# Cash Transfers*

Paul Niehaus† Tavneet Suri‡

May 30, 2024

**Abstract**

Cash transfers are now a ubiquitous tool for social protection in low- and middle-income countries, and have generated a correspondingly large program evaluation literature. From the point of view of optimal public finance, however, some decisive issues remain unclear. We emphasize, in particular, the incidence of transfers within as well as across households; the incremental value and optimal design of transfers in the presence of financial market frictions; and their external effects. Work on these issues will be particularly valuable if it is distributed less unevenly across space than past research.

*Thanks to Adrien Auclert, Milkah Chebii, Han Shen Chia, Abdourahmane Cissé, Mohamed Fuaad Daboh, Ugo Gentilini, Daniel Handel, Amrik Heyer, Anton Heil, Niale Kaba, Adama Kamara, Santiago Levy, Nicholas Li, Ben Olken, Peter Ombasa, Hannah Max-Kyne Sao-Kpato, and staff at Financial Sector Deepening for helpful conversation and suggestions, and to Muhammad Karim, Wilson King, and Shruthi Ramesh for research assistance.

†University of California, San Diego. pniehaus@ucsd.edu.

‡MIT Sloan. tavneet@mit.edu.

# 1 Introduction

Cash transfers for social protection have reached very large scale, covering up to 1.4 billion people in low- and middle-income countries (Gentilini et al., 2022). This is in part because states' capacity to deliver them has grown. It was once easier to give Indian farmers free electricity than free money, so this was done despite the resulting environmental havoc. Now things are different: governments can use national identification systems and digital payments infrastructure (see Sukhtankar (2024) in this volume) including private sector innovations such as mobile money. The pandemic both illustrated and accelerated these patterns. Togo delivered cash transfers to a quarter of its adult population without "boots on the ground," selecting them based on meta-data about their mobile phone usage and enrolling and paying them over the phone (Chowdhury et al., 2022).

It is also arguably a success story for evidence-based policy-making. High-quality research on the impacts of some of the earliest cash transfer programs, such as Progressa in Mexico, featured prominently in (successful) efforts by multinational organizations to promote subsequent adoption (Leisering, 2018, Chapter 4). The volume of high-quality evidence on the causal effects of cash transfers is now very large, and its production has accelerated in recent years (see Figure 1). With nine years elapsed since the last systematic review (Bastagli et al., 2016), one might argue that the main task is to synthesize.

We offer an alternative perspective. Consider allocating a national budget to increase overall well-being. How much should go to cash transfers as opposed to other forms of redistribution, to public goods, and so on? This question lies at the heart of optimal public finance, and it was echoed in many of the preparatory discussions we held with people in policy-making roles. Their questions were often less about cash transfers per se than about getting the balance right between transfers, infrastructure, and other priorities. On its own, no amount of evidence on causal effects can answer these questions; one must take a stand on what the effects mean for well-being—on valuation, in other words. And thinking of value in the usual economic sense of equivalent variation leads to a learning agenda with priorities different from those reflected in work to date.

We focus in particular on three frontiers that matter for optimal transfer policy. One is characterizing the incidence of transfers not only across households but also within them. There is compelling evidence of meaningful inequities within some households, to the particular disadvantage of women. Deducing what this means for how much these households should be given is an open problem, as is estimating how much intra-household allocation can be influenced and at what cost through design changes such as issuing individuals their own transfers, into their own bank accounts, etc.

Another is to work through how transfers interact with imperfections in markets for savings, credit and insurance. In a world of imperfect financial markets, the value of transfers depends in part on their contribution to solving the core problems of smoothing consumption and financing investment. They might even contribute so much that there is no trade-off between equity and efficiency. But it is not obvious how to move from such venerable conceptual points to quantification. To illustrate, suppose we see treatment effects on investment and future earnings. This could be because the transfer relaxed a credit constraint, but could also simply be because this was the only way to shift consumption into the future in the absence of reliable savings vehicles. Nor is it yet clear how realistic it is to even try to target transfers to people for whom they do relax financial constraints.

Financial market failures also imply that the structure of transfers is consequential. Without well-performing financial markets, a lump sum is not the same as a stream of small payments; money that shows up when school fees are due is not the same as money that shows up at harvest time. There is an under-exploited opportunity to figure how to structure transfers in light of this fact, identifying what works well for the financial (not to mention psychological) lives of the recipients. Doing so may help bridge the somewhat artificial gap between thinking about cash transfers "for" social protection and cash transfers "for" development.

A third frontier is measuring the externalities transfers induce. These include both fiscal externalities (i.e. impacts on public revenue and public expenditure on other programs) that affect the true cost of a transfer program, and economic externalities that affect its benefits. To what degree can we expect economic stimulus via demand externalities, for example, or effects on prices? Should governments take mitigating steps—informing traders in advance, building all-weather roads, etc.? What about classic economic externalities—how does redistributing a dollar affect deforestation, carbon emissions, and so on? There is a lot of open terrain here.

Finally, we highlight an opportunity to study these issues in places that have previously been under-studied. India, for example, has 21% of the population but just 1.6% of the experimental studies within low- and middle-income countries. Such unequal coverage may reflect in part a chicken-and-egg problem: countries that do not try cash transfers generate little contextually relevant evidence about them, and so remain uncertain whether to use them. It may also reflect a degree of "physics envy," or the hope that there exist deep economic parameters that generalize across contexts. Economic logic, however, suggests that the particular quantities we emphasize (e.g. the transfer multiplier) are likely to be far more contingent. It is thus quite reasonable for decision-makers to ask, what should I expect to happen *here*?

## 2 What Questions?

Because cash transfers give recipients the freedom to choose what to do with the money, one should expect them to affect many things. Some of the money will be spent on food, some on healthcare, some on housing, some on investment projects, etc.; some recipients will adjust the time they spend on work and leisure; and so on. This is indeed the picture we get from experimental studies, which have (collectively) documented effects on a list of outcomes that includes more or less everything development economists study, from health to education to livelihoods to women's empowerment all the way to suicide rates (Christian et al., 2019).

This raises the question where to focus. Which outcomes matter, and why? One can ask this of any intervention, of course; it is just particularly salient for cash transfers. And when the goal is to say something about how much to spend on transfers, the question boils down to how to assign value to all of these outcomes. One simply cannot do program evaluation without valuation. We sometimes forget this and refer to a collection of estimated causal effects as a program evaluation, but it is not.

Methodology here has become somewhat muddled. Some studies have asked whether cash transfers are cost-effective at moving a particular outcome. But even the staunchest advocates of cost-effectiveness analysis acknowledge that it is unclear how to apply it to interventions that affect more than one outcome, which cash transfers (and arguably most other interventions) do–both because individual people use them for many things, and because different people use them differently. Asking whether unconditional transfers are a cost-effective way of increasing schooling outcomes, for example, makes little sense if recipients spend the majority of the money on something other than schooling—not because schooling is not a good and important thing, but because the calculation effectively assigns a value of zero to everything else.

Other studies collapse several outcomes into weighted-average indices—partly, in our experience, to avoid referee demands for multiple hypothesis testing adjustments. An appropriately constructed index can indeed be useful for the narrow goal of efficiently testing the null that the intervention did nothing. But indices optimized for testing are not useful for valuation: any cardinal meaning is lost when we re-weight variables by economically meaningless factors such as inverse covariances. Such indices are of little help to decision-makers who care not just whether or not all effects were precisely zero (which in any case no one believes to begin with), but about magnitudes.

The good news is that economics offers a coherent approach to valuation. Consider someone with preferences represented by $u(x, z)$ over things $x$ that they can choose (consumption, labor supply, etc.) and things $z$ that they cannot, and a budget constraint $p(x) \leq y + t$ where $y$ is exogenous non-transfer income, $t$ a cash transfer, and $p(x)$ the net price of choices $x$. Their indirect utility is

$$\begin{aligned} v(z, p, y+t) &= \max_x u(x, z) \\ \text{s.t. } & p(x) \leq y + t \end{aligned} \tag{1}$$

We can then ask how they would value any given change in their circumstances by asking what change in $y$ would yield the equivalent change in $v$. We see immediately that increasing $t$ is equivalent to increasing $y$; we have discovered the tautology that, from the point of view of the recipient, the value of \$1 is \$1. This not a deep point; it is an immediate consequence of using money as a numeraire. But it is a useful place to start, as it forces us to to be precise about why we think things are not so simple. In what follows we will adapt the problem (1) in three different ways to illustrate three senses in which \$1 is not simply \$1, and work out what tasks each then implies for empirical research.

## 2.1 Incidence

Knowing their particular circumstances in life, an agent solving (1) would value a policy that gives them \$1 at \$1, and a policy that gives \$1 to anyone else at \$0. But now suppose we place the same agent behind a “veil of ignorance” so that they do not know their exact circumstances—whether, for example, their income will be high ($\overline{y}$) with probability $\pi$ or low ($\underline{y} < \overline{y}$) with probability $1 - \pi$. Then their expected indirect utility is

$$\pi v(z, p, \overline{y} + t(\overline{y})) + (1 - \pi) v(z, p, \underline{y} + t(\underline{y})) \tag{2}$$

If they are averse to income risk, this expression will be higher for more progressive transfer policies $t(y)$.

This is one important reason to study the targeting of transfer programs. The more progressive their incidence, the more value—assessed using criteria like (2)—they create per dollar. Another is that recipients stress the importance of fair and appropriate selection processes in their feedback (Wingfield et al., 2023; Samuels et al., 2013), both per se and because they anticipate that, depending on the process, transfers can give rise either to greater cohesion or to jealousy, tension and conflict within their communities. In the longer run, targeting methods may affect the kinds of constituencies that form around a program and advocate for its continuation. Finally, targeting methods (including conditions, in the case of conditional cash transfers) may create (dis)incentives. Cash transfers are typically not means-tested and hence are unlikely to generate the kinds of mechanical disincentives to earn emphasized in rich-country public finance following Mirrlees (1971); the question is rather whether the alternative targeting methods used, such as proxy means tests, generate distortions of their own (on which see Banerjee et al. (2020b)).

For a general review of work on targeting per se, we defer to Alatas et al. (2024) in this volume. We focus here on one issue that is particularly germane to cash transfers, and less prominent in the targeting literature: the allocation of resources within households. The ultimate incidence of transfers depends both on which households receive them, and how those households divvy them up. It is notoriously hard to measure resource allocation within households; standard surveys do not even attempt to separately measure the consumption of different members, for example. (Exceptions that prove the rule include some project-specific surveys and Brazil's Household Budget Surveys.) And yet there is enough evidence to make a persuasive case that meaningful inequities exist. Overall spending patterns vary with the source of income (including from cash transfers) in ways that do not match the "income pooling" implication of the unitary model of the household (Lundberg and Pollak, 2008). Even if we attribute this mismatch to unobserved factors, it is harder to explain away inequities in directly measured allocations—the fact, for example, that some households invest more in male than in female children (Barcellos et al., 2014).

What does this mean for targeting? To be concrete, consider the finding in Brown et al. (2019) that the majority of underweight women and undernourished children in Sub-Saharan Africa are found in households that are not among the poorest in per-capita consumption terms. Standard algorithms are trained to avoid selecting these households. Should we include them? Or does the very fact that these households have divided existing resources unequally imply that incremental resources would not reach their most deprived members, making them less attractive to target? There are both "income" and "price" effects here, so to speak, pushing in different directions.

One path forward might be to target treatment effects on individual-level outcomes rather than deprivation at the household level. This approach would ask, "how do we target the households in which effects on the most deprived members are largest," as opposed to "how do we target the most deprived households." Recently developed tools for estimating heterogeneous effects, as in Haushofer et al. (2022), could be applied. But the traverse from heterogeneous effects to optimal policy would still require care. Comparing impacts on the nutrition intake of children to that of adults, for instance, would require adjustments akin to classic "equivalence scales" which are difficult to convincingly identify (Lewbel and Pendakur, 2016).

Another potential lever is selecting which household member receives the transfer. Ironically it is so common to issue transfers to female heads of households (at least nominally) that there is relatively little evidence on how much this matters. Only a few studies directly vary the rule for selecting recipient gender, and the picture that emerges from these is unclear. Some simply are not powered to rule out large differences (e.g. Haushofer and Shapiro (2016) in Kenya) or study transfers with features like conditions or "labels" attached which may have constrained how they were used (e.g. Benhassine et al. (2015) in Morocco). And among the more dispositive results there is quite a bit of variety. Consider impacts on nutrition: Somville et al. (2020) reject large differences in expenditure shares on food in India, while Armand et al. (2020) estimate that giving transfers to women significantly increased food shares in Macedonia, and Akresh et al. (2016) find that giving transfers to men had significantly larger effects on children's nutrition in Burkina Faso (as well as on some measures of investment and production). Overall our view is that there is limited evidence that women systematically use money in ways that are "better" than men, but also that this is inessential to rationalize transfers to women: there is a strong argument to give transfers to the household members with the least influence purely on a priori distributional grounds. And recipients themselves—especially women and people with disabilities—do describe receiving transfers as being empowering and making them more independent (Wingfield et al., 2023).

Assuming along these lines that most social protection transfers will continue to be issued "to women," does it matter what exactly this means—that is, how the transfers are delivered? It seems reasonable to infer this from evidence on the adjacent problems of loan disbursement (Riley, 2023) and wage payment (Field et al., 2021), where paying women directly into digital accounts which they control has made a difference. But it has not, to our knowledge, been tested in the domain of social protection specifically. Pushing even further in this direction, one could decouple transfer receipt from household membership entirely. This could have empowerment effects through changing outside options: a woman could continue to receive transfers after separating from her partner, for example. We know of one ongoing study (our own, in Kenya) that examines transfers with this feature, but that one does not compare to more conventional transfers attached to household membership.

An unfortunate feature these questions share is that they are quite demanding to study in terms of statistical power. Suppose that the average effect of a transfer on some outcome is $X\sigma$, and that the differential effect of delivering it to a woman v.s. a man is $0.2X\sigma$ (say). In economic terms, this is a substantial difference. But rejecting the null that it is zero requires 25 times the sample size as rejecting the null that the transfer itself has no impact. This harsh arithmetic may be one reason we see few studies of this type. Part of the disconnect is, of course, that testing hypotheses at the 95% level is not a sensible way to make practical decisions, but it seems likely to remain the norm. Given this, one under-utilized alternative is to elicit preferences. It will often be much easier and more informative to ask women whether they would prefer to receive \$1 herself or to have their husbands receive \$$Y$ than to power a large RCT to test for differential effects.

## 2.2 Financial Market Imperfections

Households at or near the extreme poverty line often face severe financial constraints. The instruments available to them for saving, borrowing and insuring are limited, expensive, and unreliable. They struggle to manage cash flow, accumulate large enough lump sums to make investments, and so on (c.f. Collins et al., (2009)). All this should matter, one might think, for how we value cash transfer programs. A \$1 transfer could have disproportionate value in a second-best world of imperfect financial markets if it enables the recipient to achieve a resource allocation closer to the first-best.

To illustrate this formally, we must extend the problem in (1) to make household finance, in the sense of moving money across time or across states of the world, possible. Let us consider two possible states defined by realizations of an income shock: income may be either high ($\overline{y}$) with probability $\pi$ or low ($\underline{y} < \overline{y}$) with probability $1-\pi$. If actuarially fair insurance is available then the agent can perfectly insure himself, obtaining indirect utility

$$v(z, p, \pi(\overline{y} + t(\overline{y})) + (1-\pi)(\underline{y} + t(\underline{y}))) \tag{3}$$

regardless of the state. In this case the expected value $\pi t(\overline{y}) + (1-\pi)t(\underline{y})$ of the transfer is a sufficient statistic for utility, and so increasing it by \$1 is always worth \$1 of non-contingent income, whether the transfer itself is contingent or not. But if no insurance is available the agent's indirect utility is equivalent to expression (2) above, and in this case a \$1 increase in expected transfer value can be worth more than a \$1 increase in non-contingent income if the transfer value is concentrated in the low state.

This scenario captures a core part of what was meant by "social protection" in the first place (although official definitions have come to encompass a much broader array of objectives). When a bread-winner gets sick, cannot work, and on top of this needs expensive medical treatment, for example, it may be hard to ensure that everyone gets enough to eat. Well-designed transfer programs can help prevent such routine tragedies.

The research question this motivates is how well-correlated transfers are with shocks. With respect to aggregate shocks such as natural disasters and macroeconomic crises, governments have taken various approaches to achieving this, including introducing new transfer schemes and expanding eligibility for or the generosity of existing ones, under the heading of "shock-responsive social protection" (Oxford Policy Management, 2017, Section 7). It would be useful to assess the bottom-line performance of these efforts using metrics such as the high-frequency covariance between public transfers and income from other sources in household panel data. It may also be fruitful to study the performance of the reinsurance market for the relevant aggregate risks: in theory, governments should purchase insurance contracts whose payouts are linked to their social protection programs, but such sovereign risk financing is relatively nascent.

And it could be useful to develop and test transfers whose size is linked to the prices of essential commodities, as otherwise the recipients may face non-trivial risk in terms of real consumption possibilities (Gadenne et al., 2021).

Idiosyncratic shocks are another matter. One could argue that these are better addressed using explicit insurance schemes (for health, unemployment, and so on) rather than cash transfers. But such schemes are rare (or inapplicable) in low-income countries, so the question whether cash transfers could be made more responsive to shocks remains relevant. Targeting as it is typically done—infrequently, using proxy means tests that emphasize slow-moving indicators of well-being such as assets—is unlikely to achieve this. The fortunes of low-income households are far more volatile, with many in the same communities moving both out of and into extreme poverty in a given year (Baulch and Hoddinott, 2000; Krishna, 2010). It would be useful to develop and test new approaches here. Another alternative might be to explicitly target households based on their vulnerability to future shocks, as opposed to their current situation per se, as suggested for example by (Carter and Janzen, 2018). Alternatively, households could be given the option to receive insurance in lieu of (part of) their non-contingent cash transfers, effectively using the transfers to address intertemporal barriers to insurance takeup (Casaburi and Willis, 2018).

A second, weaker thesis is that even if transfers themselves are not indexed to shocks, providing them to vulnerable households will reduce the likelihood that they have to forego necessities when the shocks hit. In this view, useful metrics for a cash transfer program are how well it targets shock-prone households, or how much it affects consumption's variance over time or covariance with shocks. Notice that these are question about how transfers affect consumption's *second* moments, not whether they affect consumption on average (a surprisingly popular question). Asfaw et al. (2017), for example, document that transfers significantly reduce the effect of weather shocks on food expenditure and caloric intake. Interpreting these kinds of results can be subtle, however. If a household is given enough money to eat well during a shock and still chooses not to, that could simply indicate that it valued some other use of the money even more.

Transfers may also affect the profile of risks that households face in the first place by inducing them to change how they earn, where they live, and so on. This idea has drawn attention particularly in relation to the risks posed by climate change. But, again, the interpretation of effects here requires some care. A household might appropriately choose to taken on *more* risk after receiving a transfer (or after learning that it will receive future transfers) in pursuit of higher returns. Banerjee et al. (2020a) find, for example, that households receiving a long-term stream of "basic income" lost more business revenue at the start of the pandemic—because they had been more likely to start a business previously, and thus had more to lose—but also suffered less from hunger.

As this example calls to mind, imperfect financial markets limit investment as well as consumption-smoothing. Transfers may then affect not just who can claim the returns on investments, but also which investments get made in the first place. What does this mean for optimal policy? One common intuition is that it makes redistribution more attractive overall. People who have less to begin with are perhaps more likely to have untapped high-yield investment opportunities, as well higher marginal value to consumption (and if this effect dominates any incentive effects then there might not even be a trade-off between equity and efficiency). This idea has played an important role in the financing of cash transfer programs—the Inter-American Development Bank began lending to support cash transfer programs in Latin America, for example, after deciding that they could be viewed as a form of investment. Many papers have studied the impacts of cash transfers on investment (in either or both of physical and human capital) with this idea in mind. (Gertler et al., 2012), for example, start from the hypothesis that transfers can increase investment in the presence of either credit constraints or uninsured risk, and then show that Progressa transfers did indeed increase agricultural investment and earnings in Mexico.

There are some subtleties to assessing such arguments. The first is simply that the distinction between "investment" and "saving" is often blurry in practice. Consider a household that is not credit constrained, but that cannot access a savings account paying a non-negative real interest rate at a financial institution it trusts. When it receives a temporary cash transfer, this household uses some of the money to increase its holdings of a liquid but relatively low-return asset (perhaps small livestock). We might easily misinterpret this as evidence that the transfers relaxed a credit constraint, when in fact it is evidence of a persistent problem in the market for savings vehicles. These are very different interpretations of the same data, and (to our knowledge) the literature has not yet grappled with how to separate them.

A second subtlety is even if some households do have under-exploited high-return investment opportunities, this may already be captured implicitly by analyses that consider only their consumption today. Suppose that each household $h$ faces an idiosyncratic rate of return $r_h$ on investment (because there is no financial intermediation) but can invest as much or as little as it wants at that rate. The standard Euler equation

$$u'(c_h^t) = \delta r_h u'(c_h^{t+1}) \tag{4}$$

may then imply that consumption $c_h^t$ today is decreasing in $r_h$. Intuitively, households with high-return opportunities will voluntarily choose to consume less today in order to take advantage of those opportunities, and will thus appear more deprived in terms of present consumption. As a result $u'(c_h^t)$ remains a sufficient statistic for the social value of an incremental transfer to $h$, just as it would be if returns were equalized ($r_h = r \forall h$). No information on investment, earnings, or future consumption $c_h^{t+1}$ is required.

This simplification breaks down when investment opportunities are lumpy. In that case first-order conditions such as (4) no longer fully characterize household behavior. A household might have a high value of $c_h^t$ not because it has no high-return investment opportunities, but because the opportunities it does have are indivisible and it does not have quite enough money to cover their cost. Treatment effects on investment, earnings, and (future) consumption may then vary independently from initial deprivation. This is what (Haushofer et al., 2022) find. They use household covariates to predict both deprivation and the impacts of cash transfers, and find that impacts vary substantially conditional on deprivation, with the result that welfare-maximizing allocations typically condition on impact as well as on deprivation.

Another piece of evidence for this view comes from studies that compare smooth v.s. lumpy payment streams, holding fixed the total amount. This distinction would be irrelevant in a world of perfect financial markets (since any payment stream could be re-engineered into any other NPV-equivalent stream), and would not matter for investment behavior if investment opportunities were not lumpy. But impact evaluations have typically found meaningful differences, with lumpy payments inducing more investment in lumpy assets (Haushofer and Shapiro, 2016; Aguila et al., 2017) and larger impacts on aggregate economic activity (Banerjee et al., 2023). And when given a choice recipients have overwhelmingly chosen relatively lumpy payment schedules (Kansikas et al., 2023), as one would expect if they find accumulating lump sums both valuable and difficult. This is consistent with the popularity of institutions such as ROSCAs (for example see the Global Findex, especially for economies in sub-Saharan Africa), which enable participants to convert income streams into larger lumps of cash, in many parts of the world. All of this suggests an important opportunity to redesign existing social protection (and humanitarian) programs—which typically deliver cash transfers in a series of small payments—in ways that increase welfare.

How big should the lumps be? In principle there is a tradeoff between giving smaller transfers to more households or larger transfers to fewer, the terms of which depend on the distribution of lumpy investment opportunities across households. One instinct might be to learn about this distribution by experimentally varying transfer size. This turns out to be difficult. One obstacle is power. Suppose that a transfer of size $2t$ has 1.6 times the impact of a transfer of size $t$. Economically, this is a large deviation from constant returns, but statistically it implies that we need 25 times the sample to detect differences between treatment groups as to detect the effect of size $t$ transfers per se. Aggregating across studies, as Kondylis and Loeser (2021) do, can increase sample size, but there is no reason to expect the distribution of lumpy investment opportunities to be similar across contexts. Another obstacle involves measurement. Different lumpy investments yield different kinds of value over different time horizons. A motorcycle yields revenue quickly; coffee or cocoa trees yield revenue only after several years; education yields higher wages (and perhaps intangible benefits) over a lifetime; a more durable metal roof yields savings on repair costs; a larger house yields a flow of use value. Putting all of these investments on a truly even footing seems very difficult, and we do not know of any attempts to do so. Overall, the optimal sizing of wealth transfers remains an important open problem.

Another margin on which transfer design can likely be improved is timing. Social protection programs typically default to a regular schedule of payments throughout the year. This would not matter if financial markets (and self-control, and foresight) were perfect: recipients could re-engineer any pattern of cash inflows into any desired pattern of outflows. However, in practice, cash flow management is a major challenge for low-income households. Transfers that show up when school fees are due, or when fertilizer must be purchased, or during the "lean season" when many go hungry, are likely to have disproportionate value. In the few cases where recipients have been asked, they often prefer irregular timing (Kansikas et al., 2023, Section 3.4). It would be useful to measure more systematically the demand for and impacts of contextually appropriate timing, and also to try pairing social protection transfers with budgeting exercises such as that in Augenblick et al. (2023).

## 2.3 Externalities

Cash transfers are likely to affect people other than those who initially receive them. In part this is simply because recipients use them to transact—buying things, putting money into financial vehicles, and so on—and these transactions affect the counterparties in some way. Notice that this feature distinguishes cash from in-kind transfers such as food, which might conceivably be consumed without immediately affecting anyone else.

Beyond the transactions themselves, cash transfers have been shown to induce a number of behavioral responses which are likely to have either pecuniary or non-pecuniary external effects. They tend to increase labor supply (Banerjee et al., 2017) and to induce substitution from wage- into self-employment, which will affect wages. They often increase schooling, which we think (though it is hard to prove) has external effects. Even consumption increases on their own will probably have some environmental consequences. And so on. Enriching the simple problem (1) we began with to reflect such possibilities would yield an indirect utility function like

$$v(z(t), p(;t), y(t) + t) \tag{5}$$

under which increasing $t$ by \$1 could evidently raise utility by more or by less than \$1 of exogenous income.

One important class of external effects are those on government budgets. Changes in consumption, labor supply, and output will typically affect tax revenue. Changes in participation in other government programs may affect public expenditure. The true cost of a cash transfer program includes these fiscal externalities as well as the accounting costs of the program itself. This point features prominently in recent work in public finance; Finkelstein and Hendren (2020), for example, begin their illustration of the Marginal Value of Public Funds concept by analyzing an increase in the generosity of a cash transfer, summarizing as follows:

> "...we start with a benchmark case of a small increase in a cash transfer that only affects its recipients, whose response to the policy is privately optimal. Under these assumptions, we show that estimates of causal effects of the policy are needed only for estimating the policy's costs, not its benefits."

Whether on the cost or benefit side, external effects have been under-emphasized in the program evaluation literature. Indeed they have even been viewed at times as a nuisance parameter. Take the common argument that spillovers are not a major concern because the share of people treated within each community or market (say) was small. This is coherent insofar as we only want estimates of treatment effects on the treated. If the expected spillover effect on any one neighbor is small, then the bias induced by using neighbors as counterfactuals is small. But because there are many affected neighbors the *total* spillover effect may be large! If we care about aggregate well-being, we should be keen to capture such effects.

In fairness, it is much easier to measure behavioral responses than the broader consequences these have for the rest of society. When transfers have far-reaching effects like those described above, how can we find any good counterfactuals? One approach is to directly estimate both direct and indirect effects using a spatial approach, pairing experiments (natural or otherwise) in which the unit of treatment assignment is large with some assumption about how space bounds the effects (Muralidharan and Niehaus, 2017). The classic example in this genre is village-level randomization paired with the assumption that spillovers are contained within villages. This is what (Angelucci and De Giorgi, 2009) did in an influential early study of Progressa in Mexico, for example, estimating that much of the benefit of the program accrued to ineligible households, whose food consumption increased by 50% as much as that of eligibles. Of course, the assumption that economic spillovers are contained within villages—or more generally within any administrative unit—is itself surely wrong to some extent. But recent work has shown that it is also possible to achieve consistent and rate-optimal estimation of the total effect of a program under the weaker assumption that spillovers decay across space at some bounded rate (Leung, 2022; Faridani and Niehaus, 2022).

What we know today about external effects comes largely from a small group of studies that take such a spatial approach, starting with Angelucci and De Giorgi. Two subsequent studies have measured impacts on aggregate output, and found substantial effects: Egger et al. (2022) estimate that output increased by \$2.5 for every \$1 of NGO transfers in rural Kenya, and Gerard et al. (2021) estimate that an expansion of social protection transfers in Brazilian municipalities increased output by 1.7%. In the former case this expansion had little detectable effect on local government finances (Walker, 2018), while in the latter tax revenue increased by 2.7%. Estimated effects on price levels have generally been modest, but with some exceptions and nuances. In rural Kenya, for example, large inflows amounting to 15% of GDP raised consumer goods prices by 0.1–0.2% (Egger et al., 2022). In Mexico, Progressa transfers did not significantly change prices on average, but did increase quantity discounts (Attanasio and Pastorino, 2020). Unconditional transfers had no significant effects on average, but raised food prices by 1.5% in more remote villages where they amounted to 10% of aggregate income (in contrast to food transfers, which lowered prices) (Cunha et al., 2019).

Collectively these studies demonstrate that external or "general equilibrium" effects can be first-order to any bottom-line assessment of a cash transfer program. But there are few of them. Statistical power remains a serious hurdle. Many studies use spatial designs and test for spillovers but report inconclusive results, with confidence intervals spanning a range of values with very different economic implications. This is not surprising: a typical design featuring both within- and between-village randomization, for example, will generally be far better-powered to detect differences within than between villages.

There are also open questions about the dynamics of these effects. Rich-country macroeconomists typically expect the stimulus effects of transfers to be short-lived, while development economists have conjectured that aggregate demand shocks could stimulate a "big push" with longer-lasting effects. The estimated effects on output above are time-averaged over periods of 27 months after one-time transfers (Egger et al., 2022) and three years after the onset of ongoing transfers (Gerard et al., 2021), respectively. It is an open question by how much (if at all) estimates of the total multiplier effect would increase with longer-term measurement.

These difficulties make it all the more important to assess how well we can predict and value the aggregate effects of transfers in advance. One could do this using a combination of general equilibrium modeling and partial equilibrium estimates of behavioral responses in domains where externalities are thought to be important, including health (Dupas and Miguel, 2017), education (see Snilstveit et al., 2015), environment, and crime (see for example Blattman et al., 2017). Estimates of the size of the externalities would also be of great value; to take a canonical example, effects on years of schooling should be multiplied by the external (not the private) returns to schooling to obtain the total external effect. Exercises like this would lend themselves to testable predictions about how transfers interact with other forms of public spending—whether constructing more roads would alter the macroeconomic effects of transfers, for example—which would speak to the higher-level budget allocation decisions that policy-makers face.

# 3 Where?

Considering the various quantities we have discussed above—the incidence of transfers within households, the covariance of transfers with shocks, the excess investment returns generated particularly by lumpy transfers, the externalities working through education, aggregate demand, public budgets, and so on—there is no reason to expect them to be the same everywhere. On the contrary, one could (and in some cases people have) build economic models to try to predict how they vary. Given this, our working hypothesis should be that it matters where future empirical research is done.

All else equal, one would ideally allocate more research effort to places where research is inexpensive; where we have weak priors; and where many people will benefit from better estimates. We do not have the data or space to explore all of these issues. But on the last point we can get some sense by examining the distribution of past research across low- and middle-income countries. To this end we searched the 3ie Development Evidence Portal for completed experimental evaluations of conditional or unconditional cash transfers, identifying 315 in total, with new studies appearing at an increasing rate from 2000 onward (Figure 1; see table notes for details of the query). This sample likely omits some studies that we would ideally include and includes some that we would omit, but we expect it to give a reasonably accurate picture of the *relative* availability of evidence across space.

That picture is very uneven. Figure 2 plots the geographic distribution of studies relative to population, illustrating some striking disparities. Kenya, for example, has 10 times as large a share of LMIC cash transfer evaluations as it has of LMIC population. Mexico has 8.5 times as many. Near the other end of the spectrum, India—with 21% of LMIC population—has just 1.6% of LMIC studies. These highly skewed distributions echo those for empirical research in economics more generally (Das et al., 2013).

We suspect that future work on cash transfers will have more value if it is done in more diverse places. There are sizeable opportunities still on the table. In India, for example, there has been much discussion about whether to replace in-kind food transfers with cash transfers, but very little evidence on the consequences this would have, particularly for food prices in different parts of the country. China's Dibao scheme is one of the largest in the world, reaching 53 million people as of 2017, and also one of the least-studied. Namibia hosted the first significant pilot of universal basic income in a low-income country, but the idea went no further due (it seems) to untested concerns that giving Namibians money would make them lazy (see

Banerjee et al., 2020a).

What of the incentives to do this kind of research? Some worry that professional returns are lower for "replications" than for clever new insights. But, in this case, there are both contexts and concepts that are understudied, and hence room to be innovative on both dimensions at once. And in any case our whole argument is that we should not expect results to "replicate." The quantities we have emphasized here are not fundamental constants. They are the sorts of things we should expect to vary, and expect economic modelling to enable us to predict—to extrapolate from what happened in one place at one time and anticipate what will happen in another place at another time. Testing how effectively we can do so is not a box-ticking replication exercise; it is the essence of the scientific method.

## References


**Aguila, Emma, Arie Kapteyn, and Francisco Perez-Arce**, "Consumption Smoothing and Frequency of Benefit Payments of Cash Transfer Programs," *American Economic Review*, May 2017, *107* (5), 430–35.

**Akresh, Richard, Damien De Walque, and Harounan Kazianga**, "Evidence from a randomized evaluation of the household welfare impacts of conditional and unconditional cash transfers given to mothers or fathers," *World Bank Policy Research Working Paper*, 2016, (7730).

**Alatas, Vivi, Abhijit Banerjee, Rema Hanna, and Ben Olken**, "Targeting," in "The Handbook of Social Protection: Evidence to Inform Policy in Low- and Middle-Income Countries" 2024.

**Angelucci, Manuela and Giacomo De Giorgi**, "Indirect effects of an aid program: how do cash transfers affect ineligibles' consumption?," *American economic review*, 2009, *99* (1), 486–508.

**Armand, Alex, Orazio Attanasio, Pedro Carneiro, and Valérie Lechene**, "The Effect of Gender-Targeted Conditional Cash Transfers on Household Expenditures: Evidence from a Randomized Experiment," *The Economic Journal*, 05 2020, *130* (631), 1875–1897.

**Asfaw, Solomon, Alessandro Carraro, Benjamin Davis, Sudhanshu Handa, and David Seidenfeld**, "Cash transfer programmes, weather shocks and household welfare: evidence from a randomised experiment in Zambia," *Journal of Development Effectiveness*, 2017, *9* (4), 419–442.

**Attanasio, Orazio and Elena Pastorino**, "Nonlinear Pricing in Village Economies," *Econometrica*, 2020, *88* (1), 207–263.

**Augenblick, Ned, Kelsey Jack, Supreet Kaur, Felix Masiye, and Nicholas Swanson**, "Retrieval Failures and Consumption Smoothing: A Field Experiment on Seasonal Poverty," 2023.

**Banerjee, Abhijit, Michael Faye, Alan Krueger, Paul Niehaus, and Tavneet Suri**, "Effects of a Universal Basic Income during the pandemic," Technical Report, UC San Diego 2020.

**_ , _ , _ , _ , and _** , "Universal Basic Income: Short-Term Results from a Long-Term Experiment in Kenya," 2023.

**_ , Rema Hanna, Benjamin A. Olken, and Sudarno Sumarto**, "The (lack of) distortionary effects of proxy-means tests: Results from a nationwide experiment in Indonesia," *Journal of Public Economics Plus*, 2020, *1*, 100001.

**Banerjee, Abhijit V, Rema Hanna, Gabriel E Kreindler, and Benjamin A Olken**, "Debunking the stereotype of the lazy welfare recipient: Evidence from cash transfer programs," *The World Bank Research Observer*, 2017, *32* (2), 155–184.

**Barcellos, Silvia Helena, Leandro S Carvalho, and Adriana Lleras-Muney**, "Child gender and parental investments in India: Are boys and girls treated differently?," *American Economic Journal: Applied Economics*, 2014, *6* (1), 157–189.

**Bastagli, Francesca, Jessica Hagen-Zanker, Luke Harman, Valentina Barca, Georgina Sturge, Tanja Schmidt, and Luca Pellerano**, "Cash transfers: what does the evidence say?," 2016.

**Baulch, Bob and John Hoddinott**, "Economic mobility and poverty dynamics in developing countries," *Journal of Development Studies*, 2000, *36* (6), 1–24.

**Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen**, "Turning a shove into a nudge? A "labeled cash transfer" for education," *American Economic Journal: Economic Policy*, 2015, *7* (3), 86–125.

**Blattman, Christopher, Julian C Jamison, and Margaret Sheridan**, "Reducing crime and violence: Experimental evidence from cognitive behavioral therapy in Liberia," *American Economic Review*, 2017, *107* (4), 1165–1206.

**Brown, Caitlin, Martin Ravallion, and Dominique van de Walle**, "Most of Africa's nutritionally deprived women and children are not found in poor households," *Review of Economics and Statistics*, 2019, *101* (4), 631–644.

**Carter, Michael R. and Sarah A. Janzen**, "Social protection in the face of climate change: targeting principles and financing mechanisms," *Environment and Development Economics*, 2018, *23* (3), 369–389.

**Casaburi, Lorenzo and Jack Willis**, "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya," *American Economic Review*, December 2018, *108* (12), 3778–3813.

**Chowdhury, Anir, Cina Lawson, Elizabeth Kellison, Han Sheng Chia, Homi Kharas, Jacquelline Fuller, Michael Faye, Michal Rutkowski, Rodrigo Salvado, and Stefan Dercon**, "Accelerating digital cash transfers to the worlds poorest," 2022.

**Christian, Cornelius, Lukas Hensel, and Christopher Roth**, "Income shocks and suicides: Causal evidence from Indonesia," *Review of Economics and Statistics*, 2019, *101* (5), 905–920.

**Collins, D., J. Morduch, S. Rutherford, and O. Ruthven**, *Portfolios of the Poor: How the World's Poor Live on $2 a Day*, Princeton University Press, 2009.

**Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran**, "The price effects of cash versus in-kind transfers," *The Review of Economic Studies*, 2019, *86* (1), 240–281.

**Das, Jishnu, Quy-Toan Do, Karen Shaines, and Sowmya Srikant**, "US and them: The geography of academic research," *Journal of Development Economics*, 2013, *105*, 112–130.

**Dupas, Pascaline and Edward Miguel**, "Impacts and Determinants of Health Levels in Low-Income Countries," *Handbook of Field Experiments*, 2017.

**Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker**, "General equilibrium effects of cash transfers: experimental evidence from Kenya," *Econometrica*, 2022, *90* (6), 2603–2643.

**Faridani, Stefan and Paul Niehaus**, "Rate-optimal linear estimation of average global effects," *arXiv preprint arXiv:2209.14181*, 2022.

**Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore**, "On her own account: How strengthening women's financial control impacts labor supply and gender norms," *American Economic Review*, 2021, *111* (7), 2342–2375.

**Finkelstein, Amy and Nathaniel Hendren**, "Welfare analysis meets causal inference," *Journal of Economic Perspectives*, 2020, *34* (4), 146–167.

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

**Gentilini, Ugo, Mohamed Almenfi, Ian Orton, and Pamela Dale**, "Social Protection and Jobs Responses to COVID-19: A Real-Time Review of Country Measures," February 2022.

**Gerard, François, Joana Naritomi, and Joana Silva**, "Cash transfers and formal labor markets: Evidence from Brazil," 2021.

**Gertler, Paul J, Sebastian W Martinez, and Marta Rubio-Codina**, "Investing cash transfers to raise long-term living standards," *American Economic Journal: Applied Economics*, 2012, *4* (1), 164–192.

**Haushofer, Johannes and Jeremy Shapiro**, "The Short-term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya*," *The Quarterly Journal of Economics*, 07 2016, *131* (4), 1973–2042.

**_ , Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael Walker**, "Targeting impact versus deprivation," Technical Report, UC San Diego 2022.

**Kansikas, Carolina, Anandi Mani, and Paul Niehaus**, "Customized cash transfers: financial lives and cash-flow preferences in rural Kenya," Technical Report, UC San Diego January 2023.

**Kondylis, Florence and John Loeser**, "Intervention Size and Persistence," 2021.

**Krishna, Anirudh**, *One illness away: Why people become poor and how they escape poverty*, OUP Oxford, 2010.

**Leisering, Lutz**, *The global rise of social cash transfers: How states and international organizations constructed a new instrument for combating poverty*, Oxford University Press, 2018.

**Leung, Michael P**, "Rate-optimal cluster-randomized designs for spatial interference," *The Annals of Statistics*, 2022, *50* (5), 3064–3087.

**Lewbel, Arthur and Krishna Pendakur**, *Equivalence Scales*, London: Palgrave Macmillan UK,

**Lundberg, Shelly and Robert A Pollak**, *Family decision-making*, na, 2008.

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

**Muralidharan, Karthik and Paul Niehaus**, "Experimentation at scale," *Journal of Economic Perspectives*, 2017, *31* (4), 103–124.

**Oxford Policy Management**, "Shock-Responsive Social Protection Systems Research: Literature review (2nd Edition)," 2017.

**Riley, Emma**, "Resisting social pressure in the household using mobile money: Experimental evidence on microenterprise investment in Uganda," 2023.

**Samuels, Fiona, Nicola Jones, and Agnieszka Malachowska**, "Holding cash transfers to account: beneficiary and community perspectives," *London: ODI*, 2013.

**Snilstveit, B., J. Stevenson, D. Phillips, M. Vojtkova, E. Gallagher, T. Schmidt, H. Jobse, M. Geelen, M. Pastorello, and J. Eyers**, "Interventions for improving learning outcomes and access to education in low- and middle- income countries: a systematic review.," 2015.

**Somville, Vincent, Ingvild Almås, and Lore Vandewalle**, "The Effect of Gender-Targeted Transfers: Experimental Evidence From India," 2020.

**Sukhtankar, Sandip**, "Digital Technology in the Delivery of Social Protection," in "The Handbook of Social Protection: Evidence to Inform Policy in Low- and Middle-Income Countries" 2024.

**Walker, Michael**, "Informal Taxation Responses to Cash Transfers: Experimental Evidence from Kenya," July 2018.

**Wingfield, Tom, Beatrice Kirubi, Kerri Viney, Delia Boccia, Salla Atkins et al.**, "Experiences of conditional and unconditional cash transfers intended for improving health outcomes and health service use: a qualitative evidence synthesis," *Cochrane Database of Systematic Reviews*, 2023, (3).

Figure 1: Cash Transfer Evaluations by Year

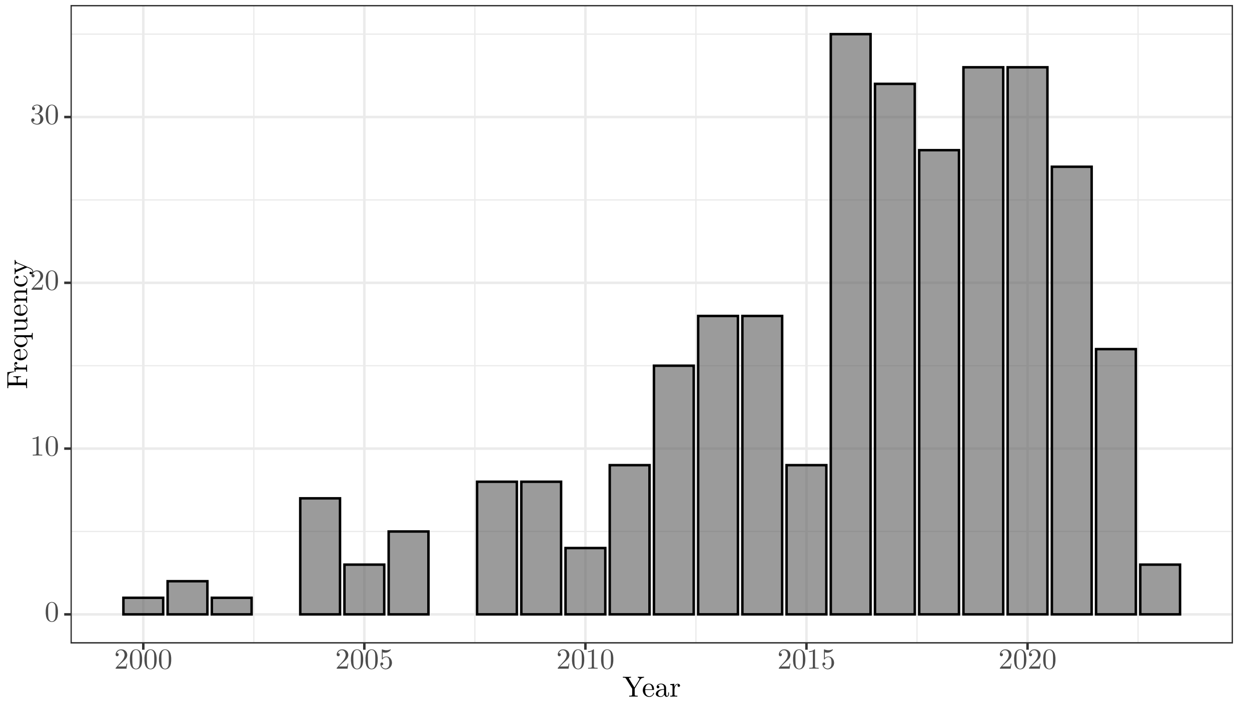


This figure plots trends in the rate of release of experimental evaluations of cash transfers as recorded in the 3ie Development Evidence Portal. To obtain this sample we first searched on the string "cash transfers" and then filtered the results to restrict attention to experimental evaluations in which one of the interventions was a conditional cash transfer or an unconditional cash transfer. We further excluded from the results studies whose status was not "completed."

Figure 2: Cash Transfer Evaluations per Million Inhabitants, by Country

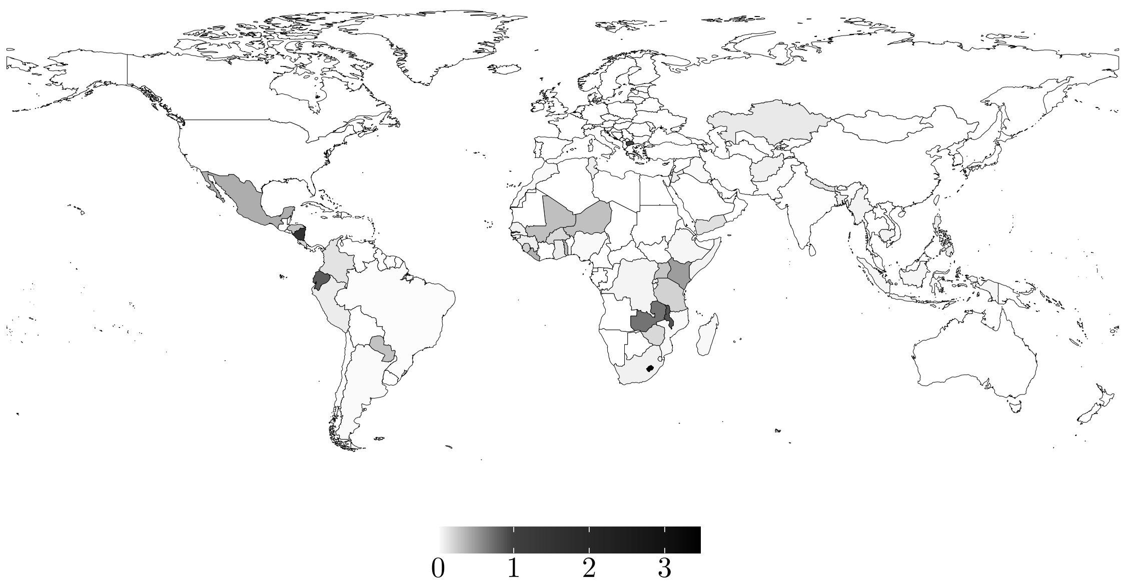


This figure plots the geographic distribution of experimental evaluations of cash transfers as recorded in the 3ie Development Evidence Portal relative to population. See notes to Figure 1 for details.