Chapter 20

Spillovers and General Equilibrium Effects of Social Protection Programs

Edward Miguel

University of California, Berkeley, National Bureau of Economic Research (NBER); CEGA

1. Introduction

Social protection (SP) programs are expanding in coverage and scale globally, including in low- and middle-income countries (LMICs), in part due to innovations in program design and growing state capacity, and these programs have been demonstrated to generate large benefits for many recipients (J-PAL and EPoD, 2023). Yet despite the rising profile of SP programs, there remains little evidence on their spillover and general equilibrium (GE) effects beyond direct recipients. This is due to both logistical constraints – estimating broader effects often requires large sample sizes to achieve adequate statistical power – and methodological limitations, as standard individual-level randomized control trial (RCT) research designs are often not well-suited to study spillovers.

This chapter is a curated review of existing research (employing rigorous designs) on program spillovers across major SP sectors and aims to provide useful directions for future investigation. Understanding the nature and magnitude of spillover and GE effects of SP programs is a first-order issue for scholarship and public policy. A key goal of most SP programs is redistribution to poor or vulnerable (e.g., elderly) individuals and households. Yet this redistribution immediately introduces oft-discussed concerns regarding possible equity versus efficiency trade-offs. For instance, expanding public cash transfers directed to low-income households could mean the state needs to raise additional taxes on productive enterprises or divert government revenue away from essential physical investments (like transport or energy infrastructure). The existence of spillovers could alter this calculation: if the cash transfer program also generated broader economic benefits for local households and firms, or increased tax revenue, it could boost such programs' cost effectiveness and blunt the hypothesized trade-offs (while the opposite would hold for negative externalities), potentially affecting public policy recommendations.

In the studies reviewed in this chapter, large-scale SP programs have generally been shown to create broader impacts for society along multiple dimensions, and it is useful for the discussion below to delineate three major categories of effects (although there is inevitably some overlap between them and some gray areas).

These include the following: (1) within-household impacts (for instance, on the extent of intimate partner violence, IPV) and inter-generational effects (on medium- to long-run child outcomes); (2) effects on other (non-recipient) local households and enterprises mediated through market outcomes (including via prices and multiplier effects); and (3) broader impacts on politics, governance, the state, norms, and culture more broadly, including through fiscal externalities and the likelihood of armed conflict.

Section 2 opens with a brief discussion of the main *econometric and statistical methods* used to estimate spillover and GE effects of SP programs in economics. In addition to covering commonly used approaches, we also mention some proposed new strategies and highlight fundamental challenges. One key limitation of most existing methods is their inability to capture general equilibrium impacts that are so pervasive that they affect all units in a study population, and thus are observationally equivalent to zero spillover effects (given that most estimators rely on comparisons across units with differing degrees of indirect exposure). This concern is most important for estimating effects in the third category listed above, since a program that affects society or politics as a whole would be conflated with secular trends or time effects.

Section 3 discusses evidence on the spillover and GE effects of *cash transfer and related programs*, the most widely studied SP program in LMICs by far (J-PAL and EPoD, 2023) and the one that this chapters dedicates the most energy to discussing.

Due to space limitations (and the excellent coverage in other chapters in this Handbook), this chapter provides a selective review of the most influential studies employing rigorous research designs, including experimental designs where possible, rather than aiming for comprehensive coverage. Taken together, the bottom line assessment of the spillover and GE effects of cash transfer programs is that they are largely positive: existing evidence indicates that cash transfers in LMICs can produce positive spillover gains within households in some cases (with little risk of increased IPV); meaningful benefits for non-recipient households and businesses, generating local economic multipliers (often with limited local price inflation); and can boost support for the political party responsible for the program, and reduce the risk of armed conflict.

Section 4 discusses evidence on the spillover and GE impacts of SP programs other than cash transfers. While the evidence from LMICs is more limited in these other program areas, we focus on two important types of programs with a small but growing evidence base, namely, *old-age pensions* and *labor market interventions*.

We lead with a discussion of old-age pension programs in 4.1. These are of first-order public policy importance in poor countries given the exploding numbers of the elderly: for instance, in Sub-Saharan Africa, the world's poorest region, the number of people over age 60 is projected to triple from 2020 to 2050 (and to double in the rest of the world) and yet only a handful of African countries have set up pensions or other elder social support and care programs with broad population coverage beyond a narrow slice of formal sector workers or government employees (see Duhon et al. 2023 for a detailed discussion on leading research issues on the economics of aging in Africa). As a result, the evidence base on direct programs impacts is limited, and on spillover effects even more so.

While there is some relevant work from Latin American settings, the most expansive body of existing work covers the South African pension program that was created in the 1990's in the aftermath of Apartheid, which was designed to have universal coverage. Several studies document large within-household impacts, in particular for children and young adults. But overall, the evidence base on the broader effects of old-age pensions in LMICs remains limited, especially for the crucial categories 2 and 3 noted above.

The chapter then turns to a discussion of the broader impacts of labor market SP programs in 4.2. These programs are found in many LMICs and aim to provide labor opportunities for poor adults, often working on public works projects.

The most influential collection of studies in this area estimates broader impacts of the famous India National Rural Employment Guarantee scheme (NREGA), which provides a certain number of guaranteed days (up to 100 typically) of low-skill labor for tens of millions of Indians annually, the largest program of its kind globally in terms of recipients. There are other related RCTs from Sub-Saharan African settings, although of smaller scale. Regarding category 2 effects, several studies document positive effects on local wages in areas with more expansive public employment programs, and these (perhaps surprisingly) often lead to higher local private employment rather than crowd-out. Separate work focuses on the category 3 impacts of NREGA and shows that its introduction can reduce the likelihood that adverse local climate shocks translate into armed conflict.

It is worth defining the scope of this chapter up front and what has been excluded. For reasons of space alone, this chapter is unable to cover the broader spillover and GE effects of health or education sector SP programs. An important literature documents externalities in both sectors. An early example in health (that also employed a since widely adopted econometric estimation approach) is Miguel and Kremer (2004); Benjamin-Chung et al. (2017) and Benjamin-Chung et al. (2018) contain a comprehensive discussion of both the methods used and key findings of spillover studies in global health. It seems obvious that the design of public health insurance systems could also have many economic and social effects beyond impacts on recipients themselves (i.e., fiscal implications). This is a critical issue but one that is challenging to study with the most common empirical methods and where there is limited evidence from LMICs. There is also high-quality evidence on the existence of large GE effects of education interventions, including from school building (Duflo 2004) and vouchers (Muralidharan and Sundararaman 2015), that are also beyond the scope of this chapter.

The conclusion summarizes and ties together the main patterns we have documented, and in particular the central finding that the spillover and general equilibrium effects of social protection programs can often be large, altering the cost effectiveness and attractiveness of many common policies. The final section also discusses key gaps in the evidence and raises other issues of interpretation that could be important for future research on this emerging topic.

2. Methodological Issues

There is a long tradition of studying the broader spillover effects of public policies and programs in economics, although it is only recently that the literature has incorporated approaches that prioritize causal inference and econometric identification. For instance, a useful discussion of the existing literature on the multiplier impacts of cash transfers is presented in Gassmann et al. (2023), who conclude that "there is scant rigorous evidence" on the topic. Of the 23 studies they identified in their systematic review, the majority (13 studies) use local economy-wide impact evaluations (LEWIE), computable general equilibrium models (CGE), or social accounting matrixes (SAM), approaches that rely on economic model structure and often on strong and untested assumptions. Yet a growing body of evidence both on cash transfer impacts and beyond utilize modern econometric methods that employ experimental or quasi-experimental variation to estimate spillover and general equilibrium effects, and they are our focus here.

An influential approach to estimating spillover effects that has been widely used in economics was developed in Miguel and Kremer (2004), in their case to estimate externalities from treating an infectious disease although the approach can be applied more broadly in health and beyond (see Benjamin-Chung et al. 2017 and Benjamin-Chung et al. 2018 for examples). The Miguel and Kremer (2004) approach is attractive because it can be applied to RCT data even in cases in which the experiment was not explicitly designed to capture such cross-unit externalities; it relies on idiosyncratic variation in local treatment saturation induced by the randomization, in most cases by geographic proximity (although other dimensions of "closeness" could also be used, such as social ties). (For an entirely different nonparametric approach to establishing a lower bound on the size of spillover effects in the setting of a traditional RCT, refer to Choi 2021.) The following estimation equation, slightly modified from Miguel and Kremer (2004), illustrates the method:

$$Y_{ij} = \alpha + \beta T_i + \gamma_d N_{di}^T + \varphi_d N_{di} + \epsilon_{ij} \qquad (1)$$

where Y_{ij} is an outcome for individual j in treatment unit i (i.e., a school in Miguel and Kremer 2004), T_i is an indicator for the treatment status of the school, N_{di}^T is the local density of treatment individuals within d km, while N_{di} is the local density of all individuals within d km, and ϵ_{ij} is the error term, clustered at the school level. The γ_d term captures the spillover effect of an additional treated unit (within the appropriate distance) on the outcome of individual j. Some applications of the method in equation 1 alternatively utilize the proportion of treated units rather than the total number (including in robustness checks in Miguel and Kremer 2004 itself, and also see Dupas 2014), though in this case the interpretation of the γ_d coefficient is obviously somewhat different.

For all its attractiveness and ease of use, the Miguel and Kremer (2004) approach has several well-known limitations. First, it can obviously only be usefully applied in cases where the randomization led to sufficiently wide variation in local treatment exposure to allow for reasonably precise estimation of the γ_d coefficient. Second, equation 1 has a simple linearly separable form , that does not immediately allow for interactions or non-linearities; these can be accommodated with a more complicated functional form, although in some designs, including the original Miguel and Kremer (2004) study, the resulting higher-order estimates are not statistically precise.

Third, and perhaps most importantly, this approach and related methods that build on it all suffer from a key blind spot: spillover or general equilibrium effects that affect the whole study population (rather than just leading to differences across units) are missed entirely and just captured in the constant term. This is related to well-known limitations of other applied econometric methods, such as difference-in-differences approaches, which also miss economy-wide effects (Roth et al., 2023); this is sometimes called the missing intercept problem in macroeconomics. This opens up the possibility that the largest spillover and GE effects – those that strongly affect an entire region, say – are precisely those that are missed entirely, while smaller and more localized effects are successfully captured by the method. At the same time, to the extent that the estimated effects appear highly localized and can be shown to fade out over a distance that is far smaller than the extent of the study area, then concerns about missing overall ambient effects (intuitively) seem less severe. Starting with Miguel and Kremer (2004), many other studies have adopted this approach to investigate the extent of treatment spillovers of different types over various distances (e.g., Dupas 2014, Egger et al. 2022, among others).

This clearly unattractive element of the approach is part of the reason why alternative methods that attempt to leverage economic theory to capture broader impacts (including the LEWIE, CGE and SAM approaches alluded to above) remain important to consider and improve on going forward. There has been important recent work that attempts to combine experimental estimates with economic models and structural econometric methods to estimate GE effects, including the all-important area-wide GE effects that are typically missed in existing econometric approaches. This is still a nascent literature and we do not survey these methods here, but Bergquist et al. (2022) is one recent step in this direction.

More recent studies have built on the Miguel and Kremer (2004) approach in equation 1 by experimentally varying the treatment saturation (or density) levels at more aggregated geographic levels – say among a collection of villages – typically using a two-stage randomization (or even three-stage approach if there is also randomization of individuals' treatment status within villages), see Baird et al. (2014) for a useful discussion and Muralidharan and Sundararaman (2015) and Muralidharan et al. (2023) for related applications, among many others. This has the immediate benefit of allowing the researcher to control the amount of variation in local saturation (the N_{di}^{T} term in equation 1 above) and hopefully ensuring that it is sufficient to enable precise estimation of the spillover effect γ_d.

The Egger et al. (2022) study of cash transfer impacts in Kenya is an influential recent example of this approach in the SP space that is discussed below. A noteworthy aspect of that study is the large geographic scale of the study area and the ability to randomize treatment saturation at the level of the geographic sublocation, which is a group of typically 10 to 15 villages in their setting. The Egger et al. (2022) study utilizes a two-stage instrumental variable (IV) approach since the actual timing and delivery of cash transfers did not always exactly line-up with experimental assignment, although the first-stage relationship is very strong and so intention-to-treat (ITT) reduced-form estimates are similar.

Faridani and Niehaus (2022) formalize the properties of the class of linearly separable spillover estimators like those discussed here, with the goal of establishing whether they provide consistent estimates of what they call the average global effect (AGE), namely the total effect of treating the entire population. They conclude that the current approach has largely favorable properties under reasonable assumptions while recommending regressing on the share of nearby treated clusters, as opposed to treated units.

3. Cash Transfers

Two of the earliest spillover studies of cash transfer programs examined impacts among non-recipients of the well-known Mexico Progresa/Oportunidades program. The research design consists of comparing outcomes among non-eligible households in the treatment villages (where many households received the conditional cash transfer, CCT) versus non-eligible households in the control villages (where no households received the transfer). The difference is interpreted as the spillover effect of the program on these households. Recall that the Progresa program provided transfers conditional on households enrolling their children in school and visiting local health clinics.

The first study to employ this approach was Bobonis and Finan (2009) and they find that school attendance among children in non-eligible households increases in the treatment villages due to some combination of the cash transfers and the conditionality. They go on to conclude that "policies aimed at encouraging [school] enrollment can produce large social multiplier effects." Angelucci and De Giorgi (2009) subsequently adopted the same approach and find that the consumption of non-eligible households in treatment villages also rises (relative to control), due to a combination of social insurance and credit market effects.

Other scholars have documented broader impacts of the Mexico program by comparing treatment versus control villages. In terms of within household "category 1" outcomes, Bobonis et al. (2013) show that spousal physical abuse rates fall by 40 percent in treatment villages, although there is an increase in violent threats with no associated action. There is also evidence that CCTs can have adverse environmental consequences: Alix-Garcia et al. (2013) show that greater cash transfer receipt through the Mexico CCT program accelerates local deforestation, especially in areas with worse transport infrastructure.

We next discuss the Egger et al. (2022) study (mentioned above) in more detail, since it was explicitly designed to estimate broader impacts of an anti-poverty cash transfer (both in terms of treatment allocation and the data and measurement strategy), and thus can shed light on many of the central issues across multiple categories of impacts. The authors designed and carried out a large-scale experiment in rural Kenya providing one-time cash transfers distributed by the NGO GiveDirectly worth approximately USD 1000 to more than 10,500 poor households within 653 villages. Figure 1 (reproduced from the article) displays the study area, highlighting treatment and control villages as well as sublocations randomized into high and low treatment saturation. The experiment was designed specifically to address questions on the aggregate impacts of cash stimulus programs at the intersection of development and macroeconomics.

Figure 1: Study Area in Egger et al (2022), reproduced from the article

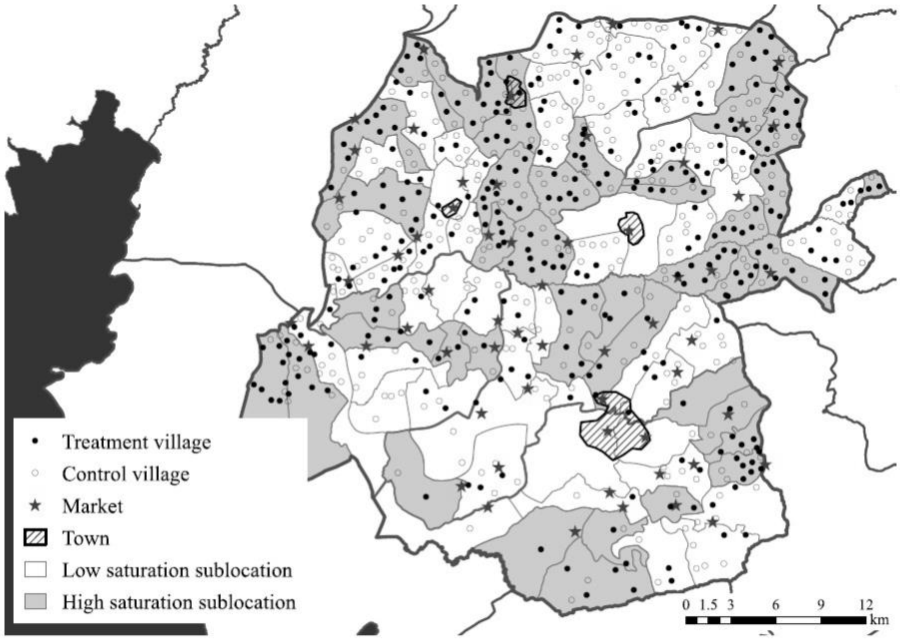

FIGURE A.2.—Study area. *Notes*: This figure plots the approximate location of study villages, sublocation boundaries, and weekly markets in the study area in Siaya County, Kenya. Control villages are denoted by hollow circles, treatment villages are denoted by solid circles, and blue stars indicate the locations of markets. High-saturation sublocations are shaded in gray, while low-saturation sublocations are those in white. Town boundaries are shaded with diagonal lines.

The program provides four main methodological advances: (i) the influx of cash is very large— more than 10M USD over 24 months, (ii) the randomization is across large units with spatial variation in exposure (at the sublocation level above the village level, as noted above), (iii) extensive measurement of outcomes for treatment and control as well as local markets and businesses, and (iv) a theoretical framework to provide structure to the results. The study took place in Siaya County in rural western Kenya, and the transfers took place from mid-2014 to early 2017, during a period of steady economic growth. At that time, GiveDirectly typically provided one-time unconditional cash transfers to poor households in LMICs (although their model has evolved over time), and required households to live in homes with thatched roofs for eligibility, a simple means-test for poverty. About 1/3 of all households met this criterion and were eligible. These households received 1000 USD (nominal) through M-Pesa made in a series of three payments over 8 months, which included a token transfer followed by two regular installments. The overall transfer was large in magnitude, in fact much larger than most large-scale government SP programs: the aggregate amount of the transfer was equivalent to approximately 16% of annual GDP in treated areas during the peak 12 months, and 24% of annual GDP during the full 24 month roll out.

The primary interest was estimating the average treatment effect (ATE) on outcomes for treated and untreated households and firms, including direct effects of village treatment as well as neighborhood effects at different radii from the village (examined empirical in 2 km bands, although in practice effects were found to be highly localized within 2 km or at most 4 km in most cases). To capture spillovers the authors estimate models closely related to equation 1 above in which a household's (or enterprise's) outcomes depended on the amount of money distributed in its own and other geographically proximate villages, as shown here (in a form that is slightly simplified relative to the original article):

$$Y_{iv} = \alpha + \beta \mathrm{Amt}_v + \sum_{r=2}^{R} \gamma_r \mathrm{Amt}_{v,r}^{\neg v} + \epsilon_{iv} \quad (2)$$

where Amt_v was the amount of cash per capita transferred to household i's own village v, $\mathrm{Amt}_{v,r}^{\neg v}$ was the amount of cash per capita transferred to villages other than v in a series of bands with inner radius $r - 2$ km and outer radius r km around the village centroid, and both amount variables depended on random assignment of the villages and sublocations to treatment and the (endogenous) share of eligible households. As a result, they instrument for them using the own-village treatment indicator and the share of eligible households in each band assigned to treatment.

In terms of main impacts, the authors estimate large positive direct impacts on consumption expenditures and durable goods for households receiving transfers, but no meaningful change in total labor supply for households in areas receiving more cash transfers. They also find that enterprises in these areas experience revenue gains and increases in sales – an important first indication of positive spillovers – without any meaningful change in investment behavior, leading them to believe expansion is largely driven from the demand side (at least in the short to medium run). Results also seem to be driven by the overall intensity of treatment in nearby communities instead of just the treatment status of the enterprise's village. Importantly, and echoing some of the findings in Angelucci and De Giorgi (2009), non-recipient households also experience large expenditure gains over the first three years after transfers were distributed. However, unlike that earlier Mexico study, these appear to be mainly driven by increases in income, including from wage labor earnings, rather than transfers among households or changing access to credit. In sum, impacts on “category 2” outcomes (using the definitions above) are largely positive.

In terms of “category 1” outcomes, Egger et al. (2022) provides evidence of generally positive or null impacts on various child outcomes in treated households (i.e., related to nutrition, education and health), and no meaningful effects – positive or negative – on measures of women’s empowerment or IPV. This latter result is in contrast to the Bobonis et al. (2013) study mentioned above showing meaningful reductions in Mexico, as do Hidrobo et al. (2016) in an RCT in Ecuador. Future work could usefully examine longer-term impacts on child life outcomes.

The ability to estimate broader “category 3” outcomes is arguably the most innovative aspect of the project. Egger et al. (2022) go on to investigate what the effects noted above for households and firms imply for the aggregate level of economic activity by computing a local transfer multiplier. This is enabled by their unusually detailed data collection – including censusing all local households and enterprises, as well as gathering detailed local public finance data – over multiple years, with the timing of data collection randomized to ensure representative data at each time point. Following the literature, they define this multiplier **M** as the cumulative effect of transfers on local real GDP, relative to the total amount T transferred, over a given time interval:

$$\mathbf{M} = \frac{1}{T}\left(\int_{t=0}^{t=\bar{t}} \Delta \mathrm{GDP}_t \, dt\right) \tag{3}$$

They obtain reassuringly similar local transfer multiplier estimates whether using income data or expenditure data, with a meaningful average multiplier estimate of 2.5.

This “experimental macroeconomic” estimate is somewhat larger than most recent fiscal multiplier estimates in rich countries like the US (Nakamura and Steinsson 2014, Chodorow-Reich 2019). Other recent studies in LMICs also present evidence of large positive multipliers of cash transfer programs. For instance, Gerard et al. (2021) finds positive local labor market effects of expansions of a household anti-poverty cash transfer program (Bolsa familia) in Brazil, which is consistent with the presence of a large positive multiplier, although they are not able to compute a local transfer multiplier in the same way as the Kenya study.

The study of local prices shows minimal price inflation in the Kenya sample, with effects typically of less than one percentage point even in peak periods of cash transfers. The authors argue that the increase in aggregate production, with minimal inflation, could be driven by increased factor productivity (labor productivity) rather than factor supply. They argue that the increase in productivity was made possible at least in part due to “slack” or under-utilization of factors in some enterprises, an issue that is worthy of further exploration. The evidence on the degree of price inflation due to cash transfer programs in other settings is mixed, with a large transfer in Mexico leading to only minimal price inflation (except for in the most remote areas, Cunha et al. 2019). Another project in the Philippines found larger impacts on the prices of certain food categories but not others, and this led to adverse child nutritional outcomes (Filmer et al. 2023).

Another key and largely unexplored dimension of large-scale SP programs is their fiscal implications. The Kenya GiveDirectly project also gathered detailed information on local tax and fee revenue collection as well as local public project spending and found minimal effects over a three-year time horizon after the distribution of the transfers (see Walker 2018 for a detailed discussion). It appears that little of the cash was "captured" by local government officials; future work should examine longer-term effects on local public finance outcomes.

Finally, the Kenya study examined a dimension of social conflict (another broader outcome), namely, crime victimization, and finds no meaningful overall impacts due to the cash transfer program. In contrast, the Crost et al. (2016) study finds that a conditional cash transfer program in the Philippines "caused a substantial decrease in conflict-related incidents in treatment villages relative to control villages in the first 9 months of the program." In terms of other work on the link between SP programs and conflict, Nunn and Qian (2014) find that an increase in US food aid to LMICs increases the incidence and duration of civil conflicts. The sharp differences in findings across these three settings are worthy of further investigation.

In terms of another dimension of broader "category 3" impacts, several studies have documented how receipt of a cash transfer can affect households' subsequent political attitudes, support and turnout. These issues are challenging to study in individual level studies due to the secret ballot, meaning that it is typically impossible for a voter to prove how they voted, and so some existing studies rely on survey self-reports. For instance, using a regression discontinuity identification strategy (comparing those just poor enough to receive assistance versus others), Manacorda et al. (2011) find that the large Uruguay PANES cash transfer program increased recipients' stated support for the government behind the program by 11 to 13 percentage points, perhaps since voters updated on the party's redistributive preference or competence. In a related finding, Conover et al. (2020) examines the impact of a conditional cash transfer program in Colombia at the polling center level and finds that voters in neighborhoods with higher program eligibility show a shift in support for the incumbent party as well as higher turnout overall. Clearly, the implementation and perceived success of large-scale SP programs can have feedback effects on political outcomes.

4. Other Social Protection Sectors

4.1 Pensions

There is a growing literature examining the effects of pension program expansion and reform on the recipients themselves (J-PAL and EPoD 2023), but as noted in the introduction, the body of research on spillover and GE effects is far more limited. It goes without saying that the design and implementation of old age pension programs – which have grown to become some of the largest public expenditure categories in wealthy countries – have major macroeconomic implications (which falls in our category 3) including on households' need to accumulate precautionary savings and for national savings rates, even if these have rarely been rigorously assessed in poor countries. Yet it has been argued that even some of the rapidly growing middle-income Asian economies, including China and India, have not yet established adequate elderly pension programs (ILO 2023, WHO 2023). We also know little about how formal social care for the elderly should be structured in LMICs. This is going to be increasingly important as these societies experience continued urbanization, migration, and growing elderly populations, with rising numbers experiencing Alzheimer's disease and related dementias (Duhon et al. 2023).

Most of the evidence on pension spillovers focuses on effects within the household (category 1) and finds evidence of heterogeneous effects as a function of the gender of both the pension recipient and children. The earliest and best-known work in this area examines the large-scale expansion of the South African old-age pension program to Black citizens in the early 1990s. South Africa is a middle-income country, at the time had the largest economy in Sub-Saharan Africa, and unlike most of its neighbors had sufficient government revenue to launch a generous program of this kind. Thirty years later, few Sub-Saharan African countries have set up similarly generous universal old-age pension programs.

An influential study by Duflo (2003) examines spillovers on the nutrition of children living in households with a pension recipient versus other households. Pensions received by men do not appear to impact child outcomes but those received by women lead to improved growth (presumably reflecting improved nutrition and health status) for girl children in the household, thus shedding light on differences in preferences by gender. Edmonds (2006) also studies the South African program and finds somewhat more positive effects when men receive pension on the schooling of boys as well as reductions in the child labor of girls and boys. Carvalho Filho (2012) examines the expansion of an old age pension in Brazil and, echoing the gender patterns in Duflo (2003), shows that pension receipt boosts 10 to 14 year old girls' schooling, and reduces girls' labor only when the pension is received by females in the household.

In terms of other within-household outcomes, receipt of the South African old age pension also appears to affect labor market decisions of other adult household members (Ardington et al. 2009), leading to higher employment due to increased labor migration, perhaps because improved household income provides more financial security for these mainly young adult job seekers.

4.2 Labor Programs

Other major non-cash SP sector programs intervene in the labor market. The most common type of program is workfare, which typically provides low-wage employment on public works projects for otherwise unemployed or underemployed workers. As noted in the introduction, the most famous and largest such program is

India's Mahatma Gandhi National Rural Employment Guarantee Act (which is also known by various acronyms though we will use NREGA here). The public hiring of workers in these programs makes them quite different from the cash transfer and pension programs discussed above. It is also not surprising that these interventions have been found to have meaningful impacts on local labor wages (Imbert and Papp 2015), especially when carried out at the scale of NREGA, employing individuals from tens of millions of households per year.

An important recent paper on the spillover and GE effects of workfare is Muralidharan et al. (2023), which focuses on the impacts of a treatment that improved implementation of and enrollment in NREGA (using biometric smart cards), which was randomized across 157 sub-districts in India. The authors had earlier found that the use of smart cards improved program targeting and led to an increase in participating workers (Muralidharan et al. 2016). This improvement in program efficiency and targeting led to direct gains for those enrolled in the program in terms of increased earnings and reductions in poverty. Importantly, these earnings gains mainly come from higher earnings from non-NREGA jobs: in fact, 86% of income gains come from their non-program earnings. Confirming earlier research, it also increased local wages as the NREGA wage provided a wage floor of sorts and increased workers' outside options (beyond the private labor market).

More relevant for this review, the intervention also significantly increased local employment, rather than reducing it as a simple price theory type analysis based on perfectly competitive labor markets might suggest (if the program effectively raises the local prevailing wage above what would occur in the absence of the program). While workers across the board, both those participating in NREGA and otherwise, appear to experience higher wages and employment in areas where NREGA was more effectively implemented, not everyone "won" from the treatment: large landowners appear to lose as land prices fall as do farm earnings per acre, presumably in large part due to higher labor costs. This detailed distributional analysis in enabled by a combination of detailed censuses of local economic activity plus rich surveys for a subsample of households, in an approach related to the Egger et al. (2022) Kenya study described above.

The study goes on to discuss the reasons for the perhaps counter-intuitive finding that a public workfare program could boost overall local employment and earning rather than simply crowding out private employment, as has been seen in some other settings (for instance, by Bertrand et al. 2021 in Ivory Coast). They find some evidence for two explanations. The first is that imposing a higher local wage can theoretically lead to more local hiring in models with monopsonistic local labor markets, which appears to be the relevant case for much of rural India (given a limited number of large landowners and high transport costs to work farther afield). Consistent with this view, they find the largest positive treatment effects on local employment in areas with higher levels of land ownership concentration. The second explanation is that NREGA expansion – and the resulting higher earnings for poor workers – constitutes a local aggregate demand shock that produces multiplier effects for local economic activity, providing further evidence for the empirical importance of multiplier effects of SP programs in LMICs, together with Egger et al. (2022) and Gerard et al. (2021), among others.

Franklin et al. (2023) also documents positive broader GE effects of a large-scale workfare SP program, in this case the Urban Productive Safety Net Program in Ethiopia, which was randomly rolled out across areas of Addis Ababa. The program provided an hourly wage that was substantially higher than the prevailing private wage, and understandably led to a shift towards public employment. At the same time, just as in Muralidharan et al. (2023), most of the welfare gains experienced by participants in the program were due to the higher local market wages that they received. The areas where the program was introduced also had improved local amenities (e.g., mainly higher quality of local sanitation, cleanliness, and sewage). Once again, the broader spillover and GE effects of the program in Ethiopia are central to understanding its impact, and a traditional individual level randomization design would greatly understate program benefits and possibly lead to different public policy recommendations.

Finally, the India NREGA employment program has also been shown to affect armed conflict, another category 3 outcome. The Fetzer (2020) study exploits the roll-out of the program (rather than a randomized design) to assess how the increased income stability provided by access to guaranteed employment affects the local sensitivity of armed conflict (mainly driven by Naxalite Maoist groups) to local rainfall shocks. There is a large body of research documenting that local weather shocks (both in terms of temperature and rainfall) can lead to greater risk of armed conflict and other forms of violence (see Hsiang et al., 2013 for an early review). The main contribution of Fetzer (2020) is to provide evidence that a large-scale SP program targeted to poor households can substantially weaken this relationship: the impact of extreme monsoon rain on conflict risk is large and statistically significant before NREGA expansion (see Panel A in Figure 2) but the slope falls to near zero after the program is introduced (Panel B). This finding points to the potential importance of SP programs in reducing conflict risk in the coming decades as most of the globe will experience increasingly severe climate and weather shocks.

Figure 2: The relationship between climate shocks and conflict, before and after NREGA, Fetzer (2020), reproduced from the article

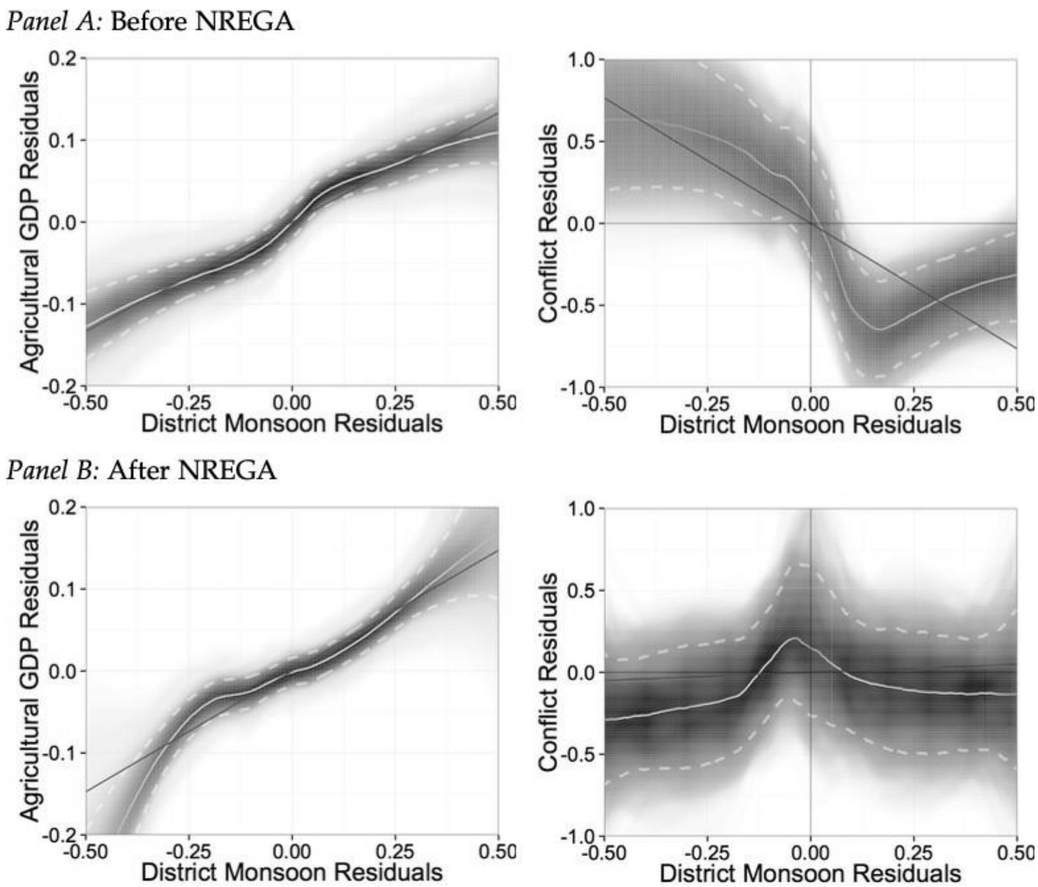

5. Conclusion

A curated survey of the existing literature finds many instances of mostly positive spillover and general equilibrium effects of social protection programs in low- and middle-income countries, including cash transfer, pension and labor market programs, with impacts on the three categories of impact laid out above, including on (1) other household members, (2) on the consumption and market earnings of other households and enterprises, and (3) large transfer multiplier effects, and reductions in armed conflict (in some but not all instances), as well as influencing political preferences and turnout.

Yet there remain many limitations and open questions. For one, the number of estimates of these broader impacts remains quite limited (especially relative to the size of the research literature investigating direct effects, as discussed in other chapters in this Handbook). This is not surprising given that relatively few studies are purposely designed to estimate these types of effects, as they typically require larger sample sizes and more complicated coordination with government authorities and NGOs than standard RCTs due to their scale.

The limited coverage of existing work raises immediate external validity concerns. For instance, the Egger et al. (2022) study of broader cash transfer impacts in Kenya finds evidence of large transfer multiplier effects that could alter conclusions about the attractiveness of large cash transfers targeted to poor households. But as they point out, there remains considerable uncertainty about what the multiplier would be if the program were targeted to other settings (i.e., urban areas) with greater density of economic activity and lower transport costs, or to households with higher living standards with a lower marginal propensity to consume the transfers. The macroeconomic environment could also matter, and their finding of minimal price inflation might not hold if a program were brought up to national scale and funded through government tax revenue (rather than external donor assistance), and so on. Similar issues are relevant for other types of SP programs: for instance, an important driver of the large GE effects in the Indian labor market documented in Muralidharan et al. (2023) appears to be the structure of local labor markets, and in particular, the extent of local employer monopsony power. Simply put, the existing estimates surveyed in this chapter may not be readily portable to other settings. Taken together, these and other results make a strong case (in the author's opinion) for making deliberate investments in more research in this area an important scientific and policy priority (see Muralidharan and Niehaus 2017 for a related discussion and call for additional research).

There are other noteworthy gaps in the existing body of evidence, including the relative lack of estimates of the fiscal externalities of large-scale SP programs (with some notable exceptions like Walker 2018), and limited research on broader impacts of pensions and health care system reforms. These issues will only increase in scholarly and policy importance as LMIC populations experience rapid aging in the coming decades. Another open question is which types of individuals who receive social protection assistance (or employment) will generate the largest spillover effects for others. In some cases, it may be that poor households are the ones who create the greatest broader benefits (for instance, if they have the highest marginal propensity to consume transfers, leading to a larger local demand shock) but in other cases it might be somewhat better-off households, especially if part of the broader gain is driven by an investment response (Haushofer et al 2022). This is a question that has rarely been tackled in existing research but one that could be valuable to explore to improve program design.

A related but distinct concern with the existing literature is the possibility of publication bias. This is a potential problem with any literature with a limited number of studies, many of which are statistically under-powered (especially due to the fact that some spillover estimates are derived from standard individual or community level RCTs that were not designed to estimate broader impacts), raising the possibility that the published estimates are those that, by chance, yielded large magnitudes and statistically significant results, but that these are not representative of the body of estimates as a whole – which are left unpublished and largely unknown to the research community if they yield null spillover estimates. Again, a deliberate research program of large-scale RCTs designed to estimate spillover and GE effects, across the major SP program types and sectors, would go a long way towards addressing these concerns.

The bottom line is that gaining a more robust understanding of the broader – spillover and general equilibrium – economic, political and social impacts of social protection programs remains a top research priority given the critical role that SP programs play in redistribution to poor and vulnerable households, in order to move low- and middle-income countries towards greater equality of opportunity and fairness. Getting the design of SP programs "right" will only become more important in the coming years and decades given the accelerating change in the global climate and its associated shocks.

Author Acknowledgements: I am grateful to Sheah Deilami for excellent recent assistance, Paul Neihaus and Michael Walker for useful conversations, and Pascaline Dupas and Ben Olken for helpful comments.

A full set of references for all Handbook chapters, including this one, can be found at this link: https://www.dropbox.com/scl/fi/9lqs2mdrawkjdrv4m648e/References-Social-Protection-Handbook.pdf?rlkey=jt0f8kute31mhdke77aoiw99d&st=kd7l8ff1&dl=0