

# Social Protection in the Developing World<sup>1</sup>

Abhijit Banerjee, Rema Hanna, Benjamin A. Olken and Diana Sverdlin-Lisker

September 2022

## 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.

---

<sup>1</sup> We thank Amy Finkelstein, Nathan Hendren, the JEL editors, and four anonymous referees for helpful comments. We thank Emily Romano, Beatriz Velho and especially Kim Lan Mallon for outstanding research assistance. Financial support from the Wellspring Philanthropic Fund and the Australian Department of Foreign Affairs and Trade is gratefully acknowledged. The views expressed in this publication are the authors' alone and are not necessarily the views of the Australian Government. Sverdlin-Lisker: 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.

## 1. Introduction

Over the past several decades, social assistance programs have been rapidly evolving in the developing world. Today, for example, 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. Safety net programs cover 2.5 billion people, of which 650 million are from the poorest quintile. In just the past decade, the number of developing countries running conditional cash transfer programs, which condition benefits on households making human capital investments in their children, has more than doubled (World Bank 2018a). These new 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 crisis led to a dramatic expansion in social protection policies, which are expected to continue to play a pivotal role in the COVID recovery and beyond.<sup>2</sup>

Designing these programs for a developing country context, however, entails a variety of challenges that differ from those faced in more developed economies. For example, one of the biggest questions in the public economics literature is who should be targeted for social assistance, and how can governments best obtain information to identify these individuals. But, unlike in developed countries, information on income is not readily observed, partly because informal work and self-employment are very widespread. For example, the bottom 70 to 90 percent of households in many low-income nations are often excluded from the income tax system, making it challenging to collect systematic data (Jensen 2022). This means that the standard frameworks used for thinking about targeting benefits to households – e.g. income eligibility thresholds, taking into account the endogeneity of labor supply to changes in effective marginal tax rates resulting from program eligibility (e.g. Mirrlees 1971; Saez 2001) – do not readily apply in a developing country setting, necessitating other approaches. These alternative targeting approaches, in turn, demand rigorous economic thinking to understand how best to design them and to guide empirical studies of their effectiveness and tradeoffs (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 highly successful development intervention is to give households a productive livestock asset—chickens, a goat, a cow, etc. —and support them in taking care of it (e.g. Banerjee, Duflo, and Sharma 2021; Banerjee et al. 2015; 2016; 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 have substantially larger long-run effects than a series of ongoing transfers with the same net present value if it pushes a household 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. For example, 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.

---

<sup>2</sup> In April 2020, 126 countries introduced or adapted social protection programs as a response to COVID-19 (Rutkowski, 2020). For a full list of social protection programs and jobs responses to COVID-19 see Gentilini et al. 2021.

This paper thus serves three distinct, but related, goals. First, we outline some theoretical frameworks for understanding the social protection challenges in low- and middle-income countries, highlighting their distinctive features relative to high-income ones. We take an expansive view of social protection, including redistribution for the permanent component of income, as well as social insurance programs that help to insure against shocks. We begin by outlining what we think of as the ‘*ex-post*’ problem; that is, after types have been realized, how should the social planner think about improving welfare through redistribution to poorer households? 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, thinking about the tradeoff between using noisy proxies and potentially slanted self-reports. We discuss conceptually how to evaluate such programs, including when there are wedges 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 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 that include the reasons private insurance markets may be incomplete (e.g. adverse selection, moral hazard, commitment issues on the part of the insurer), but then discuss how the problems and solutions may differ in countries with large informal sectors.

Second, we aim to review the growing empirical evidence on these tradeoffs, in order to help consolidate and unify the literature, and to draw connections across the literature. For example, we will discuss the growing literature on understanding targeting questions, on the tradeoffs between in-kind and cash transfers in contexts with weak market integration, on behavioral biases and take-up decisions, leakages from programs and how to combat it, the relationship between transfers and within-household bargaining, and the low demand for insurance.

Finally, in reviewing the literature, we will also 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; whether the rise of formal social protection programs will crowd in or crowd out existing informal insurance arrangements; and how the rise of digital technology and mobile payments will change options for both targeting and implementation.

## **2. How to target for redistribution once income has been realized**

### **2.1: 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. The benchmark public finance model assumes there is unobserved heterogeneity among taxpayers and the social planner would like to tax those with high income to give transfers to those with less. The planner needs to ensure that the tax system considers the incentives for high-income earners to want 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 underlying earning ability, and therefore can only base their decisions on realized

earnings. This then introduces a reason for high earners to claim lower earnings potential and undersupply effort, which is at the heart of the classic Mirrlees model (J. A. Mirrlees 1971).<sup>3</sup>

As highlighted by Saez (2001), this approach identifies the elasticity of labor supply with respect to taxes (or equivalently to the net-of-tax wage rate) and with respect to income as the key parameters in the design of welfare programs. There is an enormous 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 very high tax rate on earnings though a steep phasing out of benefits.

The available literature suggests that labor supply distortions might not be as relevant for developing countries. Two recent (unrelated) experimental studies of the wage elasticity, both from Malawi, report low values of the labor supply elasticity (Goldberg 2016; Guiteras and Jack 2018). Moreover, Banerjee et al. (2017) re-analyzed data from seven different experimental trials of large-scale, conditional cash-transfer programs where earnings were explicitly not part of the conditionality—from Mexico to the Philippines—and find no evidence of income effects on work. Baird, McKenzie, and Özler's (2018) systematic literature review comes to a similar conclusion.

In fact, there is some evidence that the income effect may be positive—being richer actually makes people work harder or more intensively. As pointed out by Benjamin (1992) this could be because of credit and labor market imperfections, but it could also be a result of the kind of psychological mechanisms highlighted by Mullainathan and Shafir (2013).<sup>4</sup> Indeed, there is some evidence that these “multiplier” effects on earnings from relatively large transfers can be both large and durable, so that a (large) temporary transfer can have an outsized effect. To the extent that this is true, it obviously has important implications for the design of social transfers.<sup>5</sup> We will return to this point below.

The fact that both income effects and the wage elasticity are small (at least for small transfers) suggests that it is less important to focus on the endogeneity of earnings in designing transfer schemes in developing countries. On the other hand, in the developing world, governments face challenges observing earnings for informal workers or those with relatively low earnings (which includes most of the poor) (Paul Glewwe 1992; A. Jensen 2022). Most governments therefore worry more about the endogeneity of *reported* earnings in the design of social transfer schemes. Specifically, they worry that there would be a lot of ineligible people claiming benefits if the program eligibility was based solely on self-reported income, with no other incentives put in place to discourage them from doing so.

For this reason, instead of targeting transfers based on income or employment, as in high-income countries, most developing country governments use some combination of three options: targeting transfers based on easily observed characteristics of households (location; ownership of relative

---

<sup>3</sup> We focus here on income as the key metric, following the public finance tradition of Mirrlees (1971). However, the society and the government may have broader metrics of utility, including physical and mental health and other broader measures of well-being, which could also be included in the social utility function.

<sup>4</sup> See, for example, Kaur et al. (2021) and Banerjee, Duflo, and Sharma (2021).

<sup>5</sup> See the discussion in Banerjee, Niehaus, and Suri (2019) for example.

lumpy and durable assets), self-selection, and community information, or some combination of these.<sup>6</sup> In the rest of this section, we discuss the theory and practice of these methods.

## **2.2: Targeting based on observable characteristics**

In a classic paper, Akerlof (1978) discusses how what he terms ‘tags’—i.e. fixed characteristics of people—can be used to improve the targeting performance of self-selection schemes.<sup>7</sup> The idea of using ‘tags’ forms the theoretical basis of what is known as proxy-means testing (more on this in Section 2.4), in which assets and other observable characteristics are used to predict people’s incomes for targeting purposes. This approach forms the basis for targeting in much of the developing world (more on this below).

Clearly, tags that are exogenously determined, such as an individual’s age, can be used for targeting purposes without inducing distortions in consumption or work. However, in order to better predict incomes, most applications of this approach also use observable characteristics that are not completely exogenous, either because people can change their decisions, or perhaps because they can conceal them from government auditors at some cost. Examples of these types of characteristics include what kind of house to live in (e.g. should it have a thatch or tile roof, how many rooms), and whether and what type of television or motorbike to own. Moreover, in practice, this information is often combined with some degree of self-selection—for example, only those who apply for the benefits get evaluated and scored based on the proxy-means test.

While these proxy-means tests allow the government to target the poor without the type of administrative data on incomes used in high-income countries, there are two potential issues. The first issue 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. The imperfect model, combined with potentially noisy data on the tags, means that there are inevitably both inclusion and exclusion errors. 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, in most rich countries the data comes from the tax system and administrative employment records so it gets updated automatically each year). Second, knowing that the targeting is imperfect, even those who are well off may apply for benefits, potentially by misrepresenting their economic status, or even if not, just applying as is and hoping that the errors in the targeting system break in their favor. This too gets worsened by the delays in updating the data. The next section presents a simple model that highlights these issues and their consequences.

## **2.3: A model of targeting**

In this section, we present a simple model that allows us to think through a number of tradeoffs in the design of redistribution schemes in developing countries. This context informs us in two important respects: first, incomes are very hard to measure precisely since much work is in the informal sector (Jensen 2022). Second, labor supply elasticities tend to be quite low, as shown by

---

<sup>6</sup> While we discuss each methodology separately, it is worth noting that many countries use some combination of the methods. For example, Pop (2015) note that many UCTs and CCTs in Africa use a combination of all three targeting methods.

<sup>7</sup> Akerlof (1978) makes it clear that tags can be helpful either if they predict the level of poverty, or if they predict people whose behavior is more or less elastic with respect to labor supply.

Goldberg (2016) and others, which may be driven by the fact that people are already working as hard as they can, given their low income levels.

Our model builds on these two insights: unlike in traditional models of redistribution design (Saez 2002), where the tradeoffs are primarily driven by the labor supply response, we will assume here that the labor supply is fixed and unaffected by redistribution. Instead, the key driver in the model is about the government’s ability to learn people’s incomes in a low-information environment where people can potentially conceal their income. The model therefore explores the design of optimal social transfers in a context in which people can mis-report incomes and the government is only able to detect these mis-reports with some noise.<sup>8</sup>

### 2.3.1: The baseline model

The problem is one of how to best redistribute to a population of low-income citizens. 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 more than some cut off,  $\hat{y}$ .<sup>9</sup> Let  $h(y)$  represent the density of the true underlying income distribution of those who might be eligible for the program, so  $\int^{\hat{y}} h(y) = 1$ .

In the model, we study the degree to which transfers should be universal (i.e., lump-sum) or income based. We assume that the government adopts a linear rule for assigning benefits, so that benefits are given by  $b = T + t(y^* - y^p)$ , where  $T$  is the lump-sum portion and  $t(y^* - y^p)$  is the targeted transfer, based on the difference between the poverty line  $y^*$  and the person’s ‘predicted income’  $y^p$ . The government is subject to a budget constraint such that the total net payouts have to add up to the allocated total budget  $B$ , i.e.  $\int^{\hat{y}} h(y)(T + t(y^* - y^p)) = B$ . We also impose that  $T \geq 0$  and  $0 \leq t \leq 1$ .<sup>10</sup>

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,  $y^a$ , where the superscript  $a$  refers to income information from an ‘audit’:

$$y^p = \alpha \tilde{y} + (1 - \alpha)y^a$$

The ‘audit’ here encompasses whatever procedure the government uses to observe an estimate of incomes, such as a proxy-means test (described in detail below) where the government proxies for incomes based on assets it can observe. The key point is that the audit data is noisy, so  $y^a = y + \epsilon$ , where  $\epsilon$  has expectation of zero and a fixed finite support. We study the choice of  $\alpha$ , i.e.,

<sup>8</sup> For related work exploring optimal targeting schemes, see also Kleven and Kopczuk (2011).

<sup>9</sup> This threshold  $\hat{y}$  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. One could of course set  $\hat{y} = \infty$  and consider the entire population.

<sup>10</sup> We allow for the possibility that those who are close to  $\hat{y}$  actually pay into the system, since  $T$  is allowed to be zero and  $\hat{y}$  is allowed to be more than  $y^*$ . Nothing essential would change if instead we imposed the constraint that  $T + t(\hat{y} - y^p) \geq 0$ .

how much weight to put on self-reported income vs. the noisy audit data. Note that while we study the choice of  $\alpha$ , we do not insist that  $\alpha$  is always set optimally.

Our modeling of the tax payer’s choice problem is inspired by Chetty (2009) who provides a simplified version of Allingham and Sandmo (1972). We work with a model where the constraint on lying too much about your income comes from a potential penalty  $F = \frac{a}{2}(y^a - \tilde{y})^2$ , based on the difference between audited income  $y^a$  and reported income. The penalty is assumed to be something that does not contribute to the program budget—so it could be a psychic cost (e.g. from being called a liar) or a bribe to a government official.

We follow Saez (2002) to introduce the government’s preference for redistributing to the most needy 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  $\int^{y^*} h(y)g(y) = 1$ . However, in addition, we allow for the possibility that the government cares 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.

One important way that we differ from traditional analyses (such as Chetty (2009) and Feldstein (1999)) is that, in our baseline model, we assume that the government’s maximand in choosing the transfer scheme does not take into account the cost that the transfer beneficiary pays as penalty for misrepresenting their income, that is, the government’s objective function does not include  $F = \frac{a}{2}(y^a - \tilde{y})^2$ . This simplifies our analysis while being fairly realistic—at least as long as the cost is primarily psychic (or a bribe), it seems unreasonable to assume that the government is concerned about the welfare consequences of its population being forced to lie about its income (whether or not it should).<sup>11</sup>

### *The welfare recipient’s problem*

Individual utility is given by

$$y + T + t(y^* - \alpha\tilde{y} - (1 - \alpha)(y + \epsilon)) - \frac{a}{2}E_\epsilon(y + \epsilon - \tilde{y})^2. \tag{1}$$

Assuming interior optima, the optimal choice of  $\tilde{y}$  is given by

$$-at + a(y - \tilde{y}) = 0$$

which tells us that

---

<sup>11</sup> We find similar results in the opposite extreme case where the penalty is financial and that the penalties help relax the government’s budget constraint, as in Feldstein (1999). We discuss in the Online Appendix how including these costs as real costs leads to an increased social cost of misreported incomes.

$$\tilde{y} = y - \frac{\alpha t}{a}. \quad (2)$$

*The government's problem in the homogenous case*

We start with the case where all welfare recipients have the same cost of misreporting income, i.e. there is a single value of  $a$  for everyone (we will relax this assumption in Section 2.3.2 below). Under this assumption, and noting as above that individuals' reported income is given by  $\tilde{y} = y - \frac{\alpha t}{a}$ , the social maximand can be written as

$$\int^{\hat{y}} g(y)h(y) E_{\epsilon} u \left( y + T + t \left( y^* - y + \frac{\alpha^2 t}{a} - (1 - \alpha)\epsilon \right) \right) dy \quad (3)$$

The government's budget constraint is given by

$$\int^{\hat{y}} h(y) E_{\epsilon} [T + t(y^* - y + \frac{\alpha^2 t}{a} - (1 - \alpha)\epsilon)] = B$$

which reduces to

$$T = B - t(y^* - \bar{y}) - \frac{\alpha^2 t^2}{a}$$

where  $\bar{y}$  is the mean income, i.e.,  $\bar{y} = \int^{\hat{y}} h(y)y$ .

Combining terms, we can write the social maximand  $W(t)$  as

$$W(t) = \int^{\hat{y}} g(y)h(y) E_{\epsilon} u(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon) dy \quad (4)$$

Several key observations follow from this expression.

*Implications*

To find the optimal level of transfers, we consider  $W'(t)$ . Differentiating equation (4) with respect to  $t$  yields

$$W'(t) = \int^{\hat{y}} g(y)h(y) E_{\epsilon} ((\bar{y} - y) - (1 - \alpha)\epsilon) u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon) dy \quad (5)$$

which can be rewritten as

$$\begin{aligned}
W'(t) = & \int_{\hat{y}}^{\hat{y}} g(y)h(y) (\bar{y} - y)E_{\epsilon}[u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)]dy \\
& - \int_{\hat{y}}^{\hat{y}} g(y)h(y) (1 - \alpha)E_{\epsilon}[\epsilon u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)]dy
\end{aligned} \tag{6}$$

The first term in  $W'(t)$  represents the welfare gains from more effective targeting of the benefits towards the poor when  $t$  is higher. This first term is always positive.<sup>12</sup> Offsetting this, however, the second term in  $W'(t)$  captures the welfare losses that come from the fact that greater targeting of benefits to the poor creates more random variation in the transfers, as a result of the noise ( $\epsilon$ ) in the government's audit function. This noise reduces social welfare so long as there is some weight on the audit term (i.e.  $\alpha < 1$ ) and the government cares about horizontal equity ( $u'' < 0$ ), so if these conditions are true, then the second term is negative.<sup>13</sup>

Several observations follow from the structure of equations (4) and (6).

First, equation (6) implies that a positive slope to the benefits function is always optimal, i.e.,  $t > 0$ . To see this, note that when  $t = 0$ , the second term vanishes, so  $W'(0) > 0$ .<sup>14</sup> Intuitively, this is because the second term captured the social losses from noise in the targeting function, and if there were no income-dependent transfers, the noise would not matter. The fact that  $W'(0) > 0$  then implies that  $t > 0$ ; that is, the social optimum in this model cannot be strictly a universal transfer; some income-dependent component of the transfer scheme is always optimal.

Second, if the government puts no weight on horizontal equity considerations, the social optimum does not feature any lump sum component. To see this, note that when the social welfare function  $u(y)$  is linear so that  $u'(y)$  is a constant, the second term in the expression for  $W'(t)$  drops out. Intuitively, this is because when  $u(y)$  is linear, the noise from the targeting formula does not create any social losses—one person's loss is another person's gain. Thus, when  $u(y)$  is linear, as long as poorer people get higher weight in social welfare ( $g'(y) < 0$ ),  $W'(t) > 0$  for all values of  $t$ . In this case, it is optimal to raise  $t$  as far as possible and set  $T = 0$ , i.e., no universal transfer.

---

<sup>12</sup> To see this, compare the value of the integrand when  $y = \bar{y} + \Delta$  with when  $y = \bar{y} - \Delta, \Delta > 0$ . Clearly, as long as  $g(y)$  is strictly decreasing,  $g(\bar{y} + \Delta)$  is smaller than  $g(\bar{y} - \Delta)$ . Alternately, if  $u(y)$  is strictly concave,  $u'$  is higher when  $y = \bar{y} - \Delta$  than when  $y = \bar{y} + \Delta$ . In other words, either of these conditions is enough to ensure that the value of the integrand is higher when  $\bar{y} - y > 0$  than when it is positive. As a result, the positive terms, i.e., the people for whom  $(\bar{y} - y) > 0$ , dominate and the full expression is positive.

<sup>13</sup> To see this mathematically, note that  $u'' < 0$  implies that  $u'$  is greater when  $\epsilon$  is positive than when it is negative. Thus, the terms where  $\epsilon$  is positive dominate, which implies that  $E_{\epsilon}\epsilon u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon) > 0$  and the entire second term will be negative.

<sup>14</sup> It vanishes because, when  $t = 0$ ,  $E_{\epsilon}\epsilon u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)$  can be rewritten as  $E_{\epsilon}\epsilon u'(y + B) = u'(y + B)E_{\epsilon}\epsilon$ . Since  $E_{\epsilon}\epsilon = 0$ , this implies that when  $t = 0$  the entire second term is 0.

Third, from the expression for  $W(t)$  in equation (4), it is clear that more noise in the audit process (a mean-preserving spread of  $\epsilon$ ) reduces  $W(t)$  as long as  $u$  is strictly convex and the government is using the audit data, i.e. when  $u'' < 0$  and  $\alpha < 1$ . This is because it increases horizontal inequity in the targeting process.

Fourth, when  $\alpha$ , the weight the government places on self-reported data, increases, the noise in the audit data matters less, and social welfare is higher. This can be seen by differentiating  $W(t)$  in equation (4) with respect to  $\alpha$ .<sup>15</sup> This implies that the optimum is to set  $\alpha = 1$  that is, the optimal scheme relies entirely on self-reports, and does not use the audit data at all. Moreover, when  $\alpha = 1$ , the second term in equation (6) vanishes, so  $W'(t) > 0$  for all values of  $t$  whenever  $\alpha = 1$ . Thus, if the government can choose both  $\alpha$  and  $T$ , it will choose  $\alpha = 1$  and  $T = 0$ ; that is, the optimal scheme is based entirely on self-reports, has no universal component, and features the maximum possible income-dependence on transfers.

Finally, note that the expression for  $W(t)$  is independent of  $a$ . In other words, it is possible to achieve the same targeting as in the case where people are fully truthful. This is despite the fact that the self-reported data is always distorted, and the distortion is greater when the tax rate is higher. The reason is that, effectively, the government understands exactly by how much people are distorting their self-reported incomes, and can work around it. (This may not be fully realistic, for reasons we will discuss below.) Thus, with  $\alpha = 1$ , the government not only makes no inclusion or exclusion errors, but it achieves the first-best level of targeting.

We summarize all these results as follows:

**Result 1:** In the baseline model with homogenous preferences over the penalty for misreporting, the optimal redistributive scheme always has the maximum feasible slope with respect to earnings and has no universal component ( $T = 0$ ); that is, the transfers go down at the maximum possible rate as income goes up. 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 their incomes, the amount by which they will distort is predictable, and therefore the optimal targeting mechanism can take that fully into account and avoid using the noisy audit data.

This 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 data, 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.

---

<sup>15</sup> Differentiating (4) with respect to  $\alpha$  yields  $\frac{\partial W}{\partial \alpha} = \int^{\bar{y}} g(y)h(y) E_{\epsilon}(t\epsilon)u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)dy$ . This is positive by the same argument given in footnote 12.

### 2.3.2: 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  $\mu$  denote the fraction of the population that are  $a_1$  types. The cost of misreporting  $a$  is unobserved by the government.

We begin with the case where both types have the same income distribution  $h(y)$ . Given these assumptions, we can rewrite equation (3) as

$$W(t) = \int^{\hat{y}} g(y)h(y) \left[ \mu E_\varepsilon u \left( y + T + t(y^* - y + \frac{\alpha^2 t}{a_1} - (1 - \alpha)\varepsilon) \right) + (1 - \mu) E_\varepsilon u \left( y + T + t(y^* - y + \frac{\alpha^2 t}{a_2} - (1 - \alpha)\varepsilon) \right) \right] dy \quad (7)$$

The government's budget constraint is given by

$$\mu \int^{\hat{y}} h(y) E_\varepsilon [T + t(y^* - y + \frac{\alpha^2 t}{a_2} - (1 - \alpha)\varepsilon] + (1 - \mu) \int^{\hat{y}} h(y) E_\varepsilon [T + t(y^* - y + \frac{\alpha^2 t}{a_2} - (1 - \alpha)\varepsilon] = B$$

which reduces to

$$T = B - t(y^* - \bar{y}) - \mu \frac{\alpha^2 t^2}{a_1} - (1 - \mu) \frac{\alpha^2 t^2}{a_2}$$

Substituting, the social maximand from equation (4) can now be rewritten as

$$W(t) = \int^{\hat{y}} g(y)h(y) \left[ \mu E_\varepsilon u(y + t(\bar{y} - y) + B + (1 - \mu)\alpha^2 t^2 A - t(1 - \alpha)\varepsilon) + (1 - \mu) E_\varepsilon u(y + t(\bar{y} - y) + B - \mu\alpha^2 t^2 A - t(1 - \alpha)\varepsilon) \right] dy \quad (8)$$

where  $A = \frac{1}{a_1} - \frac{1}{a_2} > 0$ .

There is now an additional effect in this model that was not present before: transfer schemes that rely on self-reports, in effect, redistribute from those with high values of  $a$  (those who do not misreport incomes very much) to those with low values of  $a$  (those who do). That is, the government's ability to unravel misreports, to the benefit of those who are willing to lie more and self-report a lower income.

To see this algebraically, note that for any value of  $y$ , the type  $a_1$ 's (i.e. the lower  $a_1$ , who misreport more) are getting a positive shock of  $(1 - \mu)\alpha^2 t^2 A$ , while type  $a_2$ 's getting a negative shock of  $\mu\alpha t^2 A$ . Note that the expected value of the two shocks together is zero since there are  $\mu$  fraction of type  $a_1$  and  $1 - \mu$  fraction of type  $a_2$ . This effect—redistributing towards those who mis-state their income more—occurs to some degree whenever there is heterogeneity in  $a$ , whenever there is a redistributive element to the scheme ( $t > 0$ ), and whenever the scheme relies even a bit on self-reports ( $\alpha > 0$ ). This effect is strongest in precisely the scheme that was optimal in the previous model (maximal  $t$  and complete reliance on self-reports, i.e.  $\alpha = 1$ .)

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

When both forces are present—i.e. when  $Var(\epsilon) > 0$  and  $A > 0$ —it can be optimal for the government to use the audit data, i.e. it can be the case that  $\alpha < 1$ , unlike in the previous model. Interestingly, however, it will always be optimal to rely at least a bit on self-reports, so  $0 < \alpha < 1$ . To see this, note that since the welfare loss from redistribution from  $a_2$  types to  $a_1$  types depends on  $\alpha^2$ , for  $\alpha$  small enough, raising  $\alpha$  has a second order negative effect on  $W(t)$  but first order positive effect by reducing the cost of targeting noise  $\epsilon$ . Hence it is never optimal in this model to completely ignore self-reports, though it can clearly now be optimal to combine self-reports with audit data.<sup>16</sup>

Turning next to the question of how to set  $t$ , it is easy to see that the positive effect of raising  $t$  is still present (a higher  $t$  means more of the transfers go to the poor). But, unlike in the baseline model in Section 2.3.1, now there are 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 the result that  $\alpha = 1$ , there are also losses due to noise  $\epsilon$  from the targeting process, and these are also stronger when  $t$  is higher. Hence relative to the case where  $A = 0$ , as in Section 2.3.1, the optimal value of  $t$  may be lower.<sup>17</sup> 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 audit data as part of the targeting process. The weights placed on audit data relative to self-reports is increasing in heterogeneity in misreporting costs ( $A$ ) and

---

<sup>16</sup> In practice, if there is a fixed cost of collecting the self-report data at all, and if the optimal  $\alpha$  is very small absent these fixed costs, the government may not bother paying the fixed costs and may choose to ignore self-reports altogether. More generally, though, the model suggests combining the two approaches, as both contain useful sources of information.

<sup>17</sup> 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  $t$ .

decreasing in noise in the audit process ( $Var(\epsilon)$ ). Moreover, it can now be optimal for the benefits to have a universal component, and not be entirely income-based.

### 2.3.3: When heterogeneity in misreporting is correlated with incomes

In the previous example, the fact that  $u$  is strictly concave was key to the results. With linear  $u$ , the government no longer cares about horizontal equity, so the mean preserving spread in consumption due to an increase in  $A$  does not matter to the government. However, this is because in the previous section, we assumed that the two types have the exact same income distribution, that is, that the heterogeneity in misreporting costs was uncorrelated with incomes.

To see what can happen when the income distributions are correlated with misreporting costs, consider the following (somewhat extreme) example. Assume as before that there are two types  $a_1$  and  $a_2$  in proportions  $\mu$  and  $1 - \mu$ , but now assume that incomes of the two types are  $y_1 < y^*$  (for  $a_1$ ) and  $y_2 < y_1$  (for  $a_2$ ), and consider the case when  $u$  is linear. Given these assumptions

$$W(t) = \mu g(y_1)y_1 + (1 - \mu)g(y_2)y_2 + T + \mu g(y_1) \left[ (y^* - y_1)t + \frac{\alpha^2 t^2}{a_1} \right] + (1 - \mu)g(y_2) \left[ (y^* - y_2)t + \frac{\alpha^2 t^2}{a_2} \right] \quad (9)$$

and the budget constraint is

$$B = T + \mu \left[ (y^* - y_1)t + \frac{\alpha^2 t^2}{a_1} \right] + (1 - \mu) \left[ (y^* - y_2)t + \frac{\alpha^2 t^2}{a_2} \right].$$

Substituting for  $T$  using the budget constraint, we get

$$W(t) = \mu g(y_1)y_1 + (1 - \mu)g(y_2)y_2 + \mu(g(y_1) - 1) \left[ (y^* - y_1)t + \frac{\alpha^2 t^2}{a_1} \right] + (1 - \mu)(g(y_2) - 1) \left[ (y^* - y_2)t + \frac{\alpha^2 t^2}{a_2} \right]$$

Using the fact that  $\mu(g(y_1) - 1) = -(1 - \mu)(g(y_2) - 1)$ , we can rewrite this as

$$W(t) = \mu g(y_1)y_1 + (1 - \mu)g(y_2)y_2 + (1 - \mu)(g(y_2) - 1) \left[ (y_1 - y_2)t + \frac{\alpha^2 t^2}{a_2} - \frac{\alpha^2 t^2}{a_1} \right] \quad (10)$$

Finally, note that:

$$W'(t) = (1 - \mu)(g(y_2) - 1) \left[ (y_1 - y_2) + \frac{2\alpha^2 t}{a_2} - \frac{2\alpha^2 t}{a_1} \right] \quad (11)$$

Two observations follow from this revised expression for  $W'(t)$ . First, as long as  $a_2 < a_1$  this expression will be always positive, so it makes sense to maximize  $t$  and set  $T = 0$ , just as in the homogenous case with linear social welfare function.

Second, if  $a_2 > a_1$ , then there is a clear benefit from setting  $\alpha = 0$ , since from equation (10)  $W(t)$  is decreasing in  $\alpha$  whenever  $a_2 > a_1$ . Setting  $\alpha = 0$ , it is also clear that one should set  $t$  as large as possible and set  $T = 0$ .<sup>18</sup> We summarize this as follows:

**Result 3:** If preferences are linear, it may even be optimal to ignore self-reports entirely and to use just audit data. This will be the case if those who have higher incomes have a lower cost of misreporting.

### 2.3.4: Summary

Our model, though extremely stylized, helps highlight the various tradeoffs faced by a government in designing a transfer system in the presence of noisy data. To start with, the optimal transfer scheme will typically involve both inclusion and exclusion errors, for several reasons. First, with self-reporting, beneficiaries may misrepresent their incomes, but potentially differentially so, such that some who are more deserving might get less than those who are not. Second, the fact that the government responds to the imperfections in self-reported data by using its audit data will introduce a second source of noise into the decisions, which generates a second source of inclusion and exclusion 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 heterogenous in terms of the cost of mis-reporting, 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 a substantial universal component in social transfers, because of a combination of the concern for equity (concave  $u$ ), heterogeneity in misreporting costs ( $A > 0$ ), and noise in the audit process ( $Var(\epsilon) > 0$ ). Finally, it helps us understand under what conditions noise in the audit data can be very costly.

## 2.4: Targeting on Observable Characteristics in Practice

---

<sup>18</sup> Note that if there is a lower bound on  $\alpha$  (perhaps for some institutional reason), and  $t$  is large enough that

$$(y_1 - y_2) + \frac{2\alpha^2 t}{a_2} - \frac{2\alpha^2 t}{a_1} < 0,$$

then it does not make sense to raise the tax rate further. Instead, the optimal value of  $t$  in that case is actually the one that sets

$$(y_1 - y_2) + \frac{2\alpha^2 t}{a_2} - \frac{2\alpha^2 t}{a_1} = 0$$

which is to say that it equalizes after tax incomes. In other words, the incentive to mis-report should never be so strong that the order of reported incomes reverses the order of actual incomes, at least in this two-type case.

The simplest version of characteristic-based targeting is **geographic targeting**, i.e., targeting programs to poorer regions. Its main attraction is its simplicity in selecting regions and the administrative ease of implementing it—since one does not need individual data, one can actually do the entire exercise (prediction and targeting) with a single, representative dataset. Using finer geographic areas rather than targeting larger areas can improve poverty reduction (Baker and Grosh 1994; Elbers et al. 2007). For example, one can combine survey sample data with census data to predict the poverty status of smaller geographical areas (Elbers, Lanjouw, and Lanjouw 2003), or use newer remote-sensed satellite or administration datasets that help predict smaller regions (e.g. Jean et al. 2016; Blumenstock, Cadamuro, and On 2015).<sup>19</sup> Note, however, this approach produces less accurate estimates than individual targeting, but is much less data intensive.

A more ambitious form of targeting, which forms the basis for targeting in many low- and middle-income countries, is the **proxy means test** (PMT). The real world implementation of a PMT is based on two datasets. First, a training dataset that includes measures of what the government is actually trying to target (e.g., per-capita household consumption, income, a dummy for being below the poverty line, etc.), as well as the characteristics ( $X'$ )—usually demographics and verifiable assets—that will be used for targeting. This training dataset is usually a nationally representative survey of at least several thousand households that is collected for research or general statistical purposes (such as estimating national poverty lines); since it is for research purposes, households typically have no reason to lie about their actual consumption levels or income. Using these data, one then estimates a model of the form  $y_i = X'_i\beta + \epsilon_i$ .

To target, the second step is to obtain a much larger dataset of just 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 is the step where there is often some self-selection—for example, in some cases, the government will collect the characteristics  $X_i$  only for those who apply, and hence only those who apply may be deemed eligible; in other cases, the government attempts to collect the  $X_i$  for everyone, or everything they suspect may be eligible, 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  $i$  in this population that is then used for targeting.


This type of proxy-means testing was first used in the early 1980s in Chile for the targeting of its Ficha CAS program, and has become quite common (Coady, Grosh, and Hoddinott 2004). Well-known examples of targeting systems based on proxy-means tests include the SISBEN in Colombia, SISFOH in Peru, Listahanan in the Philippines, BISP in Pakistan and Takaful and Karama in Egypt.

---

<sup>19</sup> This approach, however, depends on a homogeneity assumption that supposes that the small area that is being targeted has the same conditional distribution as the larger area used to calibrate the imputation rule, which might not always be true (Tarozzi and Deaton 2009). One also needs to think carefully about the error properties in these geographic ‘poverty maps.’

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 corresponds to the audit data in the model, are imperfect predictors of poverty. To illustrate this, Hanna and Olken (2018) simulate this type of individual targeting for two countries—Indonesia and Peru—using household survey data and set of assets and household characteristics typically used in proxy-means tests. Figure 1, reproduced from that paper, illustrates the basic challenges. It is clear that in cases the prediction equation has substantial predictive power—the  $R^2$  is between 0.53 and 0.66—but it is not perfect. A substantial fraction of households are incorrectly excluded from the program, and another set is incorrectly included. Other papers show similar findings across different contexts: for example, Brown, Ravallion, and van de Walle (2018) show, using data from nine African countries, that standard PMTs help filter out non-poor households, but also exclude some poor households as well.

**Figure 1: Predicted vs. actual per-capita consumption in test set data, from Hanna and Olken (2018)**



*Notes:* The figures above plot actual log per-capita monthly consumption against predicted log per-capital monthly, where the prediction is based on a set of household assets and characteristics typically used in proxy-means tests. See Hanna and Olken (2018) for more information.

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. Such households will be systematically penalized by the PMT algorithm. As emphasized in the model, such heterogeneity in the link between true income and PMT-predicted incomes 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. McBride and Nichols (2018), for example, examine using random forests to estimate the PMT using data from Bolivia, Timor-Leste, and Malawi, as compared to quantile regression. Overall, they find minimal reductions in targeting error using random forests as compared to the quantile regression, but modest improvements (2-17 percent) using other assessments of targeting accuracy. Baez, Kshirsagar, and Skoufias (2020) and Areias and Wai-Poi (2022) also find little additional benefit of a wide variety of machine learning approaches across countries in Africa. That said, since using machine learning instead of linear regression is essentially free, it seems useful to explore these approaches, even if they do not radically transform PMT accuracy.

Second, one can use new types of administrative data rather than collecting a door-to-door census, under the idea that you can save money and time by not having to collect a whole new dataset. For example, Blumenstock, Cadamuro, and On (2015) use anonymized data from mobile phone networks to predict the socioeconomic status of individual subscribers; similarly Abelson, Varshney, and Sun (2014) use remote sensing to differentiate between thatched and metal roofs for targeting 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 would be obtained from a more traditional full PMT.

Perhaps equally important, administrative or remote-sensed data is collected continually and automatically, unlike census data, which must be collected manually and therefore is only collected every few years at most. This means that targeting based on this type of administrative data may be much more responsive to shocks.<sup>20</sup> Some countries, from Pakistan to Togo, used this type of administrative data to target cash assistance during the 2020 COVID-19 crisis (Gentilini et al. 2020); understanding whether this was effective, and how to do so well, is an important area of ongoing and future work.

### *Endogenous characteristics?*

A further challenge, as highlighted by the model in Section 2.3, is that many characteristics used in this type of targeting are, themselves, choice variables, and this introduces an additional potential source of inefficiency, since households need to distort their consumption to remain eligible.<sup>21</sup> In one famous example, after England imposed its “window tax” in 1696—a tax levied on the number of windows a house had, which was an easily observable characteristic that proxied for wealth—people started building 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 these types of geographic targeting incentives. But, 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) conducted a nationwide randomized experiment, in cooperation with the Indonesian Government's Central Bureau of Statistics, which administers the census used for actual PMT

---

<sup>20</sup> In fact, others (e.g. Baez, Kshirsagar, and Skoufias 2020; del Ninno and Mills 2015) suggest augmenting PMT targeting system with administrative data of this sort to help capture shocks.

<sup>21</sup> Specifically, many of these, such as assets (like house type), are choices that are made infrequently and at some cost, but could still, in principle, change and lead to distortions.

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, to test whether this changed the consumption of those assets. While self-reports of flat-screen television ownership fell in treated provinces six months later, that 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 2.3, 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 not the only interpretation. First, recall that in the model,  $a$  referred to the cost of changing your reported income, not your 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, Hanna, et al. (2020) result had hints of exactly this; reported ownership of televisions declined in government data 6 months later, even though actual television ownership was unchanged. Similarly, Martinelli and Parker (2009), for example, 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 (that is,  $y^a - \tilde{y}$ , in the language of the model) when the monetary stakes are higher.

Beyond this, there is a second, important real-world consideration. Specifically, in the model, we assume that the household knows exactly what to do to claim the benefits. In practice, the fact that there are so many different variables, interacting in complex ways in the PMT formula—on top of the fact that the formula itself is typically kept secret—means that households have limited control over the outcome and therefore the returns to manipulation are small. Consistent with this, Camacho and Conover (2011) examine the proxy-means test in Colombia, and test whether there is bunching in proxy-means test scores under the eligibility threshold. With no manipulation, one would expect the distribution of scores to be continuous around the threshold, and this is indeed what they find for the first four years the PMT was in place. They do, however, detect bunching starting right after the algorithm used to construct the scores become known, and they suggest that local politicians used their knowledge of the formula to get more people in their municipalities' federal transfers. To prevent this type of manipulation, most governments keep the PMT targeting formulas secret and update them periodically, so the typical case may be closer to the first years of the Colombian program, when the formula was secret and there was little manipulation.

Though there is limited evidence so far of actual manipulation of consumption, one could imagine cases in which it may be. To the extent this becomes an issue, recent theoretical advances suggest ways of operationalizing the simple theoretical idea that is highlighted by our model—that it is possible to improve targeting by anticipating problems and taking the potential responses of agents into account (Björkegren, Blumenstock, and Knight 2020; Ball 2022). To the extent that the manipulation of particular assets is a concern in certain contexts, operationalizing these various approaches could be a useful solution. It is also worth noting the converse: simpler, more transparent rules can also make manipulation easier (Niehaus et al. 2013; Camacho and Conover 2011a).

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) show that India's public works program, MGNREGA, the Mahatma Gandhi National Rural Employment Guarantee Act—which was geographically targeted to rural areas—reduced seasonal migration to urban areas. One reason geographic targeting may create distortions, is that it is coarse, transparent, and easy to understand. Since one often expects higher productivity in urban sectors (e.g., Harris and Todaro 1970; Lagakos 2020), this type of coarse rural-specific targeting, while it may have been useful as a targeting device, runs the risk of reducing aggregate efficiency.

## 2.5: Self-selection mechanisms

An alternative way to distinguish between rich and poor households is through **self-selection mechanisms**. As already pointed out in the model section, any mechanism that involves taking a certain action to be compliant with the entitlement rules (which includes choosing certain assets to manipulate the PMT score) is in effect a self-selection mechanism. The reason why this can work is if the required action is easier for the targeted population (the poor) than for the rest (the non-poor), an idea going back at least to Nichols and Zeckhauser (1982). If so, then as described by 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, living in less nice neighborhoods, having to work in order to receive the transfer, but also ordeals like standing in line, filling out lengthy paperwork or having to go through an interview process. The disadvantage of this approach is that it imposes costs on the poor—i.e. lower quality food than they would prefer, or time wasted standing in pointless lines (this is the welfare loss discussed in the Online Appendix). The key question is whether the targeting benefits, which would allow governments to deliver more aid to the poor, outweigh these additional costs, so that both the poor individuals themselves, and society as a whole, are nevertheless better off even once these costs are taken into consideration.

*Workfare.* One of the most common ordeal mechanisms is the use of workfare—i.e. requiring beneficiaries to work in exchange for payments. These programs have a long history: they have been around as early as the 19<sup>th</sup> century in England, where transfers were granted through residence in a workhouse (Besley and Coate 1992). In the United States, they began at a large scale with the Civilian Conservation Corps and Works Progress Administration, which collectively employed millions of jobseekers to carry out public works projects during the Great Depression (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 unpleasant or difficult, this will generate self-selection: anyone who can get a more attractive job (either in terms of compensation or job conditions) will select out (Ravallion 1991; Besley and Coate 1992). But, if the wage is set too high, the workfare 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 (in the language of the model, do we add back the value of the work to the social welfare function, or not). Using workfare labor to build needed roads and irrigation canals could actually have very small efficiency loss; having workfare labor dig ditches and fill them in again would also generate self-selection, but at much larger social efficiency costs.

Perhaps the largest workfare program in the world is India's Mahatma Gandhi National Rural Employment Guarantee Act (also known as MGNREGA), which offers 100 days of paid employment per year to anyone in rural areas who is willing to do casual manual labor. MGNREGA, in turn, built on previous workfare schemes, such as Maharashtra's Employment Guarantee Scheme (Ravallion 1991). MGNREGA's wage rate is tied to the official state minimum wage, which may be above the de facto wage available to rural agricultural laborers, given the large informal economy. Nevertheless, Dutta et al. (2014) find that the participation rate for MGNREGA falls steadily from 35% of the population for the poorest wealth percentiles to around 10% for the richest, suggesting substantial self-selection, with the declines largest in the top half of the income distribution.

Given that workfare, by its nature, crowds 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. Their simulations suggest that a simple targeted transfer, using pre-existing ration cards, or a basic-income scheme would result in slightly higher poverty reductions than the workfare program, given both the cost of the workfare in terms of the government's expenditures and these forgone wages.

Similarly, Bertrand et al. (2021) examine related questions in the context of a youth workfare program in urban Cote d'Ivoire. Since the program was oversubscribed, they randomize allocation to the program. They find that while the program led to earnings gains, because it crowded out lower-wage private employment, earning gains are only 53 percent of the transfer. As discussed more in below, the program nevertheless had impacts—on composition of work and savings—but the crowd-out suggests that if the only value of the work requirement was for targeting purposes, it is not particularly efficient.

Finally, it is important to note that one common theme to workfare programs is that provided employment is often set at the formal minimum wage, which—unlike in many developed country settings—is often *above* the prevailing wage, given the large informal sector. This may be a fundamental constraint: governments, after all, typically cannot pay below their own minimum wage. But, this suggests that these programs may have more effective targeting properties in places where the formal minimum wage is lower and/or where there is less informal employment.

*Self-selection with small costs.* One challenge with ordeal mechanisms such as workfare is that they impose very large costs on beneficiaries—working a full day of hard labor under the hot sun for a day's wage, for example. Is it possible to get the gains from self-selection without imposing such large costs? Several recent studies suggest that this may be possible.

*Applications as a selection mechanism.* One possible selection mechanism is to make prospective applicants apply for a program, rather than be automatically enrolled.

Once beneficiaries apply, programs often add additional screening mechanisms, such as a PMT or additional forms of demographic and means testing.<sup>22</sup> 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. But second, and perhaps much more importantly—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 beneficiaries have imperfect information or beliefs about the screening mechanism, asking them to decide to apply or not can, in effect, induce them to reveal some of this private information to the government.

These two selection mechanisms above suggest that this approach has the potential to improve targeting and reduce program costs by reducing inclusion error, without imposing the types of large costs imposed by workfare programs.<sup>23</sup>

Alatas et al. (2016) develop a model that captures this idea, and then examine this experimentally in a field experiment in Indonesia. They examine targeting outcomes in villages where the government's conditional cash transfer program, PKH, was expanding for the first time. Villages were randomized into two groups: an automatic-enrollment group, in which statistics department officials went door-to-door to administer a proxy-means test using their normal procedures, and a self-targeting group, where people needed to apply for the program. In the self-targeting group, those who applied were verified using the same PMT formula as the automatic enrollment group, so the key difference is the application step—which entailed a few hours of an applicant's time.

Alatas et al. (2016) found that self-selection led to dramatically poorer beneficiaries—those chosen for the program in self-selection villages were about 20 percent poorer than those selected in automatic enrollment villages. They show self-selection on both observables—i.e. those who the PMT would deem ineligible are less likely to apply—and on unobservables—i.e., those who are wealthier on dimensions that would be missed by the PMT are also less likely to apply. Surprisingly, this approach *also* reduces exclusion error compared to the PMT—while the government makes its best efforts to include all relevant people in the automatic enrollment system, some of the very poor, who live more on the margins of society, can be missed; in the self-targeting group, these individuals make themselves known, increasing their take-up. This suggests that self-selection at the application stage can allow for the revelation of important information, potentially decreasing both inclusion and exclusion error.


One can see this in Figure 2 below, based on Alatas et al. (2016), which shows the probability of receiving benefits for those in automatic enrollment vs. self-selection villages, as a function of per-capita consumption. The figure shows that self-selection led to both higher probabilities of receiving benefits for the very poor and lower probabilities of receiving benefits for the very rich.

---

<sup>22</sup> For example, in the U.S., applicants apply for Unemployment Insurance and then only become eligible after their employment status is verified.

<sup>23</sup> However, this type of selection by making people apply has its own issues—if the application process ends up discouraging the poor from applying (because of stigma, or time costs, or fear of bureaucracy, or other reasons), then this type of mechanism could actually worsen targeting. We will return to this point.

**Figure 2: Probability of receiving benefits for those in automatic enrollment vs. self-selection villages, based on Alatas et al. (2016)**



Notes: The graph shows the probabilities of receiving benefits in both the automatic enrollment and self-targeting treatments described in Alatas et al. (2016) (right axis), along with a histogram of log per-capita consumption (x-axis). See Alatas et al. (2016) for more details.

## 2.6: Community-information based targeting

An alternative approach to targeting relies on the fact that there may be local information about people’s poverty status that is not captured very easily by data (Alderman 2002; Galasso and Ravallion 2005; Alatas et al. 2012). This is related to a common theme across development economics about the value of local information, from the informal insurance literature (e.g. Townsend 1994) to the literature on entrepreneurship potential (e.g., Hussam, Rigol, and Roth 2022). But, if governments try to use this information for **community-based targeting**, is the outcome better than a proxy-means test based approach?

To examine this, Alatas et al. (2012) conducted a randomized trial in cooperation with the Indonesian statistical agency across more than 600 villages, that compared proxy-means test-based targeting to community-based targeting, as well as examining a hybrid of the two approaches. To obtain an independent measure of well-being, a baseline survey was conducted to measure household consumption, as well as measures of self-reported well-being. The survey was done for research purposes only, so households would have no incentive to misreport.

The results suggest that the optimal approach may depend on the government’s objective function. On the one hand, the authors found that while the proxy-means test did somewhat 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 with consumption below PPP \$2/day, but 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. However, community targeting did substantially better in terms of identifying households who self-identify as poor. Perhaps as a result, citizens in villages randomized to community targeting reported higher satisfaction with both the targeting process and the targeting outcomes than in those with PMT-based targeting.<sup>24</sup>

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 (Bardhan and Mookherjee 2000; Acemoglu 2006; Acemoglu, Reed, and Robinson 2014).<sup>25</sup>

Empirically, however, elite capture seems to be small, at least in comparison to other sources of targeting failures. Alatas et al. (2019) find that when local elites are the ones making targeting decisions, they and their relatives are typically not more likely to receive aid programs than non-elites. Even in the cases with some elite capture, the welfare losses appear to be small, because they are proportional to the *product* of three relatively small numbers: the additional boost that relatives of elites receive in receiving transfer programs (19 percent in the very highest program considered; substantially less or zero on most programs), how much richer these elite relatives are than the general population (about 9 percent), and their proportion in the general population (about 15 percent). By eliminating elite capture entirely, they estimate that welfare gains would be less than one percent. By contrast, feasible targeting improvements—just implementing the existing PMT more reliably—could lead to welfare improvements of between 17 to 57 percent, depending on the program.

Basurto, Dupas, and Robinson (2020) find similar results studying the targeting of subsidies for both agricultural inputs and food in Malawi, which is done by local chiefs, and which they compare to hypothetical targeting done through a PMT. As in Alatas et al. (2019), they find that while nepotism exists—i.e. chiefs are somewhat more likely to target their kin—the welfare consequences of this are very small, since the kin who receive the subsidies are just about as poor as other would-be beneficiaries. They also find that the chiefs use their local knowledge to target

---

<sup>24</sup> The gain in legitimacy from community targeting may be quite important. For example, Cameron and Shah (2014) show that mistargeting in cash transfer programs in Indonesia increased crime and led to less participation in community groups.

<sup>25</sup> While it is not theoretically obvious that capture should be worse at the local level than at the national level, the idea that local government may feature more extensive elite capture goes back at least to the *Federalist Papers*, where in *Federalist Papers No. 10* James Madison asked whether “Men of factious tempers, or of local prejudices, or of sinister designs, may, by intrigue, by corruption, or by other means, first obtain the suffrages, and then betray the interests, of the people. The question resulting is, whether small or extensive republics are more favorable to the election of proper guardians of the public weal; and it is clearly decided in favor of the latter....”

those recently hit by shocks, and in the case of agricultural inputs, to target them to those with the higher returns.

While these studies suggest that community targeting has potential, especially when assessed against local objective functions (such as self-assessed welfare) or to fill in the gaps in 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.

Premand and Schnitzer (2021), for example, run a randomized trial in Niger with a design similar to Alatas et al. (2012), randomizing villages about to receive a national cash transfer to be targeted using PMT, community-based approach, or a formula that aims for the food insecure. In their study, they find that households selected by the PMT are 8 percentage points more likely to be poor, and have 15 percent lower consumption, than households selected by community-based targeting. This is a substantially larger gap than Alatas et al. (2012). Like Alatas et al. (2012), Premand and Schnitzer (2021) 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. They also find that community-based targeting is more effective at identifying households with recent shocks. Legitimacy is similar for both types of targeting regimes.

Several other studies 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-a-vis self-assessed welfare (Stoeffler, Mills, and del Ninno 2016 in Cameroon; Beaman et al. 2021 in urban Monrovia, Liberia; Dupas, Fafchamps, and Houeix 2022 in urban Cote D'Ivoire). Other studies suggest the efficacy of community targeting compared to a PMT depends on the variables included in the PMT; for example, Sabates-Wheeler, Hurrell, and Devereux (2015) experimentally show that community-based targeting performs better at identifying the poor than simple proxy measures of either dependency ratio targeting and categorical age targeting, but in simulations show that a more standard PMT could have dominated all three measures.

These papers suggest several important directions for future research. First, are there ways to make community selection more effective? For example, Alatas et al. (2012) report an intriguing fact. They had randomized the order in which households were ranked by the community in terms of income purely for fairness reasons, but using that randomization they find that households ranked early in the process are ranked much more accurately than those ranked later, suggesting that the community finds the process tiring. Finding ways to combine multiple rounds of ranking might improve the effectiveness of community ranking very substantially. More generally, it suggests that the micro-structure of the algorithm used to elicit the community ranking can affect the accuracy of the overall community ranking.

A second natural question is whether the efficacy of community targeting differs systematically. For example, one question is whether it is more effective in communities with more local information. Alatas et al. (2016), for example, studies this question, finding that at the aggregate

level, there is more information in villages with denser networks, using a variety of metrics of network density. These network characteristics matter for targeting efficacy: the community targeting treatment in Alatas et al. (2012) was more effective at identifying households with lower self-assessed welfare in areas with more diffusive networks. More generally, being able to predict in which types of locations community-based targeting is likely to be most effective is an important direction.

A third important question is whether and how 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 the community 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 transfers from a higher level of government. Designing incentive-compatible procedures for community-based targeting that reveal the level of poverty, not just who is poor, seems like an important mechanism design problem for future research.

## **2.7: Incomplete take-up, and what we can learn from it.**

The flipside of self-targeting—through which take-up decisions can be used to improve targeting—is that many people do not take up the program, even when they are entitled to it. How do we think about this lack of take up? 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?

Currie's (2004) review on social programs in the United States concludes that other costs associated with participation (such as transaction costs and complicated application processes) are more important than stigma.<sup>26</sup> But this issue has received comparatively less attention in the developing world. How important is this phenomenon, and what explains it?

### **2.7.1. How prevalent are take-up challenges?**

The fact that not everyone takes up programs that they are offered is not per se surprising. Self-selection costs vary across people—finding appropriate childcare in order to spend an afternoon signing up for a program might be more difficult for some, for example. Moreover, not everyone who we think is poor is actually as 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, a survey by the World Bank in 2014 (S. Bhattacharya et al. 2015) found that two-thirds of eligible women to a non-contributory social

---

<sup>26</sup> In the U.S., a recent study by Finkelstein and Notowidigdo (2019) who look at experimental evidence from SNAP (Supplemental Nutrition Assistance Program, commonly known as Food Stamps) suggest that stigma costs of participating would have to be 60 times larger than transactional costs of applying to explain low observed take-up. An exception to these papers is Schanzenbach (2009) who finds that clients were significantly more likely to be interested in SNAP when the word “food stamp” was not used. One reason why it might be especially complicated to detect the extent to which stigma matters for take-up is that the two are jointly determined. As more people participate in a specific program, it becomes harder to single out specific individuals for stigma because they use the program (Ewoudou, Tsimpo, and Wodon 2009).

pension for life are not enrolled in the program. The pension, given to young widows and divorced women, gives out Rs. 1,500 per month, which is equivalent to 70% of median per-capita consumption across urban areas. Given the size of the transfer, it seems unlikely that this could be rationalized by a cost-benefit calculation. Also in India, Demirguc-Kunt, Klapper, and Prasad (2017) find that only 40% of citizens apply for the goods and services they report needing from the government. Perhaps more interestingly, people miss out on benefits even after signing up for them. Hirshleifer et al. (2016) and Alfonsi et al. (2020) provide evidence from Turkey and Uganda respectively, where unemployed workers who enroll in training programs do not show up. And while our work on on-demand application described above (Alatas et al. 2016) showed that this approach found more of the poor than the government's PMT approach, which missed some of the very poorest, still only about 60 percent of the very poorest applied for benefits.

### **2.7.2. Application costs and information barriers**

The fact that so many of the poor are missing out on benefits is suggestive that some fraction of the non-take-up might be inefficient. Finkelstein and Notowidigdo (2019) argue that in the US, misperception of benefits and costs might explain a part of the lack of take up. If this is the case then it is useful to correct these misperceptions. And indeed, they find that information about benefits nearly doubles take up (from a low base), and when combined with help with signing up (which is useful because there is an unfounded fear of just how hard it is to apply) can triple take up.

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 for programs by providing enhanced take-up assistance to targeted households. A common theme from these studies is 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 households that were in the bottom 5 percent of the population, but not enrolled in government assistance programs. They provided extensive home visits by a social worker—21 visits of 45 minutes each—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. They then use the Chile Solidario eligibility cutoff and a treatment-on-treated design, which captures the local average treatment effect *for the subpopulation who took up Chile Solidario*, whom one may expect to be a particularly responsive subpopulation. They find that, 4 years later, the assistance from Chile Solidario increased the take-up of a separate government family child allowance program (“SUP”) by 17 percentage points, from a base of 53 percent. The intensive take-up assistance of Chile Solidario made a difference, 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 to which they are entitled. The bureaucracy around application is a substantial hassle: widows must fill out forms, provide supporting documentation (proof of marital status, unique ID card, proof of 5-year residency in Delhi, and bank account information), 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 obtain the pension, compared to 15 percent of control households. Viewed one way, this is a 47 percent increase; but viewed another way, it still leaves 80 percent of eligible women. Remarkably, about only  $\frac{1}{3}$  of those induced to start the application actually successfully receive the program, and the assistance particularly helped more advantaged households (among this very disadvantaged population) — e.g. literate households are more likely to benefit from the program.

Banerjee et al. (2021) examine related questions in Indonesia in the context of the government's national health insurance scheme, which many households do not take up even though it is mandatory according to the law. Universal health insurance was launched in 2014 and by 2015, the contributory portion of the program, known as JKN Mandiri, had enrolled less than 20 percent of the intended population. Banerjee et al. (2021) focused on the uninsured and examined the role of intensive application assistance, as well as subsidies. The application assistance offered households help to apply for the insurance online from home, with enumerators walking households through the online process using the enumerator's computers and internet connection. They found that while application assistance increased *attempted* enrollments substantially—by 16 percentage points for those who did not receive any financial subsidy, compared to just 1.8 percent in the control group—as in Gupta (2017), the vast majority of those who attempted to enroll did not successfully do so. As a result, the assistance alone only increased enrollment by 2.4 percentage points. This means that over 85 percent of those induced to apply by the application assistance did not successfully do so. In this case, a main reason was that there were many problems in the government's underlying administrative data, which would have required a trip to a distant government office to solve—which households did not do. The underlying administrative capacity challenges thus formed a substantial take-up barrier that simple application assistance could not solve.

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 the requirements to participate—say by pre-filling the forms using administrative data may be helpful. 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 an inconvenient trip to a government office (or maybe multiple government offices). Improving this underlying administrative capacity to allow for simpler enrollment procedures may be challenging, but important.<sup>27</sup>

### 2.7.3. Stigma

---

<sup>27</sup> It is important to note that none of these papers that address take-up—to the best of our knowledge—address the question of whether social welfare is improved or not by these marginal changes in take up. Finkelstein and Notowidigdo (2019) argue that if households are fully optimizing and rationally not signing up for programs, there is no net welfare gain to households from an intervention that reduces signup costs, even if targeting improves in the sense that more poor people participate in the program, because the households induced to sign up by a marginal reduction in signup costs were indifferent, on net, between receiving the program and not. Of course, this depends on the fact that households are fully optimizing; if they over-estimate the cost or underestimate the benefits from applying, interventions that improve targeting also improve welfare (Finkelstein and Notowidigdo 2019).

A common refrain in the developed world is that low take-up reflects low demand for government services due to the stigma associated with them (Moffitt 1983). While there is comparatively little work investigating stigma in developing countries, the research that there is suggests that ‘welfare stigma’—i.e. the idea that take-up of a program connotes poverty—does not find particular support.

In particular, Osman and Speer (2020) conduct a series of experiments in Egypt to recruit young people to labor market assistance programs, varying the message associated with recruiting to distinguish the effects of different types of stigma. Several messages designed to reduce “professional” or “social stigma” actually seemed to have primed respondents to be more alert to these concerns.<sup>28</sup> However, an explicit ‘welfare stigma’ framing—i.e. focusing on the idea that the cost of the job training program was subsidized “to help those in financial hardship” —had no effect whatsoever on program take-up.

Understanding the role of stigma in more detail—and in particular whether or not program participation is public information or not—may be an important direction for future research.

## **2.8: Targeting on Treatment effects, not levels**

The discussion thus far has focused on targeting specific characteristics of households that make them desirable from the point of view of the government. That is, we are interested in trying to identify the poor, or those with low earnings ability, in the Mirrlees (1971) formulation, and assign the program to those people. 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 the treatment effects of a program. To simplify matters, imagine that a program’s only effect is to generate income, but that 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 specific treatment effect and income. Even for programs that don’t target income, the point is that one is no longer interested in targeting the program just based on an observable characteristic (i.e. poverty  $y_i$ ), 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. D. Bhattacharya and Dupas (2012), for example, consider this approach in the context of subsidies for anti-malarial bednets, where they account for the probability of use of the bednets. More generally, recent work has shown how to use machine learning to estimate heterogeneity in treatment effects (e.g., Athey and Wager 2017; Wager and Athey 2018; Chernozhukov, Demirer, et al. 2018; Chernozhukov, Chetverikov, et al. 2018). Chernozhukov,

---

<sup>28</sup> The authors define welfare stigma as the disutility that comes from participating in a program meant for the poor or disadvantaged (the same way Moffitt 1983 describes it); social stigma is associated with believing that society, family and potential marriage partners look down on the program; professional stigma is the belief that by signing up to the entry-level jobs offered might hinder career progress; while personal stigma refers to an internal sense of disappointment associated with performing a job that is not rewarding.

Demirer, et al. (2018) for example, show how to use these approaches to identify heterogeneity in treatment effects for an immunization program. This could then be used to target the program to those locations where it is most effective. Bertrand et al. (2021) examine a workfare program in Cote D'Ivoire, and find, using these techniques, that targeting the program just to women or to youths with low predicted baseline earnings would improve the program's cost-effectiveness by 30-52 percent. Caria et al. (2021) 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, in the context of a job training program for Syrian refugees in Jordan. Haushofer et al. (2022) examine these issues in the context of a cash transfer program in Kenya, where there may be heterogeneity in how households use the cash transfer in terms of how it affects their future income, and find that targeting based on estimate of  $dy_i$  appears more effective than targeting based on  $y_i$ ; one might imagine that heterogeneity in  $dy_i$  would be even larger in programs other than pure cash transfers.

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, where there are concerns that some may not use the chlorine. They compare free distribution of chlorine vs. distribution of vouchers redeemable at nearby shops for free chlorine. Redeeming a voucher is a screening mechanism—in general, only those who intend to use the chlorine will bother redeeming the voucher. 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 bother to use it.

These examples illustrate the potential to apply these approaches to social protection more broadly. For example, the Graduation approach (discussed 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, Finkelstein, et al. 2021a). We regard targeting on treatment effects as an important dimension for future research.

## **2.9: Should we even target at all? Universal vs. targeted program**

All of the targeting methods that we have discussed thus far suffer from exclusion and inclusion errors (Robles, Rubio, and Stampini 2019; Brown, Ravallion, and van de Walle 2018). Given these challenges, another option is to make programs universal, for example universal basic income (UBI) programs or a universal in-kind program. That is, each individual receives a fixed amount of money or goods—regardless of income—which is financed through proportional or progressive taxation. Conceptually, one can think of these programs 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 budget constraint (Saez 2002).

Universal transfers have several advantages. In principle, they have no exclusion error, though in practice take-up may be far from universal even for a program with universal eligibility. As such, they also avoid horizontal inequality problems (i.e. the challenge that similar people may not

receive like benefits). UBIs also have the potential to reduce administrative costs (by not needing to collect data on households), 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 (by reducing the scope for corruption and other abuses of power) (Banerjee, Niehaus, and Suri 2019).

One challenge, however, is that for a given budget, universal programs typically transfer substantially less resources to each beneficiary than targeted programs. That is, solving the problem of eliminating exclusion error by transferring programs to all can end up being quite expensive (Hoynes and Rothstein 2019). While this is true everywhere, this tradeoff is particularly felt in developing countries because the money is not easy to tax back given the large share of the informal economy. As a result, UBIs end up giving the same *net* transfer quite high up the income distribution, which means that UBI programs are substantially less redistributive in developing countries than in developed ones (Hanna and Olken 2018). This makes UBI 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 the government gives to inclusion error, exclusion error, and per-capita benefits. Hanna and Olken (2018) calibrate these benefits using standard CRRA utility functions and find that, under a variety of assumptions, for a given budget programs targeted using standard proxy-means tests substantially outperform UBIs because they transfer substantially more resources to the poor. But, they do entail more horizontal inequity, so ultimately the tradeoff comes down to how much one values increasing welfare of the poor on average vs. reducing exclusion error 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 (M. E. Grosh and Baker 1995; Coady, Grosh, and Hoddinott 2004b; Hanna and Olken 2018), this assumes that the budget is fixed. But as the elasticity between the total budget and the number of beneficiaries increases, universal transfers may become more attractive politically (Klasen and Lange 2016). Indeed, many have argued that the durability of programs like Social Security and Medicare in the United States comes from their universal eligibility (even though they have many redistributive features), which makes them politically popular. Understanding the political response to targeted as opposed to universal programs, or programs (like Social Security) which are both universal and progressive, is an important dimension for future work.

## **2.10: 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), the institutional ability to implement different methods, and the kinds of frictions that exist that could ultimately impact take-up. Thus, understanding the context well is key to determining what method to deploy, and its likely success in practice.

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 2004a). 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 is imperfect—with community methods to fill in the gaps and reduce exclusion error.

While the research has been growing in this area, and we know a lot more than we did just ten years ago, substantial questions exist for future research, 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 does stigma matter most for take-up.

### 3. Designing transfers for redistribution

#### 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 about benefits, there is however 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 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  $v(c_1, \dots, c_n, b, l)$  where  $c_1, \dots, c_n$  is a vector of consumption goods,  $b$  is publicly provided benefit and  $l$  is labor supply. The household maximizes this utility with respect to a budget constraint that is given by

$$\sum_1^n c_i \leq y + wl + b - t(wl + y - y^*)$$

where  $y$  is non-labor income,  $w$  is the wage rate,  $b$  is the monetary value associated with the government benefit, and  $t$  is the tax/transfer rate. This 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 psychosocial 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. we can assume  $\frac{\partial^2 u}{\partial l \partial b} < 0$ ).

In addition, the household may face an additional constraint on its choice which takes the form  $(c_1, \dots, c_n, l) \in \Phi(b)$ . 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 the constraints on its ability to make the right decision and stick to it, and how they are affected by  $b$ . Denote the indirect utility function generated by the household's constrained utility maximization by  $u(y, b, t)$ .

The notation is a bit misleading in the sense that we do not explicitly index any of the variables by the identity or income of the recipient, but we will allow the government to vary  $b$  for those at a specific income level. 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:

$$W = \int^{\hat{y}} g(y)h(y)u(y, b, t)dy$$

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

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 viability of a particular intervention under these assumptions is the ratio:

$$\frac{g(y)MU(y) \cdot WTP(y)}{\phi'(y)} \tag{12}$$

As noted above, in developing country settings, the labor supply reduction in response to additional public benefits appears limited, so expression (12) captures 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 (12). So, as long as either a) labor supply does not change or b) even 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 (12) is the key expression of interest.<sup>29</sup>

---

<sup>29</sup> If, however, labor supply changes, and this affects taxes or transfers. then we do need to modify expression (12) 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  $(\phi'(b) - t \frac{dl}{db}) db$ . This is then akin to the marginal value of public funds (MVPF) discussed by Hendren and Sprung-Keyser (2020), but augmented by  $g(y)MU(y)$  to take distributional consequences into account.

### 3.1.2: Implications

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

*Implications for evaluating cash transfers.* The usual assumption with cash transfers is that the willingness to pay for \$1 is always \$1 (i.e.  $WTP(y) = 1$ ), 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. did a household buy food versus a television) is not first order in understanding a program’s welfare impacts.<sup>30</sup> Second, and related to this, is that from a welfare perspective, it does not really matter if cash transfers cause people to work less, unless, as noted above, doing so has implications for taxes, i.e. that the transfers affect work, which in turn affects the 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.

On the other hand, the thing that does matter is how poor are the program beneficiaries. 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.<sup>31</sup> What also matters is the program’s governance and administration that we discuss below in Section 3.3, i.e. did the poor actually get their full cash transfer or did a percentage get lost to corruption. But, how beneficiaries choose to spend \$1 of benefits—on nicer food or a new television or a new floor—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 hugely important to the political sustainability and messaging of these programs, from a pure welfare perspective, this matters less. Unless, of course, the three conditions outlined above fail—which they often do, as we discuss more in detail below.

---

<sup>30</sup> If we do not know how poor the household is we may sometimes use the household’s consumption patterns to infer their poverty levels.

<sup>31</sup> Ravallion (2009) notes, in the context of China’s DiBao program, that some standard heuristic measures of targeting, which discretize targeting into, say, the share of transfers going to the poor do not predict the poverty impact of the program. This discrepancy comes from the fact these heuristic measures do not necessarily capture social marginal utility, i.e.  $\int h(y)b_i g(y_i)MU(y_i)$ , which is what is needed to calculate the marginal social impact properly.

*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. At this point it is worth acknowledging that even for cash transfers it is not always clear that \$1 is \$1. For example, a savings-constrained household may be willing to pay a premium to get a lump sum of money in return for a set of smaller payments that have the same present value from the point of the government. That said, most cash transfers programs offer small periodic transfers rather than a large lump sum, and for those the value of \$1 is probably close 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 estimate, directly, 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). Or, one can experimentally vary 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. This could be done either as a take-it-or-leave it option or by using a Becker-DeGroot-Marchak mechanism (Berry, Fischer, and Guiteras 2020).

Second, if one cannot measure willingness-to-pay directly, one can impute it. Hendren and Sprung-Keuser (2020) give examples of how to do this for 133 programs in the United States using the envelope theorem: the willingness-to-pay for a program is equal to the change in the N.P.V. of a household's income from receiving the program. There are two important caveats that should be made here: First, the relevant concept here is *income holding labor supply constant*; that is, ideally one wants to capture the exogenous change in income due to the program (for example, because the program provides a cash or cash equivalent, or because it raises the wage). If the labor/leisure decision changes, that is a consumption change like any other consumption choice, and so by the envelope theorem the welfare effects are second-order.<sup>32</sup>

Second, this presumes that there are no direct utility benefits or costs of a government program. For example, in some programs, beneficiaries have to wait in long lines to apply for the cash transfers, or walk very far to access the transfer from an ATM, or may only be given food of a type they do not like. For these households the net utility from the program may be less than the cash value of the transfer. Valuing willingness-to-pay only by effects on income negates these other utility costs or benefits.

*Limitations to this approach: wedges, spillovers, and frictions*

The willingness-to-pay approach is a benchmark. As described above, there are three important assumptions required 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

---

<sup>32</sup> The point that one should examine changes in income holding labor supply constant, rather than total changes in income, is also clarified in Egger et al. (2019), who helpfully articulate the conceptual limits on using changes in expenditure to measure changes in welfare.

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.*

One key assumption in the framework is that the social planner maximized the sum of each household's utility. There are two important ways in which this assumption may be an important over-simplification.

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 how to spend money and make other choices, and this may or may not aggregate into a well-defined utility function (Chiappori 1988; 1992; Browning and Chiappori 1998). Even if it does aggregate in a formal sense, the within-household welfare weights may not be the same as those that the social planner may want to choose. A household where the husband makes all decisions ignoring his wife's preferences may have a well-defined utility function, but the social planner may not 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, and 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 have preferences that welfare recipients do not spend cash assistance on 'temptation goods' such as alcohol or cigarettes (see Currie and Gahvari 2008 for a discussion). More generally one could imagine that voters would prefer a household receiving cash assistance bought food rather than a television, even if television is an extremely cost-efficient form of entertainment that may generate a lot of utility. Similarly, voters seem to prefer that those who receive assistance work. Therefore, while, as a benchmark, we adopt the view that households should be free to choose how to spend their transfers, we discuss cases below in which policy makers or voters may have particular preferences on the composition of beneficiaries' consumption decisions.

### *Spillovers.*

A second issue is spillovers. Households that receive benefits may spend it, creating Keynesian multipliers that can have positive spillovers to other households (Sadoulet, Janvry, and Davis 2001; Angelucci and De Giorgi 2009; Egger et al. 2019). In principle, the theoretical framework above can incorporate spillovers if one is careful to measure them, that is, if in the social planners' problem above, one considers the impact of a policy change  $b$  on all households, not just the households that receive it. However, in the presence of such spillovers, measuring willingness-to-pay for beneficiaries alone will be insufficient.

A related issue is if the transfers are large enough to cause prices to rise (Cunha, De Giorgi, and Jayachandran 2019; Egger et al. 2019). 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.

*Optimization frictions.*

The third potential limitation of this framework is the failure of the sophistication assumption so that the household's choices do not reflect its true preferences. One important constraint is information: household's may not understand what the benefit really is or the value that they will receive from a particular benefit, and so they may make mistakes in forecasting their willingness-to-pay (Banerjee, Duflo, and Hornbeck 2018). In fact, some may not really understand value until they actually experience the benefit, and so their willingness-to-pay before receiving the benefit may not reflect their willingness-to-pay after. For example, Banerjee et al. 2021 find evidence for this: a temporary subsidy government-provided health insurance in Indonesia leads to increased demand for the insurance even after the temporary subsidies have been lifted.

Behavioral frictions may also affect willingness-to-pay. For example, if benefits have endowment effects, it would imply that the 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. 2019).

A third friction comes from credit constraints. This is primarily an issue for the estimation of the willingness-to-pay. When we estimate people's willingness-to-pay from observational data, we need to worry about the fact that they may not have had enough money at hand to express their preferences and this could lead to an underestimation of the willingness-to-pay. Consider an investment with a co-payment of \$10 that would pay a return of \$100 tomorrow. The willingness-to-pay for this investment (above the co-payment) should be \$90. But a household who only has \$5 on hand, and who cannot borrow, will only have a stated willingness-to-pay of \$5,<sup>33</sup> and this is what one would find if one elicits willingness-to-pay directly. But, if one takes the second approach discussed above—measuring the impact of the program on future incomes, and asserting that the willingness-to-pay must be at least as high as the (discounted) impacts on future impacts, then measuring effects on incomes will get the right answer even in the presence of credit constraints. Moreover, one can 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.

A more subtle issue comes up if the poor face obstacles navigating the bureaucratic process of obtaining benefits, as we discussed above. 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 or believe that to be the case.

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

---

<sup>33</sup> In fact, the worry here is that many poor households are too poor to access transfers programs, whose benefits could help push them over the poverty line. In some sense, this creates another form of poverty trap, wherein poor households stay poor since they are too poor to access assistance. Relaxing credit constraints could help alleviate this.

willingness-to-pay for the program, again multiplied by the marginal utility of income, again provides a useful benchmark. That said, one should also 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. Thus, as we turn to reviewing the empirical evidence on design choices of transfer program below, we do so through the lens of these ideas outlined above.

### **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—and then organize the subsequent empirical sections based on the key program design choices that one can make beyond this benchmark. In doing so, we focus on whether and how those choices affect welfare in the senses discussed here:

- affecting future income above and beyond the transfer itself (i.e. by raising wages, or allowing the household to overcome an important credit constraint)
- 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
- 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 that have evaluated these types of programs in this space. We refer interested readers to several recently completed exercises that do exactly this, such as Bastagli et al. (2016) and Ralston, Andrews, and Hsiao (2017). Instead, our aim is to capture key themes from these many studies, to relate them to the framework above, and to highlight where additional evidence is needed.

#### **3.2.1: The benchmark: welfare effects of unconditional cash transfers.**


We begin by considering the simplest programs: 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, from a welfare perspective, the basic question of an unconditional cash transfer program is who receives the benefits, i.e., the targeting of the program. This is because, as described in the framework above, a household's willingness-to-pay for \$1 of these types of improvements should be \$1 (subject to the caveats already made), with the social value of this inflated by a measure of  $g(y)MU(y)$  capturing how poor the household is. To first order, the *key* thing for evaluating the impact of unconditional cash transfers is therefore the targeting of the transfers, which affects  $g(y)MU(y)$ .

Several randomized studies consider cash-transfers directly to evaluate their impacts, such as the Haushofer and Shapiro (2016; 2018) and Egger et al. (2019) studies of the GiveDirectly program in Kenya, or consider it in comparison to other programs, such as Baird, McIntosh, and Özler (2011) study of transfers in Malawi. There are many other non-randomized studies as well; indeed

Bastagli et al. (2016) provide a systematic review of 201 cash transfer studies, of which 25 percent were unconditional cash transfer programs.

Figure 3 plots the percentage of countries at different income levels with unconditional cash transfer programs 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 essentially only found in the wealthiest countries. By the end of the period, these programs have become much more common—and almost equally likely to be found in rich or poor countries alike (and in fact are most likely in upper middle-income countries).

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



*Notes:* This figure plots the percentage of countries that enacted UCTs covering at least 1% of the population, by each year considered (1980, 1990, 2000, 2010 and 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 UCT as a government-implemented, large-scale program that enforces no conditions on recipients and is means-tested. 135 countries are in the 1980 sample, 166 countries are in the 1990 sample, 192 are in the 2000 sample, 203 countries are in the 2010 sample and 200 countries are in the 2020 sample. *Source:* GDP per capita and population data from the World Bank (2021a; 2021b).

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 (2016) report 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. Many of these effects persist even several years after all transfers ended (Haushofer and Shapiro 2018). 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, with no to limited effects on unintended outcomes such as labor supply or fertility. There is less systematic attention to the differential targeting of these programs; the framework above suggests this is an important direction for future work in this area.

*Effects on adult labor supply.* As described in the framework above, changes in labor supply themselves are not welfare-relevant; the labor/leisure choice should be considered one of many consumption choices, and a household choosing to work less in response to a transfer should 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).<sup>34</sup>

*Effects on consumption of so-called ‘temptation goods.’* A common critique of cash transfers is that households will use them for ‘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 (indeed, many people are known to enjoy a drink after a long day’s work). But, voters appear to be less enthusiastic about funding transfer programs if they were to lead to consumption of these goods, either because of fears that these choices reflect short-term temptations (Banerjee and Mullainathan 2010) rather than rational decisions, because of intra-household issues (e.g., an addict may grab the funds to fund his or her addiction), or simply because voters disapprove of this type of consumption. A systematic review by Evans and Popova (2016) examines 19 cash transfer studies (a mix of unconditional and conditional cash transfers) and finds no systematic 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 liquidity or other credit constraints (though if the household understands them, they may still be reflected in a properly measured willingness-to-pay). In the presence of these constraints, households are not inter-temporally optimizing, so one cannot use the envelope theorem to argue that future income increases are second-order. Instead, one should properly count future income increases in addition. Given the ample evidence that households in developing country settings are credit constrained (e.g., Banerjee and Duflo 2010), these effects may be substantial.

---

<sup>34</sup> Specifically, Banerjee et al (2017) examine 7 different cash transfer programs in 6 countries (Honduras, Indonesia, Morocco, Mexico, Nicaragua, and the Philippines – a mix of unconditional and conditional programs), all evaluated through randomized trials. Across all 7 programs, they find no systematic impacts on labor supply, for either men or women. Hanna et al (2018) examine 8 unconditional cash transfer programs in 7 African countries (Ethiopia, Ghana, Kenya, Lesotho, Malawi, Zambia, and Zimbabwe), evaluated using a mix of experimental (5 programs) and non-experimental (3 programs) matching methods. They find evidence of a shift in *types* of labor—away from agricultural wage labor, and towards labor in own farming activities and on own non-farm enterprises—but not a systematic change in overall labor engagement.

Indeed, there are a number of studies that suggest that households use cash transfers to make future 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.<sup>35</sup> 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 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 as an example. One approach is to assume households are optimizing, and value a transfer of \$1 at \$1. In this case, 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 so one should not also include the investment income as this would be double-counting. Alternatively, a second approach is to assume that households are not fully intertemporally optimizing beforehand due to credit or liquidity constraints. In this case, one should value the initial consumption (i.e. \$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 about 12 months after the transfer. The net welfare effect is therefore given by  $0.74 + PDV(0.016 \text{ per month})$ . For example, assume a discount rate of 10% per year and assume the investment effects started 1 year after the transfer and lasted 10 years post transfer, the net welfare effect of this transfer would be \$1.8 instead of \$1. While the precise numbers depend substantially on the assumptions made, the point is that, if one considers these effects separately from a welfare perspective, investment effects like this can make a big difference.

This discussion thus far is just about how households choose to invest a stream of transfers. Beyond this, several studies study alternate forms of transfer schemes (e.g. asset-transfers), or alternate delivery mechanisms (e.g. a large lump-sum payment rather than a smaller stream of payments), with the idea that these particular delivery forms may be more effective in allowing households to invest, and potentially enough to escape a poverty trap. We discuss these in more detail in Section 3.2.6.

*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 De Giorgi (2009), for example, document that non-beneficiary households in Mexican villages in which the Progresa program was introduced report higher consumption than comparable households control villagers. They show that this comes from ineligible households receiving more

---

<sup>35</sup> Note that Oportunidades is a conditional cash transfer, as are several other papers we discuss in this section. Here, we discuss aspects of these programs related to just the cash transfer component; we discuss the aspects specifically related to conditionality in Section 3.2.2 below.

loans and transfers, presumably from eligible households who received Progresa benefits. Of course this does not always take place—Evans and Kosec (2020), by contrast, 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 financial or in-kind assistance to other households.

Prices: A second channel through which cash transfer programs can affect welfare of non-beneficiaries is through price changes. However, since those receiving benefits tend to spend the benefits on a wide variety of different items, the demand shock for any given item is likely to be small. This, combined with the fact that supply of those items may be at least sometime elastic, suggests that in practice price effects for pure cash transfers may be small. Indeed, several studies, such as Cunha, De Giorgi, and Jayachandran (2019) and Egger et al. (2019) find either zero or economically very small impacts of generalized cash transfers on prices.<sup>36</sup>

One important counter-example is a study by Filmer et al. (forthcoming) of a cash transfer program in the Philippines. Filmer et al. (forthcoming) study a cash transfer program in which many, but not all, households in treated areas received benefits—on average 65 percent of households were treated, and in some villages 90 percent or more. As a result of this high saturation, the program raised aggregate incomes in treated villages by about 15 percent. They show this led to an increase in prices of protein-rich perishable foods (e.g. eggs, fresh fish) by 6-8 percent, with price changes happening more in villages where the share treated was high and in remote areas where supply is presumably 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 potential 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 those who received the transfers, who then in turn spend some of their income locally, and so on.<sup>37</sup>

Egger et al. (2019) use the fact the transfer program that they study was randomized at the village level to estimate a village level ‘fiscal multiplier’ of this sort, which they estimate to be between 2.5–2.8; that is, each dollar transferred leads to 2.5-2.8 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, one-time transfers (US 1,000) that they study in this study. Note, however, that since transfers are paid for externally, this estimate does not capture the net general equilibrium ‘multiplier’ if one was to impose such a policy nationally. That is, if one were to create transfer programs as national policy, the transfers would need to be paid for somehow, and those taxed *also* could have multiplier effects (e.g., people who are taxed to pay for the transfers consume less, and this also has multiplier effects). Understanding the full general equilibrium implications of these types of tax-and-transfer schemes in developing country settings is an important area for future work.

---

<sup>36</sup> By contrast, Cunha, De Giorgi, and Jayachandran (2019) and Banerjee et al (2021) 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 below.

<sup>37</sup> Chodorow-Reich (2019) provides a conceptual overview of this literature based in the US.

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 Familia). Unlike other programs, Bolsa Familia is explicitly means-tested (i.e. those with formal sector income above a certain level are ineligible), so one may be particularly concerned about labor market effects of such a program. However, consistent with the presence of Keynesian multipliers, they find that expansions in Bolsa Familia substantially increased labor market participation, including among workers who were never part of the program. Since they do not observe consumption or everyone's income, they do not formally estimate a Keynesian multiplier, but the results are consistent with cash transfer programs having strong local multipliers.

\* \* \*

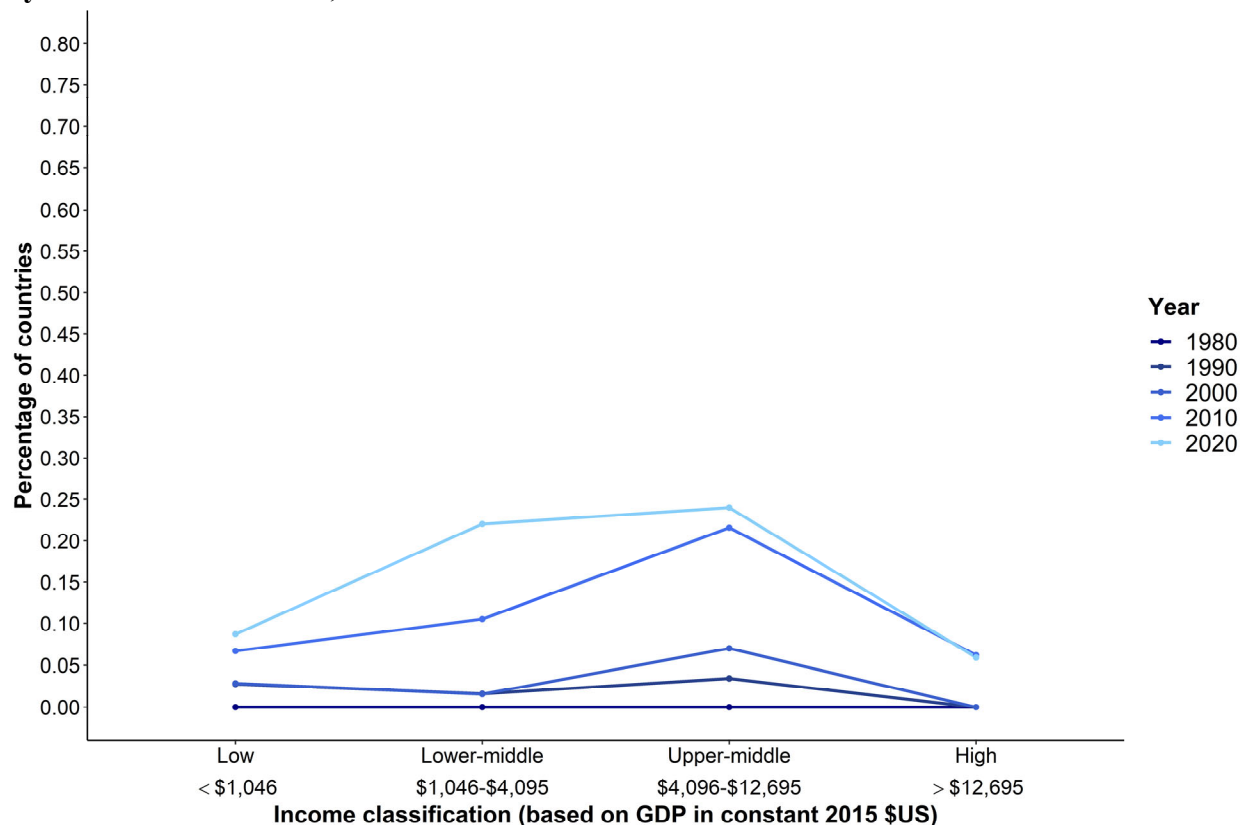
In addition to these above effects, unconditional cash transfer programs may have important welfare effects within households that are not captured by this simple framework, e.g., on education and health on children (which may, or may not, be fully counted in a household utility function, depending on whether parents are completely altruistic towards their children), and on the relative bargaining weights and outcomes among spouses. We address these effects in Section 3.2.2 and 3.2.5 below where we discuss interventions that systematically target these issues.

### **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 (e.g., monthly) cash transfer as in many UCTs, but then *condition* the transfer on the household fulfilling a set of criteria, usually related to investments in human capital for children. For example, pregnant mothers may be required to get ante-natal and post-natal care and have professionally-assisted childbirth; young children are required to be vaccinated and monitored for regular growth, and school-age children are required to be enrolled and attend school. Some programs also require nutritional counseling and participation in other child health services such as de-worming; the precise conditions vary a bit depending on the context, with some programs placing more weight on the cash transfer and others placing more weight on encouraging human capital investments. In Appendix Table 1 we present the conditions of sixty-seven major CCT programs from forty-five countries. We find that the most common conditions attached to CCTs are: school enrollment and attendance (fifty-two CCTs), health checkups (thirty-eight CCTs), complete vaccination schedule (twenty-two CCTs) and attendance of training sessions or workshops (twenty CCTs).

CCTs were pioneered in in the 1990s, with Brazil's Bolsa Familia, Mexico's Progresa, and Bangladesh's Female Secondary School Assistance Project being early examples. Since then, over 60 countries have started their own programs of this type (World Bank 2018a). Figure 4 shows the percentage of countries with large-scale CCT programs (i.e. covering at least 1 percent of the population), by quintile of GDP per capita (in constant 2015 \$US). It demonstrates the rise in popularity of CCTs over recent decades, from being essentially non-existent in 1980, to between 10% and 35% of countries in the bottom four GDP per capita quintiles having large-scale CCTs by 2020. Unlike with UCTs shown in Figure 3, however, these programs are much less common in the wealthiest countries.

**Figure 4: Percentage of countries that enacted CCTs covering at least 1% of the population by income classification, over time**



*Notes:* This figure plots the percentage of countries that enacted CCTs covering at least 1% of the population, by each year considered (1980, 1990, 2000, 2010 and 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. 135 countries are in the 1980 sample, 166 countries are in the 1990 sample, 192 are in the 2000 sample, 203 countries are in the 2010 sample and 200 countries are in the 2020 sample. *Source:* GDP per capita and population data from the World Bank (2021a; 2021b).

The theory behind these programs is that the cash transfers provide assistance to the current generation, while a combination of cash and the requirements for human capital incentives aim to break the inter-generational cycle of poverty.<sup>38</sup> Implicit in these requirements is also a statement about welfare: that the government values the child human capital investments at a greater weight than the household decision maker would absent the incentives (if not, then there would be no need for explicit 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.

Mexico’s Progresa program was evaluated using a phase-in design. Surveys were conducted in 320 treated municipalities as well as 186 control municipalities, which received the program about

<sup>38</sup> For example, Santiago Levy, who helped create the CCT model with the Mexican PROGRESA program in the 1990s, argued, “clearly achieving good health is a cumulative process, and temporary investments in nutrition are of little help. The same is true of education: children must be supported year after year.... [PROGRESA’s] central effects will gradually occur through the accumulation of human capital” (Levy 2006).

18 months later (Skoufias 2005). This phase-in design allows researchers to understand if the programs led to the desired human capital investments in health (Gertler 2004), education (Schultz 2004), as well as a variety of other outcomes (Skoufias 2005). Moreover, when CCTs were instituted in countries worldwide throughout the early 2000s, a number of them were experimentally evaluated during early stages. Indeed, there are now randomized evaluations of CCT programs in Colombia (Barrera-Osorio, Linden, and Saavedra 2019), Honduras (Benedetti, Ibararán, 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. forthcoming; Kandpal et al. 2016) and Tanzania (Evans et al. 2014; Evans, Holtemeyer, and Kosec 2019). Moreover, many CCTs are targeted using proxy-means tests with strict cutoffs, which also enables impact evaluation using RD methods (e.g., Buddelmeyer and Skoufias 2004).

These evaluations often find impacts of CCTs on human capital accumulation, both in terms of child health and schooling. In the short run, there is evidence that these programs led 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 cognitive outcomes for children (Macours, Schady, and Vakis 2012).<sup>39</sup> They also show reductions in child labor (Barrera-Osorio et al. 2011; Schady and Araujo 2006; Benedetti, Ibararán, and McEwan 2016).

*Impacts of conditional cash transfers over longer time horizons.* Several studies, using a variety of empirical approaches, suggest that these human capital benefits may persist. Cahyadi et al. (2020) experimentally measure the impacts of a CCT program in Indonesia after six years. They find evidence that human capital gains persisted: 6 years after the start of the program, under-15 children who were not in school fell by half, child labor fell by almost half, and stunting—which requires cumulative investments in health—fell by 23 percent.

Molina Millán et al. (2020) also exploit experimental variation in their study of the CCT in Honduras. In their case, the program ran for 5 years (2000-2005) in treatment municipalities, but the program was never run in the control ones. They examine outcomes 13 years after the program began (i.e. 8 years after it ended). They find large increases in education—particularly secondary and university enrollment—for non-indigenous groups, but that the effects are more muted for indigenous groups. They also find evidence that the CCT doubled the probability of international migration for men who experienced the program as older cohorts. They find little robust evidence on wages.<sup>40</sup>

---

<sup>39</sup> Parker and Todd (2017) provide an in-depth discussion of the many studies of the Progresa program.

<sup>40</sup> Dervisevic et al (2021) also exploit experimental variation to examine the impacts of short-run (1.5 years) CCT exposure for adolescents 10 years later. While the control group was treated, due to the age requirements for entry into the program, they can identify a cohort who was never treated because they had aged out when the program was expanded. While they found impacts on marriage and fertility for women, they found no impacts on educational or labor market outcomes. However, it is important to note that as the cohort studied were between the ages of 12.5 and 14 years of age at the time of CCT exposure, it is not capturing impacts of a CCT at early childhood. A 10-year follow-up of the Ecuadorian CCT rollout, using both a phased-in strategy and a regression discontinuity based on eligibility, finds small (about 2 percent) increases in education completion rates 10 years later (Araújo, Bosch, and Schady 2017).

A second empirical strategy uses the fact that some programs started earlier in some places than others; for example, treatment locations in Mexico's Progresa received the program about 18 months earlier than those in controls. Behrman, Parker, and Todd (2011) use this strategy, comparing those who received an 'extra' 18 months of the program, after five and a half years of exposure in treatment areas, as well as comparing both treatment and control municipalities to additional never-treated areas using a matching design. They find that increases in schooling attainment for those who received the program earlier, as well as lower rates of labor force participation. Using the same identification strategy, Gertler, Martinez, and Rubio-Codina (2012) find that those who were treated 18 months earlier also have 5.6 percent higher consumption than those who were treated later, even 4 years after the controls were treated.

Barham, Macours, and Maluccio (2013; 2018a; 2018b) use a related design in Nicaragua, using a randomized design and comparing early treatment locations (who received CCTs from 2000-2003) with late treatment locations (who received CCTs from 2003-2005), and comparing outcomes in 2010. In Barham, Macours, and Maluccio (2013), they 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 (2018a), they focus on girls ages 9-12 at the start of the program, and compare girls treated at ages 9-12 with those treated at ages 11-14, and examine outcomes when these young women were 19-22 years old. They find girls with early exposure to the program had higher educational attainment, higher earnings, later age of menarche, age of initiation of sexual activity, and lower BMI compared to late treatment girls. They also had lower fertility by ages 18-21. In Barham, Macours, and Maluccio (2018b), they examine boys of the same ages, and find similar education and labor market effects.

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.<sup>41</sup> Parker and Vogl (2018) use this strategy to analyze the long run impacts of Progresa. They find that childhood exposure to Mexico's Progresa in primary school (i.e. before the transition to secondary school) leads to increased educational attainment, greater labor force participation, greater propensity for migration, and higher earnings for women. They also find modest impacts on education and migration for men, though the labor market results are less robust.

In fact, Araujo and Macours (2021) actually follow the Progresa cohorts *about 20 years later* and find that educational attainment increases, but no effects on marriage or childbirth. For the children who were older when the program started (i.e. enough time has gone by to look at labor market effects), labor income is higher, particularly for the top quarter of the income distribution. Those in the treatment group were more likely to migrate, particularly to the U.S., which could in part account for the higher incomes.

*Does conditionality matter? Can it be a two-edged sword?*

---

<sup>41</sup> Kugler and Rojas (2018) use a similar strategy, but focused more narrowly on the ~500 villages in the early vs. late Progresa impact evaluation sample.

These human capital improvements provide evidence that that CCTs 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, without relying on explicitly incentivizing this behavior. Or could it be that simply suggesting greater human capital investment without enforcing it might be enough. How can we tell? Do the conditions themselves matter?

Several studies have found evidence that households respond to conditionality on the targeted outcomes. Brollo, Kaufmann, and La Ferrara (2017), for example, study Bolsa Familia, the Brazilian CCT, and examine how households respond to warnings that they are not complying with school attendance requirements. Under Bolsa Familia, households receive up to 5 warnings before being terminated from the program, with increasing penalties for each warning. The date of these warnings is essentially random, based on the last digit of the household head's ID number. They find increases in attendance right after households receive the warnings. Moreover, these effects are larger for later warnings, which carry larger penalties (and for which *subsequent* violations would carry increasingly large penalties). These results confirm that households are responsive to conditionality.

Macours, Schady, and Vakis (2012) examine the role of cash vs conditionality in these programs using an alternative approach. In their study of the Nicaraguan CCT, they compare the basic CCT program to a variant of the program in which households who submitted a business plan were given an additional lump-sum cash transfer, which raised the total amount of cash received from 15 percent of annual per-capita expenditures to 26 percent of annual per-capita expenditures in one year. They find that while the basic CCT improved an index of child cognitive, socio-emotional, and health outcomes for children (such as language, memory, motor skills, and anthropometrics), there were no additional increases observed for those households who received the additional cash, and in fact, they can rule out that changes in child outcomes are linear in the amount of increase in household consumption. They interpret this as evidence that it was not just the cash transfer, but other aspects of the program (e.g., the conditionality, labeling effects, other aspects of the basic program design, etc.) that was crucial for these outcomes.

But could there be downsides to conditionality? Baird, McIntosh, and Özler (2011) provide perhaps the most direct test of the differences between these two programs in Malawi, where they randomize communities into one of two distinct interventions: one arm which gives unconditional cash transfers (UCTs) to households with school-age girls, and a second arm which gives conditional cash transfers (CCTs), conditioned on school attendance, to a similar group of households, in addition to a control group. The results show that conditionality matters, but also entails risks of its own. On the one hand, for targeted indicators, the CCT had much larger effects, implying that conditionality *per se* may improve targeted outcomes: while both unconditional and conditional transfers led to a reduction in dropout rates, the reduction was more than twice as large in the CCT compared to the UCT. 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 drop out of school. This suggests that the success of the conditionality comes 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.

Akresh, de Walque, and Kazianga (2013) conduct a similar randomized trial in Burkina Faso, where villages were randomized into receiving a conditional cash transfer program with typical growth monitoring and enrollment/attendance conditions, or a similar cash transfer program without conditions. For boys and older children, who tend to be prioritized by parents, both the conditional and unconditional programs led to increases in enrollments, by about 20 percent in both programs. However, they find that the conditional program led to large increases in enrollments—20.3 percent for girls and 37.3 percent for younger children—who they identify as more ‘marginal’ for education while the unconditional transfer had no impacts on enrollment for these groups.

Combined, all these studies suggest that conditions may matter for human capital outcomes for the next generation, but, there can also be a potential downside that comes from cutting off assistance to those who do not comply. This presents the intriguing possibility that a *labeled* transfer program—where households are told that there are conditions, but they are not rigorously enforced—may achieve the impacts results on the next generation from conditional transfer programs, without the downsides. Such an approach could also save extra monitoring costs.

Benhassine et al. (2015) examine this in Morocco, where they run a randomized experiment giving a small cash transfer to fathers of school-aged children that is explicitly labeled as an education support program but does not actually condition on school attendance. Labeling was done by enrolling households at schools by school headmasters. They randomize school sectors into receiving this ‘labeled cash transfer’, or a regular conditional cash transfer program where the conditions are rigorously enforced. They 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.

In short, it appears that conditions can help encourage human capital outcomes—and hence welfare beyond the current household, if these effects are not fully internalized—but that rigorously enforcing these conditions does have some potential risks. The idea that one can achieve the best of both worlds through ‘labeling’ the transfers but not enforcing them is intriguing, but whether this ‘labeling’ effect persists in the longer term—i.e. after households have experience and learn that the transfers are not removed—is an important question for future research.

### **3.2.3: In-Kind vs. Cash Transfers**

*Conceptual Issues.* An alternative approach to delivering assistance is to provide households with goods or services directly, rather than cash. There are good reasons why cash could be preferred to in-kind transfers, and why cash transfers remain a benchmark against which other programs can be measured—after all, households can choose to use the cash to purchase whatever will increase their utility the most. Moreover, for households that would purchase more of the in-kind good than is provided by the transfer even in the absence of the transfer providing in-kind goods should be the same as cash, since households can just consume the in-kind transfer instead of buying the good, and can reallocate their cash elsewhere.

However, theory provides some reasons why policy makers may prefer in-kind. The first has to do with the types of wedges discussed in Section 3.1.2. If, for example, the social decision maker has

different weights over household members than the household would use in its decision making, or if the social decision maker cares explicitly about certain types of ‘merit goods’ (Musgrave 1959), such as nutrition, health care, or housing, rather than just household utility, the social decision maker may explicitly want to influence the consumption basket of recipients. 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. That would also work if households are simply not optimizing or use mental accounting, in which case there will be stickiness from the in-kind transfer to total consumption (Hastings and Shapiro 2018).

A second reason has to do with how these two different types of programs may interact with prices. There are two issues. First, in-kind transfers typically represent a positive supply shock of the good in question (unless the government is sourcing it all locally). This means that, if supply is inelastic, prices for the subsidized goods may be lower in an in-kind program than in a cash program which does not (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 set vouchers in terms of prices (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 in principle the government could adjust the quantity of in-kind benefits based on prices to hold value constant, or conversely adjust the value of vouchers to keep their real value constant, the fact that an in-kind transfer is inherently set in quantity terms and a voucher is inherently set in dollar terms means that these adjustments are less likely. Gadenne et al. (2021) demonstrate this mechanism by showing that expansions of in-kind transfers in India reduce sensitivity of household consumption to price variation.

A third potential difference is that in-kind transfers have the appeal of improving self-targeting if it increases the cost of participation more for those with higher income (Nichols and Zeckhauser 1982). For example, if the government provides subsidized bread of worse quality than the average available in the market, richer households will not bother consuming it. Lieber and Lockwood (2019), for example, examine home health-care benefits in the United States, and argue that in-kind programs are better targeted for this reason, despite the fact that households strongly prefer cash.

Finally, there are administrative differences between these programs. The infrastructure required to deliver millions of tons of food or other goods is different than the infrastructure required to run a voucher system using electronic debit cards (Banerjee, Hanna, et al. 2021), and these administrative differences may lead to important differences in low state capacity settings, as we document below.

#### *Experimental comparisons of in-kind, voucher, and cash programs.*

Several studies examine these issues by experimentally allocating areas into some combination of cash, voucher, or in-kind distribution programs. We examine several such studies, focused on food programs, to elucidate the empirical evidence.

One study was done in the context of Mexico’s food assistance program, PAL, and was conducted during the 2003 rollout of PAL to 200 rural villages. The government randomized villages into

receiving in-kind food transfers, approximately equivalent cash transfers, or a control group. Several studies have examined different aspects of this experiment. Overall, Skoufias, Unar, and Cossío (2008; 2013) find that both types of transfers have similar effects on total consumption, male labor supply, and poverty outcomes.<sup>42</sup> Cunha (2014), however, notes that the in-kind transfer was worth more when valued at local prices, and when this adjustment is made, finds that—per *real* dollar transferred, the in-kind and cash programs have similar effects on total food consumption.

The study shows two potential important differences between cash and in-kind programs. First, Leroy et al. (2010) find, and Cunha (2014) also confirms, that the in-kind program leads to increases in both total calories consumed and in micro-nutrients, through stickiness in terms of consumption choices. Second, Cunha, De Giorgi, and Jayachandran (2019) test the Coate, Johnson, and Zeckhauser (1994) hypothesis that in-kind transfers lead to pecuniary effects through prices, and find evidence that in-kind transfers cause price declines in remote areas, by about 5 percent. They argue that, in remote locations, in-kind rather than cash increases the real value of transfer 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 or made by recipient households. They find no changes of prices in more developed locations.<sup>43</sup> Jiménez-Hernández and Seira (2022) also find that direct government provision of milk in Mexico reduces prices, though they argue for a different channel: direct provision by government introduces competition, and hence is effective when there is local market power in the provision of goods.

A second study comparing cash, in-kind, 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 on the quantity and quality of food consumed in an urban areas with well-functioning markets, randomizing 145 clusters within 80 urban neighborhoods (*barrios*) into one of the three treatment groups or control. The in-kind transfer included rice, oil, lentils, and canned sardines; the vouchers were redeemable for a broader range of goods, including cereals, tubers, fruits, vegetables, legumes, meats, fish, dairy, and eggs. They find similar effects of the three programs on both food and non-food consumption, but they find differences in the composition of food. Specifically, relative to the in-kind program, the voucher program—which had a much wider range of portions—leads to higher dietary diversity. For example, voucher households consumed more vegetables, eggs, and milk and dairy than households that received the in-kind transfer, which is consistent with the fact that these items were included in the voucher list but not in the in-kind transfer. The in-kind program is also more than twice as expensive (in terms of costs per outcome achieved) than either the cash or voucher, given the much higher administrative costs associated with moving and sourcing the food.

---

<sup>42</sup>Skoufias, Unar, and Cossío (2013) also finds suggestive evidence that cash transfers lead to increases in both food and non-food consumption, whereas in-kind transfers lead just to increases in food consumption, though these difference are not statistically significant.

<sup>43</sup> This finding is consistent with Michelson et al. (2012), who argue that rural villages would have a lot more trouble responding to large cash injections stemming from less market access and less competition among suppliers, examining one rural and one urban site in Kenya.

A third study was conducted by Banerjee, Hanna, et al. (2021), comparing an in-kind rice delivery program with a voucher program that allowed recipient households to purchase an approximately equivalent value of rice and eggs from a large network of private providers. Importantly, the Banerjee et al study is an at-scale experiment: 105 districts, with an average population of 500,000 each and a combined 3.4 million beneficiary households, were randomized into either the in-kind or voucher program, with both programs implemented by the government bureaucracy as usual. They find a dramatic difference: voucher programs—which were administered through electronic debit cards recipient names pre-printed—delivered concentrated assistance among targeted households, whereas in-kind aid was spread around much more widely by local implementing officials. As a result, targeted households received 45 percent more assistance in voucher districts 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 easier to more reliably administer in low state capacity settings, ensuring that the program on the ground looks like it was conceived in theory.

The Banerjee, Hanna, et al. (2021) study also confirms several other results discussed above. First, they find evidence of small price effects associated with the move from in-kind to vouchers, but only in very remote areas. Given the dramatic increases in assistance delivered to the poor by vouchers, and the relatively small size of the price effects, they do not change the bottom-line conclusions, but this result confirms the finding that an advantage of in-kind transfers is lower prices in very isolated locations. Second, they find, as above that vouchers lead to more flexibility than in-kind: households in voucher areas consume more egg proteins than households in in-kind areas, consistent with the fact that the in-kind transfer was for rice only, whereas vouchers could be used for rice and eggs.

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 the most. To the extent that policy makers believe that households are privately not optimizing their nutrition choices or otherwise prefer food programs (e.g. for political reasons), the evidence suggests that food vouchers can have impacts on the set of foods that households consume. In-kind programs are costlier to administer, and harder to administer faithfully, but may be useful in very rural, isolated areas where supply is inelastic or non-competitive.

#### **3.2.4: Workfare programs**

Another type of transfer program is workfare programs. In these programs, the government provides public employment, typically at a low wage, for those who want it. Historically, these programs were quite common in the United States during the New Deal era (e.g. the Civilian Conservation Corps and Works Progress Administration). Today, the largest workfare program in the world is India's Mahatma Gandhi National Rural Employment Act (MGNREGA), which provides 100 days of work at the official minimum wage for anyone in rural India who wants it—providing jobs to over 88 million people per year.

A critical element of these programs is that the government is directly intervening in the labor market. This has a few implications. First, in addition to differences in workers' preferences (i.e. how much they prefer working on a workfare job compared to their alternative job in the private

market), one also needs to consider the productivity of the work they are doing. If the work is productive (i.e. building public infrastructure with the same efficiency as laborers working for private contractors), 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.

Several studies have examined these issues, largely in the context of MGNREGA. MGNREGA was rolled out at the district level in three waves from 2006-2008, so several papers use this strategy to identify its impact. Imbert and Papp (2015) use this strategy to examine the impacts of MGNREGA on the labor market, as well as using the fact that in this period, work was substantially rationed in many states. Because the statutory minimum wage of MGNREGA is higher than the de-facto prevailing minimum wage in most cases, the MGNREGA wage is above the prevailing wage in many cases. Crowd-out is substantial: while they find that MGNREGA increased the fraction of the population working on public works by 1.17 percentage points in the dry season and 0.46 percentage points in the wet season, they detect equivalent falls in the probability of private employment in the dry season, and they cannot reject that private employment falls one-for-one with public employment in either season—so there is substantial crowd-out of private employment. As described above, Bertrand et al. (2021) examine related questions in the context of a youth workfare program in urban Cote d’Ivoire, and also find substantial evidence of crowd-out.

Consistent with MGNREGA paying a higher wage, and this having overall equilibrium labor market impacts, Imbert and Papp (2015) find that dry season wages increase by 4.7 percent. They find that between 22 percent and 42 percent of the total welfare gains from the program for the three poorest quintiles come from the wage change on average. Berg et al. (2018) and Azam (2012) use similar identification strategies, also finding impacts on wages. Aggregating across all these studies, it seems as though an important impact of the MGNREGA program was through its impact on wages in general, not just for those who were employed in the program.

Muralidharan, Niehaus, and Sukhtankar (2020) study the general equilibrium market effects of the MGNREGA program, using a different source of variation: a randomized experiment at the sub-district level that improved the program administration (biometric smartcards, which reduced leakage, reduced payment delays, and generally made the program function better; more on these in Section 3.3 below). They find wages increased, with treated areas having wage increases of 6.5%, as well as detecting positive spillovers of wages to nearby control locations (so that, on net, they estimate a 10% increase in wages). They find that the reform raised the income of low-income households by 13%, and that 90% of the gain came from the 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. (2021) studies 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, and thus perhaps unsurprisingly led households to shift from private to public employment. The shift in labor supply, however, led to an increase in private market wages that accounted for the lion’s share of the increase in welfare of program

beneficiaries. Moreover, as discussed above, whether society is better off from these types of workfare schemes also depends on whether workfare has spillovers on local amenities, i.e. whether workfare is productive or not. Franklin et al. (2021) also experimentally show that the introduction of workfare increased the amenities in treated neighborhoods, suggesting that the investments they study were productive.

### **3.2.5: Programs targeting specific individuals within the household**

Another decision margin for a transfer program is who should be the recipient. If the household behaved like the unitary model predicts (i.e. the household solves their optimization as a single unit) then this decision is irrelevant: no matter who received the income, the money would be spent in the same way (Chiappori 1988; 1992; Browning and Chiappori 1998).

However, there is reason to think that who the is recipient matters, and indeed, this is one of the frictions that we discuss in Section 3.1.2 which moves us away from the benchmark case. Robinson (2012), for example, experimentally tested whether households are able to pool risk efficiently in response to small transfers. The field experiment consisted of weekly lotteries where each individual in a married couple could win 150 Kenyan shillings—a day and a half’s worth of work for men and a full week’s worth of work for women—with a 50 percent probability. Information about the lottery’s outcome was public so everyone knew who won. Robinson finds husbands increase their private consumption in weeks in which they win but have no change in consumption when the wife is the one who receives the winnings, implying a rejection of Pareto efficiency.

If money is spent differently in the household depending on who receives it, then who is the recipient matters for policy. The empirical evidence on the degree to which this matters in a real-world policy context, however, is mixed.

In a well-known study, Duflo (2003) studies the impact of an expansion of South Africa’s social pension program. She compares nutritional outcomes of children in households where the recipient of the transfer was a woman to those where it was a man. She finds a large improvement in outcomes for girls when the recipient is a woman but little effect when the man is the one who receives the pension. Bertrand, Mullainathan, and Miller (2003) study the same program and find additional evidence of imperfect pooling: they find a sharp drop in the working hours of working-age men when a woman in the household becomes eligible for a pension, but no sharp drop when a man does.

However, a number of subsequent studies explicitly randomize whether transfers are given to men or women in low- and middle-income countries, finding little effects on average, building on Lundberg, Pollak, and Wales's (1997) quasi-experimental study in the United Kingdom. Akresh, de Walque, and Kazianga (2016), in their study of a conditional transfer in Burkina Faso, randomize whether the transfer is given to mothers or fathers. 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. Benhassine et al. (2015) also randomize the gender of the recipient of a transfer, and find little difference on average. Likewise, Haushofer and Shapiro (2016), in their study of unconditional cash transfers in Kenya, randomize the recipient to be men or women, and

again find few differences on average. They do, however, find that giving transfers to women (as opposed to men) leads to higher reported psychological well-being and greater female empowerment, though they find no differences on many other dimensions mentioned (e.g. food security, health, education). Haushofer and Shapiro note that a higher-powered experiment may be better able to detect effects.

However, recent evidence from Field et al. (2021) show how female empowerment that stems from control over one's transfer could have real effects: studying the government workfare program in India discussed above (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' accounts), as well as training on how to use the account, worked more in both the public sector program, as well as in private sector jobs. Three years later, this even shifted perceptions of community norms of female employment.

Finally, note that another serious gender related concern is whether cash transfers themselves, by causing discord within households over spending preferences (particularly if women receive the program), would lead to increases in intimate partner violence (IPV). Buller et al. (2018) review fourteen quantitative and eight qualitative studies and find little evidence, however, that transfers drive IPV. In fact, they find that most of the studies demonstrated evidence of decreased IPV, with only two studies showing overall mixed or adverse impacts. For example, Bobonis, González-Brenes, and Castro (2013) studies the Mexican Oportunidades program and shows that beneficiary women are 40 percent less likely to be the victims of physical abuse (although they are more likely to receive violent threats). Helping women communicate why they want to spend their funds in particular ways may also help women manage tricky aspects of household negotiations: Roy et al. (2019) examine a randomized experiment in which women receive cash or food, and find no effect on IPV after the six to ten months post program. However, women who were also randomized to receive an additional training on communications experienced 26% less IPV.

### 3.2.6: Poverty Traps and Lumpy Transfers

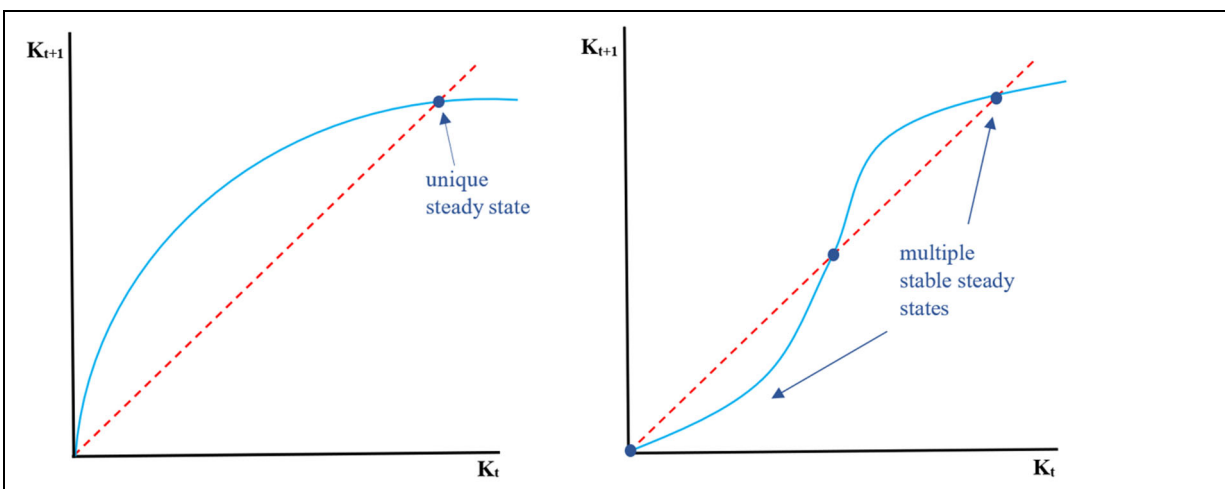
*Conceptual issues:* An important decision when designing a transfer program is the size of the transfer and the frequency with which it is handed out. Thinking about this question in the developing world context is important given that extreme poverty together with saving constraints may be more likely to create poverty traps (Dasgupta and Ray 1986; 1987; Banerjee and Newman 1993; Galor and Zeira 1993). This suggests a role for lumpy transfers for some types of people rather than streams of payments to enable investments.

To fix ideas, consider a very simple model where individuals save a fixed fraction of consumption, and invest the rest. The capital stock in period  $t$  is  $k_t$ . Each period individual produces  $f(k_t)$  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  $\delta$  is the depreciation rate of capital. The steady state level of capital is given by setting  $k_{t+1} = k_t$ , which yields  $\frac{f(k)}{k} = \frac{s}{\delta}$ .

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 unique steady states. Figure 5 illustrates two versions of this simple model with unique and non-unique steady states:

**Figure 5: unique steady state vs. multiple steady states in this model**



Note how a world with poverty traps—i.e. the right-hand panel of Figure 5—illustrates the role for large, lumpy transfers. If an individual is at the lower steady state, a small transfer will increase the consumption a bit, but will not change the steady state. But, a large, lumpy transfer could be enough to push the household past the critical point, so they could end up at the higher steady state. For example, one could imagine all kinds of lumpy capital investments—i.e., buying a non-divisible asset such as a cow, or a sewing machine—that could allow the household to move to the higher steady state. Note also that this model requires credit constraints or some other intertemporal imperfection—otherwise, a household with a profitable business opportunity could borrow and achieve the higher steady state.

From a social protection perspective, which model we are in matters substantially. If we are in the world with a unique steady state, then the goal of a transfer is to raise consumption levels. If we are in the world with multiple steady states, a second goal may be possible: boosting the household enough so they can escape poverty. This comes back 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 the returns it entails.

*Empirics:* The first question to ask is whether poverty traps exist. Balboni et al (2022) examine a randomized trial in Bangladesh that provided assets to a random subset of poor households. They sort people based on their initial wealth and then use the shock of the capital transfer to estimate an empirical analogue of Figure 5, comparing assets at baseline ( $k_t$ ) with assets 4 years later ( $k_{t+1}$ ). They find that, in control villages, there is a unique steady state—akin to the left panel of Figure 5. But, they find, transferring productive assets creates multiple steady states—so that not only does the transfer temporarily increase incomes in the short run, but for some households, it

can be enough to tip them into a substantially higher steady state.<sup>44</sup> Similarly, in the context of microfinance Banerjee et al. (2019) also find important heterogeneity in the effects of microcredit. Businesses that had already been running before microfinance entered, showed persistent benefits that increased over time. Both of these results support the existence of poverty traps—confirmed by their structural estimates of the underlying model of investment and returns—suggesting that poverty could be reduced with large one-time transfers.

More generally, a number of studies have examined big-push ‘graduation’ programs. These types of programs were pioneered by BRAC in Bangladesh and are now present in over 40 countries and make transfers to over 3 million of their poorest households (Banerjee et al. 2020). 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 both capital constraints that households might face, as well as skills constraints to eventually push them out of poverty. Several studies have found substantial 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 year 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).<sup>45</sup> In the long run (10 year follow up), there appears to be no additional growth, but still persistent effects for those who received the program (Banerjee, Duflo, and Sharma 2021). While the programs show substantial long-run effects, they are also expensive, and so whether the returns actually 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 in the other countries, where the discount rate matters.

Two questions that come from these programs are: (1) can the government implement these programs? and (2) to make the programs cheaper and more cost effective, can you scale down some components of the program? Botea et al. (2021) study the government of Zambia’s streamlined version of the BRAC graduation model. Within the large-scale launch, they randomized households into a pure control group, productive asset and savings tools only, training only, or complete graduation package (productive asset, savings tools, and training). First, they show that governments can run the programs: they show increases in consumption, assets and mental health as a result of the full package. In fact, 16-18 months after the first grant disbursement, consumption was 20 percent higher in the complete graduation group compared to the control group, which they note is similar to the results seen in Banerjee et al. (2015). But, importantly, they show similar results between the complete graduation group and the group that receive the productive assets and savings tools; the training on its own had no effect between the treatment and control groups.

---

<sup>44</sup> Lybbert et al. (2004) and Carter and Lybbert (2012) also provide empirical evidence suggesting convexities and showing the existence of poverty traps.

<sup>45</sup> Bandiera et al. (2017) in Bangladesh; Banerjee et al. (2016) in West Bengal look at medium run effects of big push programs.

In another RCT in Ghana, Banerjee et al. (2022) find that neither just giving people assets nor just the savings tool have a positive long-run impact, while the whole package does. This is consistent with the result from Botea et al. (2021)'s Zambia study, if we assume that both the asset and the savings tool are important for the success of the program but the training is not. Combined, these results suggest that perhaps some of the additional human capital investments in these programs can be stripped off these graduation programs, reducing the cost and receiving the same impacts.

Blattman, Fiala, and Martinez (2014; 2020) study a different type of lump-sum transfer program, the Youth Opportunities Program, a one-time unconditional grant to young adults in Uganda. In order to receive the transfer, people formed groups and submitted a grant proposal for vocational training or business start-ups. Funding was randomized and groups received \$382 per member (about \$7,500 transfer per group), but they were not supervised on how they spent the funds. 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, mostly due to the fact that 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 (as suggested by the theory), or the other 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 nine month installments. The first evaluation at 9 months showed evidence that monthly transfers increased food security, while lump-sum transfers increased assets. After 3 years of the transfer there are no differences between households who received lump-sum payments and those with monthly transfers. It is important to note that the monthly payments were only spread over nine months and thus may understate differences when compared to a steady-state transfer program. Understanding this question, and whether even larger lump sums could matter, is an important direction for future work.

The other important question is whether the targeting of the program matters. In the original experimental implementations of the program reported in Banerjee et al. (2015) and the Bandiera et al. (2017) the program was targeted towards the poorest of the poor, though the extent towards which this was adhered to varied across countries. Other implementations, like Blattman, Fiala, and Martinez (2020) are targeted to a broader low-income population, which may be more self-sufficient and able to self-pilot, given the opportunity. In relation to the more recent findings that training itself is not essential, it would be worth exploring to what extent training is critical at least for the most disadvantaged population.

### **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 the benefits get delivered to the targeted individuals.

In many countries, this is a non-trivial problem. In India's MGNREGA workfare program, for example, documents a 74 to 86 percent leakage rate in Odisha in 2007-2008, and Muralidharan,

Niehaus, and Sukhtankar (2016) find a 30.7 percent leakage rate in Andhra Pradesh in 2012. In Indonesia's subsidized rice program, which allows households to buy 15 kilograms of heavily subsidized rice per month, Olken (2006) estimated that at least 18 percent of the assistance never reaches any beneficiaries, and Banerjee et al. (2018) calculate that eligible households only receive a third of the intended subsidy, with substantial amounts also going to the ineligible in addition to some of the transfers being totally unaccounted for. Of course, these are particular examples, and researchers often study leakage precisely in those cases where it is thought to be high, so one should not take these as unbiased estimates of the overall rates of leakage from these programs in the developing world. And, much of the 'leakage' does go to other households who are nevertheless still near poor. But these illustrative findings do suggest that meaningful improvements to programs can be made by improving the efficacy of service delivery.

In addition to general leakage, there is also evidence of manipulation of eligibility rules by households and local leaders. Camacho and Conover (2011) estimate cheating on PMT formulas by documenting bunching in PMT scores just below the eligibility threshold, which emerges over time as the PMT formula becomes better known. 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.

The overall risk of corruption can also distort program choices, as people may prefer programs with a lower risk of corruption. For example, in an Indonesian study, Kyle (2018) shows that in general, 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 and there is a greater fear of corruption, the 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). But recent evidence suggests that reforms to payment systems, private sector involvement, improvements in bureaucratic quality, and information provision can also help. We will discuss the existing evidence of these three approaches.

*Digital payment systems.* One potential issue is in the authentication of beneficiaries, and the links to payment systems. In many developing countries, the system is manual: there is a list of beneficiaries, and some local official needs to verify who receives benefits. This is hard for the central government to enforce—both in terms of who gets access to the list, and how many actual beneficiaries know they are eligible and receive their benefits.

Several studies have examined the implications of digitizing this process. A study by Muralidharan, Niehaus, and Sukhtankar (2016) looks at the randomized rollout of "Smartcards"—biometrically authenticated payments—for receiving benefits from India's workfare program (MGNREGA) and pension programs (SSP) in the Indian state of Andhra Pradesh. They find substantial program improvements associated with smartcards: it reduced payment delays, and reduced leakage by 41 percent. Importantly, the reform does not adversely affect who has access to the program: in fact, the proportion of households who report having worked on MGNREGA increases by 17 percent. Their interpretation for this second result is that 'ghost' beneficiaries fall, which frees up funding for actual beneficiaries.

As described in Section 3.2.2 above, Banerjee et al. (2021) study the conversion of an Indonesian food subsidy program from in-kind to digital vouchers. The vouchers were redeemable using a digital debit card at a network of bank agents, whereas the in-kind program was administered manually by local officials. The voucher program concentrated benefits among the eligible and prevented diversions of benefits to the ineligible, so that delivered 45 percent more subsidy on net to eligible households. As a result, for those in the poorest 15 percent at baseline, poverty fell by 20 percent. The authors hypothesize that the indivisibility of the debit cards—you either receive a debit card with your name printed on it, or you do not—may have been an important mechanism—in-kind transfers were often split among a number of eligible and ineligible households, whereas vouchers were not. Interestingly, this occurred even without biometric identification of any type, suggesting that the digital card, rather than the biometrics, may be important in this case.

It is important to note, however, that care must be taken when implementing these types of systems. Muralidharan, Niehaus, and Sukhtankar (2020) study a reform in the Indian state of Jharkhand where a subsidized food program began to require biometric authentication. Importantly, the state rolled out the reform in areas where at least 75 percent of beneficiaries had linked their biometrically authenticated identification number to their food subsidy account. This meant that some households had not made this link, and so when the biometrics became required, they were unable to receive benefits. While some were able to make this linkage, several percent of households—amounting nearly 300,000 households in a large state like Jharkhand—were unable to claim benefits under the new system. The study emphasizes the need for access to these systems to be really universal before they are required, and/or the need to have some type of robust verification system.

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, but it was the opposite in Ghana where mobile money is less widespread: here, nearly half of recipients preferred electricity transfers and many are willing to forgo high levels of subsidy in order to receive electricity instead of mobile money.

Banerjee et al. (2021) also found related challenges in their study of take-up of Indonesian health insurance. One aspect of the study was to offer people help enrolling at home via the government’s website, in addition to subsidies. They found that more than 80 percent of those who attempted to enroll on the government’s website were unable to do so, largely because of errors in the underlying administrative data (e.g. on the names of who was in their family) that couldn’t be corrected without visiting an office in person. These challenges underscore the importance of investing in the ‘infrastructure’ of social protection, such as reliable national registries, in order to make the transition to more digital processes successful.

*Back-end payment systems.* A related type of reform is linking the back-end payment systems to delivery accounts. Banerjee et al. (2020) study this in the case of MGNREGA in Bihar, which simplified the payment process between the central government and the local one also showed important reduction in leakages. The two biggest changes to the payment system were 1) removing

several layers of bureaucracy and 2) replacing the status quo which involved advanced payments to the local government with “just in time” payments that required listing the beneficiaries before receiving the money. The changes reduced program expenditure by 24 percent while maintaining the same payments to workers. While successful overall, the program also delayed payments but that might be the result of the fact that the administrative changes were recent.

It is worth noting, however, that these reforms must be done with care. Muralidharan, Niehaus, and Sukhtankar (2020) note in their study of Jharkhand that when the government transitioned to such a system where payments were only made against biometrically authenticated payments, it ended up underpaying in some areas where there was not enough of a historical record, which resulted in lower payments in those areas. While these payment systems are important, such transitions must therefore be managed with care.

*Information provision.* If some of the “leakage” arises from local officials not implementing the program as envisioned by the central government, then one approach to fix this is to directly provide information to eligible households. Banerjee et al. (2018) set out to test this in the context of Raskin (the rice subsidy program) in Indonesia. They ran a randomized trial in which the Indonesian government mailed “Raskin identification cards” to beneficiaries in randomly selected villages to inform them of their eligibility and the quantity of rice they were entitled to. This increased the subsidy received by eligible households by 26 percent. The increase came both from an increase in quantity bought and a decrease in the copay price. The benefit that information brings is not just through increasing knowledge but also by creating common knowledge and thus creating accountability of local official’s actions.

*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 can provide lower quality than the government. Empirically, Banerjee et al. (2019) 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 compared to the information provision process. 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 the assignment of these officials to local villages, they show that in villages with these college graduate officials, more households are registered as poor, more are registered as having disabilities and, most importantly, more poor households benefit from subsidies targeted to the poor.

### **3.4 Summing up**

The basic welfare framework that we introduced in Section 3.1 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 (and hence what the marginal utility value of the \$1 transfer

to the recipient is), rather than the details of the form the \$1 in transfers takes. This suggests that understanding the targeting of programs, outlined 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 intact.

However, the flipside of this is that while the cash equivalent of a program, combined with who receives it, is a useful benchmark for evaluating a program, there are a number of circumstances where there may be wedges such that this does not hold, and other design features may be important. 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. This may be valuable above and beyond the value of the cash to the recipient households. Likewise, within-household optimization frictions may mean that program design details, such as *who* within the household receives the transfer, may be also welfare-relevant. The presence of spillovers to other households and poverty traps can also change the welfare calculus, with implications for program design.

The literature discussed in this section also suggests a number of important directions for additional research. First, the heterogeneity in the presence of these frictions suggests that different programs may be appropriate for different types of people. For example, conditional cash transfers may work well if households are near enough the margin of taking up the incentivized behavior that they may actually be induced to do it; for households far from the threshold who cannot comply, imposing conditions may make them worse off by denying them program access (Baird, McIntosh, and Özler 2011). Analogously, some households near the poverty trap threshold can have transformative effects of a big-push program (Balboni et al. 2022), but for households further from the threshold, or who just want a steady job rather than a small business, or do not know how to manage an animal or other asset well, other programs may work better. Figuring out how to assign the right programs to the right people is an important area of work.<sup>46</sup>

On a related theme, if households are not fully optimizing – for example, due to behavioral or cognitive frictions, perhaps made worse by poverty (e.g., 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 for the program. In this context, understanding how best to leverage behavioral nudges to improve program outcomes could be important.

A third area we wish to highlight for future research concerns poverty traps. The literature discussed in Section 3.2.6 suggests that there is the possibility of sustained, long-run changes in people's outcomes from programs such as the so-called 'Graduation approach' discussed above. But this is clearly not for everyone—the standard 'Graduation approach' model may work well in

---

<sup>46</sup> This idea is related to work by Haushofer et al. (2022), who study the question of targeting on the marginal impact of programs in the context of cash transfers, but the problem is likely even more first-order when one considers programs other than cash, where heterogeneity in impacts is likely substantially larger.

a rural area for someone with some ability in animal husbandry, but this may not work for people in urban sectors, or who may want something else in terms of employment. One way of thinking about it is that the studies reviewed above are 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 the program has many different components that may or may not matter for a given person. Untangling this knot and figuring out approaches that work for others is an important area for future work.

A fourth area that we wish to highlight is the interactions with labor markets. Programs that provide low-wage work, such as MGNREGA in India, discussed above are one approach that links work to social protection goals. Likewise, in some high-income countries, cash assistance programs are also tied to work requirements, though the job often needs to come from the private sector rather than being provided by the government. This can have targeting benefits as we discussed above, and it can also have utility benefits to the extent that people receive utility or dignity from work per se (Hussam et al. 2021). On the flip side, it also has targeting challenges, for example, if those who are very much in need (e.g. single mothers, the disabled, etc.) cannot easily work. In evaluating these kind of workfare programs, one can also think about other related labor market policies that can also be thought of as part of an integrated social protection system, such as minimum wages, job training, job matching programs, and so on, and this of how these kinds of programs compare or complement workfare programs. Though a full review of these types of labor market policies is outside the scope of this review, thinking about these links may be an important direction for program design.

Finally, the programs discussed in this section have primarily been conceptualized as responses to long-run poverty. But, to the extent eligibility for these programs is dynamic, they can also help provide insurance against shocks. We tackle the question of how to provide insurance against shocks more systematically in the next section; we return in Section 4.4 to the question of how the two types of programs can be interlinked in a comprehensive social protection system.

#### **4. Mitigating risks**

Poverty is not a static concept: someone may belong to the middle class one day, but a bad health shock leading to a loss of income, or the 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 is often difficult or expensive. Moreover, for whatever reason, even at 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, and indeed there is a vast literature on the challenges of private insurance markets, and the need for and consequences of social insurance, in developed countries. But these problems end up being in many ways more severe in the developing world. For many, life in many low- and middle-income countries is particularly risky—for example, 78 percent of poor people around the world reside in rural areas and mostly subsist on agriculture (World Bank 2014), so they face much more income risk than employees. And there is, on net, much less insurance against these shocks, so households end up

having to take much more severe actions to smooth shocks. Chetty and Looney (2007), for example, show how households in Indonesia respond to an unemployment shock by pulling children out of school and by increasing labor supply of other family members; by contrast, they argue, in the United States, social insurance schemes help households smooth consumption, and there is little impact on either education or other household labor supply.

In this section, we broaden our framework to think about the constraints faced by governments when creating social insurance policies. In Section 4.1, we first begin by summarizing the basic theoretical challenges of privately providing insurance, discussing the basic rationale for these programs. We then briefly discuss the additional challenges that may exacerbate the ability to provide private insurance in developing countries, such as the observed low demand for insurance and the challenges with insurance supply. Finally, as many low income and emerging economies are characterized with informal insurance, we discuss the limitations of these kind of markets and the places where formal insurance could help fill in the gaps.

We then review the current state of knowledge on social insurance schemes in developing countries. We focus on programs designed to address two key types of shocks. In Section 4.2, we examine programs that provide insurance against *income* shocks, including unemployment insurance, agricultural insurance, disability insurance, and maternity benefits. In Section 4.3, we discuss programs that are designed to cushion *expenditure* shocks, including health insurance, property insurance, pensions/annuities and funeral insurance.

Note that while the literature on social insurance in developed countries is extensive, this is a topic area where the evidence for low-income ones is just starting to emerge. This is, in part, because there are a lot of differences across countries in the scope of government-provided social insurance, with social insurance programs more common in the middle and emerging economies than the lowest income ones.<sup>47</sup> We therefore also discuss directions for future research in light of this.

## **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 studied a model where the government redistributes between households with different realized incomes. The premise of insurance markets, by contrast, is that a lot of this 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 high ability type (and thus will earn higher income  $y_{hi}$ ) and with probability  $(1 - p_i)$  she will be a low ability type (and earn lower income  $y_{li}$ ). If the probabilities are public information, then the individual can purchase actuarially fair insurance that pays  $y_h - y_l$  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 of this type, individuals will insure themselves fully against risk to smooth their

---

<sup>47</sup> For example, for these reasons, social insurance programs are, for example, much more highly developed in countries in Latin America. See Levy and Norbert Schady (2013) for a nice review, including a discussion of the role of social assistance and other redistributive programs in reducing poverty and inequality in the region.

consumption across states by purchasing insurance. In this world, the private market will provide all sorts 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 comes from the fact that the individual has private information about their distribution of outcomes that the insurer can't 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. The insurance company would want to charge a higher premium to those with a worse distribution of outcomes, but it can't easily identify them. If it prices its contract to at least break even with the high-risk types, then those with lower risk won't buy into the product offered. In an extreme case, the market can completely unravel, and no insurance contracts exist (Stiglitz and Weiss 1981).

The second source of market failure is moral hazard, i.e. the idea that once you insure individuals against adverse events, you can encourage adverse behavior. For example, if an individual is insured against unemployment spells, they might put less effort in finding a new job. In the model, this can be included 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).

Market failures stemming from either adverse selection or moral hazard already provide justifications for some government involvement in insuring against risk, as private markets may not adequately provide for insurance in these cases. In the following sections, we also discuss additional challenges to providing insurance.

#### **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, anyone with risk-aversion would buy it. And indeed, private insurance markets in developed countries are enormous: the private insurance market accounts for about 9 percent of GDP across the OECD, for example (OECD 2020). By contrast, it 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). Why?

This section explores several possible explanations for low insurance *demand*, i.e. the idea that faced with even actuarially fair insurance, many consumers in developing countries would not want to buy it: information, credit constraints, basis risk, and trust.<sup>48</sup>

##### *Information, trust, and experience*

One challenge, in low-income countries in particular, is that people may not fully understand insurance products, and therefore may be unwilling to purchase them. These factors may also be related to the fact that an insurance purchase involves a large degree of trust: you need to pay money now in the hope that sometime in the future, if you have a problem, the insurer will have the financial resources to cover your loss, and that they will actually do so rather than deny the

---

<sup>48</sup> The interested reader is also referred to Cole and Xiong (2017), who discuss many related issues in the particular context of agricultural insurance in developing countries.

claim. Many developed countries have complex systems of insurance regulation and legal processes to ensure that this is the case; even there, denial of insurance claims is 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.

Several papers provide evidence along these lines, typically by inducing people to purchase insurance (e.g. through a temporary subsidy) and showing that this leads to greater insurance demand. One challenge, however, is that insurable shocks are often rare—after all insurance works best against small-probability large risks—so it can be difficult to learn from experience. Cai, de Janvry, and Sadoulet (2020) investigate this issue in the context of weather insurance in rural China, that is provided by the People’s Insurance Company of China (PICC), and show that the *combination* of both a) having experienced a shock, and b) receiving education so that people use the experience of the shock to better understand how insurance works, leads to greater increased insurance demand in the long run. Cai, De Janvry, and Sadoulet (2015) and Cai, de Janvry, and Sadoulet (2020) also show that one does not need to experience the shock themselves; they find broadly similar effects when they examine shocks (and payouts) received by a member of the potential insuree’s social network. Cole, Stein, and Tobacman (2014) also delve deeper into the types of experiences that matter for insurance demand, and find that personal experience with insurance has longer-lasting effects than experience of others in one’s social network.

Of course, it is worth noting that this is not only a problem in private insurance markets. Experience also matters in government provided insurance programs that require co-payments. Both Asuming, Kim, and Sim (2021) and Banerjee, Finkelstein, et al. (2021) find that temporary subsidies can lead to longer-run insurance demand through experience. In contrast, simply providing information alone, often may not be enough, as a number of papers that attempt to simply explain to people the benefits of insurance find little effect for either public or private types of insurance (e.g., Dercon, Gunning, and Zeitlin 2019; Banerjee, Finkelstein, et al. 2021).

One factor that may reconcile these two effects—that close experience with insurance products can influence demand, but merely explaining the value of insurance does not—is trust. If you see insurance actually pay out, you may be more likely to believe it will pay out again in the future. However, this may be a slow road to increase insurance demand if payouts are rare. Cole et al. (2013), for example, ask whether changing the identity of the educator can make a difference here. In particular they randomize the identity of the educator when studying demand for weather insurance, and they find that demand is 36 percent higher when the educator is endorsed by a trusted local representative (in their case, a microfinance agency with long experience in area).

#### *Credit constraints and timing issues*

One important challenge in insurance demand is that premiums are typically paid for up front. If households are credit constrained—which is widely considered a challenge in many developing country contexts—they may be less likely to take up insurance against income shocks if they need to pay the insurance premiums before income is realized. Put another way, in the language above, if there are credit constraints, there is a big difference between having to pay a premium  $m$  in period 1 in order to receive a payment of  $y_h - y_l$  in period 2 if income realization is low, vs. having to pay nothing in period 2, but instead signing a contract in which you agree to pay  $m$  in period 2 if income is high in return for a promise of receiving  $y_h - y_l - m$  in period 2 if income is low.

These are equivalent contracts without credit constraints, but with credit constraints, the second contract is strongly preferable, because the individual only needs to pay a premium in the time and state when they have the liquidity to do so.

Casaburi and Willis (2018) develop this argument and test it in the context of agricultural insurance demand in Kenya. Specifically, they run an experiment with over 600 farmers in Kenya, and randomize the timing of the insurance premium  $m$ , either in advance or at harvest time. While demand for the standard contract that required payment upfront was only 5 percent, payment of the exact same contract (including adding interest to  $m$  to the payments are the same in net present value terms) increased take-up to 72 percent. By contrast, cutting the upfront premium by 30 percent increased take-up from 5 percent to 6 percent.

Of course, the challenge with a contract of this type is it requires commitment from the insured to actually pay the premium in the low state. In some contexts, this can be enforced—for example, if it can be garnished from wages or from agricultural sales in a context where there is a single buyer, or perhaps in other contexts governments can be involved in ensuring that payments are made. In the agricultural context, one promising option could be to sell insurance contracts for next year's harvest at the time of this year's harvest (as is done in the United States with the government's Federal Crop Insurance program, for example) —allowing payments to be made in advance, but while individuals have liquidity.

#### *Basis risk*

Basis risk provides another reason why demand for insurance may be low, particularly for agricultural insurance contracts (discussed in more detail in Miranda and Farrin 2012, Carter et al. 2017, Cole and Xiong 2017, and 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*, where the prediction takes into account only exogenous determinants of risk. In developing countries, weather (e.g. rainfall) is the most important such predictor used for insurance contracts. The challenge with these contracts is that the prediction is imperfect. This is known as 'basis risk,' i.e., the difference between losses actually incurred and the losses insured based on index values, and can reduce demand for agricultural insurance products substantially. We discuss this in more detail in Section 4.2.3 below.

#### *Summing up*

In short, factors of experience, trust, credit constraints and basis risk may stifle demand for insurance products in developing countries, even if actuarially fair. While there is some evidence that government provided programs are also not immune to these issues, especially when they include co-pays, it is an open question whether government can help solve some of these issues that private insurers may not be able to fully overcome, either through providing social insurance directly or working in conjunction with private providers to help address the concerns and challenges of individuals that squash demand.

### **4.1.3: Insurance supply issued in developing countries**

The literature on the supply challenges for insurance in low- and middle-income countries is less developed, but the problem can typically 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. Here, we give two examples, but these are largely speculation; we suggest this as an important area for future research. The first comes from agricultural insurance. In developed countries, in addition to weather-based 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 (perhaps through the government) needs to be able to observe the output or profits of all nearby farmers. In contexts with detailed data collection for tax or other purposes; 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 second example comes from property insurance. Globally, property and casualty insurance comprises about 1/3 of total insurance premiums; 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. We suspect that the small size of this market is not particularly well understood. But beyond the demand factors listed above, there are reasons to think that supply-side factors may be an important constraint here. For example, property insurance relies on the insurer (and the insured) knowing the market value of the property. In developed countries, property prices are often well recorded and public information; the same is not true for most houses in developing countries where property markets are often less formalized and where record keeping is not nearly as good.

Second, challenges in contract enforcement, the legal system, and the regulatory environment may also impede insurance provision. As described above, insurance requires trust—that the insurer will deliver on the contract and not renege if a claim is made. Part of this can be enforced through reputation, but most developed countries do not stop there: there is typically often a regulatory role for the government to make sure insurers have sufficient capital to pay typical claims (and reinsurance to cover unexpected losses); and there is a functioning legal system through which people can sue if they are denied a claim to which they believe they are entitled. These systems are substantially less developed in most developing countries.

Whether providing insurance through government systems can help, in part, solve these supply challenges remains an important dimension for research. Surely, government systems—either purely public, or public-private partnerships—will face many of these same information and regulatory challenges as well, and so understanding how they can be improved within the context of government provided social insurance remains an important direction for future research.

#### **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 arrangements. For example, in the context of rural India, Rosenzweig and Wolpin (1993) show that Indian farmers purchase bullocks not just as a means of production but also as a nonland asset that is easy to transact with in periods of bad weather. Other risk-mitigating strategies are storing grains from one season to the next, borrowing (or lending) from others in the village, diversifying occupations within the family or gifts and transfers in family or community

networks (Morduch 1999), and remittances from family members who have moved to other cities or other countries (Mazzucato 2009; Millán 2020).

### *Limits of Informal Insurance and the Need for Formal Insurance*

There is evidence that villages are better at smoothing idiosyncratic consumption shocks than one would have expected given the absent formal insurance markets (Townsend 1994), suggesting that informal insurance mechanisms may have a lot of value. But, from the perspective of a potential social insurance designer, it is important to note that just because informal insurance exists does not mean that there is not a need for formal insurance schemes.

One key concern with relying on informal insurance mechanisms is that they are sustained in relatively closed networks (see, e.g. the model of 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 more formal insurance products may increase with development.

A second key challenge is that informal insurance systems can “trap” people in areas where they may have fewer opportunities. For example, 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). Understanding the degree to which introducing or improving formal social insurance arrangements allows people to take advantage of migration opportunities is an important dimension for additional research.

Third, informally insuring each other within a village or locality does not necessarily provide the level of coverage that households need since income shocks have a strong spatial correlation, especially in areas that rely heavily on agriculture. Ideally, risk-pooling would be done across villages where shocks are independent, but these arrangements are difficult to uphold in practice since monitoring costs to prevent moral hazard are too high.<sup>49</sup> These spatially correlated shocks may be a particular challenge with natural disasters or conflict, where an entire area may be affected, necessitating large scope disaster relief programs.

All of these arguments suggest that some forms of formal insurance are likely to be very useful even where informal mechanisms are present and working. This then makes it necessary to think through the potential interactions between (formal) social insurance and informal insurance.

### *Linkages between formal and informal insurance*

When thinking about the design of public programs, it is also important to consider how formal insurance will interact with existing informal insurance systems, especially as there is evidence that even as economies experience growth, many households keep relying on informal networks (see, e.g., Besley and Levenson 1996 for a related argument on informal credit). Ultimately, there are two key questions: first, does offering formal social insurance crowd-out informal insurance; and second, what are the welfare consequences of the interaction between the two.

---

<sup>49</sup> One popular mechanism to get around this problem is marriage, if the marriage takes place to someone from another village. This allows two households to insure each other with low monitoring costs, since the new household has ties to two villages with perhaps less correlated shocks (see, for example, Rosenzweig and Stark (1989) and Rosenzweig (1998)).

The evidence is mixed regarding the relationship between social insurance and informal risk sharing. Albarran and Attanasio (2003) look at individual transfers following the introduction of Progresa in Mexico and find crowd out of private transfers especially in villages where the variance of income is smaller. Strupat and Klohn (2018) document large crowd-out effects after the implementation of the National Health Insurance Scheme in Ghana. This crowd-out might be desirable as it relieves (poor) families of the burden of carrying others. On the other hand, Takahashi, Barrett, and Ikegami (2019) study this empirically and find little relationship between formal and informal insurance. They study in the context of southern Ethiopia where an index-based livestock insurance was introduced. They find no detectable crowd-out of informal risk management arrangements. While they find weak evidence of crowd-in of insurance (consistent with Dercon et al. (2014)'s model, which shows, theoretically, that the presence of basis risk in index insurance makes informal risk sharing a complement), they conclude that the crowd-in results are not robust, and suggest this is an important dimension for additional research.

In an important theoretical paper, Arnott and Stiglitz (1991) conclude that whether the existence of formal and informal insurance systems is welfare improving depends on how well informal networks monitor others to reduce the moral hazard problem. 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. In the other extreme case, welfare decreases since the formal insurer knows that there is no monitoring going on and thus has to adjust premiums and payouts. Attanasio and Ríos-Rull (2000) also show how even a well-intentioned policy might induce a reduction of welfare. The intuition behind their result is that the introduction of the policy leads to a destruction of the social fabric which leaves household's more unprotected in terms of the risk that is not covered by formal insurance.

Another way in which the provision of formal insurance have positive welfare effects through its interactions with informal insurance is if people take costly actions to maintain access to informal networks. If, for example, people do not migrate to cities (which is productive) in order to maintain access to rural informal insurance networks, the provision of formal insurance could allow them to make these more productive migration decisions (Banerjee and Newman 1998; Munshi and Rosenzweig 2016).

In more recent work, Huang and Zhang (2021) finds no crowd-out of private transfers in the context of China's rural pension scheme (which is discussed more in detail below). This is in contrast to (Jensen (2004) which finds large crowd out effects of South Africa's public pension to the elderly. Huang and Zhang argue that one reason that they find less crowd-out in China than South Africa is that the pension benefits are much smaller.

#### **4.2: Insurance Against Income loss**

We next turn to understand the challenges for specific types of risk. This section covers risks of income loss; Section 4.3 below covers expenditure risks.

#### 4.2.1. Insurance against involuntary job loss

Unemployment insurance is designed to help people smooth consumption between jobs. However, a key challenge is that unemployment benefits typically continue until the worker resumes work, and there is an extensive literature on the degree to which this conditioning of benefits on future employment discourages job search (Baily 1978; Chetty 2006). In most high-income countries, the literature has focused on the tradeoffs between the welfare gains from the insurance component of these programs against these potential disincentive effects, and whether the different policy design choices made within these programs (i.e. the level and the duration of benefits, etc.) can help improve overall welfare. There is also an additional effect whereby the unemployment benefits allow households to spend more time searching for a new job, potentially resulting in a better job-match than if they did not receive the unemployment benefits, though not all studies find such an effect (Nekoei and Weber 2017; Farooq, Kugler, and Muratori 2020). Examples of papers estimating these tradeoffs in low- and middle-income countries include Cunningham (2000), van Ours and Vodopivec (2008), Huneus, Leiva, and Micco (2012), and Gonzalez-Rozada and Ruffo (2016; 2022).<sup>50</sup>

In low- and middle-income countries, these same policy questions around unemployment insurance exist, but it is further complicated by another factor: many workers are employed in a large informal sector, and the government cannot monitor job entry and exit in this sector. This has two important implications for policy. First, the government cannot provide benefits that start conditional on unemployment for those in the informal sector, and so many—and, in particular, the most vulnerable—may not be actually covered by insurance. Second, the government also has challenges conditioning benefits on re-employment because it can only monitor whether one accepts a new job in the formal sector. Thus, workers receiving unemployment insurance could choose to seek employment in the informal sector in order to continue receiving their benefits. Therefore, there is a worry that these kinds of programs may also disincentivize formalization, which has consequences for productivity, taxation, workplace safety, etc.

Gerard and Gonzaga (2021) examine the relationship between informality and unemployment insurance in Brazil, using rich administrative data. They find, perhaps counterintuitively, that the presence of a large informal sector *reduces* the efficiency consequences of moral hazard in UI, and indeed, they estimate that the efficiency costs of UI are 5 times lower in Brazil than in the United States. This is 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 getting an informal job, this reduces the overall efficiency consequences of the UI tax on re-employment. Of course, this comes at a cost, because informal jobs have lower earnings than formal jobs, a finding echoed in Liepmann and Pignatti (2021)'s study in Mauritius. These papers suggest that understanding the optimal UI design in the presence of a large informal sector (e.g. would it be better to switch to lump-sum payouts?) is an important question for future research.

---

<sup>50</sup> A number of recent papers also examine the spillover effects of unemployment insurance both within a household and across households. Within the household, Bhalotra et al. (2021) shows that the mass layoffs in Brazil led to large increases in domestic violence, but that unemployment insurance for men did not mute this observed effect. Further, eligible men were then more likely to commit domestic violence when their benefits are expired. Examining spillover effects on crime levels, Britto, Pinotti, and Sampaio (forthcoming) finds that individuals subjected to mass layoffs in Brazil are more likely to be prosecuted after being laid off, but that receiving unemployment insurance helps offset this effect.

Given the challenge of implementing unemployment insurance strategies with a large informal sector, one often observes alternative policies to help to insure workers against job loss. One very common policy worldwide is to mandate severance pay; that is, a lump-sum payment given to workers at the termination of a labor relationship. This payment is not conditioned on future employment, and thus does not distort future employment decisions. But, these programs are not without their own administrative challenges: someone needs to adjudicate that the worker was, indeed, working, and was terminated. If firms are supposed to pay directly, the government needs to make sure they do so; if the government collects taxes and pays severance, it needs to collect the taxes to do so.

Weaker quality of institutions may further exacerbate these challenges. For example, severance is paid based on involuntary separations only; otherwise, workers could voluntarily quit to obtain the benefit. This requires a legal system that can adjudicate of disputes as to whether the separation was voluntary. Sadka, Seira, and Woodruff (2018) examine this in the context of Mexico, where despite a strong severance policy on the books, many workers do not receive their full entitlement, and delays and misinformation in the court system means that advocating for their rights is a challenge. They show that providing information about likely court outcomes substantially shortens settlement times, which in turn helps reduce the probability that the worker does not have enough food to eat or is unable to pay for basic services. In short, the paper illustrates that role of broader institutional quality in providing access to insurance against job loss.

While severance payments may reduce moral hazard for workers (since they receive a lump-sum payment, which does not distort future labor supply decisions), the need to pay large severance payments in the event of a termination may serve as a disincentive for firms to hire in the first place, and a single lump-sum means that the worker bears the risk of being unemployed longer than average. Moreover, studying administrative data from São Paulo, Brazil, Gerard, Naritomi, and Silva (2021) show that workers appear to over-spend from the lump-sum severance payment relative to their optimal consumption profile. Importantly, this was not driven by durable good purchases (i.e. using the large severance payment for investment in productive assets and/or new businesses). They suggest that a stream of payments (which could, of course, also be unconditional) may provide better consumption smoothing than a lump-sum.

A different alternative that has been widely discussed (first proposed, to the best of our knowledge, by Feldstein and Altman (1998)) has been to create unemployment savings accounts, where workers having an individual account that they can access in case of job loss. The accounts can be funded through workers and/or firms being required to contribute to the accounts for each month the worker is employed. This system helps align incentives and reduces the scope for moral hazard. However, it does not have the risk pooling and redistribution feature of traditional insurance; although, in cases where the individual account does not have any money, or has too little, the government could supplement the funds guaranteeing a minimum. Countries such as Colombia, Chile, Indonesia, and Mexico have implemented this type of system. Kugler (2005) and Nagler (2013) examine the impacts of the shift from severance to these types of linked accounts in Colombia and Chile, respectively, on ex-ante labor conditions.

The challenges of the large informal sector, and the effective inability to condition payments on continued unemployment, thus suggests that the design of unemployment insurance schemes may be fundamentally different in low- and middle-income countries from what one is used to in high-income countries. It also suggests several avenues for future research. In particular, it suggests that policies that are not conditioned on re-employment (e.g. individual unemployment-triggered savings accounts, severance payments paid by employers, lump-sum unemployment insurance paid by the government) may be an important way forward, but each of these programs has its own design questions—such as the optimal level of benefits, the degree to which employers should be experience-rated, and how to verify unemployment triggers. Moreover, these policies themselves only help workers already employed in the formal sector—but the presence of the large informal sector means that the many workers employed informally would not themselves be protected. Are there mechanisms, perhaps using self-targeting techniques, to provide insurance to those in the formal sector? For example, one could consider a program that requires job training as a function of providing unemployment assistance. Understanding both how to design programs for those in the formal sector given the presence of the informal sector as an alternative option, and how to design programs for those already in the informal sector, remains an important area for additional research.

#### **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 the challenge of providing unemployment insurance, the large informal sector hinders governments from providing universal disability coverage or workplace accident insurance through employers.

Therefore, one often observes alternative policies to provide financial assistance to the disabled. For example, many cash transfer programs, such as Argentina's *Programa de Ciudadanía Porteña*, Chile's *Subsidio Unico Familiar*, Indonesia's *Program Keluarga Harapan* and Tanzania's *Productive Social Safety Net* give additional weight to having a disabled household member in the eligibility criteria. This, however, only provides insurance for those near the threshold of poverty—if a person who is middle class experiences an income loss due to disability, but not enough to become poor, they may not qualify for benefits.

Moreover, if the targeting is done infrequently, the sign-up period may not actually align with when people need assistance. In high-income contexts, moral hazard is often a concern—i.e. disability programs usually preclude work while receiving payments (e.g. Autor et al. 2016), so there is a disincentive for work for people on the margin. In low- and middle-income countries, partly because of the possibility of informal work, this effect may be weaker but to the best of our knowledge there is no evidence on this issue from these countries and indeed this whole area is relatively unexplored.

A related issue is workplace accident insurance. This is often separate from disability insurance, in part because of an attempt to link the risks back to the employers. While some middle-income countries are trying to provide this type of insurance at least to those workers in the formal sector, to the best of our knowledge this remains an area that is comparatively underexplored in the economics literature.

Finally, we also want to touch briefly on life insurance. The death of a primary income earner is a substantial income loss, and indeed, throughout many low- and middle-income countries, there is a recognition that widowhood is often 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 indeed much more of this risk remains uninsured (“Life Insurance Volume - Country Rankings” 2019). Moreover, the relative paucity of formal pensions—which in high-income countries, often come with survivors’ benefits which help insure spouses against the income loss of their primary earner—further leaves more of this risk uninsured. As with disability, widow status is often considered in the PMT formulas, but similarly, this only provides insurance to the extent that households are near the poverty threshold. Understanding how to better insure the spouses of the diseased is an important area for additional research.

#### **4.2.3: Insurance against agricultural loss**

Agriculture plays a large role in low- and middle-income countries—in fact, agriculture is the main source of income and employment for about 60 percent of the population who live in low-income countries (ILOSTAT 2019). And 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, and 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 risks. Some of this involves around providing technologies that help ensure more stable yield and prevent crop loss, be it investments in irrigation or weather resistant seeds. Other policies revolve around providing agricultural insurance, to help farmers smooth consumption during periods of crop loss.

Formal agricultural insurance programs can be classified into two broad categories: direct indemnity-based insurance that bases payouts on farm-specific realizations, and index-based insurance that pays out based on an objective, reliably measured index such as rainfall, or the yields of other farmers in nearby areas. Both approaches have strengths and limitations.

In the direct approach, verification of loss at a plot level is costly. And even with direct verification, moral hazard and adverse selection present important challenges: farmers insure the worst land or riskiest crops, and may not put in as much effort avoiding crop failures if there is insurance. For example, Gunnsteinsson (2020) illustrates both of these challenges in a field experiment in the Philippines, working with a government-owned insurance company. Comparing plots of land for farmers that are randomly chosen for insurance with those that are not, he shows that farmers take fewer precautions to prevent pest and diseases on insured plots, consistent with moral hazard. Moreover, when he gives farmers the choice to rank which plots they would prefer to be insured, farmers select plots that are prone to flooding and crop disease, consistent with adverse selection.

Given private information concerns, instead of insuring against loss directly, agricultural insurance contracts are often written to pay out as a function of an index of *predicted losses*, where the prediction takes into account only exogenous determinants of risk, and hence farmers have no

incentive to reduce their own effort. In developing countries, weather (e.g. rainfall) is the most important such predictor used for insurance contracts. In other contexts, insurance payouts are a function of yields from other farmers in the same area. The literature on index agricultural insurance has been extensively reviewed elsewhere (Miranda and Farrin 2012; Cole and Xiong 2017; Carter et al. 2017; Jensen and Barrett 2017); we just touch on a few issues here.

Index-based schemes may also reduce transaction costs since payouts are linked to an external or aggregate indicator, curtailing the need to verify actual losses. Given these potential advantages over a more traditional policy, most programs in the developing world are now index-based. According to a 2020 survey (GIZ 2021) out of the 54 programs surveyed that cover 265 million people, 80 percent of them were index based (of which about 30 percent were area-yield based and the rest were weather based).

The challenge with these contracts is that the prediction is imperfect. This is known as ‘basis risk,’ i.e., the difference between losses actually incurred and the losses insured based on index values. For areas with substantial microclimatic variation—so that rainfall in one village does not exactly predict weather in a nearby village—this can be a major issue.

There is some empirical evidence that basis risk lowers private insurance demand: for example, Mobarak and Rosenzweig (2013) test this experimentally by randomly varying whether a rainfall gauge used for insurance is placed in the village, or further away—with the idea that the closer rainfall gauge has lower basis risk. They find that for castes that do not provide informal insurance (discussed in more detail in Section 4.1.4 above), demand is lower when there is more basis risk. Giné, Townsend, and Vickery (2008) also report lower take-up when basis risk is higher, and Karlan et al. (2014) provide evidence that farmers are concerned about this, though neither of these studies varies this experimentally as Mobarak and Rosenzweig (2013) do. Moving forward, using advances in machine learning to reduce basis risk could be a promising additional dimension of research, as well as developing contracts based on realized output of nearby farmers (as is done in the United States, for example) for developing country settings.

Low adoption of agricultural insurance may then further be exacerbated by limited trust in providers, a lack of understanding of products, and liquidity constraints, further necessitating the need for governments to play a role in the agricultural insurance provision. We refer the interested reader to nice reviews by Cole and Xiong 2017; Ali, Abdulai, and Mishra 2020; Nshakira-Rukundo, Kamau, and Baumüller 2021 for additional reviews of this topic.

#### **4.2.4: Parental Benefits**

Parental benefits are designed to help ensure 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, as well as to provide employment protection to protect against termination during this period of life. One common type of paternal benefit provides a mechanism in which workers can take time off around pregnancy and a child’s birth, by mandating allowed paid time off and/or subsidizing firms for the worker’s time off. A second parental benefit takes the form of subsidizing child-care to allow parents to re-enter or stay in the workforce.

While virtually all countries, except seven, have some sort of maternity leave policy on the books (World Bank 2018c), the benefits levels in low and emerging countries are often very low for mothers and not necessarily well-enforced in practice, paternity leave is often nonexistent, and many of those in the informal sector—large fractions of the workforce in low-income countries—do not necessarily have access to any of these formal benefits.

There is an extensive literature on parental benefits in high-income countries. Looking first at maternal paid leave, despite its insurance properties, one frequent question is how it affects women's subsequent careers; for example, if maternal leave leads more women to leave the labor force following childbirth, are there then consequences for their long-run labor supply and careers? A number of recent studies have shown, however, little long-run effect: while increased paid leave results in longer delays in returning to work in the short-run (as expected), they generally find little to no effect or no impact on longer-term female labor force participation, experience, employee performance ratings, or earnings (see, for example, Lalive and Zweimüller 2009; Lalive et al. 2014; Schönberg and Ludsteck 2014; Dahl et al. 2016; Bana, Bedard, and Rossin-Slater 2020; Bartel et al. 2021; Kleven et al. 2020).

Similarly, there is also an extensive literature on where free or subsidized childcare improves female labor force participation in high income countries, again helping to prevent income loss. Here, the evidence appears more mixed. For example, a number of recent papers have shown that expanded or subsidized childcare studies increase female labor force participation (Lefebvre and Merrigan 2008; Bettendorf, Jongen, and Muller 2015; Nollenberger and Rodríguez-Planas 2015). On the other hand, Kleven et al. (2020) shows little impacts of similar reforms on the earnings of women relative to men, and Cascio (2009) finds positive impacts only for single mothers.

In contrast to the extensive literature in high-income countries, there is relatively little empirical research on paid maternity or paternity leave in lower- and middle-income countries (Amin and Islam 2019; Fallon, Mazar, and Swiss 2017). One recent exemption is Vu and Glewwe (2022), which shows that in Vietnam, a more generous maternity leave shifted more women who are potentially eligible away from informal work to formal work, and in particular public sector jobs where they can access the benefits. The lack of research in this area may be, in part, due to more limited enforcement of these types of policies, particularly in lower-income countries where many workers are in the informal sector, or government capacity for enforcement is weaker. More research needs to be done to understand both how pregnancy affects consumption smoothing around the time of birth, as well as the career trajectory of women; and how to design effective leave policies given the presence of these informal labor markets.

In contrast, free or subsidized child-care has become increasingly common in many developing countries, particularly middle-income countries and a number of studies have tried to examine impacts. For example, exploiting a triple difference design, Calderón (2014) in Mexico shows that access to childcare increased the likelihood of working for women and reduced the likelihood of earning zero income. While most men worked, they found a small negative effect on male work (perhaps men substituting leisure for some of the additional female income). Similarly, Halim, Johnson, and Perova (2022) examined the expansion of public pre-school expansions on women's work in Indonesia, which has a large informal sector. Relying on a triple difference approach, they find positive effects on women's work participation, but it is driven by increases in unpaid family

work. They argue that because the provided childcare was only operating for 3 hours a day, it makes it unlikely that women could find jobs outside the home with longer time commitments. Finally, several recent experimental studies (Barros et al. 2011; Martínez and Perticarà 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 the gains to family income from childcare lead to at least as large a gain as an equivalent cash transfer, while also improving child development. Specifically they find that the increase in earnings with childcare comes from increased productivity rather than longer hours, consistent with the theoretical model of Banerjee and Mullainathan (2008). How these policies affect female labor-force participation in the longer-term remains an open question.

As is clear from this literature, there are still many open questions regarding benefits, particularly given the nature of the labor markets observed in low- and middle-income countries. For example, how should parental benefits be paid and who should pay for them? Should they be universal or dependent of labor status, especially in countries where the level of informal firms may make it hard to mandate benefits? How does job-protected, paid leave affect labor force participation, particularly for women? And, given the presence of the informal sector, does mandating maternity benefits result in women being concentrated in less good, informal jobs? 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 risks—i.e. risks of unexpected large outlays. Prominent examples of this are health insurance, which insures against the cost of unexpected medical expenses; property insurance, which insures against the cost of damage to property due to fire, flood, or other causes; and funeral insurance, which covers the cost of a funeral. We discuss each of these in turn.

#### **4.3.1: Health Insurance**

Government-led health insurance systems are increasingly becoming common in many low-income and middle-income countries, as governments aim to increase health care utilization, improve health outcomes, and help households manage income and consumption risks that may arise from a severe health shock. Nearly 190 countries around the world have some sort of public health insurance system.<sup>51</sup>

The literature on health insurance in low- and middle-income countries tends to fall into two common buckets. The first bucket examines the impact of health insurance on households who have this insurance, including the ability to smooth out the economic consequences that may arise from health shocks as well as impacts on health utilization and, potentially, health. The second bucket then explores the challenges in provision and design of insurance, particularly around how

---

<sup>51</sup> 188 countries in the WHO's Global Health Expenditure database report positive, non-zero spending on government health schemes in 2019 (World Health Organization 2019). These schemes are defined as noncontributory, publicly funded spending arrangements (World Health Organization 2021).

to design and fund public insurance systems given the constraints coming from informal employment, adverse selection, moral hazard, information failures and trust.

### *Impacts of Health Insurance*

We start with the first bucket first, reviewing the studies that aim to measure the impacts of health insurance. In particular, we first examine a number of recent studies that experimentally test for the health impacts of insurance by randomizing households into a treatment group where they are offered a highly subsidized or free insurance and comparing them against a control group that is not offered the subsidy (or offered a cash transfer in an equivalent amount to the insurance), as well as quasi-experimental studies where the government has rolled out public health insurance over time.

Before beginning, it is worth noting, conceptually, that the primary outcome one would expect health insurance to affect is consumption smoothing. That is, health insurance is primarily a *financial* product, that provides payments in the event of an expenditure shock, and so should make the financial consequences of bad 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, giving people a lump-sum cash transfer when they are diagnosed with illness, which would help with the financial consequences without changing prices on the margin), one would expect utilization to increase. This increased health care utilization could, in turn, potentially affect health, but one would expect this only the increased health care utilization induced by the health insurance (which may be the lower-value health care utilization, since it is what is most price sensitive) in turn affects health.

*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. For example, King et al. (2009) randomly assigned 50 of 100 communities in Mexico to receive the government-run Seguro Popular health insurance at least a year earlier than expected. They find that having health insurance reduced out of pocket expenses for medical care, as well as reduced the fraction of catastrophic medical expenditures in the respondent's spending. In a different context, Gruber, Lin, and Yi (2021) examines China's public health insurance for rural households, the National Cooperative Medical Scheme (NCMS). To evaluate impacts of the program, they take advantage of two key types of variation, specifically that it was rolled out to different areas across time and that the intensity of the program was also proportional to the share that is living in rural areas. They show that there were sizeable reductions in out-of-pocket spending for those with access to insurance in China, particularly at the top of the expense distribution. Levine, Polimeni, and Ramage (2016) randomize a subsidy for NGO-provided insurance in rural Cambodia, and also find that insured households were less likely to take on new debt to a health shock and they were less likely to pay for care by selling assets.

More broadly, health insurance may also reduce financial stress. Haushofer et al. (2020) compare the offer of free health insurance in Nairobi to an unconditional cash transfer of the same value and a control group (with a total sample size of about 789 individuals, and about 640 at endline). While they find no effects on economic outcomes (and no effect on health care utilization, either), they found suggestive evidence that the provision of health insurance reduced levels of

self-reported stress and the stress hormone cortisol relative to cash and a pure control group. Thus, insurance may help provide peace of mind for people that they could handle a future shock. Note, however, that the estimates do not survive when they impose conservative bounds for attrition and multiple inference. They conclude that these results are just suggestive and that more work needs to be done in this area.

#### *Health care utilization and health.*

Health insurance also has the potential to change health care utilization, since it typically reduces the price of health care, which could in turn generate health impacts.

In some contexts, 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 health insurance 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 tend to think many households are, particularly in developing country contexts), 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.

*Experimental evidence.* Does health insurance even change health consumption, and health, in low- and middle-income contexts? The evidence on the impact of health care utilization and health is more mixed, though one challenge may be that some of the experimental studies on this may not be sufficiently powered to detect economically meaningful, but small effects. For example, the King et al. (2009) study of Mexico's Seguro Popular found no effects on medication spending, health outcomes, or utilization. Similarly Haushofer et al. (2020)'s randomization in Nairobi of free health insurance or unconditional cash transfer (with a total sample size of about 789 individuals, and about 640 at endline) do not find effects of the insurance on self-reported health or health care utilization; they argue that they were not powered to measure small effects.

A recent study by Malani et al. (2021) highlighted one issue in particular with health insurance use. They randomly assigned about 10,900 household in India to one of four groups: the offer to purchase health insurance, the offer of free health insurance, the opportunity to purchase health insurance combined with a cash transfer of an equal size and a control group. While they find some increases in households using insurance for payment, they found that many of the beneficiaries had challenges in using their insurance—e.g. having problems with cards, not knowing how to use them, forgetting the cards. In the end, they do not find systematic evidence in favor of increases in health outcomes, perhaps for these reasons.

In short, most of the experimental studies to date find little overall impact of health insurance on health outcomes.<sup>52</sup>

---

<sup>52</sup> Rather than exploring impacts on households in general, who may or may not have health needs during the course of the experiment, an alternative experimental approach is to focus attention on individuals who need specific on-going, immediate care. Bronsoler, Gruber, and Seira (2021) take this approach in Mexico, focusing on diabetics. They find striking results, with participating individuals having lower blood sugar levels and fewer diabetes

*Larger scale changes.* A number of other studies try to improve on statistical precision by focusing on quasi-experimental variation from large-scale public health insurance reforms, and then using large administrative datasets on mortality and national samples to ensure large enough sample size needed to be powered to detect changes in health outcomes.<sup>53</sup> For example, the Gruber, Lin, and Yi (2021) study discussed above examines China’s public health insurance for rural households, and uses both large-scale administrative records to examine mortality (the China Death Surveillance Point Dataset) as well as data from the Chinese Longitudinal Healthy Longevity Survey—a large-scale, survey dataset of the elderly that had seven waves during 1998-2014 across 22 provinces in China. They find a significant decline in aggregate mortality, which they argue could explain 78% of the entire increase in life expectancy in China during this period. Using the survey data, they find large effects on health care utilization and a host of other health outcomes.<sup>54</sup>

Similarly, Gruber, Hendren, and Townsend (2014) examine Thailand’s 2021 health care reform, known as the “30 Baht” program. Prior to this reform, the poor, young and old were given free care in the public health system, but the system was seen as chronically underfunded. About 20 percent of households were covered under formal insurance schemes, while the remaining part of the population was uninsured. The reform changed the system in two ways. First, it provided universal access to public facilities at a co-pay of 30 Baht (or \$0.75) per visit. The co-pay was waived for the poor, so in practice this reform largely led to reduced health access costs for the informal, non-poor. Second, it replaced the old financing system and instead provided provincial hospitals with a universal capitation payment of 1,200 Baht (~\$35) based on population in the province; this had the feature of moving away from user fees funding hospital systems to more generous government financing. Using administrative mortality records, they find large impacts on infant mortality. Specifically, they show that prior to the reform, the infant mortality rate was related to how rich your province was; after the reform, resources were equalized across provinces leading to this correlation disappearing.

Importantly, both Gruber, Lin, and Yi (2021) and Gruber, Hendren, and Townsend (2014) argue that part of why large health effects were seen is that the insurance reforms, by increasing financing by the government for health, increased the supply of health services, as has been shown in developed countries (e.g., Finkelstein 2007). This suggests that in evaluating health insurance systems experimentally, it may be important to randomize across a hospital catchment area or health market, rather than an individual level or village, in order to estimate the full general equilibrium impacts that could arise due to an insurance expansion or enhancement.

del Valle (2021) argues that health insurance also provides another economic benefit beyond reducing or smoothing health expenditures: studying the staggered roll-out of the Seguro Popular

---

complication, even in the short run. This is more about access to care than ex-ante insurance, since the need for care is already clear.

<sup>53</sup> We discuss two key studies here for brevity, but please see Gruber, Lin, and Yi (2021), who provide a nice overview of other studies that do so as well.

<sup>54</sup> These findings are consistent with findings from large-scale expansion of health insurance in the United States, which also find impacts on reducing mortality (Miller, Johnson, and Wherry 2019; Goldin, Lurie, and McCubbin 2021).

program in Mexico, they show that by reducing the severity of health shocks, Seguro Popular reduced the degree to which other household members needed to drop out of the labor force to directly provide care. Thus, health insurance improved labor force participation of informal female workers.

### *Challenges with Health Insurance Design*

The second bucket of work falls 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 what to do about informal workers (for whom premia cannot be collected from employers) who are not poor enough to receive direct coverage from the government budget. There are, broadly speaking, three options. First, there is the option of no health insurance for non-poor, informal workers. This, of course, would mean that these workers do not receive the benefits of insurance, and remain vulnerable to the economic impacts of a health emergency. Moreover, if firms want to evade paying benefits to formal workers and/or workers do not fully value the insurance, there is a concern that mandating insurance benefits for formal workers could increase informality among both workers and firms (Levy and Schady 2013).

A second policy option—that many countries have instituted—is to attempt to mandate that non-poor informal workers contribute to insurance on their own. This is very difficult to enforce, however, particularly in countries with limited administrative capacity. Most of the relevant population is exempt from income-taxes, so tax penalties are not an option—and therefore there is relatively little other available options for sanctioning those who do not pay. As a result, many of the non-poor opt-out.

In addition to worries about increases in informality described above, there is a worry that this will lead to adverse selection problems that will financially strain the insurance system. These adverse selection problems may be particularly more severe in low-income countries, particularly if households lack information on insurance or have limited trust that the system will pay out.

Finally, the third option is to decide that it is too difficult to collect contributions from informal workers, and extend free insurance to most, if not all, informal workers. Once again, concerns about encouraging informality arise in this case.

In what follows, we first discuss some of the evidence on movement to informality when governments increase formal contributions and/or benefits for formal workers or increase subsidized options for informal workers. We then discuss the evidence on the adverse selection, as well as some of the policy tools that have been used to try to mitigate it. We then conclude with open questions for research.

### *Health Insurance and Informality*

Ex-ante, the effect of health insurance provided through the formal sector on informality is ambiguous. In some cases, if a) workers value the insurance and b) there is a cost-savings (from the workers' perspective) of gaining it by becoming a formal employee as opposed to paying an individual premium, this could lead to more formalization. On the other hand, if workers do not value the insurance at cost or if it's cheaper to obtain insurance (or care) without being part of the formal sector, a requirement that workers pay for insurance out of their formal sector wages could lead to increased informality.

Bergolo and Cruces (2014) find evidence of both effects. They examined a large-scale policy reform in Uruguay's social insurance administration (SIA) that both increased benefits and contributions for formal workers. Prior to reform, a health insurance benefit was in place in which the social insurance administration covered the employees' insurance premiums from a private health care provider. The 2008 reform extended the coverage funded by the SIA for dependent children, which made the program more attractive. At the same time, reform also increased the payroll tax contribution deducted from employees' salaried earnings, which could cause an increase in underreporting of wages for formal employees. The authors find both effects. Comparing those without eligible children prior to and after the reform, the greater benefits drew people into the system: the number of individuals in benefit-eligible, registered employment increased by 1.6 percentage points with respect to the comparison group, about 5% higher than the pre-reform level. However, in addition to the accompanying increase in the rate of contributions due to higher formal employment, there was also an increase in misreporting wages. 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, and they find some evidence that it increased informality. Formal workers were expected to make mandatory contributions to the public health insurance system through their employers. 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—creating an incentive to opt out of the contributions by becoming informal. Using a combination of survey and administrative data, they find the reform increased informal employment by about 3-4 percentage points.

A number of papers have also examined the impact of the expansion of Seguro Popular in Mexico on formalization. Prior to the rollout, insurance in Mexico used to be tied to payroll contributions, and so many people who worked in informal jobs were uncovered; Seguro Popular aimed to provide universal coverage so informal workers could be covered as well. Analyzing the staggered roll-out of the program and survey data, several papers including Aterido, Hallward-Driemeier, and Pagés (2011), 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 of Seguro Popular on social security data, Bosch and Campos-Vazquez (2014) showed that it slowed the registration of employers and employees in small and medium firms (up to 50 employees) into social security, reducing revenues paid into the social security system but also VAT taxes.

In sum, the evidence suggests that the informality margin— either complete informality, or underreporting of wages—can respond to requirements that formal sector workers obtain health insurance.

*Adverse Selection, Information and Trust with Imperfectly Enforced Mandates*

As discussed above, many countries have tried to mandate that non-formal, non-poor workers purchase insurance, but limited enforceability remains a challenge. Several studies have examined the implications of limited enforceability of mandates in the context of health insurance in developing countries, and what can be done to mitigate this issue.

*Subsidies.* Two papers, for example, ask whether time-limited subsidies can be a substitute for unenforceable mandates.<sup>55</sup> Asuming, Kim, and Sim (2021) study this in the context of the Ghanaian health insurance scheme, randomizing offers of 1/3, 2/3, and full subsidies of premiums for one year, and tracking results over the subsequent three years.<sup>56</sup> They find an increase in take-up as a result of the subsidies that persist even after the subsidies are done. They also find that those who respond to the subsidy are more ill (which is perhaps surprising, since they presumably value insurance less) and utilize baseline services more than those who did not enroll regardless of the subsidy, but that the level of selection does not vary by subsidy level.

Banerjee, Finkelstein, et al. (2021) study this in the context of the Indonesian national health insurance, randomizing offers of 1/2 and full subsidies of premiums for one year (and a control), and tracking results over the subsequent 20 months. They find that larger subsidies bring in healthier individuals into the health care system, consistent with adverse selection; Banerjee, Finkelstein, et al. (2021) also document that those who enroll with the subsidies are healthier than those who enroll on their own during the same period, and document a pattern of dynamic selection, where those who enroll in the no-subsidy condition are much more likely to immediately file large claims. In the Banerjee, Finkelstein, et al. (2021) case, because the subsidies attract healthier individuals who then pay premia in the post-subsidy period, they allow the government to cover substantially more individuals at the same total cost. The results suggest that these types of temporary subsidies, while by no means a panacea, can help ameliorate adverse selection to some degree.

*Bundling.* A second potential approach to reduce adverse selection is bundling: by tying the purchase of health insurance to that of another product, one can potentially limit the ability of households to buy insurance only for the sick. Fischer, Frölich, and Landmann (2018) conduct a similar experiment to Asuming, Kim, and Sim (2021) and Banerjee, Finkelstein, et al. (2021) in Pakistan, working with an NGO, the National Rural Support Program of Pakistan. As in the previous two studies they also experimentally vary premia, and find that adverse selection is higher when premia are higher. In addition, however, they also study an additional experimental variation, investigating whether people can choose to enroll individuals, or if they choose to enroll, must

---

<sup>55</sup> Two other important papers worth noting. Thornton et al. (2010) show that 6-month subsidies led to more take-up of health insurance in Nicaragua, but that very few people stayed beyond the subsidy. Wagstaff et al. (2016) show that providing an information kit and a 50% subsidy increase take-up of health insurance in the Philippines, but that hassle costs were still a significant issue.

<sup>56</sup> The original study design and short-run results are described in Asuming (2013).

enroll the entire household. The theory is that the bundled household contract may reduce adverse selection, if health shocks are idiosyncratic across the household. They also examine a community bundled contract, where at least fifty percent of the community must sign up in order to activate the insurance. They find that 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 leveraging households or community organizations as an alternative in developing countries where many workplaces are informal and may be unwilling or unable to provide group health insurance.

The downside of bundling, however, is that if demand for insurance is low, it can reduce demand for the other, bundled product. Banerjee, Duflo, and Hornbeck (2018) examined demand for a bundled health insurance and microfinance product, and found that this backfired: a substantial fraction of microcredit clients were apparently willing to forgo renewing their microcredit credits just so that they did not have to purchase required bundled health insurance. As Banerjee, Duflo, and Hornbeck (2018) write, the design of the bundled product “while standard, seems to have unwittingly created a much larger problem in attempting to solve a small one. [The provider] was clearly successful in avoiding adverse selection in the take-up for their product. But what they did not initially forecast is that this was because no one seemed to demand insurance...” concluding that, “there cannot be adverse selection if there is no demand.”

*Information and Trust.* In addition to selection issues, lack of information (or mis-information) and limited trust in the system may be another reason for low take-up of insurance; these factors may be particularly relevant and important factors in low institutional capacity settings. Starting with information, the evidence is mixed: Giles et al. (2021) experimentally find that information on health insurance benefits and processes for enrolling in health insurance leads to an increase in enrollment by those that do not have a contract (i.e. those who likely did not have health insurance). Similarly, Asuming (2013) also finds impacts of similar information on enrollment in public health insurance in Ghana. In contrast, Banerjee, Finkelstein, et al. (2021) does not find impacts of information experiments designed to explain the benefits of public health insurance to citizens in Indonesia. Moreover, Dercon, Gunning, and Zeitlin (2019) explore a composite health insurance product for tea farmers in Kenya, and find little impact on financial training of the insurance product (beyond the basic marketing of insurance) on take-up.

In addition to studying information provision, Dercon, Gunning, and Zeitlin (2019) also examines issues of trust. They measure trust using a modified BDM’s trust game at baseline, and then offer the composite health insurance to individuals at either no discount, 10% discount or 20% discount. They find that low trust (as well as risk aversion) is negatively associated with insurance demand, and that the purchase decisions of individuals with low trust are significantly more sensitive to price. They then use their estimates to calculate the welfare loss relative to counterfactual products that would be perceived as more trust-worthy and credible, showing large welfare losses particularly among the poor due to a lack of trust of the products.

In sum, despite the growing literature on provision of health insurance in low- and middle-income countries, doing so still remains a challenge. The combination of lack of demand for private insurance (even before one deals with adverse selection issues), strong informal sector (so one

cannot compel insurance through employers), and the relative challenges of enforcing a mandate to purchase insurance, means that it is challenging to ensure widespread insurance coverage. Many countries, faced with this, run a public-sector system with low prices, with those who want better coverage opting out at their own expense. However, doing so still leads to substantial uninsured risk, and working out how to improve take-up of more comprehensive insurance schemes, and hence moving towards more comprehensive coverage, remains an important direction for future research.

#### **4.3.2: Property insurance and insurance against climate-related shocks**

Risk from property damages and loss poses a challenge everywhere, whether it be from flooding, earthquake, fire, crime, or other perils. In wealthy countries, private insurance tends to cover many of these idiosyncratic risks: in the US, for example, over 90 percent of homeowners have property insurance (Insurance Information Institute 2020) though for certain types of perils where there are aggregate shocks—floods, for example—coverage remains low. Yet the fraction of people in developing countries who insure these risks is tiny, and these shocks can have important effects (Anttila-Hughes and Hsiang 2013). As global climate change increases the frequency of these shocks, these issues are likely to become more important.

*Idiosyncratic shocks.* It is important to distinguish between idiosyncratic perils (e.g., fire) and aggregate perils (e.g. floods, earthquakes). For idiosyncratic perils, private insurance markets should work well in principle. However, in developing country contexts, where information is worse, 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<sup>57</sup>. Moreover, like other forms of insurance, low demand often prevails due to trust in the institutions to pay out (for example, using a choice experiment, Reynaud, Nguyen, and Aubert (2018) show that demand for flood insurance is low due to lack of trust). Understanding the limits of this market seems an important question for future research.

*Aggregate shocks.* For aggregate shocks, such as floods, earthquakes, hurricanes, and so on, the issues are a bit different. Perhaps most important, governments ex-post tend to intervene in these crises and provide some amounts of emergency relief to households, though the degree to which they do so is imperfect.<sup>58</sup> For example, Gignoux and Menéndez (2016) study earthquakes in Indonesia, and find that affected households have negative effects from the shocks, but that these dissipate after 2-5 years. They find substantial external (e.g. government) aid flows following earthquakes, which may help mitigate the shocks’ effects.

However, given that there is some chance that governments may come in, ex-post, to provide assistance, households may be reluctant to purchase insurance themselves. That is, the inability of governments to commit *not* to provide assistance in the event of a negative aggregate shock

---

<sup>57</sup> With the exception of TransUnion South Africa’s similar *Insurance Claims and Policy* database.

<sup>58</sup> Aggregate shocks are harder for private insurers to cover, though in principle this can be feasible using reinsurance or other related schemes. In the Caribbean, for example, many small countries face risks from hurricanes, but they do not know ex-ante which country will be hit. They have therefore developed (with Assistance from the World Bank and multi-lateral aid agencies) the Caribbean Catastrophic Risk Insurance Facility to provide mutual insurance to help mitigate these macro shocks.

means that the private market is unlikely working. If governments anticipate they will need to bail out households ex-post, governments may be interested in formalizing this commitment so that households can rely on it, so that the coverage they provide can be more effective—and perhaps so that governments can recoup some of the costs in the form of mandatory premiums. One of the few examples of such a scheme we know of in middle-income countries is the Turkish compulsory<sup>59</sup> earthquake insurance scheme, established in 2000. Although coverage increased gradually in the 20 years since the program was rolled out, currently there are almost 11 million covered policies and 53.2% of housing units insured (Natural Disaster Insurance Institution 2022).

There is some limited evidence that these types of government-run insurance schemes can make an important difference. del Valle, de Janvry, and Sadoulet (2020) study Mexico's indexed disaster relief scheme, known as Fonden, which (according to the authors) is the only indexed disaster relief fund in the world. Fonden provides transfers to a municipality if rainfall exceeds a pre-specified threshold. Using a regression discontinuity design based on the rainfall cutoff, they show that the insurance payments lead to substantially more economic activity, as measured by night-lights, but cannot trace out effects on individual households.

#### **4.3.3: Pensions and Annuities**


Old age is often associated with poverty. Health typically falters as people age, so the elderly are less able to work, and this is particularly true for lower-income people for whom work often entails substantial physical labor. Savings may be difficult, particularly for those outside the formal sector. And even for those who do save, ageing entails risk: in particular, the risk that one will outlive one's savings. Traditionally, these challenges were borne by family members (e.g., children taking care of their parents), but these networks may be incomplete, and may break down with urbanization and the transition to the more modern economy.

Accordingly, many governments around the world have public pension programs, with a mix of non-contributory and contributory systems. Appendix Table 2 describes the types of pension systems that countries have in place: 181 have contributory pension systems, 102 have non-contributory systems and 96 have both systems in place. Governments often develop these systems as countries emerge into middle-income hood, perhaps coincident with the rise of the more modern economy, increased life expectancies, and the breakdown of the traditional family-based model of elderly care. This is shown in Figure 6, which shows a positive relationship between GDP per capita and likelihood of having a non-contributory pension. That said, Figure 6 also reveals that conditional on a level of real GDP per capita, these programs are becoming more prominent in middle-income and even low-income countries in recent decades.

---

<sup>59</sup> Proof of earthquake insurance are required when individuals take out home loans, for title deed transactions and when changing electricity and water contracts, and this is primarily how the mandatory insurance scheme is enforced (Republic of Turkey 2012).

**Figure 6: Percentage of countries that enacted non-contributory pension systems by income classification, over time**



*Notes:* This figure plots the percentage of countries that enacted non-contributory pension systems by each year considered (1960, 1980, 2000 and 2020) against quintile of GDP per capita (in constant 2015 \$US). GDP per capita quintiles are calculated as an absolute scale across the time periods considered. 85 countries are in the 1960 sample, 129 countries are in the 1980 sample and 179 countries are in the 2000 and 2020 samples. *Source:* GDP per capita data is from the World Bank (2021a).

These programs typically have three, related goals, though the emphasis on each varies from country to country: a) to help (or enforce) individuals to save for their old-age, either individually or through a taxes-and-transfer scheme, b) to provide annuities that insure against the risk of living too long 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 their old age.

A substantial literature has studied the effects of pension transfers on recipient households, primarily from the perspective of labor supply and consumption choices. There is less of a literature on the *design* of these programs in middle-income countries, however. We review these each in turn.

*Effects on consumption and well-being.* One strain of research is focused on understanding the impact of pensions on the well-being of beneficiaries (e.g. consumption, food security, health). For example, Huang and Zhang (2021) use a difference-in-difference design to examine the rollout of China’s New Rural Pension Scheme (NRPS). They find that the rollout of the scheme increased eligible households’ incomes, food expenditures. It also improves health of the elderly

and reduces mortality. Similarly, Galiani, Gertler, and Bando (2016) evaluate the impact of Mexico's *Adultos Mayores* Program (Older Adults Program), a non-contributory cash transfer provided to rural adults aged 70 years and older. Using geographic variation and age-related cutoffs, they find increases household consumption, and reductions in depression.

Consumption effects need not be limited to pensioners themselves, as consumption resources will be 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.<sup>60</sup> Edmonds (2006), studying the same program, 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. In both cases, the income shocks were anticipated, suggesting that households were liquidity constrained. de Carvalho Filho (2012) studies an unexpected pension reform in Brazil, and also shows that pension eligibility leads to an increase in school enrollment for girls. Combined, these papers suggest important inter-generational spillovers within the household from pension receipt.

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

*Labor Supply Effects.* A second set of papers tries to understand 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 quirks in 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 reductions in labor supply associated with pension receipt. de Carvalho Filho (2008), for example, studies the reform to both the eligibility policy and benefit level in the Brazilian Social Security System discussed above, and shows that these reforms 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's (2016) study of Mexico's *Adultos Mayores* Program (Older Adults Program) also finds that paid work declines, but that those who stop doing paid work switch to family businesses. And, Huang and Zhang (2021)'s study of the rollout of the China's New Rural Pension Scheme (NRPS), find that labor supply decreased for rural people older than 60 (i.e. the eligible), particularly for farm-work.<sup>61</sup>

---

<sup>60</sup> Case and Deaton (1998) provide a rich description of the South Africa's non-contributory pension program, as well as a nice analysis of the targeting properties of the program.

<sup>61</sup> The income effects from a transfer are not limited to the program recipient; they may be shared among the recipients immediate or extended family, as are other consumption decisions. A series of papers around South Africa's pension program examined spillovers effects in the labor supply of other household members. Examining the age-based discontinuity in eligibility, 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 women. Posel, Fairburn, and Lund (2006) and Ardington, Case, and Hosegood (2009) argue, however, that the pensions also drive migration among the prime age workers, i.e. the pensions help relieve credit constraints to migrate. Including the migrants in the analysis, Ardington, Case, and Hosegood (2009) they find small positive increases in work among prime-age adults.

*Program design.* There is comparatively less work, however, on program design questions in middle-income countries, but we highlight several areas where work has been done that suggests some ways in which middle-income contexts may be somewhat different from higher-income contexts.

One important question comes is about contributions: do firms actually report the correct social security contributions, and pay tax accordingly, or do they under-report? Kumler, Verhoogen, and Frías (2020) show that the design of the pension system can change workers' incentives to accurately monitor firms. In particular, they show that when the pension system switched to a system in which pensions became largely a function of contributions, payments for younger workers (who had more of an incentive to make sure their wages were reported accurately) increased. This suggests that there are important linkages between the design of pension benefits and workers ex-ante incentives to make sure the taxes to fund these systems are remitted accurately on their behalf.

A second question is about investment choices. Some contributory systems allow individuals choice in 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—the loads (i.e. the fee charged by investment manager on new inflows into the accounts) averaged 23% of the amount deposited into the account, as well as annual fees that averaged 63 basis points on annual balances. 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 is that workers are not particularly price sensitive. As a result, instead of competing on fees, firms advertise heavily, and other non-price attributes seem to dominate workers decisions of who should manage their investments, leading to high prices in equilibrium. This suggests that it can be challenging to structure market incentives so that competition can lead to low prices in these markets, and 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. This can be particularly a problem in middle-income countries where there is a large informal sector, which does not pay in to the contributory system. A country may therefore want to add a substantial component to cover the informal systems, i.e., a minimum pension floor, but the challenge is to do this without discouraging savings. Attanasio, Meghir, and Otero (2011), 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 that challenges with paperwork and dealing with the bureaucracy hinder 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 we suggest there is substantially more to be done in this area. 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 with funding sustainability of the system? And as these systems grow over time, what will the implications be for the host of other economic decisions – marriage, investments in children, savings – that are related to how people plan for their old age? Finally, does the large and fluid household system in many low- and middle-income countries change how one thinks about annuitization, and in particular, couples-based annuity systems?

#### **4.3.4. Funeral Insurance**

Funerals are often a large, not always foreseen cost that can devastate households, particularly at 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). For example, in the South Africa case, for example, about 28 percent of their sample 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) discusses, many of these informal insurance or burial clubs have sets of rules and institutions to promote fairness. Nonetheless, a question is whether certain types of households can be excluded through networks, etc. and whether there is a role for government provided funeral insurance to complement many of these informal structures.

#### **4.4: Summing up, and Broader Implications for Social Protection System Design**

Life in developing countries is risky. Even in developed countries, adverse selection and moral hazard mean that there is an important justification for government involvement to help households smooth this risk. But the additional challenges highlighted above—low insurance demand, challenges of trust, liquidity constraints, challenges of state verification, and informality—mean that demand for even if individuals were offered contracts priced fairly for the entire population, take-up may be low. Understanding how to solve these issues—and to what degree appropriate government schemes may be effective—remains an important direction for future work.

Moreover, the fact there are 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 shown in Figure 5 is that for a household in the 'good' equilibrium, an uninsured negative shock can push them across the threshold and 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 types of programs discussed in Section 3. Indeed, Ikegami et al. (2018) and Janzen, Carter, and Ikegami (2021) argue that, for this reason, in a dynamic sense governments may be better off diverting some of their standard targeted social

protection funds to providing additional social insurance. Even to the extent these programs do not solve the entire financial shock, the conditionalities in some of these programs can direct households to smooth them in ways that are less socially costly, i.e. by encouraging them to keep children in school (de Janvry et al. 2006). This can have similar inter-generational protective 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 also provide some insurance. That is, a challenge with many of the targeting approaches discussed in Section 2 is that the targeting list is updated infrequently, and when targeting it is done, it often focuses the on the permanent component of income (e.g. assets). This makes it not responsive to shocks.

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

At the aggregate level, governments can use the types of programs discussed in Section 3 to respond to aggregate shocks by changing the eligibility thresholds, or by increasing the amounts of transfers. For example, the government of Kenya explicitly relaxes eligibility rules for its food security program in time 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.

During the COVID-19 crisis, when this was an extreme issue (because so many people experienced extreme shocks), 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 the eligibility rules and expanded the amount of transfers in its food voucher program, as well as created a new community-targeted cash transfer program in all rural villages. Pakistan, for example, built on its flagship cash-transfer program for women, expanding eligibility and using cell-phone metadata to help determine eligibility (Lone et al. 2021). These examples were forged rapidly in a crisis, presumably, with advance planning, governments can do much more to make their existing redistribution programs also function to smooth shocks. We regard making these programs more responsive as an important area for future research.

## **5. The Politics of Social Protection Systems**

To run a social protection system or program, one cannot abstract away from the politics. The politics affects different aspects of how programs run, as voters make decisions about the level of redistribution they want and the form of it. Differences in who has political power and access may further determine how programs are designed and who ultimately benefits from them. And, as with any government program, there are interesting dynamics on how politicians think about these programs—do they design platforms on social protection design around addressing voter views

and needs? Or, do incumbents manipulate the programs by, for example, changing programs to shore up support with certain groups prior to elections?

While we cannot comprehensively review the entire politics literature on this topic here, we highlight a few of the issues below. We refer the interested reader to the review by Golden and Min (2013) for a discussion of related issues from the political science perspective.

### 5.1: Voters

The existence or receipt of social protection programs may affect voter behavior—either positively or negatively. For example, some voters may reward parties or politicians that introduce or improve these programs due to a stated preference towards greater redistribution—regardless of whether they receive benefits or not. Moreover, voters who receive these programs may reward the parties that implement them either because they are happy with the services and help that they are receiving from the government, or because they are dependent on the help and would want the programs to continue. But others may vote against those that implement these programs, either due to an inherent preference against redistribution, because they are not benefiting directly from these programs, or because they believe that the quality of the programs is poor.

A number of studies have looked at whether the introduction or expansions of social protection can induce political participation and/or change voting patterns. This is empirically hard to disentangle, as voting may induce the introduction of social protection as well. A number of different empirical strategies have been used to understand these issues.<sup>62</sup>

Several studies use regression discontinuity designs that compare those who are just above the poverty eligibility cutoff with those just below to understand the impact of receiving a transfer versus missing out. These papers have typically concluded that *directly receiving benefits* leads to increased political engagement and political support for those who designed or implemented the program. One such study is Manacorda, Miguel, and Vigorito (2011), which studies PANES, a large targeted temporary cash program in Uruguay. Using survey data on voting outcomes, the authors find that beneficiaries were more likely to favor the government that implemented the program. This was true even after the program itself ended, suggesting that it was less about people voting based on their current receipt of benefits but perhaps due to a change in belief about the party's beliefs on redistribution.

A second study employing these methods was Pop-Eleches and Pop-Eleches (2012) that studied a \$200 coupon to poor families for the purchase of a computer. They also found that beneficiaries were more likely to support the incumbent government coalition, driven by both high mobilization and also party-switching. But, interestingly, the higher trust was only to the

---

<sup>62</sup> One conceptual issue which plagues all of these studies is that there is a tension between empirical identification and what one would expect from rational voters. In order to generate causal identification of the impact of a policy on electoral outcomes, one needs random or quasi-random allocation of programs to some areas or people and not to others. But, if voters understood this, they would know that the only reason that a treatment area or person received the program and a control did not was due to random chance; that is, it was *not* due to any action of the politician. So, a rational voter, even one who would reward politicians who help them, should not reward a politician for having been randomly allocated to a treatment group.

local government officials who administered the program, but not the central government that designed and funded the program. A final example is from Colombia's CCT program, *Familias en Acción*. Conover et al. (2020) explored discontinuities in program eligibility and variation in program enrollment across voting booths and find that the program increased the beneficiaries' probability to register to vote, especially for women, who were the direct recipients of the program.

A series of other papers have taken advantage of experimental variation in the roll-out of transfers programs to look at the effect of the roll-out on political outcomes of everyone in an area, *regardless of whether they received the program*. The results are mixed. Labonne (2013) shows that a CCT program in the Philippines also led to increased vote share for the incumbent, but this effect was only evident in municipalities where there were high levels of political competition. Blattman, Emeriau, and Fiala (2018), however, examine a randomized allocation of grants to youth to fund entrepreneurship activities in Uganda, and find no effects on support for the ruling party.<sup>63</sup> Likewise, Imai, King, and Velasco Rivera (2020) find no electoral impact of a large-scale randomized trial in Mexico which randomized health insurance to selected areas (discussed above).

Brollo, Kaufmann, and La Ferrara (2020) show that beneficiaries may react to specific program features, not only the program as a whole. Using random variation in the timing of when beneficiaries learned about penalties for noncompliance with Bolsa Família's conditionalities around the 2008 municipal elections, they find a lower vote share for candidates aligned with the president in areas where more beneficiaries received penalties shortly before (as opposed to shortly after) the elections.

## 5.2: Politicians

A key question is how politicians develop social protection programs and policies based on political incentives. The best case, but overly simplistic, scenario is that voters have preferences over redistribution and the design of such policies and programs, they make their voices heard through activism and voting and politicians respond by providing the types of programs that citizens need. However, there are many challenges here—as citizen voices are not always aggregated perfectly through the voting booth, with those who may be most vulnerable often excluded from the systems. To gain support with particular voters or groups, politicians may also change spending patterns or manipulate rules or programs to confer benefits to particular groups.

A number of papers aim to understand whether politicians strategically time spending around elections, which could have implications for general macroeconomic conditions (e.g. too much spending in good times, and thus limitations in available budget to increase spending in recessions), following on the work of Nordhaus (1975). For example, Khemani (2004) found

---

<sup>63</sup> There is also a literature examining the electoral consequences of Progresa in Mexico. The challenge in doing so, as pointed out by Green (2006) and discussed in our discussion of Progresa above, is that Progresa treated the control areas too, and so the randomized trial compares areas that received Progresa for ~30 months prior to the election to areas that received it 6 months prior. Imai, King, and Velasco Rivera (2020) show that once data errors in the original De La O (2013) study of the electoral consequences of the Progresa RCT are corrected, there is no effect of Progresa on electoral vote share. Imai, King, and Velasco Rivera (2020) also use an alternative, discontinuity-based design and again find no impact of Progresa on vote shares.

that, in India, public investment increased more before scheduled investments, but then contracted in other times to keep the net balance unchanged. Composition of spending changed too, as Khemani (2004) found that resources shifted to narrow interest groups (e.g. tax breaks provided to small groups of producers) rather than broad-based consumption spending. Drazen and Eslava (2010) also found, for Colombia, that the composition of spending changed before elections, particularly around targeted expenditures.

In addition to changing spending patterns, politicians can also adjust the rules or implementation of existing programs to target particular groups and voters. For example, Camacho and Conover (2011), discussed above, finds evidence of manipulation of the targeting rules before elections, while Brollo, Kaufmann, and La Ferrara (2020), also discussed above, finds that enforcement of CCT rules become more lax before elections in municipalities where the incumbent associated with the program could run for election.

Given the political context, can improving representation and voice improve incentives for politicians? Can it shift policies and programs towards the previously unrepresented groups? A number of papers implies that it can. For example, Pande (2003) and Chattopadhyay and Duflo (2004) find that improving representation of minorities and women in India through political reservations led to spending allocations that better mirrored citizens' preferences. Similarly, Fujiwara (2015) found that increasing the enfranchisement of less educated citizens through electronic voting led to increased spending towards health care, which benefits low-income populations.

Note, however, that spending decisions—even if it aligns well with voters—may come at a cost in terms of other human capital investments. For example, Bursztyn (2016) found that governments invested less in public education because lower-income decisive voters preferred them to allocate resources mostly toward redistributive programs, such as cash transfers. This could potentially be welfare enhancing if, for example, public systems are poor and households decide how to invest in their child's education through private systems that improve education. But, it could also come at a cost if parents do not fully take into account the full benefits of education for their children and underinvest in schooling, or if other challenges—pressure from other family members for funds, other immediate spending needs—also lead to an underinvestment in education. It also suggests that social protection spending needs to be examined through the overall budget lens, and not just through individual components in making decisions on overall human capital investment needs.

## **6. 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, and so countries that were recently quite poor countries are now middle income, and with substantially larger tax bases are able to fund more sophisticated systems. But, 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—such as the large informal sector, and the deep absolute poverty level of the poor in these economies—means that the way these programs are designed differs in fundamental ways from the way 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 the income data is third-party reported through the tax system. With the large informal sector, this isn't 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.

A second important difference comes in the design of programs to assist the poor. Many low- and middle-income countries have in recent years dramatically expanded unconditional cash transfers, which are similar to the form income support often takes in richer countries and which remain a benchmark. But many other programs are common as well. For example, given the relatively low levels of human capital investment, conditional cash transfers, which tie cash payments to making sure 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, and have been shown to have substantial effects on the incentivized behaviors. Given the low levels of absolute poverty, there is also a focus on finding program designs that can help break poverty traps, whether within a single generation or across generations. The substantial informal sector also has implications for design of social insurance schemes. For example, unemployment insurance cannot be easily 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 numbers 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 it increase informality rates. Many advances have been made, but understanding how to design effective social insurance schemes in this context remains an important area for additional work.

The three pieces we discuss—identification of beneficiaries, design of programs to help the poor, and design of social insurance—operate together to create a unified social protection system. The increasing presence of these systems in low- and middle-income countries 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." In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1563–72. KDD '14. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/2623330.2623335>.
- Acemoglu, Daron. 2006. "A Simple Model of Inefficient Institutions." *The Scandinavian Journal of Economics* 108 (4): 515–46. <https://doi.org/10.1111/j.1467-9442.2006.00460.x>.
- Acemoglu, Daron, Tristan Reed, and James A. Robinson. 2014. "Chiefs: Economic Development and Elite Control of Civil Society in Sierra Leone." *Journal of Political Economy* 122 (2): 319–68. <https://doi.org/10.1086/674988>.
- Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E. Blumenstock. 2022. "Machine Learning and Phone Data Can Improve Targeting of Humanitarian Aid." *Nature* 603 (7903): 864–70. <https://doi.org/10.1038/s41586-022-04484-9>.
- Aizer, Anna, Shari Eli, Adriana Lleras-Muney, and Keyoung Lee. 2020. "Do Youth Employment Programs Work? Evidence from the New Deal." Working Paper 27103. NBER Working Papers. National Bureau of Economic Research. <https://doi.org/10.3386/w27103>.
- Akerlof, George A. 1978. "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning." *The American Economic Review* 68 (1): 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." Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-6340>.
- . 2016. "Evidence from a Randomized Evaluation of the Household Welfare Impacts of Conditional and Unconditional Cash Transfers Given to Mothers or Fathers." Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-7730>.
- Alatas, Vivi. 2011. "Program Keluarga Harapan: Main Findings from the Impact Evaluation of Indonesia's Pilot Household Conditional Cash Transfer Program." 72506. World Bank Office Jakarta. <http://documents1.worldbank.org/curated/en/589171468266179965/pdf/725060WP00PUBL01uation0Report0FINAL.pdf>.
- Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken. 2016. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *American Economic Review* 106 (7): 1663–1704. <https://doi.org/10.1257/aer.20140705>.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi. 2016. "Self-Targeting: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 124 (2): 371–427.
- . 2019. "Does Elite Capture Matter? Local Elites and Targeted Welfare Programs in Indonesia." *AEA Papers and Proceedings* 109 (May): 334–39. <https://doi.org/10.1257/pandp.20191047>.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias. 2012. "Targeting the Poor: Evidence from a Field Experiment in Indonesia." *American Economic Review* 102 (4): 1206–40. <https://doi.org/10.1257/aer.102.4.1206>.
- Albarran, Pedro, and Orazio P. Attanasio. 2003. "Limited Commitment and Crowding out of Private Transfers: Evidence from a Randomised Experiment." *The Economic Journal* 113 (486): C77–85.

- Alderman, Harold. 2002. “Do Local Officials Know Something We Don’t? Decentralization of Targeted Transfers in Albania.” *Journal of Public Economics* 83 (3): 375–404. [https://doi.org/10.1016/S0047-2727\(00\)00145-6](https://doi.org/10.1016/S0047-2727(00)00145-6).
- Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2020. “Tackling Youth Unemployment: Evidence From a Labor Market Experiment in Uganda.” *Econometrica* 88 (6): 2369–2414. <https://doi.org/10.3982/ECTA15959>.
- Ali, Williams, Awudu Abdulai, and Ashok K. Mishra. 2020. “Recent Advances in the Analyses of Demand for Agricultural Insurance in Developing and Emerging Countries.” SSRN Scholarly Paper 3716725. Rochester, NY: Social Science Research Network. <https://doi.org/10.1146/annurev-resource-110119-025306>.
- Allingham, Michael G., and Agnar Sandmo. 1972. “Income Tax Evasion: A Theoretical Analysis.” *Journal of Public Economics* 1 (3): 323–38. [https://doi.org/10.1016/0047-2727\(72\)90010-2](https://doi.org/10.1016/0047-2727(72)90010-2).
- Amin, Mohammad, and Asif Islam. 2019. “Paid Maternity Leave and Female Employment: Evidence Using Firm-Level Survey Data for Developing Countries.” Working Paper 8715. World Bank Policy Research Working Paper. Washington, D.C.: World Bank. <https://openknowledge.worldbank.org/handle/10986/31183>.
- Angelucci, Manuela, and Giacomo De Giorgi. 2009. “Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles’ Consumption?” *American Economic Review* 99 (1): 486–508. <https://doi.org/10.1257/aer.99.1.486>.
- Anttila-Hughes, Jesse, and Solomon Hsiang. 2013. “Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster.” [https://web.worldbank.org/archive/website01589/WEB/IMAGES/JESSE\\_-2.PDF](https://web.worldbank.org/archive/website01589/WEB/IMAGES/JESSE_-2.PDF).
- Araújo, María C., Mariano Bosch, and Norbert Schady. 2017. “Can Cash Transfers Help Households Escape an Inter-Generational Poverty Trap?” Inter-American Development Bank. <https://doi.org/10.18235/0000618>.
- 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 (1): 22–48. <https://doi.org/10.1257/app.1.1.22>.
- Areias, Ana, and Matthew Wai-Poi. 2022. “Big Data, Machine Learning and Prediction in the Targeting of Social Protection Programs.” In *Revisiting Targeting in Social Assistance: A New Look at Old Dilemmas*, edited by Margaret Grosh, Phillippe Leite, Matthew Wai-Poi, and Emil Tesliuc. Washington, DC: World Bank.
- Arnott, Richard, and Joseph E. Stiglitz. 1991. “Moral Hazard and Nonmarket Institutions: Dysfunctional Crowding Out of Peer Monitoring?” *The American Economic Review* 81 (1): 179–90.
- Asuming, Patrick Opoku. 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. New York: Columbia University. [http://www.columbia.edu/~poa2102/jmp\\_asing\\_2012.pdf](http://www.columbia.edu/~poa2102/jmp_asing_2012.pdf).
- Asuming, Patrick Opoku, Hyuncheol Bryant Kim, and Armand Sim. 2021. “Selection and Behavioral Responses of Health Insurance Subsidies in the Long Run: Evidence from a Field Experiment in Ghana.” 05/21. CDES Working Paper Series. Centre for Development Economics and Sustainability at Monash University. [https://www.monash.edu/\\_\\_data/assets/pdf\\_file/0004/2782903/WP2021n05\\_V2.pdf](https://www.monash.edu/__data/assets/pdf_file/0004/2782903/WP2021n05_V2.pdf).

- Aterido, Reyes, Mary Hallward-Driemeier, and Carmen Pagés. 2011. “Big Constraints to Small Firms’ Growth? Business Environment and Employment Growth across Firms.” *Economic Development and Cultural Change* 59 (3): 609–47. <https://doi.org/10.1086/658349>.
- Athey, Susan, and Stefan Wager. 2017. “Efficient Policy Learning.” Working Paper No. 17-031. Stanford Institute for Economic Policy Research Working Paper. [https://drive.google.com/file/d/1RzDpRXeIgi\\_GEtIWqmggwmWEPWUVb\\_v4/view?usp=drivesdk&usp=embed\\_facebook](https://drive.google.com/file/d/1RzDpRXeIgi_GEtIWqmggwmWEPWUVb_v4/view?usp=drivesdk&usp=embed_facebook).
- Attanasio, Orazio, Luis Carlos Gómez, Patricia Heredia, and Marcos Vera-Hernández. 2005. “The Short-Term Impact of a Conditional Cash Subsidy on Child Health and Nutrition in Colombia.” *Familias 03*. Institute for Fiscal Studies Report Summary. Institute for Fiscal Studies. <https://ifs.org.uk/publications/3503>.
- Attanasio, Orazio, Costas Meghir, and Andres Otero. 2011. “Pensions, Work and Informality: The Impact of the 2008 Chilean Pension Reform.” Unpublished paper. <https://www.consejoprevisional.cl/documentos/articulos/pensiones-trabajo-informalidad-attanasio-meghir-otero.pdf>.
- Attanasio, Orazio, and José-Víctor Ríos-Rull. 2000. “Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?” *European Economic Review* 44 (7): 1225–58. [https://doi.org/10.1016/S0014-2921\(00\)00034-9](https://doi.org/10.1016/S0014-2921(00)00034-9).
- 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 (4): 1799–1848. <https://doi.org/10.1093/qje/qjaa015>.
- Autor, David H., Mark Duggan, Kyle Greenberg, and David S. Lyle. 2016. “The Impact of Disability Benefits on Labor Supply: Evidence from the VA’s Disability Compensation Program.” *American Economic Journal: Applied Economics* 8 (3): 31–68. <https://doi.org/10.1257/app.20150158>.
- Azam, Mehtabul. 2012. “The Impact of Indian Job Guarantee Scheme on Labor Market Outcomes: Evidence from a Natural Experiment.” Discussion Paper 6548. IZA Discussion Papers. Bonn: Institute of Labor Economics. <https://www.iza.org/publications/dp/6548/the-impact-of-indian-job-guarantee-scheme-on-labor-market-outcomes-evidence-from-a-natural-experiment>.
- Azuara, Oliver, and Ioana Marinescu. 2013. “Informality and the Expansion of Social Protection Programs: Evidence from Mexico.” *Journal of Health Economics* 32 (5): 938–50. <https://doi.org/10.1016/j.jhealeco.2013.07.004>.
- Baez, Javier E., Varun Kshirsagar, and Emmanuel Skoufias. 2020. “Adaptive Safety Nets for Rural Africa: Drought-Sensitive Targeting with Sparse Data.” Brief. Washington, DC: World Bank. <https://doi.org/10.1596/34301>.
- Baily, Martin Neil. 1978. “Some Aspects of Optimal Unemployment Insurance.” *Journal of Public Economics* 10 (3): 379–402. [https://doi.org/10.1016/0047-2727\(78\)90053-1](https://doi.org/10.1016/0047-2727(78)90053-1).
- 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): 1–43. <https://doi.org/10.1080/19439342.2014.890362>.
- Baird, Sarah, Craig McIntosh, and Berk Özler. 2011. “Cash or Condition? Evidence from a Cash Transfer Experiment.” *The Quarterly Journal of Economics* 126 (4): 1709–53. <https://doi.org/10.1093/qje/qjr032>.

- 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 (1): 22. <https://doi.org/10.1186/s40176-018-0131-9>.
- Baker, Judy L., and Margaret E. Grosh. 1994. “Poverty Reduction through Geographic Targeting: How Well Does It Work?” *World Development* 22 (7): 983–95. [https://doi.org/10.1016/0305-750X\(94\)90143-0](https://doi.org/10.1016/0305-750X(94)90143-0).
- 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. <https://doi.org/10.1093/qje/qjab045>.
- Ball, Ian. 2022. “Scoring Strategic Agents.” Working Paper. <http://arxiv.org/abs/1909.01888>.
- Bana, Sarah H., Kelly Bedard, and Maya Rossin-Slater. 2020. “The Impacts of Paid Family Leave Benefits: Regression Kink Evidence from California Administrative Data.” *Journal of Policy Analysis and Management* 39 (4): 888–929. <https://doi.org/10.1002/pam.22242>.
- Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. 2017a. “Labor Markets and Poverty in Village Economies.” *The Quarterly Journal of Economics* 132 (2): 811–70. <https://doi.org/10.1093/qje/qjx003>.
- . 2017b. “Labor Markets and Poverty in Village Economies.” *The Quarterly Journal of Economics* 132 (2): 811–70. <https://doi.org/10.1093/qje/qjx003>.
- Banerjee, Abhijit, and Esther Duflo. 2010. “Giving Credit Where It Is Due.” *Journal of Economic Perspectives* 24 (3): 61–80. <https://doi.org/10.1257/jep.24.3.61>.
- Banerjee, Abhijit, Esther Duflo, Raghavendra Chattopadhyay, and Jeremy Shapiro. 2016a. “The Long Term Impacts of a ‘Graduation’ Program: Evidence from West Bengal.” <https://economics.mit.edu/files/12031>.
- . 2016b. “The Long Term Impacts of a ‘Graduation’ Program: Evidence from West Bengal.” <https://economics.mit.edu/files/12031>.
- Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry. 2015a. “A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries.” *Science* 348 (6236): 1260799. <https://doi.org/10.1126/science.1260799>.
- . 2015b. “A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries.” *Science* 348 (6236): 1260799. <https://doi.org/10.1126/science.1260799>.
- 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 (340): 671–700. <https://doi.org/10.1111/ecca.12271>.
- Banerjee, Abhijit, Esther Duflo, Clément Imbert, Santhosh Mathew, and Rohini Pande. 2020. “E-Governance, Accountability, and Leakage in Public Programs: Experimental Evidence from a Financial Management Reform in India.” *American Economic Journal: Applied Economics* 12 (4): 39–72. <https://doi.org/10.1257/app.20180302>.
- Banerjee, Abhijit, Esther Duflo, and Garima Sharma. 2021a. “Long-Term Effects of the Targeting the Ultra Poor Program.” *American Economic Review: Insights* 3 (4): 471–86. <https://doi.org/10.1257/aeri.20200667>.
- . 2021b. “Long-Term Effects of the Targeting the Ultra Poor Program.” *American Economic Review: Insights* 3 (4): 471–86. <https://doi.org/10.1257/aeri.20200667>.
- Banerjee, Abhijit, Amy Finkelstein, Rema Hanna, Benjamin A. Olken, Arianna Ornaghi, and Sudarno Sumarto. 2021a. “The Challenges of Universal Health Insurance in Developing

- Countries: Experimental Evidence from Indonesia’s National Health Insurance.” *American Economic Review* 111 (9): 3035–63. <https://doi.org/10.1257/aer.20200523>.
- . 2021b. “The Challenges of Universal Health Insurance in Developing Countries: Experimental Evidence from Indonesia’s National Health Insurance.” *American Economic Review* 111 (9): 3035–63. <https://doi.org/10.1257/aer.20200523>.
- Banerjee, Abhijit, Rema Hanna, Gabriel E. Kreindler, and Benjamin A. Olken. 2017a. “Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs.” *The World Bank Research Observer* 32 (2): 155–84. <https://doi.org/10.1093/wbro/lkx002>.
- . 2017b. “Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs.” *The World Bank Research Observer* 32 (2): 155–84. <https://doi.org/10.1093/wbro/lkx002>.
- Banerjee, Abhijit, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto. 2018. “Tangible Information and Citizen Empowerment: Identification Cards and Food Subsidy Programs in Indonesia.” *Journal of Political Economy* 126 (2): 451–91. <https://doi.org/10.1086/696226>.
- . 2019. “Private Outsourcing and Competition: Subsidized Food Distribution in Indonesia.” *Journal of Political Economy* 127 (1): 101–37. <https://doi.org/10.1086/700734>.
- Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, Elan Satriawan, and Sudarno Sumarto. 2021. “Food vs. Food Stamps: Evidence from an At-Scale Experiment in Indonesia.” Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w28641>.
- Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, and Sudarno Sumarto. 2020. “The (Lack of) Distortionary Effects of Proxy-Means Tests: Results from a Nationwide Experiment in Indonesia.” *Journal of Public Economics Plus* 1 (January): 100001. <https://doi.org/10.1016/j.pubecp.2020.100001>.
- 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 (March): 102781. <https://doi.org/10.1016/j.jdeveco.2021.102781>.
- Banerjee, Abhijit, and Sendhil Mullainathan. 2008. “Limited Attention and Income Distribution.” *American Economic Review* 98 (2): 489–93. <https://doi.org/10.1257/aer.98.2.489>.
- . 2010. “The Shape of Temptation: Implications for the Economic Lives of the Poor.” Working Paper 15973. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w15973>.
- Banerjee, Abhijit, and Andrew Newman. 1998. “Information, the Dual Economy, and Development.” *The Review of Economic Studies* 65 (4): 631–53.
- Banerjee, Abhijit, and Andrew F. Newman. 1993. “Occupational Choice and the Process of Development.” *Journal of Political Economy* 101 (2): 274–98.
- Banerjee, Abhijit, Paul Niehaus, and Tavneet Suri. 2019. “Universal Basic Income in the Developing World.” *Annual Review of Economics* 11 (1): 959–83. <https://doi.org/10.1146/annurev-economics-080218-030229>.
- Bardhan, Pranab K., and Dilip Mookherjee. 2000. “Capture and Governance at Local and National Levels.” *American Economic Review* 90 (2): 135–39. <https://doi.org/10.1257/aer.90.2.135>.

- 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): 467–71. <https://doi.org/10.1257/aer.103.3.467>.
- . 2018a. "Are Conditional Cash Transfers Fulfilling Their Promise? Schooling, Learning, and Earnings after 10 Years." Discussion Paper DP11937. CEPR Discussion Paper. London: Center for Economic Policy Research. <https://papers.ssrn.com/abstract=2941523>.
- . 2018b. "Experimental Evidence of Exposure to a Conditional Cash Transfer During Early Teenage Years: Young Women's Fertility and Labor Market Outcomes." CEPR Discussion Paper DP13165. Discussion Paper Series. London: Center for Economic Policy Research. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3247237](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3247237).
- 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 (3): 611–21. <https://doi.org/10.1016/j.jhealeco.2008.12.010>.
- 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 (2): 167–95. <https://doi.org/10.1257/app.3.2.167>.
- 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 (3): 54–91. <https://doi.org/10.1257/app.20170008>.
- Barros, Ricardo Paes de, 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." Working Paper. 2011 World Bank Economists' Forum. Washington, D.C.: World Bank.
- Bartel, Ann P., Maya Rossin-Slater, Christopher J. Ruhm, Meredith Slopen, and Jane Waldfogel. 2021. "The Impact of Paid Family Leave on Employers: Evidence from New York." Working Paper 28672. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w28672>.
- Bastagli, Francesca, Jessica Hagen-Zanker, Luke Harman, Valentina Barca, Georgina Sturge, Tanja Schmidt, and Luca Pellerano. 2016. *Cash Transfers: What Does the Evidence Say? A Rigorous Review of Programme Impact and of the Role of Design and Implementation Features*. <https://doi.org/10.13140/RG.2.2.29336.39687>.
- Basu, Kaushik. 1996. "Relief Programs: When It May Be Better to Give Food Instead of Cash." *World Development* 24 (1): 91–96. [https://doi.org/10.1016/0305-750X\(95\)00110-X](https://doi.org/10.1016/0305-750X(95)00110-X).
- 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 (May): 104047. <https://doi.org/10.1016/j.jpubeco.2019.07.006>.
- Bau, Natalie. 2021. "Can Policy Change Culture? Government Pension Plans and Traditional Kinship Practices." *American Economic Review* 111 (6): 1880–1917. <https://doi.org/10.1257/aer.20190098>.
- Beaman, Lori, Niall Keleher, Jeremy Magruder, and Carly Trachtman. 2021. "Urban Networks and Targeting: Evidence from Liberia." *AEA Papers and Proceedings* 111 (May): 572–76. <https://doi.org/10.1257/pandp.20211061>.

- Bedoya, Guadalupe, Aidan Coville, Johannes Haushofer, Mohammad Isaqzadeh, and Jeremy Shapiro. 2019a. “No Household Left Behind: Afghanistan Targeting the Ultra Poor Impact Evaluation.” Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-8877>.
- . 2019b. “No Household Left Behind: Afghanistan Targeting the Ultra Poor Impact Evaluation.” Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-8877>.
- 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 (1): 93–122.
- Benedetti, Fiorella, Pablo Ibararán, 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 (4): 759–93. <https://doi.org/10.1086/686583>.
- Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen. 2015. “Turning a Shove into a Nudge? A ‘Labeled Cash Transfer’ for Education.” *American Economic Journal: Economic Policy* 7 (3): 86–125. <https://doi.org/10.1257/pol.20130225>.
- Benjamin, Dwayne. 1992. “Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models.” *Econometrica* 60 (2): 287–322. <https://doi.org/10.2307/2951598>.
- Berg, Erlend. 2018. “Funeral Insurance: An Inter-Generational Commitment Device?” *Journal of African Economies* 27 (3): 321–46. <https://doi.org/10.1093/jae/ejx037>.
- 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 (March): 239–54. <https://doi.org/10.1016/j.worlddev.2017.10.027>.
- Bergolo, Marcelo, and Guillermo Cruces. 2014. “Work and Tax Evasion Incentive Effects of Social Insurance Programs: Evidence from an Employment-Based Benefit Extension.” 8198. *IZA Discussion Papers*. IZA Discussion Papers. Institute of Labor Economics (IZA). <https://ideas.repec.org/p/iza/izadps/dp8198.html>.
- 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.” Working Paper 29086. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w29086>.
- Berry, James, Greg Fischer, and Raymond Guiteras. 2020. “Eliciting and Utilizing Willingness to Pay: Evidence from Field Trials in Northern Ghana.” *Journal of Political Economy* 128 (4): 1436–73. <https://doi.org/10.1086/705374>.
- Bertrand, Marianne, Bruno Crepon, Alicia Marguerie, and Patrick Premand. 2021. “Do Workfare Programs Live Up to Their Promises? Experimental Evidence from Côte d’Ivoire.” Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-9611>.
- Bertrand, Marianne, Sendhil Mullainathan, and Douglas Miller. 2003. “Public Policy and Extended Families.” *World Bank Economic Review* 17 (1): 27–49. <https://doi.org/10.1093/wber/lhg014>.
- Besley, Timothy, and Stephen Coate. 1992. “Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs.” *The American Economic Review* 82 (1): 249–61.

- Besley, Timothy, and Alec R. Levenson. 1996. “The Role of Informal Finance in Household Capital Accumulation: Evidence from Taiwan.” *The Economic Journal* 106 (434): 39–59. <https://doi.org/10.2307/2234930>.
- Bettendorf, Leon J. H., Egbert L. W. Jongen, and Paul Muller. 2015. “Childcare Subsidies and Labour Supply — Evidence from a Large Dutch Reform.” *Labour Economics* 36 (October): 112–23. <https://doi.org/10.1016/j.labeco.2015.03.007>.
- Bhalotra, Sonia, Diogo G. C. Britto, Paolo Pinotti, and Breno Sampaio. 2021. “Job Displacement, Unemployment Benefits and Domestic Violence.” Working Paper 9186. CESifo Working Paper. Rochester, NY: Social Science Research Network. <https://doi.org/10.2139/ssrn.3887250>.
- Bhattacharya, Debopam, and Pascaline Dupas. 2012. “Inferring Welfare Maximizing Treatment Assignment under Budget Constraints.” *Journal of Econometrics* 167 (1): 168–96. <https://doi.org/10.1016/j.jeconom.2011.11.007>.
- 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 (14): 60–67.
- Björkegren, Daniel, Joshua E. Blumenstock, and Samsun Knight. 2020. “Manipulation-Proof Machine Learning.” *ArXiv:2004.03865 [Cs, Econ]*, April. <http://arxiv.org/abs/2004.03865>.
- Bjorvatn, Kjetil, Denise Ferris, Selim Gulesci, Arne Nasgowitz, Vincent Somville, and Lore Vandewalle. 2022. “Childcare, Labor Supply, and Business Development: Experimental Evidence from Uganda.” Discussion Paper DP17243. <https://portal.cepr.org/discussion-paper/19994>.
- Blattman, Christopher, Mathilde Emeriau, and Nathan Fiala. 2018. “Do Anti-Poverty Programs Sway Voters? Experimental Evidence from Uganda.” *The Review of Economics and Statistics* 100 (5): 891–905. [https://doi.org/10.1162/rest\\_a\\_00737](https://doi.org/10.1162/rest_a_00737).
- Blattman, Christopher, Nathan Fiala, and Sebastian Martinez. 2014. “Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda.” *The Quarterly Journal of Economics* 129 (2): 697–752. <https://doi.org/10.1093/qje/qjt057>.
- . 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. <https://doi.org/10.1257/aeri.20190224>.
- Blumenstock, Joshua, Gabriel Cadamuro, and Robert On. 2015. “Predicting Poverty and Wealth from Mobile Phone Metadata.” *Science (New York, N.Y.)* 350 (6264): 1073–76. <https://doi.org/10.1126/science.aac4420>.
- Bobonis, Gustavo J., Melissa González-Brenes, and Roberto Castro. 2013. “Public Transfers and Domestic Violence: The Roles of Private Information and Spousal Control.” *American Economic Journal: Economic Policy* 5 (1): 179–205.
- Bosch, Mariano, and Raymundo M. 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 (4): 71–99. <https://doi.org/10.1257/pol.6.4.71>.
- Botea, Ioana, Markus Goldstein, Corinne Low, and Gareth Roberts. 2021. “Supporting Women’s Livelihoods at Scale: RCT Evidence from a Nationwide Graduation Program.” Working Paper. <https://corinnelow.github.io/SWL.pdf>.
- Britto, Diogo G. C., Paolo Pinotti, and Breno Sampaio. forthcoming. “The Effect of Job Loss and Unemployment Insurance on Crime in Brazil.” *Econometrica*.

- Brollo, Fernanda, Katja Kaufmann, and Eliana La Ferrara. 2020. "The Political Economy of Program Enforcement: Evidence from Brazil." *Journal of the European Economic Association* 18 (2): 750–91. <https://doi.org/10.1093/jeea/jvz024>.
- Brollo, Fernanda, Katja Maria Kaufmann, and Eliana La Ferrara. 2017. "Learning About the Enforcement of Conditional Welfare Programs: Evidence from Brazil." Discussion Paper IZA DP No. 10654. IZA Discussion Paper Series. IZA. <http://ftp.iza.org/dp10654.pdf>.
- Bronsoler, Ari, Jonathan Gruber, and Enrique Seira. 2021. "The Impact of a Private Supplement to Public Health Care: The Mexico Diabetes Experiment." w28913. National Bureau of Economic Research. <https://doi.org/10.3386/w28913>.
- Brown, Caitlin, Martin Ravallion, and Dominique van de Walle. 2018. "A Poor Means Test? Econometric Targeting in Africa." *Journal of Development Economics* 134 (September): 109–24. <https://doi.org/10.1016/j.jdeveco.2018.05.004>.
- Browning, M., and P. A. Chiappori. 1998. "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests." *Econometrica* 66 (6): 1241–78. <https://doi.org/10.2307/2999616>.
- Buddelmeyer, Hielke, and Emmanuel Skoufias. 2004. "An Evaluation of the Performance of Regression Discontinuity Design on PROGRESA." Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-3386>.
- 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 (2): 218–58. <https://doi.org/10.1093/wbro/lky002>.
- Bursztyn, Leonardo. 2016. "Poverty and the Political Economy of Public Education Spending: Evidence from Brazil." *Journal of the European Economic Association* 14 (5): 1101–28.
- Cahyadi, Nur, Rema Hanna, Benjamin A. Olken, Rizal 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 (4): 88–110. <https://doi.org/10.1257/pol.20190245>.
- Cai, Jing, Alain De Janvry, and Elisabeth Sadoulet. 2015. "Social Networks and the Decision to Insure." *American Economic Journal: Applied Economics* 7 (2): 81–108. <https://doi.org/10.1257/app.20130442>.
- Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet. 2020. "Subsidy Policies and Insurance Demand." *American Economic Review* 110 (8): 2422–53. <https://doi.org/10.1257/aer.20190661>.
- Calderón, Gabriela. 2014. "The Effects of Child Care Provision in Mexico." Working Paper 2014–07. Mexico City: Banco de México. <http://hdl.handle.net/10419/100113>.
- Camacho, Adriana, and Emily Conover. 2011a. "Manipulation of Social Program Eligibility." *American Economic Journal: Economic Policy* 3 (2): 41–65. <https://doi.org/10.1257/pol.3.2.41>.
- . 2011b. "Manipulation of Social Program Eligibility." *American Economic Journal: Economic Policy* 3 (2): 41–65. <https://doi.org/10.1257/pol.3.2.41>.
- Camacho, Adriana, Emily Conover, and Alejandro Hoyos. 2014. "Effects of Colombia's Social Protection System on Workers' Choice between Formal and Informal Employment." *World Bank Economic Review* 28 (3): 446–66. <https://doi.org/10.1093/wber/lht028>.

- 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 (2): 381–415. <https://doi.org/10.1086/674102>.
- Caria, Stefano, Grant Gordon, Maximilian Kasy, Simon Quinn, Soha Shami, and Alexander Teytelboym. 2021. "An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan." Working Paper. <https://www.cesifo.org/en/publikationen/2020/working-paper/adaptive-targeted-field-experiment-job-search-assistance-refugees>.
- Carneiro, Pedro, Emanuela Galasso, and Rita Ginja. 2019. "Tackling Social Exclusion: Evidence from Chile." *The Economic Journal* 129 (617): 172–208. <https://doi.org/10.1111/eoj.12594>.
- Carney, Kevin, Michael Kremer, Xinyue Lin, and Gautam Rao. 2019. "The Endowment Effect and Collateralized Loans." Working Paper.
- 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 (1): 421–38. <https://doi.org/10.1146/annurev-resource-100516-053352>.
- 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 (2): 255–64. <https://doi.org/10.1016/j.jdeveco.2012.02.003>.
- Carvalho Filho, Irineu Evangelista de. 2008. "Old-Age Benefits and Retirement Decisions of Rural Elderly in Brazil." *Journal of Development Economics* 86 (1): 129–46. <https://doi.org/10.1016/j.jdeveco.2007.10.007>.
- . 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. <https://doi.org/10.1086/662576>.
- Casaburi, Lorenzo, and Jack Willis. 2018. "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya." *American Economic Review* 108 (12): 3778–3813. <https://doi.org/10.1257/aer.20171526>.
- Cascio, Elizabeth U. 2009. "Maternal Labor Supply and the Introduction of Kindergartens into American Public Schools." *The Journal of Human Resources* 44 (1): 140–70.
- Case, Anne, and Angus Deaton. 1998. "Large Cash Transfers to the Elderly in South Africa." *The Economic Journal* 108 (450): 1330–61. <https://doi.org/10.1111/1468-0297.00345>.
- Case, Anne, Anu Garrib, Alicia Menendez, and Analia Olgati. 2013. "Paying the Piper: The High Cost of Funerals in South Africa." *Economic Development and Cultural Change* 62 (1): 1–20. <https://doi.org/10.1086/671712>.
- Chattopadhyay, Raghavendra, and Esther Duflo. 2004. "Women as Policy Makers: Evidence from a Randomized Policy Experiment in India." *Econometrica* 72 (5): 1409–43.
- Chaudhury, Nazmul, Jed Friedman, and Junko Onishi. 2013. "Philippines Conditional Cash Transfer Program Impact Evaluation 2012." 75533-PH. The World Bank. <https://assessments.hpc.tools/attachments/23e5fbd8-5b17-4d16-b25d-62b583cdd0c4/Philippines%20Conditional%20Cash%20Transfer%20Program%2C%20Impact%20Evaluation%202012.pdf>.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. "Double/Debiased Machine Learning for Treatment and Structural Parameters." *The Econometrics Journal* 21 (1): C1–68. <https://doi.org/10.1111/ectj.12097>.
- Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2018. "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments,

- with an Application to Immunization in India.” Working Paper 24678. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w24678>.
- Chetty, Raj. 2006. “A General Formula for the Optimal Level of Social Insurance.” *Journal of Public Economics* 90 (10): 1879–1901. <https://doi.org/10.1016/j.jpubeco.2006.01.004>.
- . 2009. “Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance.” *American Economic Journal: Economic Policy* 1 (2): 31–52.
- Chetty, Raj, and Adam Looney. 2007. “Income Risk and the Benefits of Social Insurance: Evidence from Indonesia and the United States.” In *Fiscal Policy and Management in East Asia*, 99–121. University of Chicago Press. <https://www.nber.org/books-and-chapters/fiscal-policy-and-management-east-asia/income-risk-and-benefits-social-insurance-evidence-indonesia-and-united-states>.
- Chiappori, Pierre-André. 1988. “Rational Household Labor Supply.” *Econometrica* 56 (1): 63–90. <https://doi.org/10.2307/1911842>.
- . 1992. “Collective Labor Supply and Welfare.” *Journal of Political Economy* 100 (3): 437–67.
- Chodorow-Reich, Gabriel. 2019. “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?” *American Economic Journal: Economic Policy* 11 (2): 1–34. <https://doi.org/10.1257/pol.20160465>.
- 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 (4): 1247–72. <https://doi.org/10.1007/s13524-019-00793-3>.
- Coady, David, Margaret Grosh, and John Hoddinott. 2004a. “Targeting of Transfers in Developing Countries : Review of Lessons and Experience.” Washington, DC: World Bank. <https://doi.org/10.1596/0-8213-5769-7>.
- . 2004b. “Targeting Outcomes Redux.” *The World Bank Research Observer* 19 (1): 61–85. <https://doi.org/10.1093/wbro/lkh016>.
- Coate, Stephen, Stephen Johnson, and Richard Zeckhauser. 1994. “Pecuniary Redistribution through In-Kind Programs.” *Journal of Public Economics* 55 (1): 19–40. [https://doi.org/10.1016/0047-2727\(94\)90078-7](https://doi.org/10.1016/0047-2727(94)90078-7).
- Cole, Shawn A., and Wentao Xiong. 2017. “Agricultural Insurance and Economic Development.” *Annual Review of Economics* 9 (1): 235–62. <https://doi.org/10.1146/annurev-economics-080315-015225>.
- Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery. 2013. “Barriers to Household Risk Management: Evidence from India.” *American Economic Journal: Applied Economics* 5 (1): 104–35. <https://doi.org/10.1257/app.5.1.104>.
- 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 (5): 284–90. <https://doi.org/10.1257/aer.104.5.284>.
- Cole, Shawn, and Wentao Xiong. 2017. “Agricultural Insurance and Economic Development.” *Annual Review of Economics* 9 (1): 235–62. <https://doi.org/10.1146/annurev-economics-080315-015225>.
- Conover, Emily, Román A. Zárate, Adriana Camacho, and Javier E. Baez. 2020. “Cash and Ballots: Conditional Transfers, Political Participation, and Voting Behavior.” *Economic Development and Cultural Change* 68 (2): 541–66. <https://doi.org/10.1086/701211>.

- Cunha, Jesse M. 2014. "Testing Paternalism: Cash versus In-Kind Transfers." *American Economic Journal: Applied Economics* 6 (2): 195–230. <https://doi.org/10.1257/app.6.2.195>.
- Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran. 2019. "The Price Effects of Cash Versus In-Kind Transfers." *The Review of Economic Studies* 86 (1): 240–81. <https://doi.org/10.1093/restud/rdy018>.
- Cunningham, Wendy V. 2000. "Unemployment Insurance in Brazil: Unemployment Duration, Wages, and Sectoral Choice," 64.
- Currie, Janet. 2004. "The Take Up of Social Benefits." Working Paper 10488. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w10488>.
- Currie, Janet, and Firouz Gahvari. 2008a. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature* 46 (2): 333–83. <https://doi.org/10.1257/jel.46.2.333>.
- . 2008b. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature* 46 (2): 333–83. <https://doi.org/10.1257/jel.46.2.333>.
- Dahl, Gordon B., Katrine V. Løken, Magne Mogstad, and Kari Vea Salvanes. 2016. "What Is the Case for Paid Maternity Leave?" *The Review of Economics and Statistics* 98 (4): 655–70. [https://doi.org/10.1162/REST\\_a\\_00602](https://doi.org/10.1162/REST_a_00602).
- Dasgupta, Partha, and Debraj Ray. 1986. "Inequality as a Determinant of Malnutrition and Unemployment: Theory." *The Economic Journal* 96 (384): 1011–34. <https://doi.org/10.2307/2233171>.
- . 1987. "Inequality as a Determinant of Malnutrition and Unemployment: Policy." *The Economic Journal* 97 (385): 177–88. <https://doi.org/10.2307/2233329>.
- De La O, Ana L. 2013. "Do Conditional Cash Transfers Affect Electoral Behavior? Evidence from a Randomized Experiment in Mexico." *American Journal of Political Science* 57 (1): 1–14.
- 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.
- Demirguc-Kunt, Asli, Leora Klapper, and Neeraj Prasad. 2017. "Measuring the Effectiveness of Service Delivery: Delivery of Government Provided Goods and Services in India." Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-8207>.
- Dercon, Stefan, Joachim De Weerd, Tessa Bold, and Alula Pankhurst. 2006. "Group-Based Funeral Insurance in Ethiopia and Tanzania." *World Development* 34 (4): 685–703. <https://doi.org/10.1016/j.worlddev.2005.09.009>.
- Dercon, Stefan, Jan Willem Gunning, and Andrew Zeitlin. 2019. "The Demand for Insurance under Limited Trust: Evidence from a Field Experiment in Kenya." 2019–06. *CSAE Working Paper Series*. CSAE Working Paper Series. Centre for the Study of African Economies, University of Oxford. <https://ideas.repec.org/p/csa/wpaper/2019-06.html>.
- Dercon, Stefan, Ruth Vargas Hill, Daniel Clarke, Ingo Outes-Leon, and Alemayehu Seyoum Taffesse. 2014. "Offering Rainfall Insurance to Informal Insurance Groups: Evidence from a Field Experiment in Ethiopia." *Journal of Development Economics* 106 (January): 132–43. <https://doi.org/10.1016/j.jdeveco.2013.09.006>.
- Drazen, Allan, and Marcela Eslava. 2010. "Electoral Manipulation via Voter-Friendly Spending: Theory and Evidence." *Journal of Development Economics* 92 (1): 39–52.
- Duflo, Esther. 2003. "Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa." *The World Bank Economic Review* 17 (1): 1–25.

- Dupas, Pascaline, Marcel Fafchamps, and Deivy Houeix. 2022. “Measuring Relative Poverty through Peer Rankings: Evidence from Côte D’Ivoire.” Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w29911>.
- Dupas, Pascaline, Vivian Hoffmann, Michael Kremer, and Alix Peterson Zwane. 2016. “Targeting Health Subsidies through a Nonprice Mechanism: A Randomized Controlled Trial in Kenya.” *Science* 353 (6302): 889–95. <https://doi.org/10.1126/science.aaf6288>.
- 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. <https://doi.org/10.1596/978-1-4648-0130-3>.
- Edmonds, Eric V. 2006. “Child Labor and Schooling Responses to Anticipated Income in South Africa.” *Journal of Development Economics* 81 (2): 386–414. <https://doi.org/10.1016/j.jdeveco.2005.05.001>.
- Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael W. Walker. 2019. “General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya.” Working Paper 26600. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w26600>.
- Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. “Selection on Moral Hazard in Health Insurance.” *American Economic Review* 103 (1): 178–219. <https://doi.org/10.1257/aer.103.1.178>.
- 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 (1): 198–213. <https://doi.org/10.1016/j.jdeveco.2006.02.001>.
- Elbers, Chris, Jean Lanjouw, and Peter Lanjouw. 2003. “Micro-Level Estimation of Poverty and Inequality.” *Econometrica* 71 (1): 355–64.
- 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. <https://doi.org/10.1596/978-1-4648-0141-9>.
- Evans, David, Brian Holtemeyer, and Katrina Kosec. 2019. “Cash Transfers and Health: Evidence from Tanzania.” *The World Bank Economic Review* 33 (2): 394–412. <https://doi.org/10.1093/wber/lhx001>.
- Evans, David, and Katrina Kosec. 2020. “Do Cash Transfers Reduce Trust and Informal Transfers within Communities?” Discussion Paper 01994. 0 ed. IFPRI Discussion Paper. Washington, DC: International Food Policy Research Institute. <https://doi.org/10.2499/p15738coll2.134236>.
- 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 (2): 242–70. <https://doi.org/10.1093/wbro/lkw004>.
- Ewoudou, Jacques, Clarence Tsimpo, and Quentin Wodon. 2009. “Stigma and the Take-up of Social Programs.” Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-4962>.
- Fallon, Kathleen M., Alissa Mazar, and Liam Swiss. 2017. “The Development Benefits of Maternity Leave.” *World Development* 96 (August): 102–18. <https://doi.org/10.1016/j.worlddev.2017.03.001>.
- Farooq, Ammar, Adriana D. Kugler, and Umberto Muratori. 2020. “Do Unemployment Insurance Benefits Improve Match and Employer Quality? Evidence from Recent U.S. Recessions.” Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w27574>.

- Feldstein, Martin. 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." *The Review of Economics and Statistics* 81 (4): 674–80.
- Feldstein, Martin, and Daniel Altman. 1998. "Unemployment Insurance Savings Accounts." Working Paper 6860. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w6860>.
- Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore. 2021. "On Her Own Account: How Strengthening Women's Financial Control Impacts Labor Supply and Gender Norms." *American Economic Review* 111 (7): 2342–75. <https://doi.org/10.1257/aer.20200705>.
- Filmer, Deon, Jed Friedman, Eeshani Kandpal, and Junko Onishi. forthcoming. "Cash Transfers, Food Prices, and Nutrition Impacts on Nonbeneficiary Children." *Review of Economics and Statistics*. <https://papers.ssrn.com/abstract=3150090>.
- Finkelstein, Amy. 2007. "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare." *The Quarterly Journal of Economics* 122 (1): 1–37. <https://doi.org/10.1162/qjec.122.1.1>.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review* 109 (4): 1530–67. <https://doi.org/10.1257/aer.20171455>.
- 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 (7): 1644–68. <https://doi.org/10.1016/j.jpubeco.2007.10.005>.
- Finkelstein, Amy, and Matthew J Notowidigdo. 2019. "Take-Up and Targeting: Experimental Evidence from SNAP." *The Quarterly Journal of Economics* 134 (3): 1505–56. <https://doi.org/10.1093/qje/qjz013>.
- Fischer, Torben, Markus Frölich, and Andreas Landmann. 2018. "Adverse Selection in Low-Income Health Insurance Markets: Evidence from a RCT in Pakistan." 11751. *IZA Discussion Papers*. IZA Discussion Papers. Institute of Labor Economics (IZA). <https://ideas.repec.org/p/iza/izadps/dp11751.html>.
- Franklin, Simon, Clément Imbert, Girum Abebe, and Carolina Mejia-Mantilla. 2021. "Urban Public Works in Spatial Equilibrium: Experimental Evidence from Ethiopia." Working Paper. [https://www.dropbox.com/s/rxflgiwx9ub8mx/20210514\\_FIAM.pdf?dl=0](https://www.dropbox.com/s/rxflgiwx9ub8mx/20210514_FIAM.pdf?dl=0).
- Fujiwara, Thomas. 2015. "Voting Technology, Political Responsiveness, and Infant Health: Evidence From Brazil." *Econometrica* 83 (2): 423–64. <https://doi.org/10.3982/ECTA11520>.
- Gadenne, Lucie, Samuel Norris, Monica Singhal, and Sandip Sukhtankar. 2021. "In-Kind Transfers as Insurance." Working Paper 28507. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w28507>.
- Galasso, Emanuela, and Martin Ravallion. 2005. "Decentralized Targeting of an Antipoverty Program." *Journal of Public Economics*, Cornell - ISPE Conference on Public Finance and Development, 89 (4): 705–27. <https://doi.org/10.1016/j.jpubeco.2003.01.002>.
- Galiani, Sebastian, Paul Gertler, and Rosangela Bando. 2016. "Non-Contributory Pensions." *Labour Economics* 38 (January): 47–58. <https://doi.org/10.1016/j.labeco.2015.11.003>.
- Galiani, Sebastian, and Patrick J. McEwan. 2013. "The Heterogeneous Impact of Conditional Cash Transfers." *Journal of Public Economics* 103 (July): 85–96. <https://doi.org/10.1016/j.jpubeco.2013.04.004>.
- Galor, Oded, and Joseph Zeira. 1993. "Income Distribution and Macroeconomics." *The Review of Economic Studies* 60 (1): 35–52. <https://doi.org/10.2307/2297811>.

- 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.  
<https://www.opml.co.uk/files/Publications/a0013-evaluation-kenya-hunger-safety-net-programme/hsnp-legacy-systems-practices-lessons-learned.pdf>.
- 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." Brief. Washington, DC: World Bank. <https://openknowledge.worldbank.org/handle/10986/33635>.
- 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 (3): 167–206. <https://doi.org/10.1257/pol.20180072>.
- Gerard, François, Joana Naritomi, and Joana Silva. 2021. "Cash Transfers and Formal Labor Markets: Evidence from Brazil." Working Paper 9778. Policy Research Working Papers. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-9778>.
- Gertler, Paul. 2004. "Do Conditional Cash Transfers Improve Child Health? Evidence from PROGRESA's Control Randomized Experiment." *American Economic Review* 94 (2): 336–41. <https://doi.org/10.1257/0002828041302109>.
- Gertler, Paul, Sebastian W. Martinez, and Marta Rubio-Codina. 2012. "Investing Cash Transfers to Raise Long-Term Living Standards." *American Economic Journal: Applied Economics* 4 (1): 164–92. <https://doi.org/10.1257/app.4.1.164>.
- 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 (January): 26–44. <https://doi.org/10.1016/j.jdeveco.2015.08.004>.
- 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 (May): 102645. <https://doi.org/10.1016/j.jdeveco.2021.102645>.
- Giné, Xavier, Robert Townsend, and James Vickery. 2008. "Patterns of Rainfall Insurance Participation in Rural India." *World Bank Economic Review* 22 (3): 539–66. <https://doi.org/10.1093/wber/lhn015>.
- GIZ. 2021. "Innovations and Emerging Trends in Agricultural Insurance for Smallholder Farmers – an Update." Eschborn. [https://www.giz.de/expertise/downloads/2021%20GIZ\\_Innovations%20and%20emerging%20Trends%20in%20Agricultural%20Insurance-An%20update.pdf](https://www.giz.de/expertise/downloads/2021%20GIZ_Innovations%20and%20emerging%20Trends%20in%20Agricultural%20Insurance-An%20update.pdf).
- Glewwe, P., and Pedro Olinto. 2004. "Evaluating of the Impact of Conditional Cash Transfers on Schooling: An Experimental Analysis of Honduras' PRAF Program." USAID. [http://pdf.usaid.gov/pdf\\_docs/PNADT588.pdf](http://pdf.usaid.gov/pdf_docs/PNADT588.pdf).
- Glewwe, Paul. 1992. "Targeting Assistance to the Poor: Efficient Allocation of Transfers When Household Income Is Not Observed." *Journal of Development Economics* 38 (2): 297–321. [https://doi.org/10.1016/0304-3878\(92\)90002-Q](https://doi.org/10.1016/0304-3878(92)90002-Q).
- Goldberg, Jessica. 2016. "Kwacha Gonna Do? Experimental Evidence about Labor Supply in Rural Malawi." *American Economic Journal: Applied Economics* 8 (1): 129–49. <https://doi.org/10.1257/app.20130369>.
- Golden, Miriam, and Brian Min. 2013. "Distributive Politics Around the World." *Annual Review of Political Science* 16 (1): 73–99. <https://doi.org/10.1146/annurev-polisci-052209-121553>.

- Goldin, Jacob, Ithai Z Lurie, and Janet McCubbin. 2021. “Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach.” *The Quarterly Journal of Economics* 136 (1): 1–49. <https://doi.org/10.1093/qje/qjaa029>.
- Gonzalez-Rozada, Martin, and Hernán Ruffo. 2016. “Optimal Unemployment Benefits in the Presence of Informal Labor Markets.” *Labour Economics*, SOLE/EALE conference issue 2015, 41 (August): 204–27. <https://doi.org/10.1016/j.labeco.2016.06.001>.
- Gonzalez-Rozada, Martin, and Hernan Ruffo. 2022. “The Welfare Effects of Unemployment Insurance in Argentina. New Estimates Using Changes in the Schedule of Transfers.” arXiv. <https://doi.org/10.48550/arXiv.2202.01844>.
- Grosh, Margaret E, and Judy L Baker. 1995. “Proxy Means Tests for Targeting Social Programs: Simulations and Speculation.” Working Paper 118. Living Standards Measurement Study Working Paper. Washington, D.C.: World Bank. <https://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=33965>.
- Gruber, Jonathan, Nathaniel Hendren, and Robert M. Townsend. 2014. “The Great Equalizer: Health Care Access and Infant Mortality in Thailand.” *American Economic Journal: Applied Economics* 6 (1): 91–107. <https://doi.org/10.1257/app.6.1.91>.
- Gruber, Jonathan, Mengyun Lin, and Junjian Yi. 2021. “The Largest Insurance Program in History: Saving One Million Lives Per Year in China.” <https://www.bu.edu/econ/files/2021/03/NCMS-paper-3-15-21.pdf>.
- Guiteras, Raymond P., and B. Kelsey Jack. 2018. “Productivity in Piece-Rate Labor Markets: Evidence from Rural Malawi.” *Journal of Development Economics* 131: 42–61. <https://doi.org/10.1016/j.jdeveco.2017.11.002>.
- Gunnsteinsson, Snaebjorn. 2020. “Experimental Identification of Asymmetric Information: Evidence on Crop Insurance in the Philippines.” *Journal of Development Economics* 144 (May): 102414. <https://doi.org/10.1016/j.jdeveco.2019.102414>.
- Gupta, Sarika. 2017a. “Perils of the Paperwork: The Impact of Information and Application Assistance on Welfare Program Take-Up in India.” Job Market Paper. Cambridge: Harvard University. [https://scholar.harvard.edu/files/sarikagupta/files/gupta\\_jmp\\_11\\_1.pdf](https://scholar.harvard.edu/files/sarikagupta/files/gupta_jmp_11_1.pdf).
- . 2017b. “Perils of the Paperwork: The Impact of Information and Application Assistance on Welfare Program Take-Up in India.” Job Market Paper. Cambridge: Harvard University. [https://scholar.harvard.edu/files/sarikagupta/files/gupta\\_jmp\\_11\\_1.pdf](https://scholar.harvard.edu/files/sarikagupta/files/gupta_jmp_11_1.pdf).
- Halim, Daniel, Hillary C. Johnson, and Elizaveta Perova. 2022. “Preschool Availability and Women’s Employment: Evidence from Indonesia.” *Economic Development and Cultural Change*. <https://doi.org/10.1086/714439>.
- 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 (2): 259–98. <https://doi.org/10.1093/wbro/lky003>.
- Hanna, Rema, and Dean Karlan. 2017. “Designing Social Protection Programs: Using Theory and Experimentation to Understand How to Help Combat Poverty.” In *Handbook of Economic Field Experiments*, edited by Abhijit Banerjee and Esther Duflo, 2:515–53. Handbook of Economic Field Experiments. North-Holland. <https://doi.org/10.1016/bs.hefe.2016.07.002>.
- Hanna, Rema, and Benjamin A. Olken. 2018. “Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries.” *Journal of Economic Perspectives* 32 (4): 201–26. <https://doi.org/10.1257/jep.32.4.201>.

- Harris, John R., and Michael P. Todaro. 1970. "Migration, Unemployment and Development: A Two-Sector Analysis." *The American Economic Review* 60 (1): 126–42.
- 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 (6): 1723–61. <https://doi.org/10.3982/ECTA12302>.
- Hastings, Justine, and Jesse M. Shapiro. 2018. "How Are SNAP Benefits Spent? Evidence from a Retail Panel." *American Economic Review* 108 (12): 3493–3540. <https://doi.org/10.1257/aer.20170866>.
- Haushofer, Johannes, Matthieu Chemin, Channing 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 (C). <https://ideas.repec.org/a/eee/deveco/v144y2020ics0304387818310289.html>.
- Haushofer, Johannes, Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael W. Walker. 2022a. "Targeting Impact versus Deprivation." Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w30138>.
- . 2022b. "Targeting Impact versus Deprivation." Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w30138>.
- 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 (4): 1973–2042. <https://doi.org/10.1093/qje/qjw025>.
- . 2018. "The Long-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya." Working Paper. [https://jeremypshapiro.appspot.com/papers/Haushofer\\_Shapiro\\_UCT2\\_2018-01-30\\_paper\\_only.pdf](https://jeremypshapiro.appspot.com/papers/Haushofer_Shapiro_UCT2_2018-01-30_paper_only.pdf).
- He, Guojun, and Shaoda Wang. 2017. "Do College Graduates Serving as Village Officials Help Rural China?" *American Economic Journal: Applied Economics* 9 (4): 186–215. <https://doi.org/10.1257/app.20160079>.
- Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics* 135 (3): 1209–1318. <https://doi.org/10.1093/qje/qjaa006>.
- 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 (March): 144–56. <https://doi.org/10.1016/j.jdeveco.2013.11.009>.
- Hirshleifer, Sarojini, David McKenzie, Rita Almeida, and Cristobal Ridao-Cano. 2016. "The Impact of Vocational Training for the Unemployed: Experimental Evidence from Turkey." *The Economic Journal* 126 (597): 2115–46. <https://doi.org/10.1111/eoj.12211>.
- Hoynes, Hilary, and Jesse Rothstein. 2019. "Universal Basic Income in the United States and Advanced Countries." *Annual Review of Economics* 11 (1): 929–58. <https://doi.org/10.1146/annurev-economics-080218-030237>.
- 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 (2): 179–205. <https://doi.org/10.1257/app.20170789>.
- Huneus, Cristobal, Silvia Leiva, and Alejandro Micco. 2012. "Unemployment Insurance and Search Effort in Chile." SSRN Scholarly Paper. Rochester, NY. <https://doi.org/10.2139/ssrn.2149060>.

- Hussam, Reshmaan N., Erin M. Kelley, Gregory V. Lane, and Fatima T. Zahra. 2021. “The Psychosocial Value of Employment.” National Bureau of Economic Research.
- Hussam, Reshmaan, Natalia Rigol, and Benjamin N. Roth. 2022. “Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field.” *American Economic Review* 112 (3): 861–98. <https://doi.org/10.1257/aer.20200751>.
- Ikegami, Munenobu, Michael R. Carter, Christopher B. Barrett, and Sarah Janzen. 2018. “Poverty Traps and the Social Protection Paradox.” In *The Economics of Poverty Traps*, 223–56. University of Chicago Press. <https://www.nber.org/books-and-chapters/economics-poverty-traps/poverty-traps-and-social-protection-paradox>.
- ILOSTAT. 2019. “Employment in Agriculture (% of Total Employment) - Low Income.” International Labour Organization. <https://data.worldbank.org/indicator/SL.AGR.EMPL.ZS?locations=XM>.
- Imai, Kosuke, Gary King, and Carlos Velasco Rivera. 2020. “Do Nonpartisan Programmatic Policies Have Partisan Electoral Effects? Evidence from Two Large-Scale Experiments.” *The Journal of Politics* 82 (2): 714–30. <https://doi.org/10.1086/707059>.
- Imbert, Clément, and John Papp. 2015. “Labor Market Effects of Social Programs: Evidence from India’s Employment Guarantee.” *American Economic Journal: Applied Economics* 7 (2): 233–63. <https://doi.org/10.1257/app.20130401>.
- . 2020. “Short-Term Migration, Rural Public Works, and Urban Labor Markets: Evidence from India.” *Journal of the European Economic Association* 18 (2): 927–63. <https://doi.org/10.1093/jeea/jvz009>.
- Insurance Information Institute. 2020. “Insured Homes by State.” <https://www.iii.org/publications/a-firm-foundation-how-insurance-supports-the-economy/a-50-state-commitment/insured-homes-by-state>.
- Janvry, Alain de, 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*, Special Issue in honor of Pranab Bardhan, 79 (2): 349–73. <https://doi.org/10.1016/j.jdeveco.2006.01.013>.
- 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. <https://doi.org/10.1111/jori.12322>.
- Jean, Neal, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon. 2016. “Combining Satellite Imagery and Machine Learning to Predict Poverty.” *Science* 353 (6301): 790–94. <https://doi.org/10.1126/science.aaf7894>.
- Jensen, Anders. 2022. “Employment Structure and the Rise of the Modern Tax System.” *American Economic Review* 112 (1): 213–34. <https://doi.org/10.1257/aer.20191528>.
- Jensen, Nathaniel, and Christopher Barrett. 2017. “Agricultural Index Insurance for Development.” *Applied Economic Perspectives and Policy* 39 (2): 199–219. <https://doi.org/10.1093/aep/ppw022>.
- Jensen, Robert. 2004. “Do Private Transfers ‘Displace’ the Benefits of Public Transfers? Evidence from South Africa.” *Journal of Public Economics* 88 (1): 89–112. [https://doi.org/10.1016/S0047-2727\(02\)00085-3](https://doi.org/10.1016/S0047-2727(02)00085-3).
- Jiménez-Hernández, Diego, and Enrique Seira. 2022. “Should the Government Sell You Goods? Evidence from the Milk Market in Mexico.” Job Market Paper. [https://www.diegोजimenezh.com/assets/pdf/JimenezHernandez\\_JMP.pdf](https://www.diegोजimenezh.com/assets/pdf/JimenezHernandez_JMP.pdf).

- 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. <https://doi.org/10.3945/jn.116.233684>.
- Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. “Agricultural Decisions after Relaxing Credit and Risk Constraints.” *The Quarterly Journal of Economics* 129 (2): 597–652. <https://doi.org/10.1093/qje/qju002>.
- Kaur, Supreet, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach. 2021. “Do Financial Concerns Make Workers Less Productive?” Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w28338>.
- Khemani, Stuti. 2004. “Political Cycles in a Developing Economy: Effect of Elections in the Indian States.” *Journal of Development Economics* 73 (1): 125–54.
- King, Gary, Emmanuela Gakidou, Kosuke Imai, Jason Lakin, Ryan T. Moore, Clayton Nall, Nirmala Ravishankar, et al. 2009. “Public Policy for the Poor? A Randomised Assessment of the Mexican Universal Health Insurance Programme.” *The Lancet* 373 (9673): 1447–54. [https://doi.org/10.1016/S0140-6736\(09\)60239-7](https://doi.org/10.1016/S0140-6736(09)60239-7).
- Klasen, Stephan, and Simon Lange. 2016. “How Narrowly Should Anti-Poverty Programs Be Targeted? Simulation Evidence from Bolivia and Indonesia.” 213. *Courant Research Centre: Poverty, Equity and Growth - Discussion Papers*. Courant Research Centre: Poverty, Equity and Growth - Discussion Papers. Courant Research Centre PEG. <https://ideas.repec.org/p/got/gotrc/213.html>.
- Kleven, Henrik Jacobsen, and Wojciech Kopczuk. 2011. “Transfer Program Complexity and the Take-Up of Social Benefits.” *American Economic Journal: Economic Policy* 3 (1): 54–90. <https://doi.org/10.1257/pol.3.1.54>.
- Kleven, Henrik, Camille Landais, Johanna Posch, Andreas Steinhauer, and Josef Zweimüller. 2020. “Do Family Policies Reduce Gender Inequality? Evidence from 60 Years of Policy Experimentation.” Working Paper 28082. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w28082>.
- Kugler, Adriana D. 2005. “Wage-Shifting Effects of Severance Payments Savings Accounts in Colombia.” *Journal of Public Economics* 89 (2): 487–500. <https://doi.org/10.1016/j.jpubeco.2004.04.006>.
- Kugler, Adriana D., and Ingrid Rojas. 2018. “Do CCTs Improve Employment and Earnings in the Very Long-Term? Evidence from Mexico.” Working Paper 24248. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w24248>.
- 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): 881–96. [https://doi.org/10.1162/rest\\_a\\_00907](https://doi.org/10.1162/rest_a_00907).
- Kyle, Jordan. 2018. “Local Corruption and Popular Support for Fuel Subsidy Reform in Indonesia.” *Comparative Political Studies* 51 (11): 1472–1503. <https://doi.org/10.1177/0010414018758755>.
- Labonne, Julien. 2013. “The Local Electoral Impacts of Conditional Cash Transfers: Evidence from a Field Experiment.” *Journal of Development Economics* 104 (September): 73–88. <https://doi.org/10.1016/j.jdeveco.2013.04.006>.
- Lagakos, David. 2020. “Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities?” *The Journal of Economic Perspectives* 34 (3): 174–92.

- Lalive, Rafael, Analía Schlosser, Andreas Steinhauer, and Josef Zweimüller. 2014. “Parental Leave and Mothers’ Careers: The Relative Importance of Job Protection and Cash Benefits.” *The Review of Economic Studies* 81 (1): 219–65. <https://doi.org/10.1093/restud/rdt028>.
- Lalive, Rafael, and Josef Zweimüller. 2009. “How Does Parental Leave Affect Fertility and Return to Work? Evidence from Two Natural Experiments\*.” *The Quarterly Journal of Economics* 124 (3): 1363–1402. <https://doi.org/10.1162/qjec.2009.124.3.1363>.
- Lefebvre, Pierre, and Philip Merrigan. 2008. “Child-Care Policy and the Labor Supply of Mothers with Young Children: A Natural Experiment from Canada.” *Journal of Labor Economics* 26 (3): 519–48. <https://doi.org/10.1086/587760>.
- Leroy, Jef L., 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 (3): 612–17. <https://doi.org/10.3945/jn.109.116285>.
- Levine, David, Rachel Polimeni, and Ian Ramage. 2016. “Insuring Health or Insuring Wealth? An Experimental Evaluation of Health Insurance in Rural Cambodia.” *Journal of Development Economics* 119 (March): 1–15. <https://doi.org/10.1016/j.jdeveco.2015.10.008>.
- Levy, Santiago. 2006. *Progress Against Poverty*. Brookings Institution Press. <https://www.brookings.edu/book/progress-against-poverty/>.
- Levy, Santiago, and Norbert Schady. 2013. “Latin America’s Social Policy Challenge: Education, Social Insurance, Redistribution.” *Journal of Economic Perspectives* 27 (2): 193–218. <https://doi.org/10.1257/jep.27.2.193>.
- Liepmann, Hannah, and Clemente Pignatti. 2021. “Welfare Effects of Unemployment Benefits When Informality Is High.” 3900863. IZA Discussion Papers. Rochester, NY: Social Science Research Network. <https://doi.org/10.2139/ssrn.3900863>.
- “Life Insurance Volume - Country Rankings.” 2019. TheGlobalEconomy.Com. 2019. [https://www.theglobaleconomy.com/rankings/life\\_insurance\\_volume/](https://www.theglobaleconomy.com/rankings/life_insurance_volume/).
- 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): 209–44. <https://doi.org/10.1111/1467-937X.00204>.
- Lone, Tanya, Maheen Shakeel, Jana Bischler, and Marta Marzi. 2021. “Towards Shock-Responsive Social Protection: Lessons from the COVID-19 Response in Pakistan.” *Maintains*.
- 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 (3): 463–80. <https://doi.org/10.2307/146179>.
- Lybbert, Travis J., Christopher B. Barrett, Solomon Desta, and D. Layne Coppock. 2004. “Stochastic Wealth Dynamics and Risk Management among a Poor Population.” *The Economic Journal* 114 (498): 750–77. <https://doi.org/10.1111/j.1468-0297.2004.00242.x>.
- 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 (2): 247–73. <https://doi.org/10.1257/app.4.2.247>.
- 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 (576): 607–33. <https://doi.org/10.1111/eoj.12145>.
- 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.” The IFS. December 7, 2021. <https://doi.org/10.1920/wp.ifs.2021.4721>.
- Manacorda, Marco, Edward Miguel, and Andrea Vigorito. 2011. “Government Transfers and Political Support.” *American Economic Journal: Applied Economics* 3 (3): 1–28. <https://doi.org/10.1257/app.3.3.1>.
- Martinelli, César, and Susan Wendy Parker. 2009. “Deception and Misreporting in a Social Program.” *Journal of the European Economic Association* 7 (4): 886–908.
- Martínez, Claudia, A., and Marcela Perticará. 2017. “Childcare Effects on Maternal Employment: Evidence from Chile.” *Journal of Development Economics* 126 (May): 127–37. <https://doi.org/10.1016/j.jdeveco.2017.01.001>.
- Mazzucato, Valentina. 2009. “Informal Insurance Arrangements in Ghanaian Migrants’ Transnational Networks: The Role of Reverse Remittances and Geographic Proximity.” *World Development* 37 (6): 1105–15. <https://doi.org/10.1016/j.worlddev.2008.11.001>.
- McBride, Linden, and Austin Nichols. 2018. “Retooling Poverty Targeting Using Out-of-Sample Validation and Machine Learning.” *The World Bank Economic Review* 32 (3): 531–50. <https://doi.org/10.1093/wber/lhw056>.
- Meghir, Costas, and David Phillips. 2010. “Labour Supply and Taxes.” In *Dimensions of Tax Design: The Mirrlees Review*, edited by James A. Mirrlees and Institute for Fiscal Studies (Great Britain), 202–74. Oxford ; New York: Oxford University Press. <https://www.ifs.org.uk/mirrleesreview/dimensions/ch3.pdf>.
- 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 (3): 455–69. <https://doi.org/10.1007/s12571-012-0177-0>.
- Millán, Teresa Molina. 2020. “Regional Migration, Insurance and Economic Shocks: Evidence from Nicaragua.” *The Journal of Development Studies* 56 (11): 2000–2029. <https://doi.org/10.1080/00220388.2019.1703956>.
- Miller, Sarah, Norman Johnson, and Laura R. Wherry. 2019. “Medicaid and Mortality: New Evidence from Linked Survey and Administrative Data.” Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w26081>.
- Miranda, Mario J., and Katie Farrin. 2012. “Index Insurance for Developing Countries.” *Applied Economic Perspectives and Policy* 34 (3): 391–427. <https://doi.org/10.1093/aep/paps031>.
- Mirrlees, J. A. 1971. “An Exploration in the Theory of Optimum Income Taxation.” *The Review of Economic Studies* 38 (2): 175–208. <https://doi.org/10.2307/2296779>.
- Mobarak, Ahmed Mushfiq, and Mark Rosenzweig. 2013. “Informal Risk Sharing, Index Insurance, and Risk Taking in Developing Countries.” *American Economic Review* 103 (3): 375–80. <https://doi.org/10.1257/aer.103.3.375>.
- Moffitt, Robert. 1983. “An Economic Model of Welfare Stigma.” *The American Economic Review* 73 (5): 1023–35.
- Molina Millán, Teresa, 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 (March): 102385. <https://doi.org/10.1016/j.jdeveco.2019.102385>.
- Morduch, Jonathan. 1999. “The Microfinance Promise.” *Journal of Economic Literature* 37 (4): 1569–1614. <https://doi.org/10.1257/jel.37.4.1569>.

- Morris, Saul S., Rafael Flores, Pedro Olinto, and Juan Manuel 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.” *Lancet (London, England)* 364 (9450): 2030–37. [https://doi.org/10.1016/S0140-6736\(04\)17515-6](https://doi.org/10.1016/S0140-6736(04)17515-6).
- Mullainathan, Sendhil, and Eldar Shafir. 2013a. *Scarcity: Why Having Too Little Means so Much*. Scarcity: Why Having Too Little Means so Much. New York, NY, US: Times Books/Henry Holt and Co.
- . 2013b. *Scarcity: Why Having Too Little Means so Much*. Macmillan.
- Munshi, Kaivan, and Mark Rosenzweig. 2016. “Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap.” *American Economic Review* 106 (1): 46–98. <https://doi.org/10.1257/aer.20131365>.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. “Building State Capacity: Evidence from Biometric Smartcards in India.” *American Economic Review* 106 (10): 2895–2929. <https://doi.org/10.1257/aer.20141346>.
- . 2020. “Identity Verification Standards in Welfare Programs: Experimental Evidence from India.” Working Paper 26744. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w26744>.
- Murgai, Rinku, Martin Ravallion, and Dominique van de Walle. 2016. “Is Workfare Cost-Effective against Poverty in a Poor Labor-Surplus Economy?” *World Bank Economic Review* 30 (3): 413–45. <https://doi.org/10.1093/wber/lhv038>.
- Musgrave, Richard A. 1959. *The Theory of Public Finance: A Study in Public Economy*. New York: McGraw-Hill.
- Nagler, Paula. 2013. “How Unemployment Insurance Savings Accounts Affect Employment Duration: Evidence from Chile.” *IZA Journal of Labor & Development* 2 (1): 9. <https://doi.org/10.1186/2193-9020-2-9>.
- Natural Disaster Insurance Institution. 2022. “Policy Production.” DASK. 2022. <https://dask.gov.tr/en/policy-production>.
- Nekoei, Arash, and Andrea Weber. 2017. “Does Extending Unemployment Benefits Improve Job Quality?” *American Economic Review* 107 (2): 527–61. <https://doi.org/10.1257/aer.20150528>.
- Nichols, Albert L., and Richard J. Zeckhauser. 1982. “Targeting Transfers through Restrictions on Recipients.” *The American Economic Review* 72 (2): 372–77.
- Niehaus, Paul, Antonia Atanassova, Marianne Bertrand, and Sendhil Mullainathan. 2013a. “Targeting with Agents.” *American Economic Journal: Economic Policy* 5 (1): 206–38. <https://doi.org/10.1257/pol.5.1.206>.
- . 2013b. “Targeting with Agents.” *American Economic Journal: Economic Policy* 5 (1): 206–38. <https://doi.org/10.1257/pol.5.1.206>.
- Ninno, Carlo del, and Bradford Mills. 2015. “Safety Nets in Africa: Effective Mechanisms to Reach the Poor and Most Vulnerable.” Washington, DC: World Bank. <https://openknowledge.worldbank.org/handle/10986/21369>.
- Nollenberger, Natalia, and Núria Rodríguez-Planas. 2015. “Full-Time Universal Childcare in a Context of Low Maternal Employment: Quasi-Experimental Evidence from Spain.” *Labour Economics* 36 (October): 124–36. <https://doi.org/10.1016/j.labeco.2015.02.008>.
- Nordhaus, William D. 1975. “The Political Business Cycle.” *The Review of Economic Studies* 42 (2): 169–90.

- 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): 605–31. <https://doi.org/10.1017/S1355770X21000085>.
- Oates, Wallace E., and Robert M. Schwab. 2015. “The Window Tax: A Case Study in Excess Burden.” *Journal of Economic Perspectives* 29 (1): 163–80. <https://doi.org/10.1257/jep.29.1.163>.
- OECD. 2020. “Insurance Spending.” <http://data.oecd.org/insurance/insurance-spending.htm>.
- Olken, Benjamin A. 2006. “Corruption and the Costs of Redistribution: Micro Evidence from Indonesia.” *Journal of Public Economics* 90 (4): 853–70. <https://doi.org/10.1016/j.jpubeco.2005.05.004>.
- Olken, Benjamin A. 2007. “Monitoring Corruption: Evidence from a Field Experiment in Indonesia.” *Journal of Political Economy* 115 (2): 200–249. <https://doi.org/10.1086/517935>.
- Osman, Adam, and Jamin Speer. 2020. “Stigma and Take-Up of Labor Market Assistance: Evidence from Two Field Experiment.” Working Paper.
- Ours, Jan C. van, and Milan Vodopivec. 2008. “Does Reducing Unemployment Insurance Generosity Reduce Job Match Quality?” *Journal of Public Economics* 92 (3): 684–95. <https://doi.org/10.1016/j.jpubeco.2007.05.006>.
- Pande, Rohini. 2003. “Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India.” *American Economic Review* 93 (4): 1132–51.
- Parker, Susan W., and Tom Vogl. 2018. “Do Conditional Cash Transfers Improve Economic Outcomes in the Next Generation? Evidence from Mexico.” Working Paper 24303. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w24303>.
- Pop, Lucian Bucur. 2015. “Options for Improving the Targeting of Safety Nets in Ghana.” In *Safety Nets in Africa: Effective Mechanisms to Reach the Poor and Most Vulnerable*, 67–106. Africa Development Forum. The World Bank. [https://doi.org/10.1596/978-1-4648-0435-9\\_ch4](https://doi.org/10.1596/978-1-4648-0435-9_ch4).
- Pop-Eleches, Cristian, and Grigore Pop-Eleches. 2012. “Targeted Government Spending and Political Preferences.” *Quarterly Journal of Political Science* 7 (3): 285–320.
- 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*, Ninth Annual Conference on Econometric Modelling for Africa, School of Economics, University of Cape Town, 2004, 23 (5): 836–53. <https://doi.org/10.1016/j.econmod.2005.10.010>.
- 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 (4): 892–920. <https://doi.org/10.1093/wber/lhaa019>.
- Ralston, Laura, Colin Andrews, and Allan Hsiao. 2017. “The Impacts of Safety Nets in Africa: What Are We Learning?” Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-8255>.
- Ravallion, Martin. 1991. “Reaching the Rural Poor through Public Employment: Arguments, Evidence, and Lessons from South Asia.” *The World Bank Research Observer* 6 (2): 153–75.
- . 2009. “How Relevant Is Targeting to the Success of an Antipoverty Program?” *The World Bank Research Observer* 24 (2): 205–31. <https://doi.org/10.1093/wbro/lkp009>.
- Republic of Turkey. 2012. *Catastrophe Insurance Law No. 6305*. <https://dask.gov.tr/en/law>.

- 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 (3): 593–617. <https://doi.org/10.1007/s10018-017-0207-4>.
- Robinson, Jonathan. 2012. "Limited Insurance within the Household: Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics* 4 (4): 140–64.
- Robles, Marcos, Marcela G. Rubio, and Marco Stampini. 2019. "Have Cash Transfers Succeeded in Reaching the Poor in Latin America and the Caribbean?" *Development Policy Review* 37 (S2): O85–139. <https://doi.org/10.1111/dpr.12365>.
- Rosenzweig, Mark, and Oded Stark. 1989. "Consumption Smoothing, Migration, and Marriage: Evidence from Rural India." *Journal of Political Economy* 97 (4): 905–26.
- Rosenzweig, Mark, and Kenneth I. Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *Journal of Political Economy* 101 (2): 223–44.
- Rosenzweig, Philip. 1998. "Managing the New Global Workforce:: Fostering Diversity, Forging Consistency." *European Management Journal* 16 (6): 644–52. [https://doi.org/10.1016/S0263-2373\(98\)00041-3](https://doi.org/10.1016/S0263-2373(98)00041-3).
- Roy, Shalini, Melissa Hidrobo, John Hoddinott, and Akhter Ahmed. 2019. "Transfers, Behavior Change Communication, and Intimate Partner Violence: Postprogram Evidence from Rural Bangladesh." *The Review of Economics and Statistics* 101 (5): 865–77. [https://doi.org/10.1162/rest\\_a\\_00791](https://doi.org/10.1162/rest_a_00791).
- 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): 1521–45. <https://doi.org/10.1002/jid.3186>.
- Sadka, Joyce, Enrique Seira, and Christopher Woodruff. 2018. "Information and Bargaining through Agents: Experimental Evidence from Mexico's Labor Courts." Working Paper 25137. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w25137>.
- Sadoulet, Elisabeth, Alain de Janvry, and Benjamin Davis. 2001. "Cash Transfer Programs with Income Multipliers: PROCAMPO in Mexico." *World Development* 29 (6): 1043–56. [https://doi.org/10.1016/S0305-750X\(01\)00018-3](https://doi.org/10.1016/S0305-750X(01)00018-3).
- Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *The Review of Economic Studies* 68 (1): 205–29. <https://doi.org/10.1111/1467-937X.00166>.
- . 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *The Quarterly Journal of Economics* 117 (3): 1039–73. <https://doi.org/10.1162/003355302760193959>.
- 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. <https://doi.org/10.1257/jel.50.1.3>.
- Schady, Norbert, and Maria Caridad Araujo. 2006. "Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador." Working Paper 3. Impact Evaluation Series. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-3930>.
- Schanzenbach, Dianne Whitmore. 2009. "Experimental Estimates of the Barriers to Food Stamp Enrollment." Discussion Paper DP1367-09. Institute for Research on Poverty Discussion Paper. Institute for Research on Poverty. <https://www.irp.wisc.edu/resource/experimental-estimates-of-the-barriers-to-food-stamp-enrollment/>.

- Schönberg, Uta, and Johannes Ludsteck. 2014. "Expansions in Maternity Leave Coverage and Mothers' Labor Market Outcomes after Childbirth." *Journal of Labor Economics* 32 (3): 469–505. <https://doi.org/10.1086/675078>.
- Schultz, T. Paul. 2004. "School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program." *Journal of Development Economics*, New Research on Education in Developing Economies, 74 (1): 199–250. <https://doi.org/10.1016/j.jdeveco.2003.12.009>.
- Skoufias, Emmanuel. 2005. "PROGRESA and Its Impacts on the Welfare of Rural Households in Mexico:" 139. *Research Reports*. Research Report. Washington, D.C.: International Food Policy Research Institute (IFPRI). <https://ideas.repec.org/p/fpr/resrep/139.html>.
- Skoufias, Emmanuel, Mishel Unar, and Teresa Gonzalez de Cossío. 2008. "The Impacts of Cash and In-Kind Transfers on Consumption and Labor Supply : Experimental Evidence from Rural Mexico." Working Paper 4778. Policy Research Working Paper. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-4778>.
- . 2013. "The Poverty Impacts of Cash and In-Kind Transfers: Experimental Evidence from Rural Mexico." *Journal of Development Effectiveness* 5 (4): 401–29. <https://doi.org/10.1080/19439342.2013.843578>.
- Stiglitz, Joseph E. 1974. "Incentives and Risk Sharing in Sharecropping." *The Review of Economic Studies* 41 (2): 219–55. <https://doi.org/10.2307/2296714>.
- Stoeffler, Quentin, Bradford Mills, and Carlo del Ninno. 2016. "Reaching the Poor: Cash Transfer Program Targeting in Cameroon." *World Development* 83 (July): 244–63. <https://doi.org/10.1016/j.worlddev.2016.01.012>.
- Strupat, Christoph, and Florian Klohn. 2018. "Crowding out of Solidarity? Public Health Insurance versus Informal Transfer Networks in Ghana." *World Development* 104 (April): 212–21. <https://doi.org/10.1016/j.worlddev.2017.11.004>.
- 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): 672–91. <https://doi.org/10.1093/ajae/aay042>.
- 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): 773–92.
- 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* 25 (December): 650–62. <https://doi.org/10.1002/hec.3291>.
- Todd, Jessica E., 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 (3): 549–81. <https://doi.org/10.1086/658347>.
- Townsend, Robert M. 1994a. "Risk and Insurance in Village India." *Econometrica* 62 (3): 539–91. <https://doi.org/10.2307/2951659>.
- . 1994b. "Risk and Insurance in Village India." *Econometrica* 62 (3): 539–91. <https://doi.org/10.2307/2951659>.
- . 1995. "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies." *The Journal of Economic Perspectives* 9 (3): 83–102.

- Valle, Alejandro del. 2021. “The Effects of Public Health Insurance in Labor Markets with Informal Jobs: Evidence from Mexico.” *Journal of Health Economics* 77 (May): 102454. <https://doi.org/10.1016/j.jhealeco.2021.102454>.
- Valle, Alejandro del, 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): 164–95. <https://doi.org/10.1257/app.20190002>.
- Vu, Khoa, and Paul Glewwe. 2022. “Maternity Benefits Mandate and Women’s Choice of Work in Vietnam.” *World Development* 158 (October): 105964. <https://doi.org/10.1016/j.worlddev.2022.105964>.
- Wager, Stefan, and Susan Athey. 2018. “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Journal of the American Statistical Association* 113 (523): 1228–42. <https://doi.org/10.1080/01621459.2017.1319839>.
- 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): 663–74. <https://doi.org/10.1002/hec.3293>.
- World Bank. 2014. “For Up to 800 Million Rural Poor, a Strong World Bank Commitment to Agriculture.” November 12, 2014. <https://www.worldbank.org/en/news/feature/2014/11/12/for-up-to-800-million-rural-poor-a-strong-world-bank-commitment-to-agriculture>.
- . 2018a. *The State of Social Safety Nets 2018*. Washington, DC: World Bank. <https://doi.org/10.1596/978-1-4648-1254-5>.
- . 2018b. *The State of Social Safety Nets 2018*. Washington, DC: World Bank. <https://doi.org/10.1596/978-1-4648-1254-5>.
- . 2018c. *Women, Business and the Law 2018*. Washington, DC: World Bank. <https://doi.org/10.1596/978-1-4648-1252-1>.
- . 2021a. “GDP per Capita (Constant 2015 US\$).” World Bank DataBank. <https://data.worldbank.org/indicator/NY.GDP.PCAP.KD>.
- . 2021b. “Population, Total.” World Bank DataBank. <https://data.worldbank.org/indicator/SP.POP.TOTL>.
- World Economic Forum. 2019. “GCI 4.0: Insurance Premium.” The World Bank. [https://tcdata360.worldbank.org/indicators/hfc07f7af?country=BRA&indicator=41436&viz=line\\_chart&years=2017,2019](https://tcdata360.worldbank.org/indicators/hfc07f7af?country=BRA&indicator=41436&viz=line_chart&years=2017,2019).
- World Health Organization. 2019. “Global Health Expenditure Database.” <https://apps.who.int/nha/database/Select/Indicators/en>.
- . 2021. “Global Expenditure on Health: Public Spending on the Rise?” <https://www.who.int/publications-detail-redirect/9789240041219>.