

Social Protection in the Developing World¹

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.

¹ 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

Social assistance programs have been rapidly evolving in the developing world. Today, safety net programs cover about 2.5 billion people. More than 120 low- and middle-income countries run cash transfer programs for poor families, and more than 70 of them run social pension programs. 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 programs have increasingly taken the place of broad-based subsidies of basic commodities, promising increased efficiency and an enhanced ability to redistribute income to the poor. The COVID-19 pandemic further expanded these policies, suggesting their increased role in policy in the future (see Gentilini et al. 2021).

Designing these programs for a developing country context, however, entails challenges that differ from those faced in more developed economies. For example, one of the biggest questions in the public economics literature is how to target for social assistance programs. But, unlike in developed countries, information on income is hard to observe, partly because informal work and self-employment is common, resulting in the bottom 70 to 90 percent of households in many low-income nations often being excluded from the income tax system (Jensen 2022). This means that the standard frameworks used for thinking about targeting benefits – e.g. income eligibility thresholds, endogenous labor responses to effective marginal tax rates (e.g. Mirrlees 1971; Saez 2001) – do not readily apply in a developing country setting. This has led to the rise of other targeting approaches in developing countries, which in turn, has necessitated research designed specifically to understand these contexts (e.g., Alatas et al. 2012; Hanna and Karlan 2017).

Similarly, program design may also need to be radically different in a developing country context. For example, one 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 push households into a higher steady-state income level. While poverty traps could occur in more developed settings as well, the particular nature of the interventions to relieve these shocks may differ: an intervention that provides productive livestock may make sense in low-income countries given that 63 percent of the population is engaged in agriculture (ILOSTAT 2019) —but would make less sense in a developed nation where the vast majority of the poor and near-poor are not in agriculture.

This paper thus serves three distinct, but related, goals. First, we outline some theoretical frameworks for understanding the social protection challenges in developing countries. We focus both on redistributive goals and social insurance programs that help to insure against shocks. We begin by outlining the ‘*ex-post*’ problem; that is, after types have been realized, how should the social planner think about improving welfare through redistribution? 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, outlining 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, 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. 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, 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; 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, considering 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 earning ability, and therefore bases their decisions on realized earnings. This then introduces a reason for high earners to claim lower earnings potential and undersupply effort (Mirrlees 1971).²

As highlighted by Saez (2001), this approach identifies the elasticity of labor supply with respect to taxes and with respect to income as the key parameters in the design of welfare programs. There is a large literature on this point for developed countries (e.g. Saez, Slemrod, and Giertz 2012; Meghir and Phillips 2010; Currie and Gahvari 2008) motivated partly by the fact that many US welfare programs (especially prior to the 1996 welfare reform bill) implicitly put a high tax rate on earnings through a steep phasing out of benefits.

The available literature suggests that labor supply distortions might not be as relevant for developing countries. Two recent experimental studies of the wage elasticity, both from Malawi, report low values of 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 and find no evidence of income effects on work. Baird, McKenzie, and Özler's (2018) systematic literature review comes to a similar conclusion. The fact that both income effects

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

and the wage elasticity are small suggests that it is less important to focus on the endogeneity of earnings in designing transfer schemes in developing countries.

On the other hand, in the developing world, governments face challenges observing earnings for informal workers or those with relatively low earnings (Glewwe 1992; Jensen 2022). Most governments therefore worry more about the endogeneity of *reported* earnings. That is, they worry that if program eligibility was based solely on self-reported income, many non-poor people would simply report a low income to claim benefits.

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; asset ownership), self-selection, and community information, or some combination of these. 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. The idea of using ‘tags’ forms the theoretical basis of what is known as proxy-means testing, in which assets and other observable characteristics are used to predict people’s incomes for targeting 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 without inducing distortions in consumption or work. However, in order to better predict incomes, governments also use observable characteristics that are not completely exogenous (e.g. does the house have a tile roof, motorcycle ownership), either because people can change their decisions, or perhaps because they can conceal them from government auditors at some cost. 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 is mistargeting: targeting in such a system is based on the statistical model relating the tags to incomes. This model, like all statistical models, is imperfect, and the data on the tags is often noisy, leading to both inclusion 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, most rich countries the data use frequently updated administrative data). 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. The next section presents a simple model that highlights these issues and their consequences.

2.3: A model of targeting

We present a simple model to think through a number of tradeoffs in the design of redistribution schemes in developing countries. Our model accounts for two features of developing countries: first, incomes are hard to measure precisely since much work is in the informal sector (Jensen

2022). Second, as discussed above, labor supply elasticities tend to be low. Unlike in traditional models (Saez 2002), where the tradeoffs are primarily driven by the labor supply response, we will assume that labor supply is fixed and unaffected by redistribution. Instead, the key driver in our model is the government’s ability to learn people’s true incomes in a low-information environment, potentially with noise, in an environment where people can potentially mis-report their incomes when asked to do so.³

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 above some cut off, \hat{y} .⁴ Let $h(y)$ represent the density of the true underlying income distribution of those who might be eligible for the program, so $\int^{\hat{y}} h(y) = 1$.

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 $0 \leq t \leq 1$.⁵

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 audit data is noisy, so $y^a = y + \epsilon$, where ϵ has expectation of zero and a fixed finite support. We study the choice of α , i.e., how much weight to put on self-reported income vs. the noisy audit data.

If people mis-report incomes, there is a penalty $F = \frac{a}{2}(y^a - \tilde{y})^2$, based on the difference between audited income y^a and self-reported income. The penalty does not contribute to the program budget—so it could be a psychic cost or a bribe to a government official.

We follow Saez (2002) to introduce the government’s preference for redistributing to the most

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

⁴ 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.

⁵ We allow for those close to \hat{y} to pay into the system, since T can be zero and \hat{y} is allowed to be more than y^* . Nothing essential would change if instead we imposed the constraint that $T + t(\hat{y} - y^p) \geq 0$.

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 $\int^{y^*} h(y)g(y) = 1$. In addition, we allow for the government to care about equity *among those who have the same true income* (that is, we allow for 'horizontal equity' considerations). We introduce this by assuming that the government's utility function for those who have a post-transfer income of z is $u(z)$, where u is a (potentially strictly) concave function.

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 does not take into account the beneficiary's disutility from misrepresenting their income (i.e. does not include $F = \frac{a}{2}(y^a - \tilde{y})^2$). This simplifies our analysis while being fairly realistic—if 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.⁶

The welfare recipient's problem

Individual utility is given by

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

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

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

which tells us that

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

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 below). Under this assumption, and noting that self-reported income is $\tilde{y} = y - \frac{at}{a}$, 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

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

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

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

⁶ 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); see the Online Appendix.

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

$$W'(t) = \int^{\hat{y}} g(y)h(y) (\bar{y} - y)E_{\epsilon}[u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)]dy \\ - \int^{\hat{y}} g(y)h(y) (1 - \alpha)E_{\epsilon}[\epsilon u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)]dy \quad (6)$$

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.⁷ Offsetting this, however, the second term in $W'(t)$ captures the welfare losses from the fact that greater targeting of benefits to the poor creates more random variation in the transfers, as a result of the noise (ϵ) in the government's audit function. This noise reduces social welfare so long as there is some weight on audits (i.e. $\alpha < 1$) and the government cares about horizontal equity ($u'' < 0$), and if so, then the second term is negative.⁸

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 $W'(0) > 0$.⁹ 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, 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 social welfare weight ($g'(y) < 0$), $W'(t) > 0$ 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.

Third, from the expression for $W(t)$ in equation (4), it is clear that more noise in the audit process

⁷ 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.

⁸ To see this mathematically, note that $u'' < 0$ implies that u' is greater when ϵ is positive than when it is negative. Thus, the terms where ϵ 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.

⁹ 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.

(a mean-preserving spread of ϵ) 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 α , the weight the government places on self-reported data, increases, the noise in the audit data matters less, and social welfare is higher, so $\frac{\partial W}{\partial \alpha} > 0$.¹⁰ 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 α 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 ($T = 0$). The optimal scheme relies entirely on self-reports ($\alpha = 1$). Moreover, this scheme delivers the first best level of targeting.

This stark result has a simple intuition. As long as all potential beneficiaries have the same cost of misreporting incomes, the amount by which they distort is predictable, and the optimal targeting mechanism can take that fully into account and 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.

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 μ 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

¹⁰ Differentiating (4) with respect to α yields $\frac{\partial W}{\partial \alpha} = \int^y g(y)h(y) E_\epsilon(t\epsilon)u'(y + t(\bar{y} - y) + B - t(1 - \alpha)\epsilon)dy$. This is positive by the same argument given in footnote 7.

assumptions, we can rewrite equation (3) as

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

The government's budget constraint is given by

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

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

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

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

When $A > 0$, there is now an additional effect: 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 can no longer fully unravel misreports, which helps those who are willing to mis-report more.

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 μ 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 α), 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 ϵ 's in the audit process. The net effect of lowering α depends on the relative sizes of these effects. In the limit, holding other parameters fixed, as the audit process becomes better and better ($Var(\epsilon) \rightarrow 0$), 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 before. 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 α^2 , for α small enough, raising α has a second order negative effect on $W(t)$ but first order positive effect by reducing the cost of targeting noise ϵ . Hence in this model one will not completely ignore self-reports, though it can clearly now be optimal to combine self-reports with audit data.¹¹

Turning next to t , as before, a higher t still means more transfers go to the poor. But, unlike in the previous model, 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 $\alpha = 1$, there are also losses due to targeting noise ϵ , and these are also stronger when t is higher. Hence relative to the case where $A = 0$, as in Section 2.3.1, the optimal value of t may be lower.¹² 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 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 μ 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 and using the fact that $\mu(g(y_1) - 1) =$

¹¹ In practice, if there is a fixed cost of collecting the self-report data, and if the optimal α is very small absent these fixed costs, the government may not bother paying the fixed costs and may choose to ignore self-reports altogether.

¹² 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 .

$-(1 - \mu)(g(y_2) - 1)$, we get

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

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 α whenever $a_2 > a_1$. Setting $\alpha = 0$, it is also clear that one should set t as large as possible and set $T = 0$. 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 occur if those with higher incomes have a lower misreporting cost.

2.3.4: Summary

Our model, though stylized, helps highlight the tradeoffs in designing a transfer system in the presence of noisy data. The model highlights that the key challenge is heterogeneity in misreporting costs, which prevents the government from ‘backing out’ people’s true incomes from their self-reports. Given this, even the optimal transfer scheme will typically involve both inclusion and exclusion errors. The fact that the government responds to the imperfections in self-reported data by using noisy audit data generates additional 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 some 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 ($Var(\epsilon) > 0$).

2.4: Targeting on Observable Characteristics in Practice

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 (Baker and Grosh 1994; Elbers et al. 2007). To improve prediction, 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). This approach produces less accurate estimates than individual targeting, but is much less data intensive.

A more ambitious form of targeting is the **proxy means test** (PMT). The 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, etc.), as well as the characteristics (X')—usually demographics and assets—that will be used for targeting. This training dataset is usually a nationally representative household survey that is collected for research or general statistical purposes; as such, households typically have no reason to lie about their actual consumption or income. Using these data, one then estimates a model of the form $y_i = X' \beta + \epsilon_i$.

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

This type of proxy-means testing was first used in the 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 proxy-means tests include the SISBEN in Colombia, SISFOH in Peru, Listahanan in the Philippines, BISP in Pakistan and Takaful and Karama in Egypt.

The way proxy-means tests are implemented connects to the challenges highlighted by the model in section 2.3. 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. They show that the prediction equation has substantial predictive power—the R^2 is between 0.53 and 0.66—but it nonetheless has errors of both inclusion and exclusion. Other papers show similar findings (see, e.g., Brown, Ravallion, and van de Walle 2018).

There is also no reason to think that the relationship between underlying income and the types of assets used in the PMT algorithm is the same for everyone. The model used to predict PMT scores, makes no allowance for differences in preferences. Someone might end up in a bigger house than what would normally correspond with their income because, for example, they feel obligated to provide temporary housing to their extended family. As in the model, such heterogeneity in the link between true income and PMT-predicted 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. However, most papers find little to no improvements using these techniques relative to the common methods currently used (e.g. McBride and Nichols 2018; Baez, Kshirsagar, and Skoufias 2020; Areias and Wai-Poi 2022).

Second, one can use new types of administrative data rather than collecting a door-to-door census, which allows more frequent, and hence up-to-date, data to be used. 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 in Kenya’s GiveDirectly program. Using data from Togo, Aiken et al. (2022) find that using mobile phone metadata for targeting has substantive predictive power, though it is not as accurate as a more traditional full PMT. The fact that administrative data such as this is continually and automatically updated suggests that targeting using these types of administrative data may be more responsive to shocks.¹³ In fact, some countries, from Pakistan to Togo, used this approach to target assistance during the 2020 COVID-19 crisis (Gentilini et al. 2020); understanding whether this was effective is an important area of ongoing and future work.

Endogenous characteristics?

A further challenge 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. For example, after England imposed a “window tax” in 1696—a tax levied on the number of windows a house had, which proxied for wealth—people built new houses with very few windows, leading to many dark houses (Oates and Schwab 2015). Even geography is not an immutable characteristic: households could choose where to live in response to geographic targeting incentives. 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 targeting. To test for endogenous distortions, the statistics bureau randomly added questions on flat-screen televisions and/or the number of cell phone SIM cards owned to the targeting census in some randomly selected provinces, but not in others. While self-reports of television ownership fell in treated provinces six months later, 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 et al. (2020) result had hints of this; reported ownership of televisions declined in government data 6 months later, even though actual television ownership was unchanged. Similarly, Martinelli and Parker (2009) examine self-reported versus enumerator-verified PMT data in Mexico, and find more under-reporting of assets in the self-reported data compared to the enumerator-verified data (that is, $y^a - \tilde{y}$) when the monetary stakes are higher.

Beyond this, there is a second, important real-world consideration. Specifically, in the model, we

¹³ Others (e.g. Baez, Kshirsagar, and Skoufias 2020; del Ninno and Mills 2015) suggest augmenting PMT targeting system with remote sensed and/or administrative data of this sort to help capture shocks.

assume that the household knows exactly what to do to claim the benefits. In practice, the fact that there are many different variables, interacting in complex ways in the PMT formula means that households have limited control over the outcome and therefore the returns to manipulation are small. In fact, governments typically keep the formulas secret for this reason: studying the proxy-means test in Colombia, Camacho and Conover (2011) find that local politicians were only able to manipulate the data after the formula became public in order sign up more of their constituents for federal government transfers.¹⁴

Finally, in contrast to PMTs, there is some evidence that geographic targeting can lead to real distortions. In particular, place-based policies—which in contrast to PMTs are easy to understand—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.

2.5: Self-selection mechanisms

Another targeting approach is a **self-selection mechanism**. The idea is to pair benefits with an action, where the required action is relatively less costly for the targeted population (the poor) than for the rest (the non-poor) (Nichols and Zeckhauser 1982). If so, then as in the model, the selection rule allows the government to make an inference about people’s income (perhaps noisily) from their self-chosen actions. The required action can include buying low quality food, work requirements, and administrative ordeals (e.g. standing in line, filling out lengthy paperwork or 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 are nevertheless better off even once these costs are taken into consideration.

Workfare. One of the most common ordeal mechanisms is workfare—i.e. requiring beneficiaries to work in exchange for payments. These programs have a long history, dating back at least to the 19th century in England, where transfers were granted through residence in a workhouse (Besley and Coate 1992), and to the United States’s Civilian Conservation Corps and Works Progress Administration in the 1930s (Aizer et al. 2020).

Workfare programs are a classic example of an ordeal mechanism because work takes time. If the wages for the workfare job are low, or the tasks unpleasant or difficult, this will generate self-selection: those who can get a more attractive job will select out (Ravallion 1991; Besley and Coate 1992). But, if the wage is set too high, the program runs the risk of *crowding out* more productive private sector work. The net efficiency consequences also depend on whether the work being done is productive (in the model, do we add back the value of the work to the social welfare function). Using workfare to build roads may have little efficiency loss; having workfare labor dig ditches

¹⁴ Though there is limited evidence so far of actual manipulation of consumption, one could imagine cases in which it may be. If so, 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).

no one needs would also generate self-selection, but at much larger social efficiency costs.

Perhaps the largest workfare program in the world is India's 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'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 substantial self-selection in program take-up.

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 cash transfer would result in modestly higher poverty reductions than the workfare program, given both the government's costs to run the program and these forgone wages. Similarly, Bertrand et al. (2021) experimentally examine related questions in the context of a youth workfare program in urban Cote d'Ivoire. They find that while the program led to earnings gains, earning gains are only 53 percent of the transfer due to the crowd-out of private employment. Both programs set the workfare wage at the formal minimum wage, which—unlike in many developed country settings—may be *above* the prevailing wage, given the large informal sector. This suggests that workfare programs may have more effective targeting properties in places 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—e.g. working a full day of hard labor under the hot sun. Is it possible to get the gains from self-selection without imposing such large costs?

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

Alatas et al. (2016) develop a model that captures this idea, and then experimentally test a self-targeting mechanism (application with verification done by PMT) against automatic enrollment based on a PMT done by census enumerators within the context of the expansion of the government's real conditional cash transfer program. They found that self-targeting led to dramatically poorer beneficiaries, driven in large part by those who are wealthier on dimensions that would be missed by the PMT being less likely to apply. Surprisingly, this approach *also* reduces exclusion error compared to the PMT—while the government makes its best efforts 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 apply (see Appendix Figure 1. This suggests that self-selection at the application stage can allow for the revelation of important information, potentially decreasing both inclusion and exclusion error.

2.6: Community-information based targeting

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

To examine this, Alatas et al. (2012) experimentally compared proxy-means test-based targeting to community-based targeting. 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 below PPP \$2/day, although the households whose identification switched were sufficiently close to the poverty line that both approaches would perform similarly in terms of most social welfare functions. 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.¹⁵

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

Empirically, however, elite capture seems to be small, at least in comparison to other sources of targeting failures. In a field experiment in Indonesia, Alatas et al. (2019) find that increasing elite decision making over targeting does not substantially change outcomes. More generally, while they do find some evidence of elite preference in some programs, it is small in magnitude, so that the gains one would achieve by reducing elite capture entirely are dwarfed by the gains one could make by improving the quality of PMT data.¹⁶ Similarly, Basurto, Dupas, and Robinson (2020) study the targeting of subsidies for both agricultural inputs and food in Malawi, which is done by local chiefs, and compare it to a hypothetical targeting done through a PMT. While they find that nepotism exists, the welfare consequences are small, again because who receive preferential treatment are also relatively poor; plus there are potential gains as chiefs use their local knowledge to target those recently hit by shocks, and in the case of agricultural inputs, to target them to those with the higher returns.

¹⁵ 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.

¹⁶ The welfare losses from elite capture are essentially the *product* of a) how much richer elites are (9 percent), b) how much more likely elites are to obtain benefits (19 percent in the worst program considered), and c) elites' proportion in the general population. Multiplying yields a welfare cost less than 1 percent.

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 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 than those selected by community-based targeting. 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, and at identifying households with recent shocks.

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 (e.g. Sabates-Wheeler, Hurrell, and Devereux 2015).

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) found that households that were randomized to be ranked early in the process are ranked more accurately than those ranked later, suggesting that the community finds the process tiring. This suggests that the micro-structure of the algorithm used to elicit the community ranking can affect the accuracy of the overall community ranking, and understanding this better can help improve the information elicitation process. Second, does the efficacy of community targeting differ systematically? For example, Alatas et al. (2016), for example, finds that the community targeting treatment in Alatas et al. (2012) was more effective at identifying households with in areas with more diffusive interpersonal network structures. 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?

2.7.1. How prevalent are take-up challenges?

The fact that not everyone takes up programs is not, per se, surprising. Self-selection costs vary across people—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 (Bhattacharya et al. 2015) found that two-thirds of women eligible for a non-contributory social pension for life are not enrolled in the program. The generosity of the pension (the equivalent of 70% of median per-capita consumption across urban areas) makes it unlikely that 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.

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; they find that information about benefits nearly doubles take up (from a low base), and when combined with help with signing up 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. These studies typically find that while take-up assistance can meaningfully increase program take-up compared to not having assistance, it does not close the majority of the take-up gap.

For example, Carneiro, Galasso, and Ginja (2019) study a program known as Chile Solidario, that targeted extremely poor households that were not enrolled in government assistance programs. They provided extensive home visits by a social worker over a two-year period, including guidance on how to apply for government programs, as well as a financial stipend. The program itself had imperfect take-up—only about 20 percent of those eligible for Chile Solidario enrolled. For those who availed themselves of Chile Solidario assistance, 4 years later take-up of a separate government family child allowance program (“SUF”) had increased 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 behind 80 percent of eligible women.

Banerjee, et al. (2021) examine related questions in Indonesia in the context of the government’s national health insurance scheme. 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 despite being officially mandatory. Banerjee, et al. (2021) focused on the uninsured and examined the role of intensive application assistance, as well as subsidies. 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—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.

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 application requirements—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 a costly trip to a government office. Improving this underlying administrative capacity to allow for simpler enrollment procedures may be challenging, but important.¹⁷

2.7.3. Stigma

A common refrain in the developed world is that low take-up reflects low demand for government services due to the stigma associated with receiving benefits (Moffitt 1983). While there is comparatively little work investigating stigma in developing countries, the research that there is suggests less concern about stigma. For example, 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. 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. Nevertheless, understanding

¹⁷ It is also important to note that improving takeup on the margin does not automatically improve welfare” if households are fully optimizing and rationally not signing up for programs given the costs of doing so (Finkelstein and Notowidigdo 2019). However, this result depends on the fact that households are fully optimizing, which is often not the case in many contexts.

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 primarily on identifying the poor and assigning programs 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 effect on income is heterogeneous by individuals, i.e. we write dy_i . Now, optimal targeting would want to find those individuals for whom $g(y_i)u'(y_i)dy_i$ is highest, which in this particular formulation is a combination of the individual specific treatment effect and income. Even for programs that do not target income, the point is that one is no longer interested in targeting the program just based on an observable characteristic (i.e. poverty 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. 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., Wager and Athey 2018; Chernozhukov et al. 2018). Several papers have used these techniques to show how they can identify heterogeneous treatment effects to improve the targeting of immunizations (Chernozhukov et al. 2018), workfare (Bertrand et al. 2021), and cash transfers (Haushofer et al. 2022). Caria et al. (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, for job training programs for Syrian refugees in Jordan.

A second option is to encourage self-selection on the basis of treatment effects. Dupas et al. (2016), for example, investigate a program that provides free chlorine in Kenya, 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 et al. 2021). We regard targeting on treatment effects as an important dimension

for future research.

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

All of the targeting methods discussed thus far suffer from exclusion and inclusion errors (Brown, Ravallion, and van de Walle 2018). Given these challenges, another option is to make programs universal, for example universal basic income (UBI) programs or a universal in-kind program, where 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 abuse of power (Banerjee, Niehaus, and Suri 2019).

One challenge, however, is that for a given budget, universal programs 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 severe 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 in developing countries (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 (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, 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.¹⁸

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.

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, 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 stigma matters 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, however, there is both a lot of discussion about the form the transfer should take, and how the household makes use of the benefits. We will argue that under a set of quite standard conditions, the expression for the social value of a transfer is actually very simple and can clarify a number of vexing issues.

Assume that the utility function of a household is given by $u(c_1, \dots, c_n, b, l)$ 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, t is the tax/transfer rate and y^* is the poverty line. This captures a whole range

¹⁸ In general, a very important topic for future research is the political economy of social protection programs. Given space constraints, we could not do it justice within this paper. However, we refer the interested read to the online appendix for an overview of this literature.

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. $\frac{\partial^2 u}{\partial l \partial b} < 0$).

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 how they are affected by b .¹⁹ Denote the indirect utility function generated by the household's constrained utility maximization by $u(y, b, t)$. We assume no spillovers so that no one else's utility is affected by the benefits going to a particular household.

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

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

Under these assumptions, the social benefit from a household with income y getting a small increase in b is given by $g(y)h(y)\frac{du}{db}db = g(y)h(y)\frac{du}{dy}\frac{db}{du}db$. The ratio $\frac{db}{du}$ 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) is the key ratio from the government's perspective. Even if their labor supply does change, the utility consequences of this are second-order by the envelope theorem, so this does not change expression (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.²⁰

¹⁹ While the government can vary b by income level, i.e. b_i , we suppress this in the exposition to simplify notation.

²⁰ 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

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. 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.²¹ What also matters is the program’s governance and administration (discussed below), i.e. did the poor actually get their full cash transfer or did some get lost to corruption. But, how beneficiaries choose to spend \$1 of benefits—on nicer food or a new television—is to first-order not welfare relevant.

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

(MVPF) discussed by Hendren and Sprung-Keyser (2020), but augmented by $g(y)MU(y)$ to take distributional consequences into account.

²¹ Ravallion (2009) notes, in studying 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. $\int h(y)b_i g(y_i)MU(y_i)$.

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 1.²² 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. If one cannot estimate willingness-to-pay directly, one can try to impute it.²³

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 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 (Browning and Chiappori 1998). And even if a household where the husband makes all decisions ignoring his wife's preferences has a well-defined utility function, the social planner may not choose to put that particular utility function in the social welfare function. Instead, it may base its decisions about the choice of benefits on some other preferences, that, for example, give equal weight to both spouses, 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

²² It is also possible that even for some cash transfers \$1 may not actually be \$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.

²³ For example, Hendren and Sprung-Keyser (2020) give examples of how to impute willingness-to-pay for 133 programs in the United States, arguing that the willingness-to-pay for a program is equal to the change in N.P.V. income from receiving the program. (Note that this is change in income holding labor supply constant, because if the labor/leisure decision changes, that is a consumption change like any other consumption choice, as discussed above.) This approach presumes that there are no direct utility benefits or costs of the program. For example, if beneficiaries have to wait in long lines to access a transfer, 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.

recipients do not spend cash assistance on ‘temptation goods’ such as alcohol or cigarettes (Currie and Gahvari 2008). More generally one could imagine that voters would prefer a household receiving cash assistance bought food rather than a television, even if television generates 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 by 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.

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, which affects the measurement of 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 properly 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, and this is what one would find if one elicits willingness-to-pay directly.²⁴

²⁴ One could in principle design experiments to estimate willingness-to-pay where the credit constraint does not bind—for example by enrolling them in a lottery with a small chance of winning enough to cover any reasonable value of their willingness-to-pay and asking them to commit to a willingness-to-pay conditional on winning the lottery—but this is not the usual approach taken to measure willingness-to-pay.

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 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.

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 in the previous section: affecting future income above and beyond the transfer itself (i.e. by raising wages, or allowing the household to overcome credit constraints); affecting future taxpaying decisions; changing *intra*-household allocations, e.g., between spouses, or between parents and children; changing patterns of consumption in ways a social planner may care about, even if the household does not; spillovers to other households, both pecuniary and otherwise; and affecting welfare by allowing households to overcome information constraints, credit constraints, or other frictions.

Note that we are not attempting a systematic meta-analysis of all papers; we instead refer interested readers to several recent reviews, such as Bastagli et al. (2016) and Ralston, Andrews, and Hsiao (2017). Instead, our aim is to capture key themes from the literature, to relate it 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. As described above, from a welfare perspective, the basic question of an unconditional cash transfer program is who receives the benefits. This is because a household's willingness-to-pay for \$1 in transfers should be \$1 (subject to the caveats above), with the social value of this inflated by a measure of $g(y)MU(y)$. To first order, the *key* thing for evaluating the impact of unconditional cash transfers is therefore the targeting of the transfers, which affects $g(y)MU(y)$.

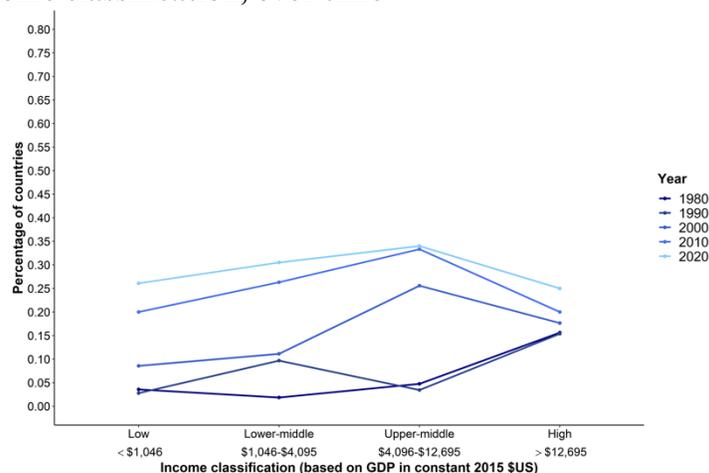
Figure 1 plots the percentage of countries at different income levels with UCTs covering at least

1 percent of the population for countries at different levels of real constant per-capita GDP. The figure shows that at the beginning of the period we study (1980) such programs were essentially only found in the wealthiest countries. By the end of the period, UCTs have become much more common—and almost equally likely to be found in rich or poor countries alike.

Several randomized studies directly evaluate the impacts of UCTs, such as the Haushofer and Shapiro (2016; 2018) and Egger et al. (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; Bastagli et al. (2016) provide a systematic review of 201 cash transfer studies (25 percent of which are UCTs).

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

Figure 1: 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 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-

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. There is less systematic attention to the targeting of these programs; the framework above suggests this is an important direction for future work.

Effects on adult labor supply. As discussed 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).

Effects on consumption of so-called ‘temptation goods.’ A common concern about cash transfers is that people will buy ‘temptation goods,’ such as alcohol and tobacco. However, the same arguments that apply to labor supply apply here: to the extent that governments’ utility function is a function of individual household utilities, and households are rationally choosing to consume more of these goods, this should not matter. Nonetheless, voters appear to be less enthusiastic about funding transfer programs 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 these goods. A systematic review by Evans and Popova (2016) examines 19 cash transfer studies (a mix of UCTs and CCTs), however, and finds no evidence that cash transfers increase expenditures on these goods.

Effects on investment and future incomes. The value of the cash transfer may understate its true welfare effects in the presence of credit constraints (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 the poor are credit constrained (e.g., Banerjee and Duflo 2010), these effects may be substantial.

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.²⁵ They then compare households who were treated 18 months earlier rather than later, and find that those treated earlier have higher incomes, even 4 years later. They interpret this 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. One approach is to assume households are optimizing, and value a transfer of \$1 at \$1. If so, the fact that the investments increase future incomes is already included in the \$1, by the envelope theorem (i.e. households were already optimizing), and 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 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 12 months after the transfer. The net welfare effect is therefore given by $0.74 + PDV(0.016 \text{ per month})$. For example, with a 10% annual discount rate, and assuming the effects started 1 year after the transfer

²⁵ Oportunidades is a conditional cash transfer, as are several other papers that we discuss in this section. Here, we discuss aspects of these programs related to just the cash transfer component; we discuss conditionality below.

and until year 10, the net welfare effect of this transfer would be \$1.8 instead of \$1. While the precise numbers depend on the assumptions, the point is that, if one considers these effects separately from a welfare perspective, investment effects can make a big difference.

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

Informal insurance: One type of spillover effect is through informal insurance arrangements in the spirit of Townsend (1994) and Ligon, Thomas, and Worrall (2002). If households are implicitly insuring one another, this could mean that if one household receives a transfer, but another household does not, they may share the proceeds of the transfer to some extent. Angelucci and De Giorgi (2009), for example, document these kinds of spillover effects in Mexico's Progresa program. 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 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.²⁶

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

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

Egger et al. (2019) use the fact that the transfer program they study was randomized at the village level to estimate a village level 'fiscal multiplier' of this sort. They find that each dollar transferred leads to 2.5-2.8 total increase in local incomes in those villages due to this local multiplier effect.

²⁶ 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.

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. 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, the transfers were paid for by increased taxes, the taxes *also* could have 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.

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 come back to these issues below.

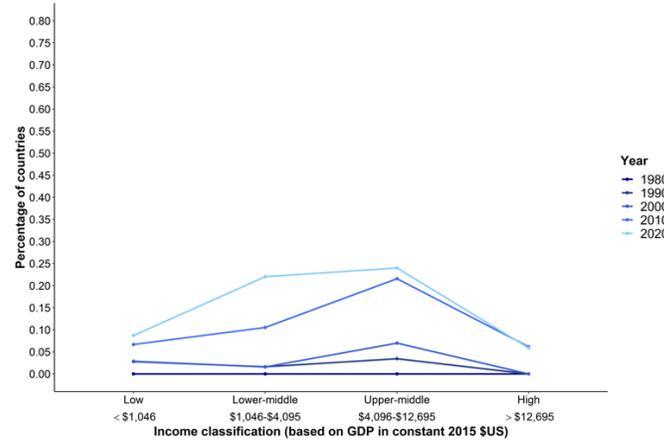
3.2.2: Conditional Cash Transfers and Intergenerational Investments

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

CCTs were pioneered in 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 CCTs (World Bank 2018a). Figure 2 shows the percentage of countries with CCT programs (i.e. covering at least 1 percent of the population), by quintile of GDP per capita (in constant 2015 \$US). CCTs went from being essentially non-existent in 1980, to between 10% and 35% of countries in the bottom four GDP per capita quintiles having CCTs by 2020. Unlike with UCTs shown in Figure 1, however, these programs are less common in the wealthiest countries.

The theory behind these programs is that the cash transfers provide assistance to the current generation, while a combination of cash and the requirements for human capital incentives aim to

Figure 2: 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

break the inter-generational cycle of poverty.²⁷ 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 ones, which received the program about 18 months later (Skoufias 2005). This phase-in design allows researchers to understand if a CCT led to human capital investments in health (Gertler 2004), education (Schultz 2004), and other outcomes (Skoufias 2005). When CCTs were instituted in other countries, a number of them were experimentally evaluated during early stages. There are now randomized evaluations of CCT programs in Colombia (Barrera-Osorio, Linden, and Saavedra 2019), Honduras (Benedetti, Ibarrarán, 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 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

²⁷ Santiago Levy, who helped create the CCT model with the Mexican PROGRESA program, 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).

(Schultz 2004; Todd and Winters 2011; Baird et al. 2014); and cognitive outcomes for children (Macours, Schady, and Vakis 2012).²⁸ They also show reductions in child labor (Barrera-Osorio et al. 2011; Schady and Araujo 2006; Benedetti, Ibararán, and McEwan 2016).

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

A second empirical strategy uses the fact that some programs started earlier in some places than others. Behrman, Parker, and Todd (2011) use this strategy, comparing those who received an ‘extra’ 18 months of Mexico’s Progresa program, after 5.5 years of exposure in treatment areas, as well as comparing these areas to additional never-treated areas using a matching design. They find that increases in schooling attainment for those who received the program earlier.²⁹

Barham, Macours, and Maluccio (2013; 2018a; 2018b) use a related design in Nicaragua, using a randomized design and comparing early treatment locations with late treatment locations. In Barham, Macours, and Maluccio (2013), 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 compare girls treated at ages 9-12 with those treated at ages 11-14, and examine outcomes when they were 19-22 years old. They find girls with early exposure to the program had higher educational attainment, higher earnings, and lower fertility compared to late treatment girls. In Barham, Macours, and Maluccio (2018b), they examine similar aged boys and find similar education and labor market effects.

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

A third empirical strategy is to use the general phase-in of programs, coupled with the differential ages of children when the CCT expands to their municipality. The variation in when the CCT enters generates a difference-in-difference in total exposure to the CCT, while controlling for

²⁸ Parker and Todd (2017) provide an in-depth discussion of the many studies of the Progresa program.

²⁹ Using the same identification strategy, Gertler, Martinez, and Rubio-Codina (2012) find that the Progresa early treatment groups have higher consumption than those that were treated later.

age.³⁰ Parker and Vogl (2018) use this strategy to analyze also study Progresa. They find that exposure to Mexico's Progresa in primary school leads to higher educational attainment, greater labor force participation, more 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.

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

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, even without the incentives. Do the conditions themselves matter? Several studies have found evidence that households respond to conditionality on the targeted outcomes. Studying random variation in the warnings for violating Brazil's Bolsa Familia's conditions on school attendance, Brollo, Kaufmann, and La Ferrara (2017) find that child attendance increases in response to the threat of penalties. Similarly, in Nicaragua, Macours, Schady, and Vakis (2012) argue that a treatment arm that gave additional cash over the basic CCT led did not to additional improvements in child outcomes, suggesting that the conditions drove the effects in the basic treatment.

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

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

Combined, these studies suggest that conditions may matter for child human capital outcomes, but, that there may be a downside from cutting off assistance to non-compliers. This presents the intriguing possibility that a *labeled* transfer program—where households are told that there are conditions, but they are not rigorously enforced—could achieve the child impacts results that could come from conditions, without the downsides. Such an approach could also save extra monitoring costs. Benhassine et al. (2015) suggests this may be the case: in Morocco, they compare a labelled cash transfer with a CCT, and find that the labeled transfer substantially reduced the drop-out rate

³⁰ 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.

and increased school enrollment among those who had dropped out at baseline, whereas the traditional conditional cash transfer had smaller positive effects. More research is needed on labelling, to better understand when it works, how long it persists, and how to structure the labels.

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—after all, households can choose to use the cash to purchase whatever will increase their utility the most. Moreover, for households that would anyways purchase more of the in-kind good than is provided by the transfer, providing in-kind goods should be the same as cash, since households can just reallocate their cash elsewhere.

However, theory provides some reasons why policy makers may prefer in-kind. The first has to do with the types of wedges discussed in Section 3.1.2. If, for example, if the social decision maker cares explicitly about certain types of ‘merit goods’, such as nutrition or health care, rather than just household utility, the social decision maker may explicitly want to influence the beneficiaries’ consumption basket. 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 programs interact with prices. There are two issues. First, in-kind transfers often 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 (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, these adjustments may not happen.

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

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

Experimental comparisons of in-kind, voucher, and cash programs. Several studies examine these issues by experimentally allocating areas into cash, voucher, or in-kind distribution programs. In Mexico, for example, the government ran an RCT that randomized villages into receiving in-kind food transfers, approximately equivalent cash transfers, or a control group. Several studies have

examined this experiment. Overall, Skoufias, Unar, and Cossío (2008) find that both types of transfers have similar effects on food consumption, total consumption, male labor supply, and poverty outcomes; Cunha (2014) also finds similar consumption effects of in-kind and cash programs.

The experiment, however, 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) 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.³¹

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

Third, Banerjee et al. (2021) examine an experiment that compared an in-kind rice program with an electronic voucher program that allowed recipient households to purchase an approximately equivalent value of rice and eggs from private providers. Importantly, the Banerjee et al. (2021) study is an at-scale experiment: 105 districts, with a combined 3.4 million beneficiary households, were included, with both programs implemented by the government bureaucracy as usual. They find a dramatic difference: e-voucher programs delivered concentrated assistance among targeted households, whereas in-kind aid was spread around much more widely. 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 more reliably administered in low state capacity settings, ensuring that the program on the ground looks like it was conceived in theory.

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,

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

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 particularly in very rural, isolated areas where supply is inelastic or non-competitive.

3.2.4: Workfare programs

In workfare programs, the government provides public employment, typically at a low wage, for those who want it. As we discussed earlier, these programs were quite common in the United States during the New Deal era (e.g. the Civilian Conservation Corps and Works Progress Administration), but are also common today. For example, India's Mahatma Gandhi National Rural Employment Act (MGNREGA) 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.

In these programs, the government is directly intervening in the labor market, which has several implications. First, in addition to differences in workers' preferences (i.e. how much they prefer a workfare job compared to their alternative private sector job), one also needs to consider the productivity of the work. If the work is productive (i.e. building public infrastructure with the same efficiency as 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), for example, use this strategy to examine the impacts of MGNREGA on the labor market. Because MGNREGA pays the statutory minimum wage, the MGNREGA wage is above the prevailing wage in many cases. Perhaps as a result, they find substantial crowd-out of private employment. On the other hand, consistent with MGNREGA paying a higher wage, and this having overall equilibrium labor market impacts, Imbert and Papp (2015) find that private-market dry season wages increase by 4.7 percent.³² As a result, between 22 percent and 42 percent of the total welfare gains from the program for the three poorest quintiles come from the average market wage effect, rather than direct program effects. As described above, Bertrand et al. (2021) experimentally examine related questions in the context of a workfare program in urban Cote d'Ivoire, and also find substantial evidence of crowd-out.

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). They find private sector wages increased. As a result, while the reform raised the income of low-income households by 13%, 90% of the gain came from 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

³² Berg et al. (2018) and Azam (2012) use similar identification strategies, also finding impacts on wages.

neighborhoods in Addis Ababa, Ethiopia. The program provided an hourly wage that was about 64 percent higher than the private market, which led households to shift to public employment. As in India, the program led to an increase in private market wages that accounted for most of the increase in welfare of program beneficiaries. Moreover, 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. On net, a common theme across these studies is that, particularly when the wage is high, the equilibrium impacts of workfare on the labor market may be larger than the direct effects of the program.

3.2.5: Programs targeting specific individuals within the household

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. However, there is reason to think that who the recipient matters (see, for example, Browning and Chiappori 1998; Robinson 2012), and indeed, this is one of the frictions that we discuss in Section 3.1.2 that moves us away from the benchmark case. If money is spent differently in the household depending on who receives it, then 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.

A number of studies explicitly randomize whether transfers are given to men or women in developing countries, finding little effects overall.³³ For example, Akresh, de Walque, and Kazianga (2016) randomize whether the transfer is given to mothers or fathers in Burkina Faso. They find no differences on child health or education outcomes on average. That said, they do find that giving transfers to mothers leads to lower levels of child labor. They also find some evidence that giving money to fathers improves outcomes in poor rainfall years, and leads to more household investment in livestock, cash crops, and housing.³⁴

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 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).

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

Finally, note that another serious gender related concern is whether cash transfers themselves, by

³³ These studies build on Lundberg, Pollak, and Wales's (1997) quasi-experimental study in the United Kingdom. Other important non-randomized studies include Duflo (2003)'s study of pensions, discussed in Section 4.3.3.

³⁴ Benhassine et al. (2015) also randomize the gender of the recipient of a transfer, and find little difference on average.

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.

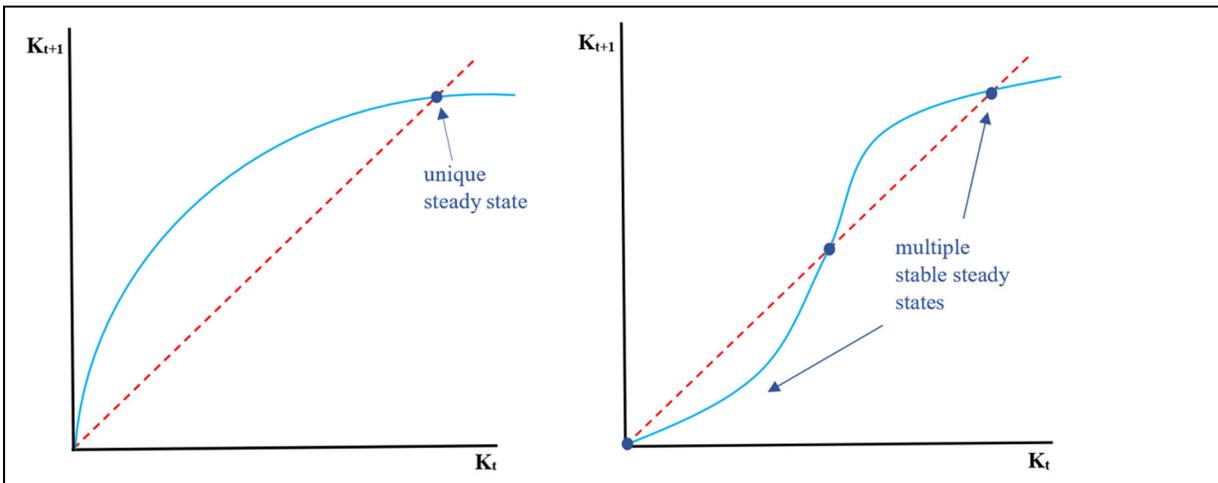
3.2.6: Poverty Traps and Lumpy Transfers

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

Conceptual issues: To fix ideas, consider a very simple 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 δ 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 3 illustrates two versions of this simple model with unique and non-unique steady states:

Figure 3: 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 3—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. From a social protection perspective, which

model we are in matters substantially.³⁵

Empirics: Which model are we in? In an RCT in Bangladesh that provided assets to a subset of poor households, Balboni et al (2022) shows the existence of poverty traps. They sort people based on their initial wealth and then use the shock of the capital transfer to estimate an empirical analogue of Figure 3, comparing assets at baseline (k_t) with assets 4 years later (k_{t+1}). They find that there is a unique steady state in control villages, but that transferring assets creates multiple steady states—so not only does the transfer temporarily increase incomes in the short run, but for some, it can tip them into a substantially higher steady state.³⁶

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 poor 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 a household’s capital constraints, as well as their skills constraints, to “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). 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? Recent experimental evidence from Botea et al. (2021) in Zambia answers yes, to both questions. First, they show that governments can run the programs: they show increases in consumption, assets, and mental health as a result of the full government-run package. But, importantly, they show similar results between a complete graduation treatment and the group that received the asset (and a savings tool) but no training.

In Ghana, however, Banerjee et al. (2022) find that neither just giving people assets nor just the savings tool have a positive long-run impact, while the whole package does. The Zambia and Ghana papers’ results are consistent if both the asset and the savings tool are important for the

³⁵ This also 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.

³⁶ Lybbert et al. (2004), Carter and Lybbert (2012), and Banerjee, et al. (2019) also provide evidence of the existence of poverty traps.

success of the program but the training is not. While the Zambia results suggest that perhaps some of training can be stripped off these graduation programs, reducing the cost, the Ghana results suggest that further testing may be useful in unbundling these programs.

Blattman, Fiala, and Martinez (2014; 2020) study a different type of lump-sum transfer program—upon submitting a business plan, young adults form groups in Uganda and receive a one-time grant. The authors find substantial increases in investment, work and income four years later. However, in the nine year long-run follow up (Blattman, Fiala, and Martinez 2020) they see the gains dissipate, 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. At 9 months, the monthly transfers increased food security, while lump-sum transfers increased assets. After 3 years of the transfer, there are no differences between the lumpy and monthly transfer groups. 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.

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, Muralidharan, Niehaus, and Sukhtankar (2016) find a 30.7 percent leakage rate in Andhra Pradesh in 2012. In Indonesia's subsidized rice program, Banerjee et al. (2018) calculate that eligible households only receive a third of the intended subsidy. Camacho and Conover (2011) find cheating on the PMT formulas, and Niehaus et al. (2013) estimate that 70 percent of ineligible households in India have below-poverty line cards, and many households report paying small bribes. Of course, these are particular examples, and researchers often study leakage where it is thought to be high, so one should not take these as unbiased estimates of the overall leakage rates. And, much of the 'leakage' does go to other households who are nevertheless still near poor. But, these facts do suggest that meaningful improvements to programs can be made by improving the efficacy of service delivery.

Corruption risk 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 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, more 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 in the social protection context.

Digital payment systems. In many developing countries, the authentication of beneficiaries and payments systems are manual: there is a list of beneficiaries, and some local official needs to verify who receives benefits, and cash is given. This makes it hard for the central government to monitor whether targeted beneficiaries get their full payment, or get a payment at all.

Several studies have examined the implications of digitizing this process. Muralidharan, Niehaus, and Sukhtankar (2016) examines the randomized rollout of “Smartcards” —biometrically authenticated payments—for receiving benefits from India’s workfare program (MGNREGA) and pension programs (SSP) in Andhra Pradesh, India. The Smartcards worked: they reduced payment delays, reduced leakage by 41 percent, and quite importantly increased program takeup by poor households. A second study, Banerjee et al. (2021) (described in more detail above) studied the conversion of an Indonesian food subsidy program from in-kind program administered by local officials to digital vouchers (a debit card) redeemable at a network of bank agents. The digital vouchers led targeted households to receive substantially more benefits, which in turn bumped many poor households over the poverty line. These effects occurred even without biometric identification, suggesting that the digital card, rather than the biometrics, may be important in this case.

A word of caution though: Muralidharan, Niehaus, and Sukhtankar (2020) showed that reform in Jharkhand, India, that began to require biometric authentication for the subsidized food program had the unfortunate side effect of dropping many poor people from the system who had not linked their biometric information to their food benefits account. Banerjee et al. (2021) similarly finds that many people could not enroll in Indonesia’s health insurance because of errors in the underlying identification data systems. These studies emphasized the need for universal access and robust administrative data before attempting these types of reforms.

Finally, it is worth noting that the degree to which people prefer digital systems may depend on the overall level of adoption of those systems in the country more broadly. Berkouwer et al. (2021) show that during the COVID-19 crisis, people in Kenya—where there is high mobile money adoption—preferred mobile money over electricity subsidies, 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.

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 changes reduced program expenditure by 24 percent while maintaining the same payments to workers. However, it is worth noting that these programs must also be done with care. For example, the reform studied by Banerjee et al. (2020) also delayed payments to beneficiaries. Similarly, Muralidharan, Niehaus, and Sukhtankar (2020) note in their study of Jharkhand that when the government transitioned to 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. While these payment systems are important, such transitions must therefore be managed with care.

Information provision. If “leakages” arise from local officials not following the national government’s rules, one fix is to provide information directly to beneficiaries so that they can potentially demand their full benefits from the local officials. Banerjee et al. (2018) experimentally tests this hypothesis, in the form of identification cards for a food transfer program in Indonesia, and find that simply mailing out identification cards dramatically increases the subsidy eligible households receive. They find additional benefits from posting information publicly, creating creating “common knowledge” and thus increasing 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 discussed 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, 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, there are some circumstances where this result may 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. The presence of spillovers to other households and poverty traps can also change the

welfare calculus.

The discussion here also suggests important directions for additional research. First, different programs may be appropriate for different types of people. For example, conditional cash transfers may work well if households are near the margin of take up for the incentivized behavior; for households far from the threshold, imposing conditions that they cannot comply with could make them worse off by denying them program access (Baird, McIntosh, and Özler 2011). 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.

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

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

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

Finally, the programs discussed in this section have primarily been conceptualized as responses to long-run poverty. But, to the extent 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, but these problems end up 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 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 taking much more severe actions to smooth shocks (see, for example, Chetty and Looney 2007).

In this section, we broaden our framework to think about social insurance policies. In Section 4.1, we first begin by summarizing the basic theoretical challenges of privately providing insurance (particularly in developing country settings), discussing the basic rationale for these programs. Next, as many low income and emerging economies are characterized by informal insurance, we discuss their limitations and the places where formal insurance could help fill in the gaps.

We then review the current state of knowledge on social insurance schemes designed to address two key kinds of shocks in developing countries. In Section 4.2, we examine those that insure against *income* shocks, e.g. unemployment insurance and paternal benefits. In Section 4.3, we discuss those that help cushion *expenditure* shocks, e.g. health insurance and pensions/annuities.

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 income economies than low income ones.³⁷ 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 (earning higher income y_{hi}) and with probability $(1 - p_i)$ she will be a low ability type (earning lower income y_{li}). If the probabilities are public information, then the individual can purchase 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, individuals will fully smooth their consumption across states by purchasing insurance. In this

³⁷ 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.

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 cannot observe and thus cannot properly price; that is, either the probabilities p_i or the risk ($y_{hi} - y_{li}$) is at least partially private information. If the insurance company prices its contract to at least break even with the high-risk types, then those with lower risk may not 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 large: the private insurance market accounts for about 9 percent of GDP across the OECD (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). 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.

Information, trust, and experience

One challenge is that people may not fully understand insurance products, and therefore may be unwilling to purchase them. Buying insurance requires a large degree of trust: you need to pay money now in the hope that sometime in the future, if you have a problem, the insurer will cover your loss (and not deny the claim). Many developed countries have 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. For example, several papers show that lack of trust hinders insurance take-up by exploring how experience effects (either from yourself or your network) can help improve insurance take-up (Cole, Stein, and Tobacman 2014; Cai, De Janvry, and Sadoulet 2015; Cai, de Janvry, and Sadoulet 2020).

It is worth noting that this is not only a problem in private insurance markets. Experience with insurance matters for government provided insurance programs that require co-payments. Both Asuming, Kim, and Sim (2021) and Banerjee et al. (2021) find that temporary subsidies can lead to longer-run insurance demand through experience. Providing just information on the benefits of insurance often is not enough (e.g., Dercon, Gunning, and Zeitlin 2019; Banerjee 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.

Credit constraints and timing issues

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

Basis risk

Basis risk provides another reason why demand for insurance may be low, particularly for agricultural insurance contracts (see 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 accounts for only exogenous determinants of risk (e.g. weather). 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.

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. For example, in developed countries, in addition to weather-based insurance, agricultural insurance contracts are often written as a function of output or profits from nearby farms. This type of contract may have lower basis risk, but requires good information: the insurer needs to be able to observe the output or profits of all nearby farmers

(perhaps through tax or other data collection). By contrast, there is little data on the output of most subsistence farmers, so these contracts are not practical to write in most developing country contexts. A similar example comes from property: in the US, for example, 93 percent of homeowners have property insurance (Insurance Information Institute 2020). Yet, the fraction of people in developing countries who insure these risks is likely to be tiny. One reason why this may be different in developing countries may be a supply side constraint: property insurers need to know the market value of a house to insure it. The less formalized property markets in developing countries may make this challenging.

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. In developed countries, there is typically 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 feel they are wrongly denied a claim. These systems are substantially less developed in most developing countries.

4.1.4: Informal insurance

A third complication for formal insurance provision in developing countries, compared to the developed world, is how it interacts with complex pre-existing, informal insurance arrangements. Many poor households in developing countries—while not formally insured, engage in various forms of risk sharing arrangements. Indeed, villagers are better at smoothing idiosyncratic consumption shocks than one may have expected given the absent formal insurance markets (Townsend 1994). 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 (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. Moreover, 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).

Even when informal insurance exists, it does not necessarily provide the level of coverage that households need, since income shocks have a strong spatial correlation, especially in agricultural areas. Ideally, risk-pooling would be done across villages where shocks are independent, but these arrangements are difficult in practice since monitoring costs to prevent moral hazard are too high.³⁸ These spatially correlated shocks may be a particular challenge with natural disasters or conflict, necessitating large scope disaster relief programs.

All of these arguments suggest that some forms of formal insurance are likely to be very useful

³⁸ One way to reduce this problem is marriage across villages (Rosenzweig and Stark 1989).

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, especially as there is evidence that even as economies experience growth, many households keep relying on informal networks. 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.

The evidence on crowd-out is mixed (Albarran and Attanasio 2003; Strupat and Klohn 2018; Takahashi, Barrett, and Ikegami 2019). The level of crowd-out may vary by characteristics of the types of insurance: for example, Huang and Zhang (2021) finds no crowd-out of private transfers in the context of China's rural pension scheme, in contrast to Jensen (2004), which finds large crowd out effects of South Africa's public pension. Huang and Zhang argue that that they find less crowd-out in China than South Africa since the pension benefits are much smaller.

Turning to the second question, Arnott and Stiglitz (1991) theorize that whether the existence of formal and informal insurance systems is welfare improving depends on how well informal networks monitor others to reduce 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 reduce 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.

Finally, another way in which formal insurance may have positive welfare effects over 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).

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 (UI) is designed to help people smooth consumption between jobs. However, a key challenge is that unemployment benefits typically continue until the worker resumes work. There is an extensive literature on the degree to which this conditioning of benefits on future employment discourages job search in high income countries, focusing on (1) 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 duration of benefits) can improve welfare, and (2) whether the benefits allow for better job matching (e.g., Chetty 2006; 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), Gonzalez-Rozada and Ruffo (2022).³⁹

In many low- and middle-income countries, these same policy questions exist, but it is further complicated by the large informal sector, where the government cannot monitor job entry and exit. This has two important implications. First, the government cannot provide benefits that start conditional on unemployment for those in the informal sector, and so informal workers may not be actually covered by UI. Second, the government also has challenges conditioning benefits on re-employment (since it only observes formal employment), and thus workers receiving UI could choose to seek employment in the informal sector in order to continue receiving their benefits. There is therefore also a worry that these kinds of programs may also disincentivize formalization, which has consequences for productivity, taxation, workplace safety, and so on.

Gerard and Gonzaga (2021) examine the relationship between informality and unemployment insurance in Brazil. They find, perhaps counterintuitively, that the presence of a large informal sector *reduces* the efficiency 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 further understanding optimal UI design in the presence of a large informal sector is important for future research.

Given these challenges, one often observes alternative policies to UI to help to insure workers against job loss. A common policy is mandated severance pay; that is, a lump-sum payment upon termination of a labor relationship. This payment is not conditioned on future employment, and thus does not distort future employment decisions.

Severance programs are not without challenges. First, someone needs to adjudicate that the worker was, indeed, working, and was terminated. If firms are supposed to pay directly, the government needs to ensure that they do; if the government collects taxes and pays severance, it needs to collect the taxes to do so. Weak institutions may exacerbate the enforcement challenges. For example, 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 abound in the court system. They show improving institutional quality could help: providing information about likely court outcomes substantially shortens settlement times, improving the food security of workers.

Severance payments have other challenges: 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 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.

³⁹ A number of recent papers also examine the spillover effects of unemployment insurance on domestic violence (Bhalotra et al. 2021) and crime (Britto, Pinotti, and Sampaio (forthcoming)).

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 (first proposed, to our knowledge, by Feldstein and Altman 1998) is to create unemployment savings accounts, where workers have an individual account that they can access in case of job loss. The accounts can be funded through mandatory contributions by workers and/or firms. The idea is that individual accounts help align incentives and reduce the scope for moral hazard. However, they do not have the risk pooling features 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 wages and job duration.

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

4.2.2: Insurance against Disability and Death

Disability insurance is designed to help insure workers against accidents or illnesses that may either temporarily or permanently remove them from the labor market. However, similar to UI, the large informal sector hinders governments from providing universal disability coverage 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 has two challenges: (1) this only provides insurance for those near the threshold of poverty (and not the middle class) and (2) if targeting is done infrequently, the sign-up period may not align with when people need assistance.

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, this remains an area that is comparatively underexplored in the economics literature.

Finally, we touch briefly on life insurance. The death of a primary income earner is a substantial income loss, and indeed, throughout many developing countries, 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 deceased 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. There are a number of excellent reviews on these topics and so we refer the interested reader to Cole and Xiong (2017), Ali, Abdulai, and Mishra (2020), and Nshakira-Rukundo, Kamau, and Baumüller (2021).

4.2.4: Parental Benefits

Parental benefits are designed to help 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 2018b), the benefits levels in developing countries are often very low for mothers, paternity leave is often nonexistent, and many of those in the large informal sector do not necessarily have access to any of these formal benefits.

In contrast to the extensive literature on these policies in high-income countries, there is relatively little micro-empirical research on paid maternity or paternity leave in lower- and middle-income countries.⁴⁰ 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

⁴⁰ For cross-country difference-in-difference analysis, see Fallon, Mazar, and Swiss (2017).

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. 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, Calderón (2014) in Mexico shows that access to childcare increased the likelihood of working for women and reduced the likelihood of earning zero income. Similarly, Halim, Johnson, and Perova (2022) examined public pre-school expansions on women's work in Indonesia, and find positive effects on women's work, but it is driven by increases in unpaid family work. They argue that limited day care hours made it unlikely that women could find jobs outside.

Recent experimental studies (Barros et al. 2011; Martínez and Peticar 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 income gains from childcare lead to at least as large a gain as an equivalent cash transfer, while also improving child development. Interestingly, they find that the earnings increase comes from higher productivity, largely in self-employment, rather than longer hours.

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

4.3: Insurance Against Expenditure loss

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

4.3.1: Health Insurance

Government-led health insurance systems are increasingly becoming common in many developing countries, as governments aim to increase health care utilization, improve health outcomes, and help households manage income and consumption risks that may arise from a health shock. Nearly 190 countries have some sort of public health insurance system.⁴¹

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

⁴¹ 188 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).

bucket then explores the challenges in provision and design of insurance, particularly around how 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, reviewing the studies that aim to measure health insurance impacts. 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 only expect this if the increased utilization induced by the health insurance 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 (King et al. 2009; Gruber, Lin, and Yi 2021; Levine, Polimeni, and Ramage 2016). del Valle (2021) argues that health insurance also provides another economic benefit: they show that by reducing the severity of health shocks, Seguro Popular in Mexico reduced the degree to which other household members needed to drop out of the labor force to directly provide care. There is also suggestive evidence that health insurance may also reduce financial stress above and beyond its impact on finances per se (Haushofer et al. 2020).

Health care utilization and health. Health insurance also has the potential to change health care utilization, since it 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 believe many households in developing country contexts are), households may under-consume health care without insurance. Understanding how to help households achieve their optimal, distortion-free level of health consumption remains an important challenge.

Does health insurance even change health consumption, and health, in developing countries? The experimental evidence on the impact of health care utilization and health is 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 (King et al. 2009; Haushofer et al. 2020). A recent study by Malani et al. (2021) highlighted another important policy issue: in their experiment, 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—and thus perhaps unsurprisingly, they

observe no health effects. In short, most of the experimental studies to date find little overall impact of health insurance on health outcomes.

Given the need for large samples sizes to measure impact, a number of other studies focus on quasi-experimental variation from large-scale public health insurance reforms, and then use large administrative datasets on mortality and national samples.⁴² For example, Gruber, Lin, and Yi (2021) examines the roll-out of China’s public health insurance for rural households, which covered as many as 800 million people. They find a significant decline in aggregate mortality, which they argue could explain 78% of the entire increase in life expectancy in China during this period. Using survey data, they find large effects on health care utilization and a host of other health outcomes.

Gruber, Hendren, and Townsend (2014) study an alternative method of health insurance: the expansion of free or heavily subsidized care at public facilities. They examine Thailand’s 2001 health care reform, known as the “30 Baht” program. Prior to this reform, the poor, young and old were given free public health care, but the system was seen as chronically underfunded. The reform changed the system in two ways. First, it provided universal access to public facilities at a co-pay of 30 Baht (or \$0.75) per visit. Since the co-pay was already waived for the poor, in practice this reform largely led to reduced health access costs for the informal, non-poor. Second, the reform provided hospitals with a universal capitation payment based on population in the province; this on net moved to more generous government financing. Using administrative mortality records, they show that prior to the reform, the infant mortality rate was related to how rich the 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.

Challenges with Health Insurance Design

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

A common challenge with these systems, however, is what to do about informal workers (for whom premia cannot be collected from employers). 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 remain vulnerable to the economic impacts of a health emergency.

⁴² See Gruber, Lin, and Yi (2021), who provide a nice overview of the literature.

Moreover, if firms want to evade paying benefits to formal workers and/or workers do not fully value the insurance, mandating insurance benefits for formal workers could increase informality.

A second policy option—which is quite common—is to mandate that non-poor, informal workers contribute to insurance. This is difficult to enforce, however, particularly in countries with limited administrative capacity, and many people do not comply. With imperfect compliance, in addition to worries about increases in informality described above, this can also lead to adverse selection problems (i.e. only signing up when sick) that can financially strain the insurance system.

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, as well as concerns about the cost.

Below, we first discuss the evidence on the interplay between health insurance and informality. 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 employer provided (or subsidized) health insurance is ambiguous. 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 increase formalization. But, if workers do not value the insurance at cost or if alternative health care arrangements are cheaper, it could increase informality.

A number of papers examine these concepts. For example, Bergolo and Cruces (2014) examine a large-scale policy reform in Uruguay's social insurance administration (SIA) that both increased benefits and contributions for formal workers, and find evidence of both effects. The reform extended the SIA's coverage for dependent children, making the program more attractive. At the same time, the reform 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: the greater benefits drew people into the system, but 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. A reform was instituted that made workers who were below a Poverty Index Score eligible for non-contributory health insurance, but those who were formally employed were ineligible regardless of their score (and needed to contribute through their employers)—creating an incentive to become informal. They find the reform increased informal employment by about 3-4 percentage points.

Finally, a number of papers have examined Mexico's Seguro Popular on formalization. Prior to the rollout, insurance used to be tied to payroll contributions, and so many people were uncovered; Seguro Popular aimed to provide universal coverage. Analyzing the staggered roll-out of the program and survey data, Aterido, Hallward-Driemeier, and Pagés (2011) and del Valle (2021), among others, find small effects on formalization, while Azuara and Marinescu (2013) find no

effect on average, but small effects for unskilled workers. More recently, examining the roll-out with social security data, Bosch and Campos-Vazquez (2014) 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 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 what can be done to mitigate this issue.

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

Banerjee et al. (2021) also study this in the context of the Indonesian national health insurance, randomizing offers of partial and full subsidies of premiums for one year (and a control), and tracking results for 20 months. Larger subsidies bring in healthier individuals into the health care system, consistent with this undoing adverse selection; they also document a pattern of dynamic selection, where those who enroll in the no-subsidy condition are much more likely to immediately file large claims. Since the subsidies attract healthier individuals who then also pay premia in the post-subsidy period, they allow the government to cover more people at the same total cost. Fischer, Frölich, and Landmann (2018) in Pakistan also experimentally vary premia, and find that adverse selection is higher when premia are higher. In short, while by no means a panacea, subsidies can help ameliorate adverse selection to some degree.⁴³

Bundling. A second approach to reduce adverse selection is bundling: by tying the purchase of health insurance to that of another product, or bundling insurance purchases for a household together, one can potentially limit the ability of households to buy insurance only for the sick. For example, in their study with an NGO in Pakistan, Fischer, Frölich, and Landmann (2018) experimentally investigate whether people can choose to enroll individuals, or whether they must enroll the entire household, under the idea that households will have to enroll all members, not the sickly. They also examine a community bundled contract, where at least fifty percent of the community must sign up in order to activate the insurance. These bundled contracts reduce "expected costs" among those who sign up (where expected costs are claims predicted from baseline covariates), suggesting that this type of group insurance may be effective. In many developed countries, workplaces are the 'group' for health insurance purposes; this study suggests other groups as an alternative in countries with high levels of informality.

The downside of bundling, however, is that if demand for insurance is low, it can reduce demand for the other, bundled product. For example, Banerjee, Duflo, and Hornbeck (2018) find that a

⁴³ Two other important papers on subsidies are Thornton et al. (2010) and Wagstaff et al. (2016).

substantial fraction of people were apparently willing to forgo renewing their microcredit just so that they did not have to purchase required bundled health insurance.

Information and Trust. Lack of information (or mis-information) about the benefits of insurance and limited trust in the system may be another reason for low insurance demand. Starting with information, the evidence is mixed: with Giles et al. (2021) and Asuming (2013) finding positive impacts, for example, and Banerjee et al. (2021) and Dercon, Gunning, and Zeitlin (2019) finding none.

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

In sum, despite the benefits of health insurance for risk smoothing, health insurance provision in developing countries remains a challenge: the combination of lack of demand for insurance, adverse selection, a large informal sector, and the relative challenges of enforcing a mandate to purchase insurance, means that it is challenging to ensure widespread insurance coverage. Many countries, 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 move 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 poses a challenge, whether it be from flooding, 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). This issue may become even more severe with global climate change.

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.⁴⁴ Moreover, like other forms of insurance, low demand often prevails due to trust in the institutions to pay out (Reynaud, Nguyen, and Aubert 2018). Understanding the limits of this market seems an important question for future research.

⁴⁴ With the exception of TransUnion South Africa’s similar *Insurance Claims and Policy* database.

Aggregate shocks. For aggregate shocks, such as floods, earthquakes, and hurricanes, the issues are a bit different. Governments *ex-post* tend to intervene and provide some amounts of emergency relief to households, though the degree to which they do so varies. For example, Gignoux and Menéndez (2016) study earthquakes in Indonesia, and find that households have negative effects from the shocks, but that these dissipate after 2-5 years. They find substantial 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 (as it is difficult to stand by and do nothing), households may be reluctant to purchase insurance themselves. 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, —and perhaps so that governments can recoup some of the costs via mandatory premiums. One of the few examples we know of in middle-income countries is the Turkish compulsory earthquake insurance scheme, established in 2000 (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 Fonden, which the authors claim is the only indexed disaster relief fund worldwide. 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.

4.3.3: Pensions and Annuities

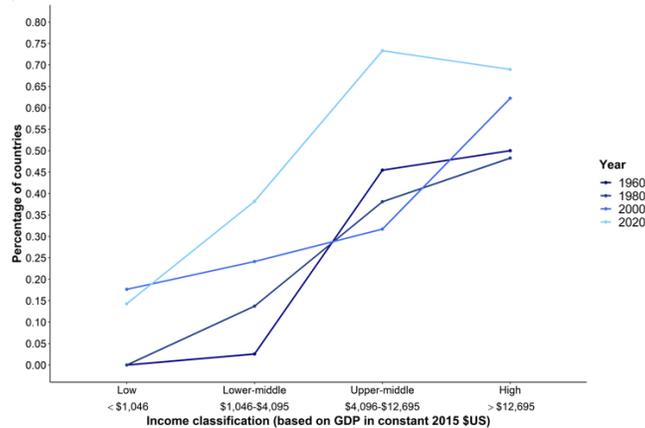
Old age is often associated with poverty. Health typically falters as people age, so the elderly are less able to work; this is particularly true for whom work often entails substantial physical labor. Saving (or saving enough) may be difficult, particularly for those outside the formal sector. And even 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 caring for parents), but these networks may be incomplete, and may break down with urbanization.

Accordingly, many governments around the world have public pension programs, with a mix of systems (Appendix Table 2 shows that 181 have contributory pension systems, 102 have non-contributory systems and 96 have both systems in place). These programs typically have three, related goals: a) to help 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 as compared to one's savings, and c) to provide some amount of redistribution so that even those with low incomes are not too poor in their old age.

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

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

Figure 4: 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. *Source:* GDP per capita data is from the World Bank (2021a).

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

Consumption effects need not be limited to pensioners themselves, as consumption resources 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.⁴⁵ 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. 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 their kids 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

⁴⁵ Case and Deaton (1998) provide a rich description of the South Africa’s non-contributory pension program.

than they otherwise would have in the absence of a pension.

Several papers find that pensions reduce labor supply. de Carvalho Filho (2008) studies the reform to 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 also finds that paid work declines, but that those who stop doing paid work switch to family businesses. Huang and Zhang (2021)'s study in China shows that labor supply decreased for rural people older than 60 (i.e. the eligible), particularly for farm-work.⁴⁶

Program design. There is comparatively less work, however, on program design questions in middle-income countries, but we highlight several papers that suggest some ways in which middle-income contexts may be somewhat different from higher-income contexts.

One important question 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 system may matter: they show that when the pension system switched from one in which pensions became largely a function of contributions, payments for younger workers (who had more of an incentive to ensure their wages were reported accurately) increased.

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. 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, and so firms compete on other non-price attributes rather than price. 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. 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 how the administrative challenges hinder take-up of a widow's pension program in India, and that

⁴⁶ The income effects from a pension could also affect the labor supply decisions of other members. Studying South Africa's pension program Bertrand, Mullainathan, and Miller (2003) find a sharp drop in the working hours of prime age men in the household, particularly when the pensioner is a women. Posel, Fairburn, and Lund (2006) argue that the pensions also drive migration among the prime age workers. Including the migrants in the analysis, Ardington, Case, and Hosegood (2009) they find small positive increases in work among prime-age adults.

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 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?

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, 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 risk. But the additional challenges highlighted above—low insurance demand, low 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 policy 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 3 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 redistributive 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 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, 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. 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—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. 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 how poor the poor can be, there is also a focus on finding program designs that can help break poverty traps, whether within a single generation or across generations.

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.

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>.
- 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).
- 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>.
- 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.
- 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_asuming_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.
- 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>.
- 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>.
- 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>.
- 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>.
- Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. 2017. “Labor Markets and Poverty in Village Economies.” *The Quarterly Journal of Economics* 132 (2): 811–70. <https://doi.org/10.1093/qje/qjx003>.
- Banerjee, Abhijit, Emily Breza, Esther Duflo, and Cynthia Kinnan. 2019. “Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs?” Working Paper 26346. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w26346>.
- 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. 2016. “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. 2015. “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. 2021. “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. 2021. “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. 2017. “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. 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, 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.
- 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.
- 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. 2019. "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>.
- 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>.
- 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.
- 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, 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>.
- 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 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>.
- 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>.
- 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. 2011. "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>.
- 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 Olgiati. 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>.
- 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>.
- 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, 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>.
- 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, and Firouz Gahvari. 2008. “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>.

- 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>.
- 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>.
- 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>.
- 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.
- 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>.
- Gardner, Clare, Caroline Riungu, Clare O’Brien, and Fred Mertz. 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>.
- 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.
- 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>.
- 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>.
- 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>.
- Gupta, Sarika. 2017. “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.
- 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>.
- 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. 2022. “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.
- 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>.
- 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>.
- 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>.
- 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).
- 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>.
- 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/gotcrc/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>.
- 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.
- 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>.
- 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/>.
- 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/.
- 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>.
- 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>.
- 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>.
- 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/pps031>.
- 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>.
- 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>.
- 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. 2013. *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>.
- 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. 2013. “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>.

- 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. 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>.
- 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>.
- 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>.
- 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.
- Rosenzweig, Mark, and Oded Stark. 1989. “Consumption Smoothing, Migration, and Marriage: Evidence from Rural India.” *Journal of Political Economy* 97 (4): 905–26.
- 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>.
- Schultz, T. Paul. 2004. "School Subsidies for the Poor: Evaluating the Mexican Progresas 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>.
- 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>.
- 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. 1994. “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. *Women, Business and the Law 2018*. Washington, DC: World Bank. <https://doi.org/10.1596/978-1-4648-1252-1>.
- 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.
- 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>.