricultural insurance contracts are often written as a function of output or profits from nearby farms. This type of contract may have lower basis risk, but requires good information: the insurer needs to be able to observe the output or profits of all nearby farmers (perhaps through tax or other data collection). By contrast, there is little data on the output of most subsistence farmers, so these contracts are not practical to write in most developing country contexts.

A similar example comes from property: in the US, for example, 93 percent of homeowners have property insurance (Insurance Information Institute, 2020). Yet, the fraction of people

in developing countries who insure these risks is likely to be tiny. One reason why this may be different in lower-income countries may be a supply-side constraint: property insurers need to know the market value of a house to insure it. The less-formalized property markets in developing countries may make such valuation challenging.

Second, challenges in contract enforcement, the legal system, and the regulatory environment may also impede insurance provision. As described in Section 4.1.2, insurance requires trust—i.e., that the insurer will deliver on the contract and not renege if a claim is made. In developed countries, there is typically a regulatory role for the government to make sure insurers have sufficient capital to pay claims (and reinsurance to cover unexpected losses). There is also a functioning legal system through which people can sue if they feel they are wrongly denied a claim. These systems are substantially less developed in most developing countries.

4.1.4 Informal insurance

A third complication for formal insurance provision in developing countries, compared to the developed world, is how it interacts with complex, pre-existing informal insurance arrangements. Many poor households in developing countries—while not formally insured—engage in various forms of risk sharing. Indeed, villagers are better at smoothing idiosyncratic consumption shocks than one may expect given the absence of formal insurance markets (Townsend, 1994). However, from the perspective of a potential social insurance designer, it is important to note that the existence of informal insurance does not preclude the need for formal insurance schemes.

One key concern with relying on informal insurance mechanisms is that they are sustained in relatively closed networks (Ligon, Thomas and Worrall, 2002). As villages become more interconnected, or as the population increasingly lives in urban environments, these systems may become less effective over time (Townsend, 1995), suggesting that the need for formal insurance products may increase with development. Moreover, informal insurance systems can 'trap' people in low-opportunity areas, and there may be inefficiencies if households need to remain in rural areas in order to take advantage of these informal networks (Banerjee and Newman, 1998; Munshi and Rosenzweig, 2016).

Even when informal insurance exists, it does not necessarily provide the level of coverage that households need, since income shocks have a strong spatial correlation, especially in agricultural areas. Ideally, risk-pooling would be done across villages, where shocks are independent. However, such arrangements are difficult in practice, given that monitoring costs to prevent moral hazard are very high. Spatially correlated shocks may be particularly challenging in the context of natural disasters or conflict, necessitating large-scale disaster relief programs, as we discuss in Section 4.4.

All of these arguments suggest that some forms of formal insurance are likely to be useful, even where informal mechanisms are present. This then makes it necessary to think through the potential interactions between (formal) social insurance and informal insurance, especially given the evidence that as economies grow, many households may no longer be able to rely as much on informal networks. Ultimately, there are two key questions: first, whether

offering formal social insurance crowds out informal insurance; and second, what welfare consequences arise from the interaction between the two.

Regarding the first question, the evidence on crowd-out is mixed (Albarran and Attanasio, 2003; Strupat and Klohn, 2018; Takahashi, Barrett and Ikegami, 2019). The level of crowd-out may vary based on characteristics of the insurance: for example, while Huang and Zhang (2021) find no crowd-out of private transfers in the context of China's rural pension scheme, Jensen (2004) find large crowd-out effects of South Africa's public pension. Huang and Zhang argue that the difference may be driven by smaller pension benefits in China.

As for the second question, Arnott and Stiglitz (1991) theorize that whether the existence of formal and informal insurance systems is welfare-improving depends on how well informal networks monitor others to reduce moral hazard. They analyze two extreme cases: one with perfect monitoring and one with no monitoring at all. In the first case, welfare improves since monitoring by informal networks helps mitigate moral hazard and thus improves risk-sharing. At the other extreme, welfare decreases since formal insurers know that there is no monitoring and adjust premiums and payouts accordingly. Building on this analysis, Attanasio and Rıos-Rull (2000) show how even a well-intentioned policy could reduce welfare. The intuition behind their result is that the introduction of formal insurance ruptures the social fabric, leaving households more vulnerable to risks not covered by formal insurance.

4.2 Insurance against income loss

We next turn to understanding the challenges of insuring against specific types of risk. In this section, we covers risks of income loss. Section 4.3 covers expenditure risks.

4.2.1 Insurance against involuntary job loss

Unemployment insurance (UI) is designed to help people smooth consumption between jobs. However, a key challenge is that unemployment benefits typically continue until the worker resumes work. There is an extensive literature on the degree to which this conditioning of benefits on future employment discourages job search in high-income countries. This literature focuses on (1) the tradeoffs between the welfare gains from insurance and their potential disincentive effects, and (2) whether the benefits allow for better job matching (e.g., Chetty, 2006; Nekoei and Weber, 2017; Farooq, Kugler and Muratori, 2020). Regarding (1), a key question is whether policy design choices made within UI programs (i.e., the level and duration of benefits) can improve welfare. Examples of papers estimating these tradeoffs in low- and middle-income countries include van Ours and Vodopivec (2008), Huneeus, Leiva and Micco (2012), and González-Rozada and Ruffo (2023).[36]

While the same policy questions exist in many low- and middle-income countries, they are complicated by the existence of a large informal sector, which prevents the government from monitoring job entry and exit. This condition has two important implications. First, the government cannot provide benefits that begin conditional on unemployment for those in

[36] A number of recent papers also examine the spillover effects of unemployment insurance on domestic violence (Bhalotra et al., 2021) and crime (Britto, Pinotti and Sampaio, 2022).

the informal sector, so informal workers may not actually be covered by UI. Second, the government may be unable to condition the end of benefits on re-employment (since it only observes formal employment). Workers receiving UI could choose to seek employment in the informal sector in order to continue receiving their benefits. In this way, UI programs may disincentivize formalization, which would have consequences for productivity, taxation, and workplace safety.

Gerard and Gonzaga (2021) examine the relationship between informality and unemployment insurance in Brazil. They find, perhaps counterintuitively, that the presence of a large informal sector *reduces* the efficiency losses from moral hazard in UI. Indeed, they estimate that the efficiency costs of UI are 5 times lower in Brazil than in the United States. This difference arises because the presence of the informal sector allows workers to keep working (albeit informally) while retaining UI benefits, so while workers do lose formal protections from taking informal jobs, informal employment reduces the overall efficiency consequences of the UI tax on re-employment. Of course, informal employment comes at a cost, because informal jobs have lower earnings than formal ones. This finding is echoed in Liepmann and Pignatti (2024)'s study in Mauritius.

Given these challenges, one often observes alternative policies to UI to help to insure workers against job loss. A common policy is mandated severance pay: lump-sum payments disbursed upon termination of a labor relationship. These payments are not conditioned on future employment, and thus do not distort future employment decisions.

Severance programs also have challenges. First, someone needs to verify that workers were, indeed, working, and that they were terminated. If firms are paying directly, the government needs to ensure that they do; if the government collects unemployment taxes and pays the severance, it needs to collect the taxes to do so. Weak institutions may exacerbate these enforcement challenges. For example, Sadka, Seira and Woodruff (2024) find that in Mexico, despite a strong de jure severance policy, many workers do not receive their full entitlement, and delays and misinformation abound in the court system. The authors show that improving institutional quality could help: providing information about likely court outcomes substantially shortens settlement times and improves workers' food security.

In addition, the need to make large severance payments in the event of termination may disincentivize firms from hiring, and a single lump-sum means that workers bear the risk of being unemployed for longer than average. Using data from São Paulo, Brazil, Gerard and Naritomi (2021) show that workers appear to over-spend from lump-sum severance payments relative to their optimal consumption profile. They suggest that a stream of payments (which could, of course, be unconditional) may lead to better consumption smoothing than a lump-sum.

Another alternative is to create unemployment savings accounts that workers could access in case of job loss, funded by mandatory contributions by workers and/or firms (Feldstein and Altman, 2007). The idea is that individual accounts could help align incentives and reduce the scope for moral hazard. These accounts would not have the risk pooling features of traditional insurance; they also have the challenge that workers who only just starting paying in may not have sufficient balances to cover an unemployment shock (though governments

could in principle supplement accounts with too little money, guaranteeing a minimum balance). Countries such as Colombia, Chile, Indonesia, and Mexico have implemented this type of system. Kugler (2005) and Nagler (2013) examine how the shift from severance to these types of linked accounts in Colombia and Chile, respectively, impacted wages and job duration.

In short, the challenges of a large informal sector mean that the design of insurance programs for unemployment in developing countries may need to be fundamentally different from those in high-income ones. These challenges also suggest avenues for future research, particularly on policies that are not conditioned on re-employment (e.g., individual unemployment-triggered savings accounts, severance, and lump-sum unemployment insurance paid by the government). Moreover, given information constraints, countries are also designing self-targeting techniques that condition benefits on costly (but potentially productive) activities. For example, Indonesia requires job training to receive benefits. Understanding the impacts of these policies on both insurance and ultimately labor market activities is important.

4.2.2 Insurance against Disability and Death

Disability insurance is designed to help insure workers against accidents or illnesses that may either temporarily or permanently remove them from the labor market. However, similar to UI, the large informal sector hinders governments from providing universal disability coverage through employers. Therefore, one often observes alternative policies to provide transfers to the disabled. For example, many cash transfer programs, such as Argentina's *Programa de Ciudadanía Porteña*, Chile's *Subsidio Único Familiar*, Indonesia's *Program Keluarga Harapa* and Tanzania's *Productive Social Safety Net*, give weight to having a disabled household member in the eligibility criteria. This has two challenges: (1) it only provides insurance for those near the poverty threshold (not the middle class) and (2) if targeting is done infrequently, the sign-up period may not align with the need for assistance.

A related issue is workplace accident insurance. This is often separate from disability insurance, in part due to attempts to link risks back to employers. While some middle-income countries are trying to provide this type of insurance to workers in the formal sector, this topic is comparatively underexplored in the economics literature.

Finally, we touch briefly on life insurance. The death of a primary income earner often implies a substantial income loss. Indeed, in many developing countries, widowhood is closely associated with poverty. Formal life insurance markets tend to be much smaller (as a share of GDP) in low- and middle-income countries compared to high-income countries, suggesting that much of this risk remains uninsured. Moreover, the relative paucity of formal pensions—which often come with survivors' benefits to insure spouses against the loss of income from their primary earner in high-income countries—leaves more of this risk uninsured. As with disability, widow status is often considered in PMT formulas, but again, this only provides insurance to the extent that households are near the poverty threshold.

4.2.3 Insurance against agricultural loss

Agriculture plays a large role in low- and middle-income countries. But, agriculture is risky: crops can fail and prices are volatile, so those engaged in agriculture as their primary occupation face much more income risk than those who work in manufacturing or other sectors. These risks are not covered by traditional unemployment insurance schemes.

For these reasons, there has been a big push among governments to develop programs that help farmers manage agricultural risk. Some programs involve providing technologies that ensure more stable yields and prevent crop loss, be it investments in irrigation or weather-resistant seeds. Other policies aim to provide agricultural insurance to help farmers smooth consumption during periods of crop loss. There are a number of excellent reviews on these topics; see, for example, Cole and Xiong (2017), Ali, Abdulai and Mishra (2020), and Nshakira-Rukundo, Kamau and Baumüller (2021).

4.2.4 Parental benefits

Parental benefits are designed to help insure families against income loss associated with pregnancy and the early period of a child's life, when parents reduce labor supply to care for the newborn child. These benefits also protect parents from employment termination during this period of life. One common type of parental benefit provides a mechanism in which workers can take time off around pregnancy and a child's birth, by mandating paid time off and/or subsidizing firms for the worker's time off. A second benefit takes the form of subsidizing child care to allow parents to re-enter or stay in the workforce.

While virtually all countries (all but seven) have some type of maternity leave policy (Iqbal, 2018), benefit levels in developing countries are often very low for mothers, and paternity leave is often nonexistent. Many of those in the large informal sector do not have access to any formal benefits.

In contrast to the extensive literature on these policies in high-income countries, there is relatively little micro-empirical research on paid maternity or paternity leave in lower- and middle-income countries. There are, however, a few notable exceptions. In Chile, Albagli and Rau (2019) show that a doubling of maternity leave (from 12 to 24 weeks) led women to take substantially more time off after childbirth, increased breastfeeding duration, and improved cognitive outcomes. Interestingly, women with the enhanced leave policy were also more likely to be in the labor force one year later. In Vietnam, Vu and Glewwe (2022) show that a more generous maternity leave shifted potentially-eligible women from informal work to formal work, and, in particular, to public sector jobs where they could access benefits. More research is needed to understand how pregnancy affects both consumption smoothing and women's career trajectories, and how effective leave policies can be designed given the presence of informal labor markets.

In contrast, free or subsidized child care has become increasingly common, particularly in middle-income countries. A number of studies have tried to examine its impacts. For example, Calderón (2014) shows that access to childcare in Mexico increased women's likelihood of working and reduced their likelihood of earning zero income. Similarly, Halim, Johnson and

Perova (2022) examine public pre-school expansions in Indonesia and find positive effects on women's unpaid family work. They argue that limited day care hours made it difficult for women to find jobs outside their homes.

Recent experimental studies (Paes de Barros et al., 2011; Martínez and Perticará, 2017; Clark et al., 2019) have also studied access to childcare in Brazil, Chile, and Kenya, respectively, with all three finding positive impacts of childcare access on women's work. In fact, Bjorvatn et al. (2022) find that income gains from childcare led to at least as large a gain as an equivalent cash transfer, while also improving child development. Interestingly, they find that the earnings increase comes from higher productivity, largely in self-employment, rather than longer hours.

There are still many open questions in this area. For example, how should parental benefits be paid and who should pay for them? Should they be universal or dependent on labor status, especially in countries where the prevalence of informal firms may make it hard to mandate benefits? How should countries think about provision of child care as larger extended family networks—which often provided this type of child care—break down? And more broadly, what are the fertility impacts of changing parental support, particularly as many of these countries are undergoing demographic transitions with rapidly falling fertility rates?

4.3 Insurance against expenditure loss

The other type of risk is expenditure risk—i.e., the risk of unexpected large outlays. In this section, we outline common social insurance programs that aim to alleviate these risks.

4.3.1 Health insurance

Government-led health insurance systems are becoming increasingly common in many developing countries, as governments aim to increase health care utilization, improve health outcomes, and help households manage income and consumption risks that may arise from health shocks. Nearly 190 countries have some sort of public health insurance system.

The literature on health insurance in low- and middle-income countries tends to fall into two buckets. The first bucket examines the impact of health insurance on insured households. This includes their ability to smooth out the economic consequences of health shocks, as well as their health care utilization and health outcomes. The second bucket explores the challenges of insurance design and provision; in particular, how to design and fund public insurance systems given the constraints of informal employment, adverse selection, moral hazard, information failures, and trust. See also Das and Do (2023) for a recent review of this literature.

Impacts of health insurance. It is worth noting that conceptually, the primary impact one would expect from health insurance is on consumption smoothing. That is, health insurance is primarily a *financial* product that provides payment in the event of an expenditure shock, and should therefore make the financial consequences of poor health less severe (Finkelstein and Mcknight, 2008). Health insurance can also affect health care consumption; since health insurance typically reduces the marginal cost of health care utilization (as opposed to, say,

lump-sum cash transfers, which would mitigate the financial consequences of illness without changing prices on the margin), one would expect utilization to increase. This increased health care utilization could, in turn, potentially affect health outcomes.

Health insurance and financial shocks. A number of papers from low- and middle-income countries suggest that health insurance does indeed provide meaningful financial insurance, and can help households manage financial shocks from health events (King et al., 2009; Levine, Polimeni and Ramage, 2016; Gruber, Lin and Yi, 2023). In fact, del Valle (2021) argues that health insurance provides another additional economic benefits: they show that, by reducing the severity of health shocks, Seguro Popular in Mexico reduced the degree to which other household members needed to drop out of the labor force to directly provide care. Other work provides suggestive evidence that health insurance may reduce financial stress above and beyond its impact on finances per se (Haushofer et al., 2020).

Health care utilization and health. Health insurance also has the potential to change health care utilization, since it can reduce the price of health care, which could in turn generate health impacts. In some cases, one may actually be concerned that this leads to *over* consumption of health care since consumers do not face the true marginal cost of the care, and indeed, in some developed countries, alternative schemes with high deductibles have been developed in an attempt to provide financial insurance while ensuring that households face the correct price of care on the margin. On the flip side, if households are liquidity-constrained (as we believe many households in developing countries are), households may under-consume health care without insurance. Understanding how to help households achieve their optimal, distortion-free level of health consumption remains an important challenge.

Does health insurance actually affect health consumption, and health, in developing countries? The experimental evidence here is mixed, though one challenge is that some of the experimental studies may not be sufficiently powered to detect small, but economically meaningful, effects (King et al., 2009; Haushofer et al., 2020). A recent study by Malani et al. (2021) highlights another important policy issue: while the authors find some increase in insurance usage in their experiment, many beneficiaries had challenges using their insurance. For example, they often had problems with their insurance cards, did not know how to use them, or forgot them. Perhaps unsurprisingly, the authors observe no health effects. In short, most of the experimental studies to date find little overall impact of health insurance on health outcomes.

Given the need for large samples sizes to measure impact, a number of other studies focus on quasi-experimental variation from large-scale public health insurance reforms, using large administrative datasets on mortality and other health outcomes. For example, Gruber, Lin and Yi (2023) examine the roll-out of China's public health insurance for rural households, which covered as many as 800 million people. They find a significant decline in aggregate mortality, which they argue explains 78% of the entire increase in life expectancy in China during this period. Using survey data, they find large effects on health care utilization and a host of other health outcomes.

Gruber, Hendren and Townsend (2014) study an alternative method of health insurance: the expansion of free or heavily-subsidized care at public facilities. They examine Thailand's

2001 health care reform, known as the "30 Baht" program. Prior to this reform, the poor, young and old were given free public health care, but the system was seen as chronically underfunded. The reform changed the system in two ways. First, it provided universal access to public facilities at a co-pay of 30 Baht (about US$0.75) per visit. Since the co-pay was already waived for the poor, in practice this reform largely led to reduced health access costs for the informal, non-poor. Second, the reform provided hospitals with a universal capitation payment based on their province's population, leading to more generous government financing on net. Using administrative morality records, the authors show that prior to the reform, infant mortality rates were related to provinces' wealth. After the reform, however, resources were equalized across provinces, leading this correlation to disappear.

Importantly, both Gruber, Lin and Yi (2023) and Gruber, Hendren and Townsend (2014) argue that part of why large health effects were seen is that insurance reforms, by increasing government financing for health, increased the supply of health services, as has been shown in developed countries (e.g., Finkelstein, 2007). They argue that this increase in supply helps explain the large health effects of insurance reforms. In evaluating health insurance systems experimentally, it may therefore be important to randomize across hospital catchment areas or health markets, rather than individuals or villages. Such designs would allow for estimation of the full general equilibrium impacts that could arise from insurance expansions or enhancements.

Challenges with health insurance design. The second bucket of work on health insurance centers around how to design and fund health insurance systems. For example, many countries choose not to universally cover health insurance premiums through the government budget, and instead set up national health insurance programs where the poor are covered directly and everyone else has to pay a mandatory contribution. Contributions are often collected for formal workers through payroll taxes remitted by employers.

A common challenge with these systems, however, is how to handle informal workers (for whom premia cannot be collected from employers). There are, broadly speaking, three options. First, the government could opt not to provide health insurance to non-poor, informal workers. This, of course, would mean that these workers would remain vulnerable to the economic impacts of health emergencies. Moreover, if firms wanted to evade paying benefits to formal workers and/or workers that do not fully value insurance, mandating insurance benefits only for formal workers could increase informality. Second, one could mandate that non-poor, informal workers contribute to insurance. This would be difficult to enforce, however, particularly in countries with limited administrative capacity. Moreover, the resulting adverse selection problems (i.e., only signing up when sick) could financially strain the insurance system. Third, one could decide that it is too hard to collect contributions from informal workers, and instead extend free insurance to most, if not all, informal workers. Once again, concerns about encouraging informality would arise, as well as cost considerations.

Below, we first discuss the evidence on the interplay between health insurance and informality. We then discuss the evidence on adverse selection, as well as policy tools that have been used to mitigate it. We conclude with open research questions.

Health insurance and informality. Ex-ante, the effect of employer-provided (or subsidized)

health insurance for formal-sector works on formality itself is ambiguous. If a) workers value the insurance and b) there is a cost savings (from the workers' perspective) of gaining insurance by becoming a formal employee, as opposed to paying an individual premium, this could increase formalization. But, if workers do not value the insurance at cost or if alternative health care arrangements are cheaper, employer-provided insurance could increase informality.

A number of papers examine these concepts. For example, Bergolo and Cruces (2014) examine a large-scale policy reform in Uruguay's social insurance administration (SIA) that increased benefits and contributions for formal workers, and find evidence of both effects. The reform extended the SIA's coverage for dependent children, making the program more attractive. At the same time, the reform increased the payroll tax contribution deducted from employees' salaried earnings, which could have caused an increase in underreporting of wages for formal employees. The authors find both effects: the greater benefits drew people into the system, but there was also an increase in wage misreporting. On net, the fiscal revenue gain from higher levels of formal employees was much larger than the loss of revenue due to under-reporting.

Camacho, Conover and Hoyos (2014) examine the introduction of subsidized insurance for informal workers on labor market choices in Colombia. A reform was instituted that made workers who were below a Poverty Index Score eligible for non-contributory health insurance, but those who were formally employed were ineligible regardless of their score (and needed to contribute through their employers), creating an incentive to become informal. They find that the reform increased informal employment by about 3-4 percentage points.

Finally, a number of papers have examined the effect of Mexico's Seguro Popular program on formalization. Prior to the rollout of Seguro Popular, insurance was tied to payroll contributions, and many people were uncovered. Seguro Popular aimed to provide universal coverage. Analyzing the staggered roll-out of the program and survey data, Aterido, Hallward-Driemeier and Carmen (2011) and del Valle (2021), among others, find small effects on formalization, while Azuara and Marinescu (2013) find no effect on average, but small effects for unskilled workers. More recently, examining the roll-out with social security data, Bosch and Campos-Vazquez (2014) show that it slowed the social security registration of employers and employees in small and medium firms (up to 50 employees), reducing revenues paid into the social security system as well as VAT payments.

Adverse Selection, Information and Trust with Imperfectly Enforced Mandates. Many countries have tried to mandate that non-formal, non-poor workers purchase insurance, but limited enforceability remains a challenge. Several studies have examined what can be done to mitigate this issue.

Subsidies. Can time-limited subsidies substitute for unenforceable mandates? In Ghana's health insurance scheme, Asuming (2013) and subsequently Asuming, Kim and Sim (2021) randomize partial and full subsidies of premiums for one year, and track the results for three years. Subsidies increase take-up, an effect that persists even after the subsidies end.

Banerjee et al. (2021) also study time-limited subsidies in the context of Indonesia's national health insurance, randomizing offers of partial and full premium subsidies (and a control) for

one year. Larger subsidies bring healthier individuals into the health care system, consistent with undoing adverse selection. They also document a pattern of dynamic selection, where those who enroll in the no-subsidy condition are much more likely to immediately file large claims. Since the subsidies attract healthier individuals who then also pay premia in the post-subsidy period, they allow the government to cover more people at the same total cost. Fischer, Frölich and Landmann (2023) in Pakistan also experimentally vary premia, and find that adverse selection is higher when premia are higher. In short, while subsidies are by no means a panacea, they can help ameliorate adverse selection.[37]

Bundling. A second approach to reduce adverse selection is bundling: by tying the purchase of health insurance to that of another product, or bundling insurance purchases for a household together, households' ability to buy insurance only for the sick can potentially be limited. For example, in their study with an NGO in Pakistan, Fischer, Frölich and Landmann (2023) experimentally vary whether people can choose to enroll individuals, or whether they must enroll the entire household, under the idea that households will have to enroll all members (not just the sick). They also examine a community bundled contract, where at least fifty percent of the community must sign up in order to activate the insurance. These bundled contracts reduce "expected costs" among those who sign up (where expected costs are claims predicted from baseline covariates), suggesting that this type of group insurance may be effective. In many developed countries, workplaces are the 'group' for health insurance purposes; this study suggests other groups as an alternative in countries with high levels of informality.

The downside of bundling, however, is that if demand for insurance is low, it can reduce demand for other bundled products. For example, Banerjee, Duflo and Hornbeck (2018) find that a substantial proportion of people were apparently willing to forgo renewing their microcredit just to avoid purchasing the required health insurance bundle.

Information and Trust. Lack of information (or misinformation) about the benefits of insurance and limited trust in the system (some warranted and some unwarranted) may be another reason for low insurance demand, as discussed in Section 4.1.2. For information provision, the evidence is mixed, with Giles et al. (2021) and Asuming (2013) finding positive impacts, for example, and Banerjee et al. (2021), and Dercon, Gunning and Zeitlin (2019) finding none.

Dercon, Gunning and Zeitlin (2019) measure trust using a trust game at baseline, and then offer composite health insurance to individuals with varying subsidies. They find that low generalized trust is negatively associated with insurance demand, and that the purchase decisions of individuals with low trust are significantly sensitive to price.

In sum, despite the benefits of health insurance for risk smoothing, health insurance provision in developing countries remains a challenge. The combination of lack of demand for insurance, adverse selection, a large informal sector, and the relative challenges of enforcing a mandate to purchase insurance means that it is challenging to ensure widespread insurance coverage. Many countries respond to these issues through public-sector systems with low prices, with those who want better coverage opting out at their own expense. However,

[37]Two other important papers on subsides are Thornton et al. (2010) and Wagstaff et al. (2016).

Figure 4: Percentage of countries that enacted non-contributory pension systems by income classification over time

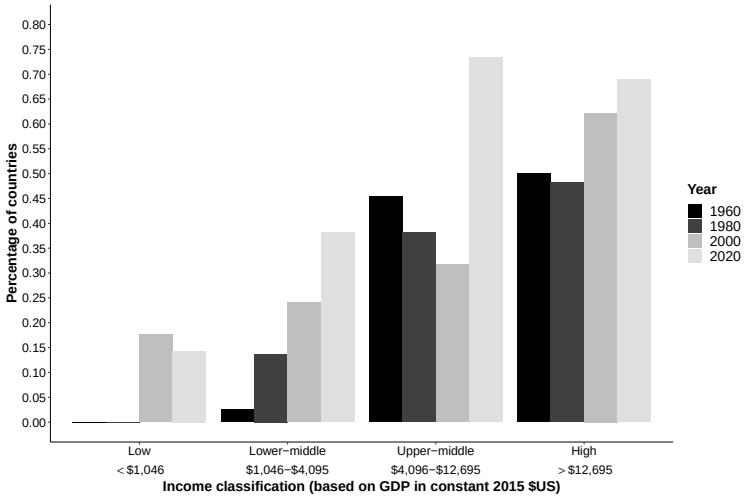

Notes: This figure graphs the percentage of countries that enacted non-contributory pension systems by each year considered (1960, 1980, 2000, 2020) against quantile of GDP per capita (in constant 2015 $US). GDP per capita quantiles are calculated as an absolute scale across the time periods considered. *Source:* GDP per capita data is from The World Bank (2021*a*).

this approach leads to substantial uninsured risk. Working out how to move towards more comprehensive coverage remains an important direction for future research.

4.3.2 Pensions and annuities

Old age is often associated with poverty. Health typically falters as people age, so the elderly are less able to work; this is particularly true for those engaged in physical labor. Saving (or saving enough) may be difficult, particularly for those outside the formal sector. Even for those who do save, aging entails risk: in particular, the risk that one will outlive one's savings. Traditionally, these challenges have been borne by family members (e.g., children caring for parents), but these networks may be incomplete and may break down with urbanization.

Accordingly, many governments around the world have public pension programs with a mix of systems. Indeed, 181 countries have some type of contributory pension system, 102 have non-contributory systems, and 96 have both (see Appendix Table 2). These programs typically have three related goals: a) to help individuals save for their old age, either individually or through a tax-and-transfer scheme; b) to provide annuities that insure against the risk of living too long as compared to one's savings; and c) to provide some amount of redistribution so that even those with low incomes are not too poor in old age.

As seen in Figure 4, governments often develop these systems as incomes rise, perhaps coincident with the rise of more modern economies and increased life expectancies. That said, Figure 4 reveals that even conditional on real GDP per capita, these programs are becoming more prominent in low- and middle-income countries.

A substantial literature has studied the effects of pension transfers on recipient households,

primarily from the perspective of labor supply and consumption choices. The literature on the *design* of these programs is less developed. We review each of these literatures in turn.

Effects on consumption and well-being. One strand of research is focused on understanding the impact of pensions on beneficaries' well-being. For example, examining the roll-out of China's New Rural Pension Scheme, Huang and Zhang (2021) find that the program increased income while reducing mortality. Similarly, Galiani, Gertler and Bando (2016) evaluate the impact of Mexico's *Adultos Mayores* Program, a non-contributory cash transfer provided to the elderly, and find increases in consumption and reductions in depression.

Consumption effects need not be limited to pensioners themselves, as consumption resources are shared within the household. Duflo (2003), for example, shows that pensions received by older women in South Africa improve the anthropometrics of female children in the household.[38] Studying the same program, Edmonds (2006) finds increases in schooling and declines in child labor for boys, as well as declines in domestic labor for girls, when men in the household become eligible. Similarly, de Carvalho Filho (2012) studies an unexpected pension reform in Brazil and shows that pension eligibility leads to an increase in school enrollment for girls. Combined, these papers suggest important inter-generational spillovers from pension receipt within the household.

This set of papers focuses on what happens when income starts flowing. Bau (2021), however, looks at ex-ante child investments. She argues, using data from Indonesia and Ghana, that the introduction of public pensions make parents less likely to invest in children's education, since they no longer need to rely on their children for old-age support.

Labor supply effects. A second set of papers examines the impact of pension receipt on labor supply and/or retirement decisions. To the extent that these are income effects (as opposed to price effects induced by features of the benefit formula), labor supply reductions associated with pension receipt are likely to be welfare-improving, as people appear to have been working longer than they otherwise would have in the absence of a pension.

Several papers find that pensions reduce labor supply. de Carvalho Filho (2012) studies the reform to the Brazilian Social Security System discussed above, and shows that it reduced the retirement age of rural men. This is an example of a pure income effect, as the transfer was not means or retirement tested. Galiani, Gertler and Bando (2016) study Mexico's *Adultos Mayores* Program and also find that paid work declines, but that those who stop doing paid work switch to family businesses. Huang and Zhang (2021) show that in China, pensions decreased labor supply for rural people older than 60 (i.e., the eligible), particularly for farm-work.[39]

Program design. There is comparatively less work, however, on program design questions in

[38] Case and Deaton (1998) provide a rich description of South Africa's non-contributory pension program.

[39] The income effects from a pension could also affect the labor supply decisions of other members. Studying South Africa's pension program Bertrand, Mullainathan and Miller (2003) find a sharp drop in the working hours of prime age men in the household, particularly when the pensioner is a woman. Posel, Fairburn and Lund (2006) argue that the pensions also drive migration among prime age workers. Including the migrants in the analysis, Ardington, Case and Hosegood (2009) find small positive increases in work among prime-age adults.

low- and middle-income countries, but we highlight several papers that suggest some ways in which low- and middle-income contexts may be somewhat different from higher-income contexts.

One important question concerns contributions: do firms report the correct social security contributions, and pay taxes accordingly, or do they under-report? Kumler, Verhoogen and Frías (2020) show that system design may matter: when the pension system became largely contributory, payments for younger workers (who had more of an incentive to ensure their wages were reported accurately) increased.

A second question concerns investment choices. Some contributory systems allow individuals to choose how their assets are invested. Hastings, Hortaçsu and Syverson (2017) document very little attention to fees among plan participants in Mexico, so that fees charged were extremely high. Combined, these fees meant that a 100-peso deposit that earned a 5% annual return would be worth only 95.4 pesos after 5 years. They document that a key constraint was that workers were not particularly price sensitive, so firms primarily competed on non-price attributes. This suggests that it can be challenging to structure market incentives so that competition leads to low prices, and that care must be paid in structuring private management of individual account systems.

Third, there are questions about the interactions of program design with savings incentives. An individual contributory system encourages private savings, but is less redistributive. A country may therefore want to add a substantial component to cover the informal system, i.e., a minimum pension floor. The challenge is to do this without discouraging savings. Attanasio, Meghir and Otero (2014), for example, document these tradeoffs in the context of Chile's 2008 reform, which introduced a minimum pension floor.

Fourth, there are questions of take-up. Many pension programs suffer from the take-up challenges described above for cash or in-kind transfer programs. Gupta (2017) experimentally shows how administrative challenges hinder the take-up of a widow's pension program in India, and that improving this process can increase the provision of benefits.

These challenges, however, are only the tip of the iceberg, and there is substantially more work to be done. A broad question concerns the informal sector: as with health insurance, if pensions are introduced in the formal sector and funded through payroll taxes, will this create a further wedge between formal and informal sectors? How does the presence of the large informal sector interact with the decision of whether pension systems should be contributory or non-contributory? And as these systems grow over time, what will the implications be for the many economic decisions—marriage, investments in children, savings—that are related to how people plan for their old age?

4.3.3 Funeral insurance

Funerals are often a large, not always foreseen cost that can devastate households, particularly in bad times. For example, analyzing data from South Africa, Case et al. (2013) finds that households spend about a year's income for an adult funeral. Funeral insurance can help manage these risks, whether through community associations or burial clubs (Dercon

et al., 2006; Case et al., 2013; Berg, 2018), religious institutions (Auriol et al., 2020), or private insurance companies (Berg, 2018). In the Case et al. (2013) study of South Africa, about 28 percent of deaths had some form of funeral insurance (either through burial clubs, funeral parlors, or private insurance). However, many more needed to rely on savings, and/or borrow (about a quarter) to help cover the costs.

As Dercon et al. (2006) discuss, many of these informal insurance or burial clubs have sets of rules and institutions to promote fairness. Nonetheless, an important question is whether certain types of households can be excluded and whether there is a role for government-provided funeral insurance to complement many of these informal structures.

4.3.4 Property damage

Risk from property damage poses a challenge, whether it be from flooding, fire, crime, or other perils. In wealthy countries, private insurance often covers these idiosyncratic risks: in the US, for example, over 90 percent of homeowners have property insurance (Insurance Information Institute, 2020). For certain types of perils where there are aggregate shocks, coverage remains low. However, the fraction of people in developing countries who insure these risks is tiny despite their important effects (Anttila-Hughes and Hsiang, 2013).

One potential reason is that adverse selection and moral hazard may be a challenge. For example, in the United States, private property insurers maintain the "CLUE" database of claims, so that insurers can price idiosyncratic risk into future premia. We are not aware of similar systems in most developing country contexts.[40]

Moreover, like other forms of insurance, low demand often prevails due to lack of trust in whether institutions will pay out (Reynaud, Nguyen and Aubert, 2018). More generally, understanding why this market is limited seems important for future research.

4.4 Natural disasters and climate-related shocks

Natural disaster and climate-related shocks, such as floods, earthquakes, droughts, and hurricanes, pose their own set of challenges. These shocks tend to be aggregate in nature, affecting a large number of people in a particular geographic area at once. This means that the informal insurance networks discussed in Section 4.1.4, which provide substantial risk-smoothing in low- and middle-income countries more generally, are not able to deal with these risks well.[41] In the absence of well-developed re-insurance markets, the private sector in low- and middle-income countries may have challenges as well.[42]

Many of these shocks are simultaneously income and expenditure shocks, destroying livelihoods as well as property. They can therefore have longer-term effects (see, for example,

[40] An exception is TransUnion South Africa's similar Insurance Claims and Policy database.

[41] While local risk-pooling networks will be ineffective, distant-and indeed even international-risk-pooling networks can respond to these types of shocks, i.e., through changes in remittances or in migration; see Yang (2008) and Mahajan and Yang (2020), respectively, for evidence on these channels.

[42] Note that there is also quite substantial international efforts as well, such as the Caribbean Catastrophe Risk Insurance Facility and Africa Risk Capacity.

Baez and Santos, 2007; Anttila-Hughes and Hsiang, 2013; Cas et al., 2014).

Ex-post emergency relief. After these shocks occur, governments tend to intervene and provide some amount of emergency relief to households, either cash or in-kind, though the degree to which they do so varies.

There is some evidence that this type of compensation can make an important difference. For example, Gignoux and Menéndez (2016) study earthquakes in Indonesia, and find that households are harmed by these shocks, but that the effects dissipate after 2-5 years. They find that government aid flows following the earthquakes may have helped mitigate the effects. Likewise, Laurito, Frankenberg and Thomas (2022) find substantial positive effects of housing assistance in Indonesia following the 2004 Indian Ocean tsunami.

However, there is evidence that if aid is not timely enough, children—who may be in a critical window when the shock occurs—can suffer long-term consequences. For example, Andrabi, Daniels and Das (2023) find that, while adults affected by the Kashmir earthquake suffered no change 4 years later (perhaps because they received transfers equal to 150 percent of their annual consumption), there were persistent shortfalls in the physical and cognitive development of children who were young at the time of the earthquake. Likewise, Dercon and Porter (2014) find long-run negative consequences for child survivors of the 1984 Ethiopian famine, despite the substantial amount of aid delivered.

Ex-ante vs. ex-post assistance. These results underscore a key challenge with ex-post emergency aid: delays in the time it takes for aid to arrive. To make this process smoother, programs can be set up in advance and triggered based on pre-specified thresholds. For example, del Valle, de Janvry and Sadoulet (2020) study Mexico's Fonden program, which provides transfers to municipalities when rainfall exceeds a pre-specified threshold. Using a regression discontinuity design based on the rainfall cutoff, they show that insurance payments lead to substantially more economic activity as measured by night-lights. Subsequent work shows that this program has also reduced mortality by restoring access to public health infrastructure more quickly (del Valle, Forthcoming). Clarke and Dercon (2016) also discuss Fonden and other programs that similarly provide index-based drought insurance (e.g., a program for pastoralists in Kenya).

More generally, there is increasing policy discussion around the idea of codifying, ex-ante, that there would be expansions in the types of social protection programs discussed in Section 3 in the event of shocks (see O'Brien et al., 2018, for a discussion). For example, the government of Kenya explicitly relaxes eligibility rules for its food security program in times of drought to make it more responsive (Gardner et al., 2017). All households in the four counties covered by the Hunger Safety Net Program were registered during the expansion of this program, regardless of eligibility status, to facilitate efforts to rapidly scale up the transfer during emergencies. This has allowed 50% to 75% of the population in these counties to receive cash transfers during droughts.

But this program was unusual in that all households were registered ex-ante, in order to facilitate expansion during times of shocks. More generally, programs must be able to rapidly identify who needs assistance because of the shock, accounting for the fact that who needs assistance may have changed since the registration period due to the shock itself.

Given the chance that governments may provide assistance ex-post (as it is difficult to stand by and do nothing), households may be reluctant to purchase insurance themselves, a classic case of a Samaritan's Dilemma. If governments anticipate the need to bail out households ex-post, they may be interested in formalizing this commitment, allowing them to recoup some of their costs via mandatory premiums. One of the few examples we know of in low- and middle-income countries is the Turkish compulsory earthquake insurance scheme, established in 2000 (Natural Disaster Insurance Institution, 2022).

A new frontier in this type of assistance is to not just commit ex-ante to deliver assistance, but to actually deliver assistance *before* shocks actually occur. A new wave of rapid detection technologies mean that for some predictable shocks (e.g., floods in the downstream part of a river), aid could be delivered before the shock is realized, allowing households to take preventative action. Pople et al. (2021) examine a pilot study of this approach in Bangladesh and find that such ex-ante assistance can make a difference. But, doing so could mean there is a risk of delivering aid to those who do not end up experiencing the shock; that is, the *realization* of the shock is also a targeting device.

4.5 Summing up, and broader implications for social protection system design

Life in developing countries is risky. While insuring against risk is a challenge everywhere, the particular challenges in low- and middle-income countries highlighted above—low insurance demand, low trust, liquidity constraints, challenges of state verification, and informality—mean that even if individuals were offered contracts priced fairly for the entire population, take-up may be low.

Moreover, the existence of these substantial uninsured risks has broader implications for the design of social protection systems as a whole. Recall the poverty traps example discussed in Section 3.2.6. A further implication of the poverty trap model is that for households in the 'good' equilibrium, an uninsured negative shock can push them across the threshold into the bad equilibrium, where they may spiral down into further poverty. This implies that solving the insurance challenges discussed here may have spillover effects, reducing the number of long-run poor who need the redistributive programs discussed in Section 3. Indeed, Ikegami et al. (2018) and Janzen et al. (2021) argue that, for this reason, governments may be better off diverting some of their standard targeted social protection funds to providing additional social insurance. Even to the extent that these programs do not entirely mitigate financial shocks, conditionality can direct households to smooth them in ways that are less socially costly, such as by encouraging them to keep children in school (de Janvry et al., 2006). This approach can have protective inter-generational effects, preventing a shock in one generation from creating a long-run cycle of poverty.

Conversely, to the extent that targeting can be made more dynamic, the same types of programs discussed in Section 3 can provide some insurance. A challenge with many of the targeting approaches discussed in Section 2 is that the targeting list is updated infrequently. When targeting is done, it often focuses on the permanent component of income (e.g., assets), making it unresponsive to shocks.

This issue can be remedied in several ways. At the individual level, to the extent that targeting can be made more responsive (for example, by using some combination of high-frequency administrative data like electricity and mobile phone usage, on-demand applications, and community-driven approaches to verify shocks), the same programs that provide assistance to the poor can *also* provide assistance to those who receive shocks. This can help fill some of the gaps in insurance highlighted in this section.

At the aggregate level, as discussed in Section 4.4, governments can use the types of programs discussed in Section 3 to respond to aggregate shocks by changing eligibility thresholds or by increasing transfer amounts. During the COVID-19 crisis, when this was an extreme issue, many countries responded in creative ways to adapt their existing programs (many of which were of the type discussed in Section 3) to address these issues. Indonesia, for example, both relaxed eligibility rules and expanded transfer amounts in its food voucher program. It also created a new community-targeted cash transfer program in all rural villages. Likewise, Pakistan built on its flagship cash transfer program for women, expanding eligibility and using cell-phone metadata to help determine eligibility (Lone, Shakeel and Bischler, 2021). These examples were forged rapidly in a crisis. Presumably, with advance planning, governments can do much more to make their existing redistribution programs function to smooth shocks.

5 Conclusion

Social protection programs are becoming increasingly prominent in low- and middle-income countries worldwide. This is partly due to the fact that countries are becoming richer, so some countries that were recently quite poor are now middle-income, and with substantially larger tax bases that can fund more sophisticated systems. It also reflects the fact that, conditional on a country's income level, countries are more likely to have systems in place that provide assistance to the poor, and social insurance of a variety of types, than they were fifty or even twenty years ago.

But the fundamental nature of these countries' economies—e.g., the large informal sector, and the deep absolute poverty level—means that the way these programs are designed fundamentally differs from how they are designed in high-income countries.

One important difference is identifying beneficiaries. In high-income countries, governments typically choose beneficiaries using income-based criteria. They can do this because income data is reported through a third party, the tax system. With the large informal sector, this is not possible in lower-income countries. Instead, countries face a tradeoff between using noisy proxies for income and using potentially biased self-reports. We discuss how to think about this tradeoff theoretically, and then review the evidence on the tradeoffs between these approaches in practice.

Once one figures out who to provide assistance to, what should the form of assistance look like? Again, the particular context of low- and middle-income countries has influenced how these programs look. For example, given the relatively low levels of human capital investment, conditional cash transfers—which tie cash payments to ensuring that family members of recipients meet a set of health and educational conditions—are far more common

in low- and middle-income countries than in high-income countries. Indeed, these programs have been shown to have substantial effects on the incentivized behaviors. Given how poor the poor can be, there is also a focus on finding program designs that can help break poverty traps, whether within a single generation or across generations.

Finally, the substantial informal sector has important implications for the design of social insurance schemes. Unemployment insurance, for example, cannot easily be conditioned on remaining unemployed if one can simply find informal employment instead. Likewise, health insurance cannot be provided through employer-based mandates, and pensions cannot be funded through payroll tax systems, for the vast number of employees who are not in the formal sector. And the fact that informality remains an option for many businesses means that governments must tread carefully before imposing too many such mandates on firms, lest they increase informality rates.

Despite the challenges involved in delivering these programs, there has been tremendous growth in both social assistance programs for redistribution, as well as social insurance, in low- and middle-income countries. This trend suggests that increasing attention to the unique challenges of developing social protection in these contexts is likely to be an important area for ongoing research in the years to come.

References

Abelson, Brian, Kush R. Varshney, and Joy Sun. 2014. "Targeting Direct Cash Transfers to the Extremely Poor." *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1563–1572.

Agüero, Jorge M., Michael R. Carter, and Ingrid Woolard. 2006. "The Impact of Unconditional Cash Transfers on Nutrition: The South African Child Support Grant." Southern Africa Labour and Development Research Unit Working Paper 06/08.

Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua Blumenstock. 2022. "Machine Learning and Phone Data Can Improve Targeting of Humanitarian Aid." *Nature*, 603: 864–870.

Aizer, Anna, Shari Eli, Adriana Lleras-Muney, and Keyoung Lee. 2020. "Do Youth Employment Programs Work? Evidence from the New Deal." National Bureau of Economic Research Working Paper 27103.

Akerlof, George. 1978. "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning." *American Economic Review*, 68: 8–19.

Akresh, Richard, Damien de Walque, and Harounan Kazianga. 2013. *Cash Transfers and Child Schooling: Evidence from a Randomized Evaluation of the Role of Conditionality.* The World Bank.

Akresh, Richard, Damien de Walque, and Harounan Kazianga. 2016. "Evidence from a Randomized Evaluation of the Household Welfare Impacts of Conditional and Unconditional Cash Transfers Given to Mothers or Fathers." World Bank Policy Research Working Paper 7730.

Alatas, Vivi. 2011. "Program Keluarga Harapan: Main Findings from the Impact Evaluation of Indonesia's Pilot Household Conditional Cash Transfer Program." World Bank Group Working Paper 72506.

Alatas, Vivi, Abhijit Banerjee, Arun Chandrasekhar, Rema Hanna, and Benjamin A. Olken. 2016*a*. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *American Economic Review*, 106: 1663–1704.

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: 1206–1240.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi. 2019. "Does Elite Capture Matter? Local Elites and Targeted Welfare Programs in Indonesia." *AEA Papers and Proceedings*, 109: 334–39.

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

Albagli, Pinjas, and Tomás Rau. 2019. "The Effects of a Maternity Leave Reform on Children's Abilities and Maternal Outcomes in Chile." *The Economic Journal*, 129(619): 1015–1047.

Albarran, Pedro, and Orazio Attanasio. 2003. "Limited Commitment and Crowding Out of Private Transfers: Evidence from a Randomised Experiment." *The Economic Journal*, 113: C77–85.

Alderman, Harold. 2002. "Do Local Officials Know Something We Don't? Decentralization of Targeted Transfers in Albania." *Journal of Public Economics*, 83: 375–404.

Ali, Williams, Awudu Abdulai, and Ashok K. Mishra. 2020. "Recent Advances in the Analyses of Demand for Agricultural Insurance in Developing and Emerging Countries." *Annual Review of Resource Economics*, 12: 411–430.

Andrabi, Tahir, Benjamin Daniels, and Jishnu Das. 2023. "Human Capital Accumulation and Disasters: Evidence from the Pakistan Earthquake of 2005." *Journal of Human Resources*, 58(4): 1057–1096.

Angelucci, Manuela, and Giacomo Giorgi. 2009. "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?" *American Economic Review*, 99: 486–508.

Anttila-Hughes, Jesse, and Solomon Hsiang. 2013. "Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster." Available at SSRN 2220501.

Araujo, Maria Caridad, and Karen Macours. 2021. "Education, Income and Mobility: Experimental Impacts of Childhood Exposure to Progresa after 20 Years." Paris School of Economics Working Paper 2021 - 57.

Ardington, Cally, Anne Case, and Victoria Hosegood. 2009. "Labor Supply Responses to Large Social Transfers: Longitudinal Evidence from South Africa." *American Economic Journal: Applied Economics*, 1: 22–48.

Areias, Ana, and Matthew Wai-Poi. 2022. "Machine Learning and Prediction of Beneficiary Eligibility for Social Protection Programs." *Revisiting Targeting in Social Assistance*, 507.

Arnott, Richard, and Joseph E. Stiglitz. 1991. "Moral Hazard and Nonmarket Institutions: Dysfunctional Crowding Out of Peer Monitoring?" *American Economic Review*, 81: 179–190.

Asuming, Patrick O. 2013. "Getting the Poor to Enroll in Health Insurance, and Its Effects on Their Health: Evidence from a Field Experiment in Ghana." Job Market Paper.

Asuming, Patrick O., Hyuncheol B. Kim, and Armand Sim. 2021. "Selection and Behavioral Responses of Health Insurance Subsidies in the Long Run: Evidence from a Field Experiment in Ghana." Centre for Development Economics and Sustainability Working Paper 05/21.

Aterido, Reyes, Mary Hallward-Driemeier, and Pagés Carmen. 2011. "Big Constraints to Small Firms' Growth? Business Environment and Employment Growth across Firms." *Economic Development and Cultural Change*, 59: 609–647.

Attanasio, Orazio, and José-Vıctor Rıos-Rull. 2000. "Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?" *European Economic Review*, 44(7): 1225–1258.

Attanasio, Orazio, Costas Meghir, and Andres Otero. 2014. "The Design of a Multi-Tier Contributory Pension System: The Distributional Impact of the 2008 Chilean Pension Reform." In *Social Insurance, Informality, and Labor Markets: How to Protect Workers While Creating Good Jobs*. 291–346. Oxford University Press.

Attanasio, Orazio, Luis Gómez, Patricia Heredia, and Marcos Vera-Hernández. 2005. "The Short-Term Impact of a Conditional Cash Subsidy on Child Health and Nutrition in Colombia." Institute for Fiscal Studies Report.

Augenblick, Ned, Kelsey Jack, Supreet Kaur, Felix Masiye, and Nicholas Swanson. 2022. "Budget Neglect in Consumption Smoothing: A Field Experiment on Seasonal Hunger." Unpublished.

Auriol, Emmanuelle, Julie Lassébie, Amma Panin, Eva Raiber, and Paul Seabright. 2020. "God Insures Those Who Pay? Formal Insurance and Religious Offerings in Ghana." *The Quarterly Journal of Economics*, 135: 1799–1848.

Azam, Mehtabul. 2012. "The Impact of Indian Job Guarantee Scheme on Labor Market Outcomes: Evidence from a Natural Experiment." IZA Discussion Paper 6548.

Azuara, Oliver, and Ioana Marinescu. 2013. "Informality and the Expansion of Social Protection Programs: Evidence from Mexico." *Journal of Health Economics*, 32: 938–950.

Baez, Javier E., and Indhira V. Santos. 2007. "Children's Vulnerability to Weather Shocks: A Natural Disaster as a Natural Experiment." Social Science Research Network.

Baez, Javier E., Varun Kshirsagar, and Emmanuel Skoufias. 2019. "Adaptive Safety Nets for Rural Africa: Drought-Sensitive Targeting with Sparse Data." World Bank Policy Research Working Paper 9071.

Baird, Sarah, Craig Mcintosh, and Berk Özler. 2011. "Cash or Condition? Evidence from a Cash Transfer Experiment." *The Quarterly Journal of Economics*, 126: 1709–1753.

Baird, Sarah, David McKenzie, and Berk Özler. 2018. "The Effects of Cash Transfers on Adult Labor Market Outcomes." *IZA Journal of Development and Migration*, 8: 22.

Baird, Sarah, Francisco H.G. Ferreira, Berk Özler, and Michael Woolcock. 2014. "Conditional, Unconditional and Everything in Between: A Systematic Review of the Effects of Cash Transfer Programmes on Schooling Outcomes." *Journal of Development Effectiveness*, 6: 1–43.

Baker, Judy, and Margaret Grosh. 1994. "Poverty Reduction through Geographic Targeting: How Well Does It Work?" *World Development*, 22: 983–995.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil. 2022. "Why Do People Stay Poor?" *The Quarterly Journal of Economics*, 137(2): 785–844.

Ball, Ian. 2023. "Scoring Strategic Agents." Working Paper.

Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. 2017. "Labor Markets and Poverty in Village Economies." *The Quarterly Journal of Economics*, 132(2): 811–870.

Banerjee, Abhijit, Amy Finkelstein, Rema Hanna, Benjamin A. Olken, Arianna Ornaghi, and Sudarno Sumarto. 2021. "The Challenges of Universal Health Insurance in Developing Countries: Experimental Evidence from Indonesia's National Health Insurance." *American Economic Review*, 111: 3035–3063.

Banerjee, Abhijit, and Andrew Newman. 1998. "Information, the Dual Economy, and Development." *The Review of Economic Studies*, 65: 631–653.

Banerjee, Abhijit, and Esther Duflo. 2010. "Giving Credit Where It Is Due." *Journal of Economic Perspectives*, 24: 61–80.

Banerjee, Abhijit, and Sendhil Mullainathan. 2010. "The Shape of Temptation: Implications for the Economic Lives of the Poor." National Bureau of Economic Research Working Paper 15973.

Banerjee, Abhijit, Dean Karlan, Robert Osei, Hannah Trachtman, and Christopher Udry. 2022. "Unpacking a Multi-Faceted Program to Build Sustainable Income for the Very Poor." *Journal of Development Economics*, 155: 102781.

Banerjee, Abhijit, Emily Breza, Esther Duflo, and Cynthia Kinnan. 2019*a*. "Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs?" National Bureau of Economic Research Working Paper 26346.

Banerjee, Abhijit, Esther Duflo, and Garima Sharma. 2021. "Long-Term Effects of the Targeting the Ultra Poor Program." *American Economic Review: Insights*, 3: 471–486.

Banerjee, Abhijit, Esther Duflo, and Richard Hornbeck. 2018. "How Much Do Existing Borrowers Value Microfinance? Evidence from an Experiment on Bundling Microcredit and Insurance." *Economica*, 85: 671–700.

Banerjee, Abhijit, Esther Duflo, Imbert Clément, Santhosh Mathew, and Rohini Pande. 2020*a*. "E-Governance, Accountability, and Leakage in Public Programs: Experimental Evidence from a Financial Management Reform in India." *American Economic Journal: Applied Economics*, 12: 39–72.

Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, Parienté William, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry. 2015. "A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries." *Science*, 348: 1260799.

Banerjee, Abhijit, Esther Duflo, Raghabendra Chattopadhyay, and Jeremy Shapiro. 2016. "The Long Term Impacts of a 'Graduation' Program: Evidence from West Bengal." Unpublished.

Banerjee, Abhijit, Paul Niehaus, and Tavneet Suri. 2019. "Universal Basic Income in the Developing World." *Annual Review of Economics*, 11: 959–983.

Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, and Sudarno Sumarto. 2020*b*. "The (Lack of) Distortionary Effects of Proxy-Means Tests: Results from a Nationwide Experiment in Indonesia." *Journal of Public Economics Plus*, 1: 100001.

Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, Elan Satriawan, and Sudarno Sumarto. 2023. "Electronic Food Vouchers: Evidence from an At-Scale Experiment in Indonesia." *American Economic Review*, 113(2): 514–547.

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

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: 101–137.

Banerjee, Abhijit, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto. 2019*b*. "Private Outsourcing and Competition: Subsidized Food Distribution in Indonesia." *Journal of Political Economy*, 127(1): 101–137.

Bardhan, Pranab, and Dilip Mookherjee. 2000. "Capture and Governance at Local and National Levels." *American Economic Review*, 90: 135–139.

Barham, Tania, and John A. Maluccio. 2009. "Eradicating Diseases: The Effect of Conditional Cash Transfers on Vaccination Coverage in Rural Nicaragua." *Journal of Health Economics*, 28: 611–621.

Barham, Tania, Karen Macours, and John A. Maluccio. 2013. "Boys' Cognitive Skill Formation and Physical Growth: Long-Term Experimental Evidence on Critical Ages for Early Childhood Interventions." *American Economic Review*, 103(3): 467–471.

Barham, Tania, Karen Macours, and John A. Maluccio. 2018*a*. "Are Conditional Cash Transfers Fulfilling Their Promise? Schooling, Learning, and Earnings After 10 Years." Centre for Economic Policy Research Discussion Paper 11937.

Barham, Tania, Karen Macours, and John A. Maluccio. 2018*b*. "Experimental Evidence of Exposure to a Conditional Cash Transfer during Early Teenage Years: Young Women's Fertility and Labor Market Outcomes." Centre for Economic Policy Research Discussion Paper 13165.

Barrera-Osorio, Felipe, Leigh L. Linden, and Juan E. Saavedra. 2019. "Medium- and Long-Term Educational Consequences of Alternative Conditional Cash Transfer Designs: Experimental Evidence from Colombia." *American Economic Journal: Applied Economics*, 11: 54–91.

Barrera-Osorio, Felipe, Marianne Bertrand, Leigh L. Linden, and Francisco Perez-Calle. 2011. "Improving the Design of Conditional Transfer Programs: Evidence from a Randomized Education Experiment in Colombia." *American Economic Journal: Applied Economics*, 3: 167–195.

Bastagli, Francesca, Jessica Hagen-Zanker, Luke Harman, Valentina Barca, Georgina Sturge, Tanja Schmidt, and Luca Pellerano. 2016. "Cash Transfers: What Does the Evidence Say?" Overseas Development Institute Research Report.

Basu, Kaushik. 1996. "Relief Programs: When It May Be Better to Give Food Instead of Cash." *World Development*, 24: 91–96.

Basurto, Maria Pia, Pascaline Dupas, and Jonathan Robinson. 2020. "Decentralization and Efficiency of Subsidy Targeting: Evidence from Chiefs in Rural Malawi." *Journal of Public Economics*, 185: 1880–1917.

Bau, Natalie. 2021. "Can Policy Change Culture? Government Pension Plans and Traditional Kinship Practices." *American Economic Review*, 111(6): 1880–1917.

Beaman, Lori, Niall Keleher, Jeremy Magruder, and Carly Trachtman. 2021. "Urban Networks and Targeting: Evidence from Liberia." *AEA Papers and Proceedings*, 111: 572–576.

Becker, Gary S., and George J. Stigler. 1974. "Law Enforcement, Malfeasance, and Compensation of Enforcers." *The Journal of Legal Studies*, 3(1): 1–18.

Bedoya, Guadalupe, Aidan Coville, Johannes Haushofer, Mohammad Isaqzadeh, and Jeremy Shapiro. 2019. "No Household Left Behind: Afghanistan Targeting the Ultra Poor Impact Evaluation." World Bank Policy Research Working Paper 8877.

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

Benedetti, Fiorella, Ibarrarán Pablo, and Patrick J. Mcewan. 2016. "Do Education and Health Conditions Matter in a Large Cash Transfer? Evidence from a Honduran Experiment." *Economic Development and Cultural Change*, 64: 759–793.

Benhassine, Najy, Florencia Devoto, and Esther Duflo. 2015. "Turning a Shove into a Nudge? A 'Labeled Cash Transfer' for Education." *American Economic Journal: Economic Policy*, 7: 86–125.

Berg, Erlend. 2018. "Funeral Insurance: An Inter-Generational Commitment Device?" *Journal of African Economies*, 27: 321–346.

Berg, Erlend, Sambit Bhattacharyya, D. Rajasekhar, and R. Manjula. 2018. "Can Public Works Increase Equilibrium Wages? Evidence from India's National Rural Employment Guarantee." *World Development*, 103: 239–254.

Bergolo, Marcelo, and Guillermo Cruces. 2014. "Work and Tax Evasion Incentive Effects of Social Insurance Programs: Evidence from an Employment-Based Benefit Extension." *Journal of Public Economics*, 117: 211–228.

Berkouwer, Susanna B., Pierre E. Biscaye, Eric Hsu, Oliver W. Kim, Kenneth Lee, Edward Miguel, and Catherine Wolfram. 2021. "Money or Power? Financial Infrastructure and Optimal Policy." National Bureau of Economic Research Working Paper 29086.

Bertrand, Marianne, Bruno Crépon, Alicia Marguerie, and Patrick Premand. 2021. "Do Workfare Programs Live Up to Their Promises? Experimental Evidence from Côte d'Ivoire." World Bank Policy Research Working Paper 9611.

Bertrand, Marianne, Sendhil Mullainathan, and Douglas Miller. 2003. "Public Policy and Extended Families." *The World Bank Economic Review*, 17: 27–49.

Besley, Timothy, and Stephen Coate. 1992. "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs." *American Economic Review*, 82: 249–261.

Bhalotra, Sonia, Diogo G.C. Britto, Paolo Pinotti, and Breno Sampaio. 2021. "Job Displacement, Unemployment Benefits and Domestic Violence." Centre for Economic Policy Research Discussion Paper 16350.

Bhattacharya, Debopam, and Pascaline Dupas. 2012. "Inferring Welfare Maximizing Treatment Assignment under Budget Constraints." *Journal of Econometrics*, 167: 168–196.

Bhattacharya, Shrayana, Maria Mini Jos, Soumya Kapoor Mehta, and Rinku Murgai. 2015. "From Policy to Practice: How Should Social Pensions Be Scaled Up?" *Economic and Political Weekly*, 50: 60–67.

Björkegren, Daniel, Joshua E. Blumenstock, and Samsun Knight. 2021. "Manipulation-Proof Machine Learning." Center for Effective Global Action Working Paper 186.

Bjorvatn, Kjetil, Denise Ferris, Selim Gulesci, Arne Nasgowitz, Vincent Somville, and Lore Vandewalle. 2022. "Childcare, Labor Supply, and Business Development: Experimental Evidence from Uganda." Centre for Economic Policy Research Discussion Paper 17243.

Blattman, Christopher, Nathan Fiala, and Sebastian W. Martinez. 2014. "Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda." *The Quarterly Journal of Economics*, 129(2): 697–752.

Blattman, Christopher, Nathan Fiala, and Sebastian W. Martinez. 2020. "The Long-Term Impacts of Grants on Poverty: Nine-Year Evidence from Uganda's Youth Opportunities Program." *American Economic Review: Insights*, 2(3): 287–304.

Blumenstock, Joshua, Gabriel Cadamuro, and Robert On. 2015. "Predicting Poverty and Wealth from Mobile Phone Metadata." *Science*, 350: 1073–1076.

Bosch, Mariano, and Raymundo Campos-Vazquez. 2014. "The Trade-Offs of Welfare Policies in Labor Markets with Informal Jobs: The Case of the 'Seguro Popular' Program in Mexico." *American Economic Journal: Economic Policy*, 6: 71–99.

Botea, Ioana, Markus Goldstein, Corinne Low, and Gareth Roberts. 2021. "Supporting Women's Livelihoods at Scale: RCT Evidence from a Nationwide Graduation Program." Unpublished.

Britto, Diogo G.C., Paolo Pinotti, and Breno Sampaio. 2022. "The Effect of Job Loss and Unemployment Insurance on Crime in Brazil." *Econometrica*, 90(4): 1393–1423.

Brollo, Fernanda, Katja Maria Kaufmann, and Eliana La Ferrara. 2017. "Learning about the Enforcement of Conditional Welfare Programs: Evidence from Brazil." IZA Discussion Paper 10654.

Brown, Caitlin, Martin Ravallion, and Dominique van de Walle. 2018. "A Poor Means Test? Econometric Targeting in Africa." *Journal of Development Economics*, 134: 109–124.

Browning, Martin, and Pierre-André Chiappori. 1998. "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests." *Econometrica*, 1241–1278.

Buera, Francisco J. 2009. "A Dynamic Model of Entrepreneurship with Borrowing Constraints: Theory and Evidence." *Annals of Finance*, 5: 443–464.

Buller, Ana Maria, Amber Peterman, Meghna Ranganathan, Alexandra Bleile, Melissa Hidrobo, and Lori Heise. 2018. "A Mixed-Method Review of Cash Transfers and Intimate Partner Violence in Low-and Middle-Income Countries." *The World Bank Research Observer*, 33: 218–258.

Cahyadi, Nur, Rema Hanna, Benjamin A. Olken, Adi Prima, Elan Satriawan, and Ekki Syamsulhakim. 2020. "Cumulative Impacts of Conditional Cash Transfer Programs: Experimental Evidence from Indonesia." *American Economic Journal: Economic Policy*, 12: 88–110.

Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet. 2020. "Subsidy Policies and Insurance Demand." *American Economic Review*, 110(8): 2422–53.

Calderón, Gabriela. 2014. "The Effects of Child Care Provision in Mexico." Banco de México Working Paper 2014-07.

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

Camacho, Adriana, Emily Conover, and Alejandro Hoyos. 2014. "Effects of Colombia's Social Protection System on Workers' Choice between Formal and Informal Employment." *The World Bank Economic Review*, 28(3): 446–466.

Cameron, Lisa, and Manisha Shah. 2014. "Can Mistargeting Destroy Social Capital and Stimulate Crime? Evidence from a Cash Transfer Program in Indonesia." *Economic Development and Cultural Change*, 62: 381–415.

Caria, A. Stefano, Grant Gordon, Maximilian Kasy, Simon Quinn, Soha Osman Shami, and Alexander Teytelboym. 2023. "An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan." *Journal of the European Economic Association*.

Carneiro, Pedro, Emanuela Galasso, and Rita Ginja. 2019. "Tackling Social Exclusion: Evidence from Chile." *The Economic Journal*, 129: 172–208.

Carney, Kevin, Michael Kremer, Xinyue Lin, and Gautam Rao. 2022. "The Endowment Effect and Collateralized Loans." National Bureau of Economic Research Working Paper 30073.

Carter, Michael, Alain de Janvry, Elisabeth Sadoulet, and Alexandros Sarris. 2017. "Index Insurance for Developing Country Agriculture: A Reassessment." *Annual Review of Resource Economics*, 9: 421–438.

Carter, Michael R., and Travis J. Lybbert. 2012. "Consumption versus Asset Smoothing: Testing the Implications of Poverty Trap Theory in Burkina Faso." *Journal of Development Economics*, 99: 255–264.

Carter, Michael R., Emilia Tjernström, and Patricia Toledo. 2019. "Heterogeneous Impact Dynamics of a Rural Business Development Program in Nicaragua." *Journal of Development Economics*, 138: 77–98.

Casaburi, Lorenzo, and Jack Willis. 2018. "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya." *American Economic Review*, 108: 3778–3813.

Cas, Ava Gail, Elizabeth Frankenberg, Wayan Suriastini, and Duncan Thomas. 2014. "The impact of Parental Death on Child Well-Being: Evidence from the Indian Ocean Tsunami." *Demography*, 51(2): 437–457.

Case, Anne, and Angus Deaton. 1998. "Large Cash Transfers to the Elderly in South Africa." *The Economic Journal*, 108(450): 1330–1361.

Case, Anne, Anu Garrib, Alicia Menendez, and Analia Olgiati. 2013. "Paying the Piper: The High Cost of Funerals in South Africa." *Economic Development and Cultural Change*, 62(1): 1–20.

Case, Anne, Victoria Hosegood, and Frances Lund. 2005. "The Reach and Impact of Child Support Grants: Evidence from KwaZulu-Natal." *Development Southern Africa*, 22(4): 467–482.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Fernández-Val Iván. 2018. "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India." National Bureau of Economic Research Working Paper 24678.

Chetty, Raj. 2006. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy*, 90: 31–52.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics*, 1(1): 451–488.

Chetty, Raj, and Adam Looney. 2007. "Income Risk and the Benefits of Social Insurance: Evidence from Indonesia and the United States." *Fiscal Policy and Management in East Asia*, 99–121.

Clarke, Daniel J., and Stefan Dercon. 2016. *Dull Disasters? How Planning Ahead Will Make a Difference.* Oxford University Press.

Clark, Shelley, Caroline W. Kabiru, Sonia Laszlo, and Stella Muthuri. 2019. "The Impact of Childcare on Poor Urban Women's Economic Empowerment in Africa." *Demography*, 56: 1247–1272.

Coady, David, Margaret Grosh, and John Hoddinott. 2004*a*. "Targeting of Transfers in Developing Countries: Review of Lessons and Experience." *The World Bank Research Observer*, 19: 61–85.

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

Coate, Stephen, Stephen Johnson, and Richard Zeckhauser. 1994. "Pecuniary Redistribution through In-Kind Programs." *Journal of Public Economics*, 55: 19–40.

Cole, Shawn, and Wentao Xiong. 2017. "Agricultural Insurance and Economic Development." *Annual Review of Economics*, 9: 235–262.

Cole, Shawn, Daniel Stein, and Jeremy Tobacman. 2014. "Dynamics of Demand for Index Insurance: Evidence from a Long-Run Field Experiment." *American Economic Review*, 104: 284–290.

Cunha, Jesse. 2014. "Testing Paternalism: Cash versus In-Kind Transfers." *American Economic Journal: Applied Economics*, 6: 195–230.

Cunha, Jesse, Giacomo Giorgi, and Seema Jayachandran. 2019. "The Price Effects of Cash Versus In-Kind Transfers." *The Review of Economic Studies*, 86: 240–281.

Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature*, 46: 333–383.

Dasgupta, Partha, and Debraj Ray. 1986. "Inequality as a Determinant of Malnutrition and Unemployment: Theory." *The Economic Journal*, 96: 1011–1034.

Das, Jishnu, and Quy-Toan Do. 2023. "The Prices in the Crises: What We Are Learning from 20 Years of Health Insurance in Low-and Middle-Income Countries." *Journal of Economic Perspectives*, 37(2): 123–152.

Dean, Emma Boswell, Frank Schilbach, and Heather Schofield. 2017. "Poverty and Cognitive Function." In *The Economics of Poverty Traps*. 57–118. University of Chicago Press.

de Carvalho Filho, Irineu Evangelista. 2012. "Household Income as a Determinant of Child Labor and School Enrollment in Brazil: Evidence from a Social Security Reform." *Economic Development and Cultural Change*, 60(2): 399–435.

de Janvry, Alain, Frederico Finan, Elisabeth Sadoulet, and Renos Vakis. 2006. "Can Conditional Cash Transfer Programs Serve as Safety Nets in Keeping Children at School and From Working When Exposed to Shocks?" *Journal of Development Economics*, 79(2): 349–373.

del Ninno, Carlo, and Bradford Mills. 2015. *Safety Nets in Africa: Effective Mechanisms to Reach the Poor and Most Vulnerable.* Washington, DC: World Bank and Agence Française de Développement.

del Valle, Alejandro. 2021. "The Effects of Public Health Insurance in Labor Markets with Informal Jobs: Evidence from Mexico." *Journal of Health Economics*, 77: 102454.

del Valle, Alejandro. Forthcoming. "Saving Lives with Indexed Disaster Funds: Evidence from Mexico." *American Economic Journal: Economic Policy.*

del Valle, Alejandro, Alain de Janvry, and Elisabeth Sadoulet. 2020. "Rules for Recovery: Impact of Indexed Disaster Funds on Shock Coping in Mexico." *American Economic Journal: Applied Economics*, 12(4): 164–195.

Dercon, Stefan, and Catherine Porter. 2014. "Live Aid Revisited: Long-Term Impacts of the 1984 Ethiopian Famine on Children." *Journal of the European Economic Association*, 12(4): 927–948.

Dercon, Stefan, Jan Willem Gunning, and Andrew Zeitlin. 2019. "The Demand for Insurance under Limited Trust: Evidence from a Field Experiment in Kenya." Centre for the Study of African Economies Working Paper 2019-06.

Dercon, Stefan, Joachim Weerdt, Tessa Bold, and Alula Pankhurst. 2006. "Group-Based Funeral Insurance in Ethiopia and Tanzania." *World Development*, 34: 685–703.

Duflo, Esther. 2003. "Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa." *The World Bank Economic Review*, 17: 1–25.

Dupas, Pascaline, Marcel Fafchamps, and Deivy Houeix. 2022. "Measuring Relative Poverty through Peer Rankings: Evidence from Côte D'Ivoire." National Bureau of Economic Research Working Paper 29911.

Dupas, Pascaline, Vivian Hoffmann, Michael Kremer, and Alix Zwane. 2016. "Targeting Health Subsidies through a Nonprice Mechanism: A Randomized Controlled Trial in Kenya." *Science*, 353: 889–895.

Dutta, Puja, Rinku Murgai, Martin Ravallion, and Dominique van de Walle. 2014. *Right to Work? Assessing India's Employment Guarantee Scheme in Bihar.* Washington, DC: World Bank.

Edmonds, Eric. 2006. "Child Labor and Schooling Responses to Anticipated Income in South Africa." *Journal of Development Economics*, 81: 386–414.

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

Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf, and Mark Cullen. 2013. "Selection on Moral Hazard in Health Insurance." *American Economic Review*, 103: 178–219.

Elbers, Chris, Jean Lanjouw, and Peter Lanjouw. 2003. "Micro-Level Estimation of Poverty and Inequality." *Econometrica*, 71: 355–364.

Elbers, Chris, Tomoki Fujii, Peter Lanjouw, Berk Özler, and Wesley Yin. 2007. "Poverty Alleviation through Geographic Targeting: How Much Does Disaggregation Help?" *Journal of Development Economics*, 83: 198–213.

Evans, David, and Anna Popova. 2016. "What Really Works to Improve Learning in Developing Countries?: An Analysis of Divergent Findings in Systematic Reviews." *World Bank Research Observer*, 31: 242–270.

Evans, David, and Katrina Kosec. 2022. "Cash Transfers, Trust, and Inter-Household Transfers: Experimental Evidence from Tanzania." International Food Policy Research Institute Discussion Paper 1994 v2.

Evans, David, Brian Holtemeyer, and Katrina Kosec. 2019. "Cash Transfers and Health: Evidence from Tanzania." *The World Bank Economic Review*, 33: 394–412.

Evans, David, Stephanie Hausladen, Katrina Kosec, and Natasha Reese. 2014. *Community-Based Conditional Cash Transfers in Tanzania: Results from a Randomized Trial.* Washington, DC: World Bank.

Farooq, Ammar, Adriana D. Kugler, and Umberto Muratori. 2020. "Do Unemployment Insurance Benefits Improve Match and Employer Quality? Evidence from Recent U.S. Recessions." National Bureau of Economic Research Working Paper 27574.

Feldstein, Martin. 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." *The Review of Economics and Statistics*, 81: 674–680.

Feldstein, Martin, and Daniel Altman. 2007. "Unemployment Insurance Savings Accounts." *Tax Policy and the Economy*, 21: 35–63.

Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Moore. 2021. "On Her Own Account: How Strengthening Women's Financial Control Impacts Labor Supply and Gender Norms." *American Economic Review*, 111: 2342–2375.

Filmer, Deon, Jed Friedman, Eeshani Kandpal, and Junko Onishi. 2023. "Cash Transfers, Food Prices, and Nutrition Impacts on Ineligible Children." *The Review of Economics and Statistics*, 105: 327–343.

Finkelstein, Amy. 2007. "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare." *The Quarterly Journal of Economics*, 122: 1–37.

Finkelstein, Amy, and Matthew Notowidigdo. 2019. "Take-Up and Targeting: Experimental Evidence from SNAP." *The Quarterly Journal of Economics*, 134: 1505–1556.

Finkelstein, Amy, and Robin Mcknight. 2008. "What Did Medicare Do? The Initial Impact of Medicare on Mortality and out of Pocket Medical Spending." *Journal of Public Economics*, 92: 1644–1668.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review*, 109: 1530–1567.

Fischer, Torben, Markus Frölich, and Andreas Landmann. 2023. "Adverse Selection in Low-Income Health Insurance Markets: Evidence from an RCT in Pakistan." *American Economic Journal: Applied Economics*, 15(3): 313–340.

Franklin, Simon, Clément Imbert, Girum Abebe, and Carolina Mejia-Mantilla. Forthcoming. "Urban Public Works in Spatial Equilibrium: Experimental Evidence from Ethiopia." *American Economic Review*.

Gadenne, Lucie, Samuel Norris, Monica Singhal, and Sandip Sukhtankar. 2021. "In-Kind Transfers as Insurance." National Bureau of Economic Research Working Paper 28507.

Galasso, Emanuela, and Martin Ravallion. 2005. "Decentralized Targeting of an Antipoverty Program." *Journal of Public Economics*, 89(4): 705–727.

Galiani, Sebastian, and Patrick J. McEwan. 2013. "The Heterogeneous Impact of Conditional Cash Transfers." *Journal of Public Economics*, 103: 85–96.

Galiani, Sebastian, Paul J. Gertler, and Rosangela Bando. 2016. "Non-Contributory Pensions." *Labour Economics*, 38: 47–58.

Gardner, Clare, Caroline Riungu, Clare O'Brien, and Fred Merttens. 2017. "Evaluation of the Hunger Safety Net Programme Phase 2: The Legacy of HSNP Phase 2: Systems, Practices and Lessons Learned." Oxford Policy Management Report.

Gentilini, Ugo, Mohamed Almenfi, Ian Orton, and Pamela Dale. 2020. "Social Protection and Jobs Responses to COVID-19: A Real-Time Review of Country Measures." World Bank Brief.

Gerard, François, and Gustavo Gonzaga. 2021. "Informal Labor and the Efficiency Cost of Social Programs: Evidence from Unemployment Insurance in Brazil." *American Economic Journal: Economic Policy*, 13: 167–206.

Gerard, François, and Joana Naritomi. 2021. "Job Displacement Insurance and (the Lack of) Consumption-Smoothing." *American Economic Review*, 111: 899–942.

Gerard, François, Joana Naritomi, and Joana Silva. 2021. "Cash Transfers and Formal Labor Markets: Evidence from Brazil." Centre for Economic Policy Research Discussion Paper 16286.

Gertler, Paul J. 2004. "Do Conditional Cash Transfers Improve Child Health? Evidence from Progresa's Control Randomized Experiment." *American Economic Review*, 94: 336–341.

Gertler, Paul J., Sebastian W. Martinez, and Marta Rubio-Codina. 2012. "Investing Cash Transfers to Raise Long-Term Living Standards." *American Economic Journal: Applied Economics*, 4(1): 164–192.

Gignoux, Jérémie, and Marta Menéndez. 2016. "Benefit in the Wake of Disaster: Long-Run Effects of Earthquakes on Welfare in Rural Indonesia." *Journal of Development Economics*, 118: 26–44.

Giles, John, Xin Meng, Sen Xue, and Guochang Zhao. 2021. "Can Information Influence the Social Insurance Participation Decision of China's Rural Migrants?" *Journal of Development Economics*, 150: 102645.

Glewwe, Paul. 1992. "Targeting Assistance to the Poor: Efficient Allocation of Transfers When Household Income Is Not Observed." *Journal of Development Economics*, 38: 297–321.

Glewwe, Paul, and Pedro Olinto. 2004. "Evaluating the Impact of Conditional Cash Transfers on Schooling: An Experimental Analysis of Honduras' PRAF program." Unpublished.

Goldberg, Jessica. 2016. "Kwacha Gonna Do? Experimental Evidence about Labor Supply in Rural Malawi." *American Economic Journal: Applied Economics*, 8(1): 129–149.

González-Rozada, Martín, and Hernán Ruffo. 2023. "The Welfare Effects of Unemployment Insurance In Argentina: New Estimates Using Changes in the Schedule of Transfers." *Journal of Human Resources*.

Gruber, Jonathan, Mengyun Lin, and Junjian Yi. 2023. "The Largest Insurance Program in History: Saving One Million Lives Per Year in China." *Journal of Public Economics*, 226: 104999.

Gruber, Jonathan, Nathaniel Hendren, and Robert Townsend. 2014. "The Great Equalizer: Health Care Access and Infant Mortality in Thailand." *American Economic Journal: Applied Economics*, 6: 91–107.

Gupta, Sarika. 2017. "Perils of the Paperwork: The Impact of Information and Application Assistance on Welfare Program Take-Up in India." Job Market Paper.

Halim, Daniel, Hillary C. Johnson, and Elizaveta Perova. 2022. "Preschool Availability and Women's Employment: Evidence from Indonesia." *Economic Development and Cultural Change*, 22.

Handa, Sudhanshu, Silvio Daidone, Amber Peterman, Benjamin Davis, Audrey Pereira, Tia Palermo, and Jennifer Yablonski. 2018. "Myth-Busting? Confronting Six Common Perceptions about Unconditional Cash Transfers as a Poverty Reduction Strategy in Africa." *The World Bank Research Observer*, 33: 259–298.

Hanna, Rema, and Benjamin A. Olken. 2018. "Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries." *Journal of Economic Perspectives*, 32: 201–226.

Hanna, Rema, and Dean Karlan. 2017. "Designing Social Protection Programs: Using Theory and Experimentation to Understand How to Help Combat Poverty." *Handbook of Economic Field Experiments*, 2: 515–553.

Hastings, Justine, Ali Hortaçsu, and Chad Syverson. 2017. "Sales Force and Competition in Financial Product Markets: The Case of Mexico's Social Security Privatization." *Econometrica*, 85: 1723–1761.

Hastings, Justine, and Jesse Shapiro. 2018. "How Are SNAP Benefits Spent? Evidence from a Retail Panel." *American Economic Review*, 108: 3493–3540.

Haushofer, Johannes, and Jeremy Shapiro. 2016. "The Short-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya." *The Quarterly Journal of Economics*, 131: 1973–2042.

Haushofer, Johannes, and Jeremy Shapiro. 2018. "The Long-Term Impact of Unconditional Cash Transfers: Experimental Evidence from Kenya." Unpublished.

Haushofer, Johannes, Matthieu Chemin, Chaning Jang, and Justin Abraham. 2020. "Economic and Psychological Effects of Health Insurance and Cash Transfers: Evidence from a Randomized Experiment in Kenya." *Journal of Development Economics*, 144.

Haushofer, Johannes, Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael Walker. 2022. "Targeting Impact versus Deprivation." National Bureau of Economic Research Working Paper 30138.

He, Guojun, and Shaoda Wang. 2017. "Do College Graduates Serving as Village Officials Help Rural China?" *American Economic Journal: Applied Economics*, 9: 186–215.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics*, 135: 1209–1318.

Hidrobo, Melissa, John Hoddinott, Amber Peterman, Amy Margolies, and Vanessa Moreira. 2014. "Cash, Food, or Vouchers? Evidence from a Randomized Experiment in Northern Ecuador." *Journal of Development Economics*, 107: 144–156.

Hoddinott, John, John A. Maluccio, Jere R. Behrman, Rafael Flores, and Reynaldo Martorell. 2008. "Effect of a Nutrition Intervention during Early Childhood on Economic Productivity in Guatemalan Adults." *The Lancet*, 371(9610): 411–416.

Hoynes, Hilary, and Jesse Rothstein. 2019. "Universal Basic Income in the United States and Advanced Countries." *Annual Review of Economics*, 11: 929–958.

Huang, Wei, and Chuanchuan Zhang. 2021. "The Power of Social Pensions: Evidence from China's New Rural Pension Scheme." *American Economic Journal: Applied Economics*, 13: 179–205.

Huneeus, Cristóbal, Silvia Leiva, and Alejandro Micco. 2012. "Unemployment Insurance and Search Effort in Chile." Inter-American Development Bank Working Paper 313.

Ikegami, Munenobu, Michael R. Carter, Christopher B. Barrett, and Sarah A. Janzen. 2018. "Poverty Traps and the Social Protection Paradox." In *The Economics of Poverty Traps*. 223–256. Chicago Press.

ILOSTAT. 2019. "Employment in Agriculture (% of Total Employment) - Low Income." International Labour Organization.

Imbert, Clément, and John Papp. 2020*a*. "Costs and Benefits of Rural-Urban Migration: Evidence from India." *Journal of Development Economics*, 146: 102473.

Imbert, Clément, and John Papp. 2020*b*. "Short-Term Migration, Rural Public Works, and Urban Labor Markets: Evidence from India." *Journal of the European Economic Association*, 18(2): 927–963.

Insurance Information Institute. 2020. "Insured Homes by State."

Iqbal, Sarah. 2018. "Women, Business and the Law 2018." World Bank Group.

Janzen, Sarah A., Michael R. Carter, and Munenobu Ikegami. 2021. "Can Insurance Alter Poverty Dynamics and Reduce the Cost of Social Protection in Developing Countries?" *Journal of Risk and Insurance*, 88(2): 293–324.

Janzen, Sarah A., Nicholas Magnan, Conner Mullally, Soye Shin, I. Bailey Palmer, Judith Oduol, and Karl Hughes. 2021. "Can Experiential Games and Improved Risk Coverage Raise Demand for Index Insurance? Evidence from Kenya." *American Journal of Agricultural Economics*, 103(1): 338–361.

Jean, Neal, Marshall Burke, Michael Xie, W Davis, David Lobell, and Stefano Ermon. 2016. "Combining Satellite Imagery and Machine Learning to Predict Poverty." *Science*, 353: 790–794.

Jensen, Anders. 2022. "Employment Structure and the Rise of the Modern Tax System." *American Economic Review*, 112: 213–234.

Jensen, Nathaniel, and Christopher Barrett. 2017. "Agricultural Index Insurance for Development." *Applied Economic Perspectives and Policy*, 39: 199–219.

Jensen, Robert. 2004. "Do Private Transfers 'Displace' the Benefits of Public Transfers? Evidence from South Africa." *Journal of Public Economics*, 88: 89–112.

Jiménez-Hernández, Diego, and Enrique Seira. 2022. "Should the Government Sell You Goods? Evidence from the Milk Market in Mexico." Job Market Paper.

Kandpal, Eeshani, Harold Alderman, Jed Friedman, Deon Filmer, Junko Onishi, and Jorge Avalos. 2016. "A Conditional Cash Transfer Program in the Philippines Reduces Severe Stunting." *The Journal of Nutrition*, 146(9): 1793–1800.

King, Gary, Emmanuela Gakidou, Kosuke Imai, Jason Lakin, Ryan T. Moore, Clayton Nall, Nirmala Ravishankar, Manett Vargas, Martha María Téllez-Rojo, Juan Eugenio Hernández Ávila, et al. 2009. "Public Policy for the Poor? A Randomised Assessment of the Mexican Universal Health Insurance Programme." *The Lancet*, 373(9673): 1447–1454.

Klasen, Stephan, and Simon Lange. 2016. "How Narrowly Should Anti-Poverty Programs Be Targeted? Simulation Evidence from Bolivia and Indonesia." Courant Research Centre Discussion Paper 213.

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

Kugler, Adriana. 2005. "Wage-Shifting Effects of Severance Payments Savings Accounts in Colombia." *Journal of Public Economics*, 89: 487–500.

Kumler, Todd, Eric Verhoogen, and Judith Frías. 2020. "Enlisting Employees in Improving Payroll Tax Compliance: Evidence from Mexico." *The Review of Economics and Statistics*, 102(5): 881–896.

Kyle, Jordan. 2018. "Local Corruption and Popular Support for Fuel Subsidy Reform in Indonesia." *Comparative Political Studies*, 51: 1472–1503.

Laurito, Maria M., Elizabeth Frankenberg, and Duncan Thomas. 2022. "Effects of Housing Aid on Psychosocial Health after a Disaster." *International Journal of Environmental Research and Public Health*, 19(12): 7302.

Leroy, Jef, Paola Gadsden, Sonia Rodríguez-Ramírez, and Teresa Gonzalez De Cossío. 2010. "Cash and In-Kind Transfers in Poor Rural Communities in Mexico Increase Household Fruit, Vegetable, and Micronutrient Consumption but Also Lead to Excess Energy Consumption." *The Journal of Nutrition*, 140: 612–617.

Levine, David, Rachel Polimeni, and Ian Ramage. 2016. "Insuring Health or Insuring Wealth? An Experimental Evaluation of Health Insurance in Rural Cambodia." *Progress Against Poverty*, 119: 1–15.

Levy, Santiago, and Norbert Schady. 2013. "Latin America's Social Policy Challenge: Education, Social Insurance, Redistribution." *Journal of Economic Perspectives*, 27: 193–218.

Liepmann, Hannah, and Clemente Pignatti. 2024. "Welfare Effects of Unemployment Benefits When Informality Is High." *Journal of Public Economics*, 229: 105032.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *The Review of Economic Studies*, 69(1): 209–244.

Lone, Tanya, Maheen Shakeel, and Jana Bischler. 2021. "Towards Shock-Responsive Social Protection: Lessons from the COVID-19 Response in Pakistan." Oxford Policy Management Report.

Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales. 1997. "Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit." *The Journal of Human Resources*, 32: 463–480.

Lybbert, Travis, Christopher B. Barrett, Solomon Desta, and D. Layne Coppock. 2004. "Stochastic Wealth Dynamics and Risk Management among a Poor Population." *The Economic Journal*, 114: 750–777.

Macours, Karen, and Renos Vakis. 2014. "Changing Households' Investment Behaviour through Social Interactions with Local Leaders: Evidence from a Randomised Transfer Programme." *The Economic Journal*, 124: 607–633.

Macours, Karen, Norbert Schady, and Renos Vakis. 2012. "Cash Transfers, Behavioral Changes, and Cognitive Development in Early Childhood: Evidence from a Randomized Experiment." *American Economic Journal: Applied Economics*, 4: 247–273.

Mahajan, Parag, and Dean Yang. 2020. "Taken by Storm: Hurricanes, Migrant Networks, and US Immigration." *American Economic Journal: Applied Economics*, 12(2): 250–77.

Malani, Anup, Phoebe Holtzman, Kosuke Imai, Cynthia Kinnan, Morgen Miller, Shailender Swaminathan, Alessandra Voena, Bartosz Woda, and Gabriella Conti. 2021. "Effect of Health Insurance in India: A Randomized Controlled Trial." National Bureau of Economic Research Working Paper 29576.

Martinelli, César, and Susan W. Parker. 2009. "Deception and Misreporting in a Social Program." *Journal of the European Economic Association*, 7: 886–908.

Martínez, Claudia, and Marcela Perticará. 2017. "Childcare Effects on Maternal Employment: Evidence from Chile." *Journal of Development Economics*, 126: 127–137.

Mcbride, Linden, and Austin Nichols. 2018. "Retooling Poverty Targeting Using Out-of-Sample Validation and Machine Learning." *The World Bank Economic Review*, 32: 531–550.

McKenzie, David. 2017. "Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition." *American Economic Review*, 107(8): 2278–2307.

Meghir, Costas, and David Phillips. 2010. "Labour Supply and Taxes." In *Dimensions of Tax Design: The Mirrlees Review.* , ed. James A. Mirrlees and Institute for Fiscal Studies (Great Britain), 202–274. New York: Oxford University Press.

Michelson, Hope, Erin C. Lentz, Richard Mulwa, Mitchell Morey, Laura Cramer, Megan McGlinchy, and Christopher B. Barrett. 2012. "Cash, Food, or Vouchers? An Application of the Market Information and Food Insecurity Response Analysis Framework in Urban and Rural Kenya." *Food Security*, 4: 455–469.

Millán, Teresa Molina, Karen Macours, John A. Maluccio, and Luis Tejerina. 2020. "Experimental Long-Term Effects of Early-Childhood and School-Age Exposure to a Conditional Cash Transfer Program." *Journal of Development Economics*, 143: 102385.

Miranda, Mario, and Katie Farrin. 2012. "Index Insurance for Developing Countries." *Applied Economic Perspectives and Policy*, 34: 391–427.

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

Moffitt, Robert. 1983. "An Economic Model of Welfare Stigma." *American Economic Review*, 73: 1023–1035.

Morris, Saul, Rafael Flores, Pedro Olinto, and Juan Medina. 2004. "Monetary Incentives in Primary Health Care and Effects on Use and Coverage of Preventive Health Care Interventions in Rural Honduras: Cluster Randomised Trial." *The Lancet*, 364: 2030–2037.

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why Having Too Little Means so Much.* New York: Macmillan.

Munshi, Kaivan, and Mark Rosenzweig. 2016. "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap." *American Economic Review*, 106: 46–98.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. "Building State Capacity: Evidence from Biometric Smartcards in India." *American Economic Review*, 106(10): 2895–2929.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2023. "Identity Verification Standards in Welfare Programs: Experimental Evidence from India." *Review of Economics and Statistics*, 1–46.

Murgai, Rinku, Martin Ravallion, and Dominique van de Walle. 2016. "Is Workfare Cost-Effective against Poverty in a Poor Labor-Surplus Economy?" *The World Bank Economic Review*, 30(3): 413–445.

Nagler, Paula. 2013. "How Unemployment Insurance Savings Accounts Affect Employment Duration: Evidence from Chile." *IZA Journal of Labor & Development*, 2: 1–25.
Natural Disaster Insurance Institution. 2022. "Policy Production."
Nekoei, Arash, and Andrea Weber. 2017. "Does Extending Unemployment Benefits Improve Job Quality?" *American Economic Review*, 107: 527–561.
Nichols, Albert, and Richard Zeckhauser. 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review*, 72: 372–377.
Niehaus, Paul, Antonia Atanassova, Marianne Bertrand, and Sendhil Mullainathan. 2013. "Targeting with Agents." *American Economic Journal: Economic Policy*, 5: 206–238.
Nshakira-Rukundo, Emmanuel, Juliet Wanjiku Kamau, and Heike Baumüller. 2021. "Determinants of Uptake and Strategies to Improve Agricultural Insurance in Africa: A Review." *Environment and Development Economics*, 26(5-6): 605–631.
Oates, Wallace, and Robert Schwab. 2015. "The Window Tax: A Case Study in Excess Burden." *Journal of Economic Perspectives*, 29: 163–180.
O'Brien, Clare, Zoë Scott, Gabrielle Smith, Valentina Barca, Andrew Kardan, Rebecca Holmes, Carol Watson, and Jenny Congrave. 2018. "Shock-Responsive Social Protection Systems Research: Synthesis Report." Oxford Policy Management Report.
OECD. 2020. "Insurance Spending."
Olken, Benjamin A. 2007. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy*, 115(2): 200–249.
Osman, Adam, and Jamin D. Speer. 2024. "Stigma and Take-Up of Labour Market Assistance: Evidence from Two Field Experiments." *Economica*, 91: 123–141.
Paes de Barros, Ricardo, Pedro Olinto, Trine Lunde, and Mirela Carvalho. 2011. "The Impact of Access to Free Childcare on Women's Labor Market Outcomes: Evidence from a Randomized Trial in Low-Income Neighborhoods of Rio de Janeiro." In *World Bank Economists' Forum*.
Parker, Susan W., and Petra E. Todd. 2017. "Conditional Cash Transfers: The Case of Progresa/Oportunidades." *Journal of Economic Literature*, 55(3): 866–915.
Parker, Susan W., and Tom Vogl. 2023. "Do Conditional Cash Transfers Improve Economic Outcomes in the Next Generation? Evidence from Mexico." *The Economic Journal*, 133(655): 2775–2806.
Pople, Ashley, Ruth Hill, Stefan Dercon, and Ben Brunckhorst. 2021. "Anticipatory Cash Transfers in Climate Disaster Response." Centre for the Study of African Economies Working Paper 2021-07.
Posel, Dorrit, James A. Fairburn, and Frances Lund. 2006. "Labour Migration and Households: A Reconsideration of the Effects of the Social Pension on Labour Supply in South Africa." *Economic Modelling*, 23(5): 836–853.
Premand, Patrick, and Pascale Schnitzer. 2021. "Efficiency, Legitimacy, and Impacts of Targeting Methods: Evidence from an Experiment in Niger." *The World Bank Economic Review*, 35: 892–920.
Ralston, Laura, Colin Andrews, and Allan Hsiao. 2017. "The Impacts of Safety Nets in Africa: What Are We Learning?" World Bank Policy Research Working Paper 8255.
Ravallion, Martin. 1991. "Reaching the Rural Poor through Public Employment: Arguments, Evidence, and Lessons from South Asia." *The World Bank Research Observer*, 6: 205–231.
Reynaud, Arnaud, Manh-Hung Nguyen, and Cécile Aubert. 2018. "Is There a Demand for Flood Insurance in Vietnam? Results from a Choice Experiment." *Environmental Economics and Policy Studies*, 20: 593–617.
Robinson, Jonathan. 2012. "Limited Insurance within the Household: Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics*, 4: 140–164.
Sabates-Wheeler, Rachel, Alex Hurrell, and Stephen Devereux. 2015. "Targeting Social Transfer Programmes: Comparing Design and Implementation Errors Across Alternative Mechanisms." *Journal of International Development*, 27(8): 1521–1545.
Sadka, Joyce, Enrique Seira, and Christopher Woodruff. 2024. "Information and Bargaining through Agents: Experimental Evidence from Mexico's Labour Courts." *The Review of Economic Studies*, 00: 1–36.
Sadoulet, Elisabeth, Alain de Janvry, and Benjamin Davis. 2001. "Cash Transfer Programs with Income Multipliers: PROCAMPO in Mexico." *World Development*, 29(6): 1043–1056.
Saez, Emmanuel. 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *The Quarterly Journal of Economics*, 117(3): 1039–1073.

Saez, Emmanuel, and Stefanie Stantcheva. 2016. "Generalized Social Marginal Welfare Weights For Optimal Tax Theory." *American Economic Review*, 106(01): 24–45.
Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz. 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature*, 50(1): 3–50.
Schady, Norbert, and Maria Caridad Araujo. 2006. "Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador." World Bank Policy Research Working Paper 3930.
Schultz, T. Paul. 2004. "School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program." *Journal of Development Economics*, 74(1): 199–250.
Skiba, Aleksandr K. 1978. "Optimal Growth with a Convex-Concave Production Function." *Econometrica*, 527–539.
Skoufias, Emmanuel. 2005. "Progresa and Its Impacts on the Welfare of Rural Households in Mexico." International Food Policy Research Institute Research Report 139.
Skoufias, Emmanuel, Mishel Unar, and Teresa González-Cossío. 2008. "The Impacts of Cash and In-Kind Transfers on Consumption and Labor Supply: Experimental Evidence from Rural Mexico." World Bank Policy Research Working Paper 4778.
Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *The Quarterly Journal of Economics*, 70(1): 65–94.
Stiglitz, Joseph E. 1974. "Incentives and Risk Sharing in Sharecropping." *The Review of Economic Studies*, 41: 219–255.
Stiglitz, Joseph E., and Andrew Weiss. 1981. "Credit Rationing in Markets with Imperfect Information." *American Economic Review*, 71(3): 393–410.
Stoeffler, Quentin, Bradford Mills, and Carlo del Ninno. 2016. "Reaching the Poor: Cash Transfer Program Targeting in Cameroon." *World Development*, 83: 244–263.
Strupat, Christoph, and Florian Klohn. 2018. "Crowding Out of Solidarity? Public Health Insurance versus Informal Transfer Networks in Ghana." *World Development*, 104: 212.
Takahashi, Kazushi, Christopher B. Barrett, and Munenobu Ikegami. 2019. "Does Index Insurance Crowd in or Crowd Out Informal Risk Sharing? Evidence from Rural Ethiopia." *American Journal of Agricultural Economics*, 101(3): 672–691.
Tarozzi, Alessandro, and Angus Deaton. 2009. "Using Census and Survey Data to Estimate Poverty and Inequality for Small Areas." *The Review of Economics and Statistics*, 91(4): 773–792.
The World Bank. 2014. "For up to 800 Million Rural Poor, a Strong World Bank Commitment to Agriculture."
The World Bank. 2021*a*. "GDP per Capita (Constant 2015 US$)."
The World Bank. 2021*b*. "Population, Total."
Thornton, Rebecca L., Laurel E. Hatt, Erica M. Field, Mursaleena Islam, Freddy Solís Diaz, and Martha Azucena González. 2010. "Social Security Health Insurance for the Informal Sector in Nicaragua: A Randomized Evaluation." *Health Economics*, 19(S1): 181–206.
Todd, Jessica, and Paul Winters. 2011. "The Effect of Early Interventions in Health and Nutrition on On-Time School Enrollment: Evidence from the Oportunidades Program in Rural Mexico." *Economic Development and Cultural Change*, 59: 549–581.
Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica*, 539–591.
Townsend, Robert M. 1995. "Financial Systems in Northern Thai Villages." *The Quarterly Journal of Economics*, 110(4): 1011–1046.
van Ours, Jan C., and Milan Vodopivec. 2008. "Does Reducing Unemployment Insurance Generosity Reduce Job Match Quality?" *Journal of Public Economics*, 92(3-4): 684–695.
Vu, Khoa, and Paul Glewwe. 2022. "Maternity Benefits Mandate and Women's Choice of Work in Vietnam." *World Development*, 158: 105964.
Wager, Stefan, and Susan Athey. 2018. "Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests." *Journal of the American Statistical Association*, 113: 1228–1242.
Wagstaff, Adam, Ha Thi Hong Nguyen, Huyen Dao, and Sarah Bales. 2016. "Encouraging Health Insurance for the Informal Sector: A Cluster Randomized Experiment in Vietnam." *Health Economics*, 25(6): 663–674.
World Bank. 2018. *The State of Social Safety Nets 2018.* Washington, DC: World Bank.

World Economic Forum. 2019. "GCI 4.0: Insurance Premium."
Yang, Dean. 2008. "Coping with Disaster: The Impact of Hurricanes on International Financial Flows, 1970-2002." *The BE Journal of Economic Analysis & Policy*, 8(1).