The Spillover Effects of Excludable Cash Transfers: What the Miracle Cure for Development Woes Means for Infant and Child Mortality[1]

Marika Csapo
marika.d.csapo@vanderbilt.edu

Center for the Study of Democratic Institutions
Vanderbilt University
August 26, 2016

[1]Acknowledgements to follow.

Abstract

Brazil's *Bolsa Família* (Family Stipend) Program is the world's flagship conditional cash transfer program. The excludable transfer is conditioned upon school attendance and the use of primary health care services—a design often lauded by the development community for its ability to produce long-term improvements in human capital in addition to providing short-term poverty relief. However, researchers have found the anticipated health benefits of the Program to be somewhat elusive. This article relies on evidence from aggregate and individual-level datasets on infant and child mortality rates to investigate the ambiguous findings regarding the role *Bolsa Família* plays in overall health outcomes of the poor. The analyses suggest that a failure to account for possible spillover effects (i.e., violations of the SUTVA assumption) of the Program with respect to the non-recipient poor could undergird the considerable variability in causal claims regarding the relationship between the *Bolsa Família* Program and health outcomes.

Keywords: Latin America; Brazil; Development; Infant Mortality; Policy Spillovers; Social Programs

1 Introduction

Within the first year after the Worker's Party (*Partido dos Trabalhadores*) gained control over the executive office in Brazil in 2003, they implemented the *Bolsa Família* (Family Stipend) Program (hereafter, BFP). The Program gives small amounts of cash directly to poor families provided that their children attend school and that the family complies with routinized preventative health care visits. The transfer of funds was designed to go directly from the federal government to the private account of the female head-of-household in recipient families. The Program has been celebrated internationally as a behavior-modifying, bottom-up development model, which simultaneously improves human capital and long-run development prospects, while helping to alleviate poverty in the short-term.

Yet the predicted human capital benefits—in particular, those related to health—have proven somewhat elusive. Some studies claim to show unequivocal health improvements while others find them to be absent. Studies do not necessarily agree on which outcomes are relevant. Perhaps more importantly, they often do not agree on which social subset is the relevant population for evaluating the Program's health effects. Many studies look exclusively at elicited, retrospective self-reports of recipients. Others rely on aggregate data and ecological inference. While some studies compare recipients to a baseline established by non-recipients, these studies typically focus on one temporal cross-section and inference is based on the assumption that the non-recipient group has not been affected by treatment externalities or spillover effects. If this assumption were incorrect, net positive externalities for the poor non-recipients would lead to attenuation bias in the estimated health effects of the BFP. Alternatively, if existing externalities were to net negative for the non-recipient poor, these cross-sectional analyses would tend to overestimate the effects of the BFP treatment on the treated (i.e., recipients) while failing to capture the negative externalities on non-recipients.

In this paper, I argue that when it comes to questions of development, one should not only care about the effect of the treatment on the treated. This is especially true if the treated group does not exhaust the universe of poor citizens, as in the case of the BFP. Any spillover effects ought to be explicitly measured and included in any assessment of the treatment effects of the Program. Applying this logic to the high-profile conditional cash transfer program in Brazil, *Bolsa Família*, in this paper I use individual data to parse the aggregate effects of the Program on both recipients of the Program and poor non-recipients. I find that while the BFP has lead to improvements in the infant and child mortality outcomes of poor recipients, the same program seems to have negative spillover effects on the infant and child mortality rates of poor families that are not BFP recipients in that it tempers improvements among this population. In the aggregate, the tempering effects appear to dominate.

Health Benefits of Conditional Cash Transfers—for Whom?

The commonly relied-upon zero spillover assumption (i.e., the assumption of SUTVA compliance) in the evaluation of the health effects of the BFP invites reasonable pause.[1] After all, several papers have shown potential spillover effects of *Bolsa Família* on various non-health outcomes such as crime (Chioda, De Mello and Soares, 2012) and economic growth and consumption (Angelucci and De Giorgi, 2009; Santos, 2010). These (and any other) non-excludable goods affected by the Program will also impact the non-recipient poor, not just the group directly treated by the Program. This logic extends to non-excludable health goods, if their provision is influenced by the Program.

The question of policy spillovers, or externalities for non-recipients, is likely to be extremely relevant in the way *Bolsa Família* affects infant mortality rates for several reasons. Though the BFP is rivalrous and excludable, it is necessarily noisily assigned (Soares, Ribas

[1]The stable unit treatment value assumption, or SUTVA, is meant to simplify individual-level inference about treatment effects. It asserts that the treatment is the same for all treated individuals and that the treatment does not alter the observed outcomes of individuals who do not receive the treatment—i.e., that there are no positive or negative spillover effects of treatment on the untreated.

and Soares, 2010). At current coverage rates, it excludes a considerable volume of eligible poor families as well as many poor ineligible families. The program's targeting is considered very good in that the number of families receiving program coverage that fall outside of the lowest two quintiles of income is very low (Lindert et al., 2007). However, it is also true that as a logical baseline, at least 37.5% of the poorest two quintiles of society are not covered, since presently, 25% of the total Brazilian population is included in the Program. In fact, due to the tendency toward income volatility among the poor and the fact that the BFP updates its income records only every two years for each family, an estimated 44% of the eligible poor were not covered even after the 2006 expansion of the Program (Soares, Ribas and Soares, 2010). This is important, since any resulting changes in the quality or amplitude of non-excludable public services, such as public primary health care services, would be shared with non-recipients.

Bolsa Família could produce such changes in the provision of primary care services if it has the effect of shifting recipient preferences over the provision decentralized health services. Such a shift in preferences among recipients of an income shock is somewhat intuitive. The transfers are often used to consume health goods previously provided to this group through public primary care. More money in-pocket to purchase medications or improve anthropometric outcomes through improved food security are among the positive but excludable benefits that accrue to *Bolsa Família* beneficiaries as a result of the Program.[2] Since the transfer benefits are excludable substitutes for non-excludable public primary care services, the program could lead to decreases in non-excludable services. This service effect would then spill over onto non-recipient families. Furthermore, the service conditions that beneficiaries face have many possible implications. For example, they are likely to result in increased referral for more advanced medical procedures. This type of outcome could actually shift recipients' demands *away* from primary care and toward more complex procedures,

[2] A report by IBASE indicates that 87% of BFP recipients report spending at least part of the transfer on food, which should greatly affect nutritional health; 22% of recipients report spending at least part of the transfer money on purchasing medications; 2% report spending it on other medical treatment (Menezes et al., 2008).

which share a budget with primary care at the municipal level.

Other unintended externalities are also plausible. It, therefore, remains unclear what the net effect of the Program should be for the welfare of non-recipients. The Program is likely to increase the institutional access of recipients by creating an institutional link between them and the local Social Assistance bureaucracy. One important source of institutional advantage that comes with BFP assignment is that many other social programs specifically target BFP recipients. These targeted programs range from preferential assignment to other social programs such as *Minha Casa, Minha Vida* (federally-provided housing for poor families) and *PRONATEC* (federally-funded vocational training for the poor), school meetings coordinated between beneficiary families and local political officials, and even the simple increase in preventative health care visits, where resultant referrals to hospitals outside the municipality require interaction with the local Secretary of Health.

This implies at least some negative externalities incurred from non-recipient status in the BFP. Since *Bolsa Família* assignment also serves as the unofficial shortcut assignment mechanism for many other anti-poverty programs, the assignment errors (even if effectively random for the BFP) are then correlated across programs. This means those eligible non-recipients (as well as poor but ineligible non-recipients) will be systematically disadvantaged with respect to the benefits of many social assistance programs. This, in turn, may have implications for the health of these non-beneficiaries.

2 Health Outcomes and Conditional Cash Transfers

The Healthcare Context of the Poor in Brazil

In Brazil, health care is public and theoretically universal, though private care is also available at a cost. About 20% of all Brazilians purchase private health care plans, but this action is highly correlated with income (Cataife and Courtemanche, 2014). The type of income that comes from the BFP would generally not be sufficient to make the difference between whether a family purchases a private plan or not. Families that are eligible for the BFP are generally far from such considerations, with their first financial concern being food security. Brandão et al, for example, estimate from self-reports that participation in the BFP reduces to 82.6% the number of recipient families that ran out of food before having money to purchase more in the previous three months, down from 87.5% prior to recipient status (Brandão et al., 2007). There also continues to exist considerable unmet health care demand among the poor (Cataife and Courtemanche, 2014).

The *Bolsa Família* Program and Health Outcomes in Brazil

In order to capture the total effects of the BFP on the health of the poor, one would need to consider both the direct effects on the explicitly treated and any spillover effects on the non-recipient poor. Ideally, one would like to be able to distinguish between these two types of effects. Existing literature examining the health impacts of the BFP in Brazil is unable to do this. Often researchers rely (explicitly or implicitly) on the SUTVA assumption to ease the high demands of research design in a context where individual-level data with temporal variation is difficult to come by.

Generally, ambiguity in inference regarding the effects of the BFP on health outcomes in Brazil is produced by three factors. First, the BFP operates through more than one mechanism, so different outcome variables may pick up different local effects. Conflating the range of health perceptions, behavior and outcome variables as substitutable proxies for health produces ambiguity of inference. Meta-studies that review the health effects of conditional cash transfer programs across contexts may tend toward ambiguous findings due to the fact that they aggregate studies over dependent variables like clinic visits, food security and anthropometric measures, and perceived health outcomes like self-reports of illness frequency (Lagarde, Haines and Palmer, 2009; Shei et al., 2014). The most consistent health results with respect to the BFP seem to be improvements in the food security perception of recipients (Menezes et al., 2008; Brandão et al., 2007; de Bem Lignani et al., 2011). Results from health outcome measures have often proven less tidy across studies.

Though the retrospective before and after studies of recipients' perceptions on food security are consistent across studies, they are symptomatic of a second pathology with respect to development studies inferences. The choice to restrict analysis to the temporal effects of the Program on only the subset of the poor that are beneficiaries is a valid measurement strategy—albeit, one that is subject to limited generalizability. Interpretation of these results as representative of the overall social impacts of the Program is equivalent to ignoring any externalities and taking the treatment effect on the treated as the whole health story. The effects of the treatment on the treated—or, arguably, the posterior self-perception of treatment effects by the treated—is an interesting quantity in and of itself. However, from a development perspective, it is not sufficient to answer the question "how does the *Bolsa Família* Program impact the overall welfare of the poor?" since "program recipients" and "the poor" are overlapping, but not completely intersecting population subsets.

Finally, studies that include both beneficiaries and non-beneficiaries of the BFP are typically cross-sectional and conditioned on a single point in time. These include studies of anthropometric measures[3] (Paes-Sousa, Santos and Miazaki, 2011), frequency of preventative health visits (Shei et al., 2014), and infant mortality rates (Rasella et al., 2013). Non-recipients are taken as a baseline, which is equivalent to assuming there are no externalities of the program for the non-recipient poor. Resulting inferences regarding treatment effects on the recipient population based on observed group differences combine treatment effects experienced by both beneficiaries and non-beneficiaries, with the entire difference between groups, and only this difference, being attributed to recipients alone.

[3]Anthropometric measures are physiological markers such as height and weight, which are often taken as measurements of overall health.

CCTs and Health in Latin America and the World

Meta-studies of health effects of conditional cash transfer programs across Latin America have led to ambiguous conclusions, probably largely due to the conflation of several different health-related dependent variables and the sole focus of the studies on recipients. For example, one review shows moderate evidence of beneficiary improvements in the frequency of preventative health care visits and nutrition and anthropometric measures (i.e., those indicators that follow directly from program conditionalities).[4] With respect to more general health outcomes including child mortality, diarrheal illnesses and respiratory infections, results were less consistent or absent, with Mexico and Colombia's programs producing somewhat clearer results than the other cases (Lagarde, Haines and Palmer, 2009). Similarly, a study of the health effects of conditional cash transfer programs spanning countries in Latin America, Africa, Asia and the Middle East finds some evidence of increased use of preventative health visits associated with conditional cash transfer programs, but generally finds "mixed" results with respect to overall health outcomes.

One potentially fruitful approach to assessing the overall health impacts of conditional cash transfers on the poor (pooling over recipients and non-recipients) is to use a catch-all measure of health outcomes. For example, Diaz-Cayeros, Estévez and Magaloni examine Mexico's conditional cash transfer program's effect on infant mortality rates (Díaz-Cayeros, Estévez and Magaloni, 2012). This study finds overall improvements in infant mortality rates geographically associated with cash transfer assignment. However the study is limited to looking at changes in infant mortality rates between two Census years (1990 and 2000), which makes the causal link fuzzy, given that the program of interest was first implemented in 1997. Their model also neglects to account for changes in the economic situation of the geographic units over this time.

Barham gets the closest to distinguishing between the general health impacts of

[4]The BFP is not included in this analysis, but an earlier Brazilian CCT, *Bolsa Alimenta cão*, is included. Most of the 10 studies are on Mexico, though Colombia, Nicaragua and Honduras are included as well.

conditional cash transfer programs on the recipient and non-recipient poor. She looks at the case of infant mortality in Mexico, as well, explicitly differentiating between the program's association with the population infant mortality rate and the program's association with the average infant mortality rates for the rural, beneficiary poor (Barham, 2011). She finds that though the program is highly associated with improved (decreased) infant mortality rates for the rural poor, the program has no association with overall mortality rates, which implies some heterogeneity of effects of the program. Similarly, this paper assesses heterogeneity in the impacts of Brazil's BFP on the health of the poor. It attempts to explicitly parse out one form of heterogeneous treatment effect which may be at the root of many of the confusing findings on conditional cash transfers and health. It supplements aggregate-level analyses with individual-level data. It is thus able to distinguish empirically between those overall health effects which pertain to recipients and those that pertain to the non-recipient poor.

3 Methodological Approach

Operationalization of the Dependent Variable

In order to assess the health and human capital impacts of the BFP, it is important to choose a broad measure of health that is sensitive to the most serious health threats to the poor—and, ideally, one with clear links to the Program's requirements. While vaccination requirements are explicitly included in the BFP, Brazil already had very high rates of vaccination coverage even before the start of the BFP. As such, changes in vaccination rates in recent years are very small and cannot relay much information on recent changes in health outcomes. There are also measurement issues related to vaccination reports since many children get innoculated multiple times against the same illness as the result of recurring vaccination campaigns. Each vaccination gets counted independently making it difficult to interpret innoculation coverage in percentage terms.

Fortunately, the BFP also requires preventative health visits for children and pre- and post-natal visits for pregnant women, so the Program also ought to be linked closely to overall infant and child mortality rates. Infant and child mortality rates serve as a strong barometer of the overall welfare of the poor (Ross, 2006). In terms of measurement, these rates are general enough to aggregate any and all grave health issues faced by the poor. This includes fatal cases of the most common diseases affecting poor children, including diarrheal and respiratory infections and waterborne illnesses (the last of which still accounts for almost a quarter of preventable illnesses in Brazil).[5] They are observable, documented, and constitute an event unlikely to be vulnerable to recall-issues with respect to self-reports. Furthermore, the fact that others have used this measure to investigate the health impacts of conditional cash transfer programs in various contexts means that carefully studying the measure in light of these previous analyses can contribute to an accumulation of knowledge including the identification of new interpretations of the current body of findings.

Empirical Identification Strategies

I employ four rigorous empirical identification strategies in this research. First, I assess the impact of the BFP on overall health outcomes (including both recipient and non-recipient outcomes), using three approaches. I use a grouped logit model of infant mortality rates with multivariate control for the lagged infant mortality rate as well as various economic, development and state intervention variables. The observations are aggregate (municipal) infant mortality rates and these are observed annually over the period 2005–2010. This method is the closest to those models used in previous published studies and so its usage is important for comparison purposes. Second, I compute a differenced OLS model for the change in infant mortality rates, conditional on the previous year's level and controls. Third,

[5]See `http://data.unicef.org/resources/pneumonia-and-diarrhoea-tackling-the-deadliest-diseases-for-the-world-s-poorest-children` and
`http://www.worldpolicy.org/blog/health-care-brazil-300-year`.

I repeat the differenced model, first subclassifying the data on the official BFP target rates used in assignment before computing the total average treatment effect.

The aggregate-level results may be vulnerable to any of a number of barriers to inference. The last method therefore uses individual-level data with cross-sections observed for two time periods to attempt to deal with problems of ecological inference, to assuage concerns of selection effects, and to identify differing (heterogeneous) recipient and non-recipient treatment effects. It consists of the analysis of the probability of infant or child mortality, given live birth, of children born to poor households for each of the two groups of interest. These analyses are based on data from a large household survey, the *Pesquisa Nacional de Amostra dos Domicílios* from 2004, and on Census data from 2010, both of which were administered by the *Instituto Brasileiro de Geografia e Estatística* and which identify the BFP status of respondents. I use child mortality data to complement the infant mortality analysis at the individual level to bolster the statistical power of the analysis, since birth is a relatively rare event in the population at large.

These methods address the three concerns I highlighted above with respect to the literature in the following ways. I deal with the multiple-mechanism phenomenon by focusing on a very general health outcome measure, rather than perceptions of health or frequency of doctor's visits or some combination. I avoid the problem of limiting inference to only the treatment effect on the treated by including both beneficiaries and non-beneficiaries in all of the analyses. First I average over the two, presenting the total treatment effects. Then I use individual data to decompose the total effect into the respective contributions of each of the groups of interest. I describe each of these methods in more detail below.

Lagged Grouped Logit Analysis of Municipal Data, Pooling over Time

The multivariate grouped logit design models the probability a certain number of infants die in a municipality in a given year, conditional on the number of live births observed that year in that municipality. The model pools over annual data from the 2005–2010 period. Grouped logit models are convenient for estimating proportions where both the numerator and denominator, and not just their ratio, are known. This is ideal since a very high observed ratio may either result from a high probability event with many trials, or from computing the observed ratio based on a very small number of trials (as may occur in many small municipalities). Grouped logit helps to make appropriate use of the variability in binomial sample (trial) sizes across units. Infants born in municipalities with a small number of births may have a different likelihood of dying than those born in large municipalities and the grouped logit helps account for this. In this model I also estimate a dispersion parameter for the quasibinomial distribution to allow for additional flexibility surrounding the assumption of independent Bernoulli trials with respect to infant deaths within a municipality.[6]

Though the proportion of households assigned to the BFP is likely correlated with development, which itself is a predictor of infant mortality, I rely on multivariate control to try to tease out the relationship of interest. I include as control variables the municipal GDP per capita, the Human Development Index score of the municipality for education,[7] the per capita size of federal government transfers to the municipal government, the population size, the percentage of that population that lives in rural conditions, the percentage of the population that is covered by the Family Health Program (a model of primary health care introduced in the 1990s that combines primary care with community health surveillance), the lagged infant mortality rate from the previous year, and the independent variable of interest, the percentage of families in the municipality that receive *Bolsa Família* transfers.[8]

[6]Since the dispersion factor was always greater than 1, I did not estimate an extended-beta binomial since the true dispersion parameter was unlikely to be negative.

[7]I include only the education component of the Human Development Index, since including the other components, GPD and health outcomes, would be akin to including two controls for GDP and to moving the left hand variable to the right hand side, respectively.

[8]The infant mortality data was gathered from the Pan-American Health Organization and is originally from the Ministry of Health in Brazil. The other variables come from the databases of IPEA, the *Instituto de Pesquisa Econômica Aplicada*, IBGE, the *Instituto Brasileiro de Geografia e Economia*, the National Treasury of Brazil, and DATASUS, the database of the public health system.

Differenced OLS Analysis of Municipal Data, Pooling over Time

This model includes the same control variables as the specification described above, however, it models the dependent variable as the difference in the infant mortality rate since the previous year. Infant mortality is a temporally slow-moving variable and its level is likely correlated with unobserved variables that mutually influence municipal BFP coverage. Looking only at the changes in infant mortality rates helps to ensure that the model only explains post-treatment (BFP assignment) variation in mortality. As the binomial specification is no longer appropriate for this model, I use Ordinary Least Squares estimation. I still include the lagged infant mortality rate, as it seems likely that the rate of change in infant mortality rates is conditioned to some degree upon the previous year's level.

Subclassified Differenced Model on Official *Bolsa Família* Target Rates

As concerns over the plausibility of the exogeneity of BFP assignment are still reasonable, I use an additional method to attempt to parse out an average causal effect. The BFP is assigned at the federal level based on household income reported during a family assessment interview and verified by official employment records. The number of stipends awarded in each municipality is a function of official target rates generated formulaically from aggregate income estimates and geographic development goals. Due to the frequency of income volatility among the poor as well as informal sector employment that is impossible to verify, the income assignment mechanism is necessarily noisy (Soares, Ribas and Soares, 2010). As the target rates represent the official assignment goals of the Program and observed variation in assignment from these target rates is generated by noise, I use subclassification on the official target rates to try to isolate the exogenous component of assignment, similar to the method proposed in Imai and Van Dyk (Imai and Van Dyk, 2004). Others have used a similar approach to try to estimate the causal effect of *Bolsa Família* assignment on voting behavior with respect to incumbent candidates (Zucco, 2013). I subclassify municipalities into 10 subgroups of roughly equal sample size based on their targeting rates. I then repeat the differenced analysis described above, including all the control variables, for each of the 10 groups separately, so that the BFP variable only captures observed municipal treatment variation between municipalities with like target rates. Finally I compute the weighted average of these 10 estimates to get the estimate of the total treatment effect.

Primary and Spillover Effects on the Poor from Time-Varying, Individual Data

The aggregate-level analyses cannot say much about how the BFP affects different segments of the population. This final analysis uses household survey data from 2004 and 2010 to estimate individual-level effects of recipient status on the health of the poor. One benefit of using self-reports for measuring income mortality is that it may be subject to fewer sources of measurement error than official hospital data. Not all babies are born or die in hospitals. Furthermore, the rates of births and deaths occuring in hospitals, rather than at home, may vary over time. Infant mortality rates also often aggregate multiple sources of information, making it possible to double-count single instances. However, the delivering mother has complete information and high recall on the birth and (with very few exceptions in the data) the death of any infants within the last year. The obvious tradeoff is that self-reports may be subject to under-reporting due to social desirability bias. However, this effect is unlikely to change much over the short time period examined in this analysis, meaning it should only affect intercept estimates for the average probability of infant mortality, but not within-group changes in mortality rates over time.

I estimate the cutoff for the poorest two household income quintiles using the household survey data on within-household per capita income. I then classify all sampled individuals falling below this cutoff who do not receive BFP stipends as the "non-recipient the poor." "Recipients" are those that self-identify as participants in the the *Bolsa Família* Program. From these two subsets I identify all individuals who report having given birth in the year (five years) before the reference date. These births are taken as the sample of infants (children) used for estimation. I use a linear probability model (i.e., I compute the average number of deaths among these children) for the two groups in 2004 and 2010 to construct a difference-in-differences estimator.

This methodology, though an improvement on those methods used to address this question in the past, is still subject to two limitations. The first is that, despite the availability of temporal variation for the two groups, the data are not panel data and the recipient status label is not necessarily static over time. The second is that it assumes (as do all infant mortality analyses, regardless of whether individual or aggregate) that the groups being compared do not differ greatly in the shape of the distribution of the time of year their children are born or the age of the mother, etc. There is some support for the assumption that the two groups are as-though randomly assigned among the poor with respect to infant mortality rates. In 2004 (the first year of the program's existence after it was created in October of 2003), the average infant mortality rates of the two groups do not differ significantly, as shown in Figure 3.[9]

4 Results

Grouped Logit Results

Table 1 below shows the results of the grouped logit analysis. The coefficient of interest on BFP coverage in a municipality (conditional on municipal income, last year's infant mortality rate, and the HDI education score, among the other control variables) is positive and

[9]Note that though *Bolsa Família* was initiated by President Lula in late 2003, a few municipalities had already implemented local versions of conditional cash transfer programs emphasizing educational conditionalities in the late 90s and President Cardoso had implemented a series of smaller national conditional cash transfer programs in 2001, including one with health requirements. The unified assignment mechanism of the BFP however, the *Cadastro Único*, was conceived as part of the consolidation and expansion of these programs in 2004 under Lula under the rubric *Bolsa Família* (Lindert et al., 2007). It is this assignment mechanism that subsequently became standardized across many social programs beyond the BFP, leading to correlated program assignment.

highly statistically significant. That is, conditional on the control variables and pooling over time, the BFP seems associated with relatively worse infant mortality outcomes.[10]

Table 1: Grouped Logit: Dependent Variable—Infant Deaths | Live Births

Variable	**Coefficient** (Std Err)
Bolsa Família Coverage (Sq Rt)	0.028***
	(0.003)
HDI Education Score	-0.579***
	(0.036)
Federal to Municipal Transfers (Log)	0.013
	(0.009)
Municipal GDP per cap (Log)	-0.039***
	(0.007)
Population (Log)	0.024***
	(0.003)
Percent Pop Living in Rural Conditions	-0.001***
	(0.000)
Family Health Program Coverage	-0.0002*
	(0.000)
Lagged Infant Mortality Rate (Log)	0.334***
	(0.007)
Intercept	-4.783***
	(0.087)
N	29,742
Quasibinomial Dispersion Parameter	2.41

***0.01 α-level; **0.05 α-level; *0.1 α-level

Not surprisingly, last year's municipal infant mortality rate is a strong predictor of this year's. Other predictable control results are that the probability an infant dies during its first year of life is a decreasing function of the municipality's education score, income, and coverage of families by the Family Health Program. Infant mortality appears to be positively associated with the size of the municipality's population, which is a plausible result given that infectious illness spreads more rapidly in densely populated regions. Infant

[10] Though the results are not included here, the bivariate grouped logit results where only BFP coverage is included on the right-hand side also produces a positive and highly significant coefficient, as does the grouped logit model of infant mortality with only BFP coverage and lagged infant mortality on the right-hand side.

deaths also seem to be negatively associated with rural living conditions, though the size of the effect is miniscule. Despite the fact that infectious diseases spread fastest in cities, this result is somewhat surprising since hospitals tend to be much less accessible to rural communities. However, the absence of hospitals may mean an increased likelihood of child delivery outside of the context of official medical care, and thus, underreporting of deaths may be more prevalent. The size of per capita transfers from the federal government to the local government, after accounting for the BFP and the Family Health Program does not produce a statistically significant result.

Figure 1 plots predicted municipal infant mortality rates (per 1,000 live births) over the range of all possible municipal coverage rates of the BFP, with a 95% confidence interval. The histogram of the actual distribution of *Bolsa Família* coverage across municipalities in the data is provided for more realistic interpretation of the range of plausible results. The predictions were computed holding all variables at their mean values, except municipal GDP, which is held at its 25th percentile value. For a municipality with these covariate values, an increase from no BFP coverage in one year to 15% BFP coverage (close to the modal coverage rate of 14.6% in the time period covered by the data) in the next year is associated with an average of about 1.7 additional infant deaths per 1,000 live births (an increase from 14.8 to 16.5).

My finding thus contrasts with that of Rasella et at., who find that the BFP is negatively associated with child mortality rates across municipalities (i.e., that the Program decreases the likelihood of deaths of children under five) (Rasella et al., 2013). One possible reason for the difference is that Rasella, et al.'s modeling strategy resulted in the elimination of nearly half of the municipalities from the dataset they used.[11] Admirably, they execute their analysis for various different dependent variables, looking not just at childhood mortality rates, but also assessing mortality from various specific causes. To maintain a consistent

[11] It is not possible to say with certainty whether selection on the dependent variable explains the difference in the results of the two studies because the authors were unable to provide the data for the purposes of replication due to an ongoing project.

Figure 1: Predicted Municipal Infant Mortality Rate

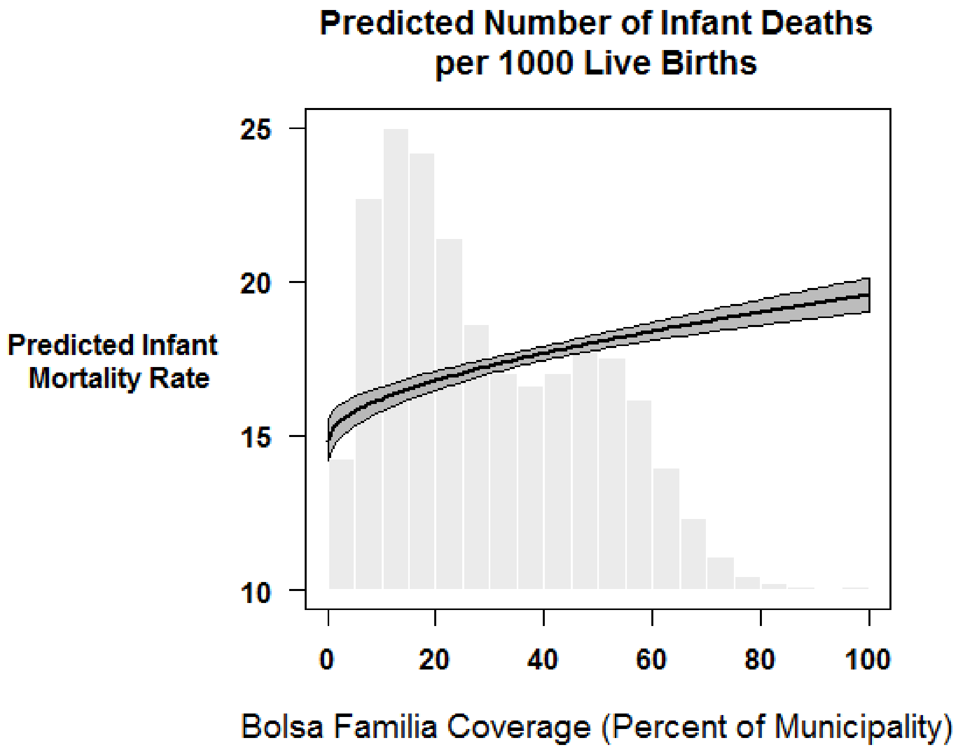

sample of cases across the different dependent variable analyses, they remove from their sample any municipality with missingness on any one of these dependent variables during the years 2004-2006 (the first three years in their dataset). As a result, they execute their analysis on 2,853 of the 5,564 municipalities in Brazil, or 51% of the cases. By contrast, Table 1 presents results from an analysis executed on a sample that includes over 90% of the universe of municipalities.

It is unlikely that the cases with imperfect health documentation are randomly distributed. In fact, research has shown that municipalities in the poorest parts of the country, especially in the north and northeastern regions, are most likely to have imperfect health data (Szwarcwald et al., 2002). I engage in a bootstrapping exercise to approximate the frequency with which one would obtain a negative coefficient (as in the Rasella, et al. analysis) using the model I have presented in Table 1 when the analysis is executed on a similarly reduced sub-sample of municipalities. Figure 2 provides the sampling distribution of coefficients on the BFP resulting from 100,000 iterations of analyses with different random samples of combinations of 2,853 municipalities chosen from among the 5,564. Less than one-half of one percent of iterations (0.49%) of randomly-selected subsets of this size produce a negative result using this model.[12]

Figure 2: Simulated Coefficient Variability: Lagged Model

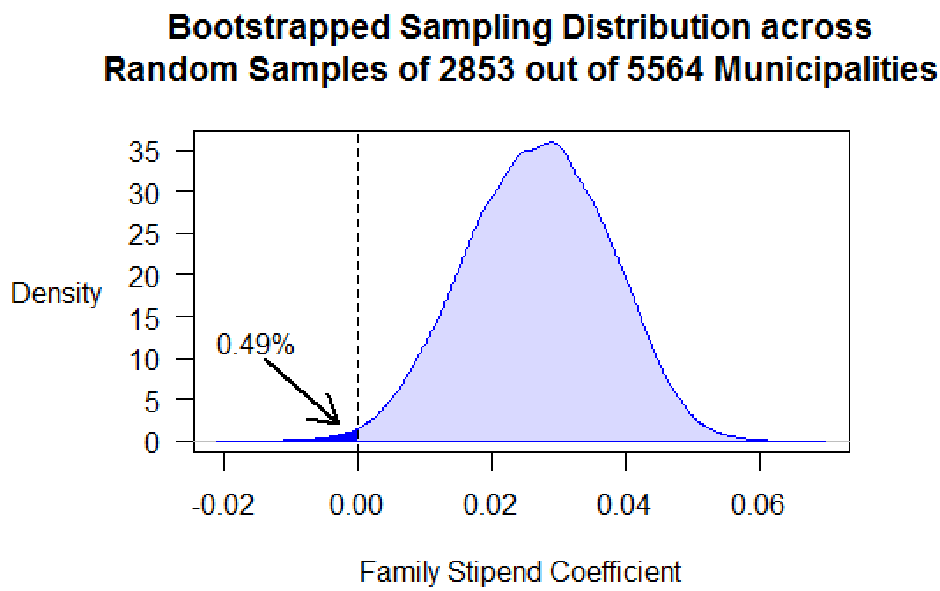

OLS Results for Change in Infant Mortality

Arguably, the ideal measurement of the dependent variable is the change in infant mortality rate associated with *Bolsa Família* treatment assignment, and not the level. One benefit

[12]Rasella, et al. use a somewhat different model specification. The most important difference is that they use municipality fixed-effects rather than the lagged infant mortality rate. Since maximum likelihood estimation of generalized linear models with geographic fixed effects is known to be inconsistent when executed on panel data with many cross-sectional cases but with few years (Nerlove, 1971), I do not take this approach. However, linearly controlling for the mean municipal infant mortality rate across years in the sample does not alter my results significantly. The statistical and substantive interpretation of the results in Table 1 are also robust to an alternative negative binomial specification with the dependent variable being the mortality rate, which is the method preferred by Rasella, et al. I do not use this model for my core specification because it does not adequately treat the varying number of "trials" (births) across municipalities of different sizes and so is most appropriate for count data with a fixed trial size, rather than ratios.

of this model specification is that it models only post-treatment movement on the dependent variable, while allowing this movement to be conditioned upon pre-treatment levels. Furthermore, in the presence of the lagged measure of infant mortality rate, this method removes a considerable amount of multicollinearity from the model. The results presented in Table 2 are from an OLS estimation procedure that uses the same independent and control variables as the model presented in Table 1, including the measure of last year's level, but substitutes the difference in infant mortality rate from the last year to this as the dependent variable. The results are comparable to those from the model presented in Table 1 in direction, magnitude and significance, despite the shift in model specification. This model also shows a positive association between the BFP and increases in infant mortality (or at least fewer decreases) relative to municipalities with fewer BFP recipients.

Table 2: OLS: Dependent Variable—Change in Infant Mortality Rate

Variable	**Coefficient** (Std Err)
Bolsa Família Coverage (Sq Rt)	0.913*** (0.242)
HDI Education Score	-21.867*** (2.714)
Federal to Municipal Transfers (Log)	-10.169*** (0.923)
Municipal GDP per cap (Log)	-0.925* (0.545)
Population (Log)	0.940** (0.389)
Percent Pop Living in Rural Conditions	-0.031** (0.014)
Family Health Program Coverage	-0.005 (0.008)
Lagged Infant Mortality Rate (Log)	-13.551*** (0.209)
Intercept	105.285*** (9.197)
N	29,742
Adjusted R^2	0.124

***0.01 α-level; **0.05 α-level; *0.1 α-level

Subclassification on Target Rates Results

Subclassifying the data on official BFP target rates for municipalities before executing the differenced model decreases the power of the BFP result in terms of statistical signficance, though the coefficient is still statistically significant at the 0.1 α level. However, the size of the BFP coefficient does not decrease, indicating a loss of efficiency is responsible for the attenuation of significance, rather than bias-correction. The substantive interpretation of the result with respect to the focal independent variable, therefore, remains the same. In municipalities with similar BFP target rates, higher BFP assignment is still associated with increases in infant mortality. These results are reported in Table 3 below.

Table 3: OLS with Subclassification: Dependent Variable—Change in Infant Mortality Rate

Variable	**Coefficient** (Std Err)
Bolsa Família Coverage (Sq Rt)	0.912* (0.478)
HDI Education Score	-12.000*** (3.436)
Federal to Municipal Transfers (Log)	-6.272*** (1.441)
Municipal GDP per cap (Log)	0.073 (0.661)
Population (Log)	2.833*** (0.492)
Percent Pop Living in Rural Conditions	-0.021 (0.015)
Family Health Program Coverage	-0.006 (0.010)
Lagged Infant Mortality Rate (Log)	-12.840*** (0.194)
Intercept	41.113*** (14.391)
N	29,742

***0.01 α-level; **0.05 α-level; *0.1 α-level

Difference-in-Differences with Individual Data

The net total of treatment effects on all members of society presented above is still difficult to interpret substantively from a development perspective. Do the total effects reflect the impact of the BFP and its externalities on the poor? Can the effects of the BFP be further partitioned into hetergeneous effects on recipients and non-recipients among the poor? These questions require individual-level data to support responsible inference. Table 4 below shows the results of a difference-in-differences estimator for the probability of infant death given live birth. The results weakly suggest that those assigned to receive the BFP transfers in the first full year of the program's existence (2004) were actually those with the highest risk of infant death among the poor. However, by 2010, the recipient group had surpassed the non-recipient poor in improvements in infant mortality.

Table 4: Diff-in-Diff: Dependent Variable—Linear Probability of Infant Mortality

Variable	**Coefficient** (Std Err)	**p-value**
Expected % Death for Poor Non-recipients in 2004	1.534*** (0.183)	0.000
Additional Expected % Death for Recipients in 2004	0.012 (0.559)	0.983
Non-recipient Decrease by 2010	-0.223 (0.186)	0.230
Additional Recipient Decrease by 2010	-0.105 (0.562)	0.852
N	173,266	

***0.01 α-level; **0.05 α-level; *0.1 α-level; N includes 4,235 and 169,031 observations in 2004 and 2010.

Since infant death is a relatively rare event and there were only 4,235 live births in the sample in 2004 (and only 453 of those belonged to BFP recipients) the results of the individual-level analysis are statistically weak. Therefore, I reproduce the analysis using individual child mortality data (estimating the probability of death before reaching five years of age). This considerably increases the statistical power.

Figure 3 below shows the frequency of child deaths per 100 living children (under the age of five) for recipients and for non-recipients in the poorest two per capita household income quintiles of society in 2004 and 2010. Using the child mortality measure increases the number of observations in 2004 to 17,032. Far more data are available for the Census year, 2010, than for 2004, which explains the differing sizes of the 90% confidence intervals for the two years. Similarly to the infant mortality results, the plot shows that in 2004—the first year after the start of the program in October of 2003—the recipient group had a somewhat more pronounced risk of child deaths.[13] The recipient group then experienced considerable improvement by 2010—in fact, converging to the trajectory for the whole country including better-off non-recipients. The non-recipient poor, despite starting off in better shape than recipients, seem to have experienced markedly slower improvements after the BFP was implemented. By the end of the 2010 period the non-recipient poor group had fallen far behind the national trajectory. This suggests the presence of negative spillover effects associated with the BFP for non-recipients.

Table 5 provides a numerical summary of the difference-in-differences results, which are statistically signficant at the 0.1 level. Taken in combination with the aggregate results, the implications of these results are important, if surprising given the conventional wisdom about the Program's immense development success. These findings may help to account for the mixed results of health assessments of the Program.

[13]Therefore, there is no evidence that non-recipient families have not registered for the Program because the parents are less competent, which would also affect child mortality, or on other such dimensions.

Figure 3: Probability of Child Deaths among Recipients and Poor Non-recipients

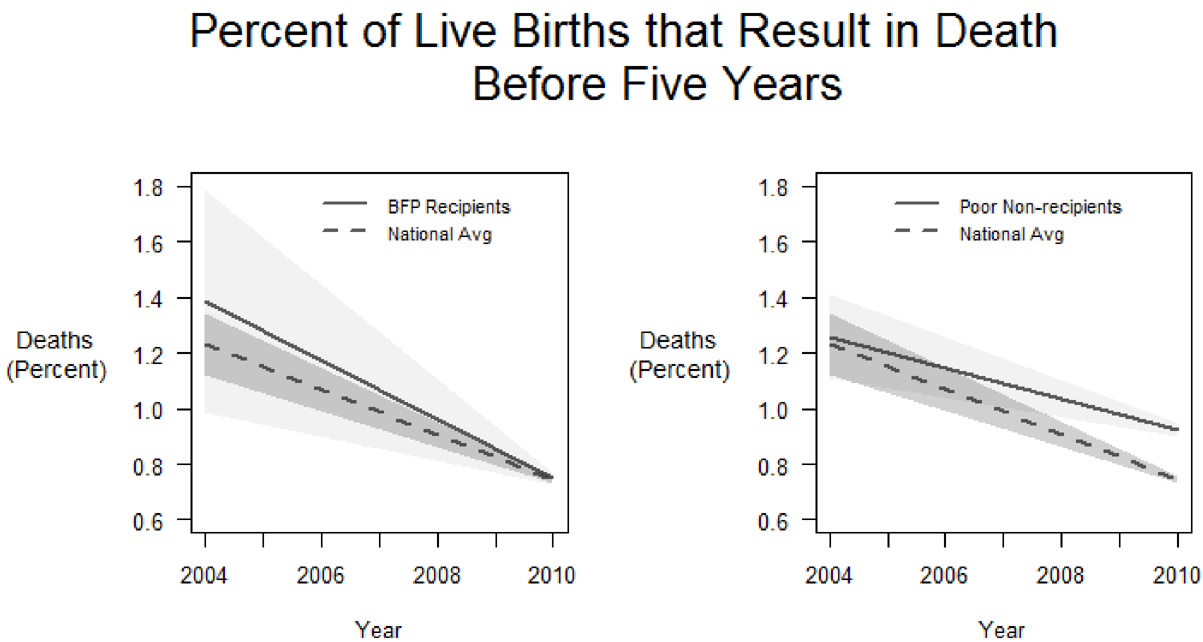

Table 5: Diff-in-Diff: Dependent Variable—Linear Probability of Child Mortality

Variable	**Coefficient** (Std Err)	**p-value**
Expected % Death for Poor Non-recipients in 2004	1.258*** (0.075)	0.000
Additional Expected % Death for Recipients in 2004	0.219 (0.211)	0.300
Non-recipient Decrease by 2010	-0.335*** (0.077)	0.000
Additional Recipient Decrease by 2010	-0.393* (0.212)	0.064
N	798,526	

***0.01 α-level; **0.05 α-level; *0.1 α-level; N includes 17,032 and 781,494 observations in 2004 and 2010.

One might question the validity of the survey instrument, since self-reports can be subject to any of a number of downward social biases, compared to official medical data. Fortunately, it is possible to compare country-wide estimates of infant mortality rates for 2001 - 2010 from the individual data with the official aggregates provided by the Ministry of Health that were used in the municipality-level analyses.[14] Figure 4 shows a comparison of the population-weighted aggregate with the estimates I calculate for the whole population from the inverse-sampling-weighted individual-level self-report data (with 95% confidence interval). The results seem to indicate that downward bias in self-reporting does exist, though at a fairly consistent rate over the 2004-2010 period included in the difference-in-differences analysis. Furthermore, the self-reports may capture year-to-year variability in deaths occurring outside of the official medical context documented by the Ministry of Health. This makes sense as the groups that do not have hospital access are likely to overlap with those whose infant mortality rates are highly sensitive to economic fluctuation as well as variability in the disease burden for any given year.

5 Discussion

The results suggest that the presence of the *Bolsa Família* Program has different effects on recipients and the non-recipient poor. The combination of the individual and aggregate data indicate that the Program may actually temper the survival prospects of poor non-recipient children—perhaps even to a greater degree than it helps poor recipient children. Given the necessarily noisy assignment of the BFP, this interpretation is consistent with the aggregate finding that, averaging over all residents including recipients and non-recipients of the Program, the BFP is associated with relatively worse(ning) infant mortality outcomes.

These results may change the way we interpret past findings regarding the health

[14]Unfortunately, the household survey used to compute these results did not regularly include a question identifying the BFP status of the family.

Figure 4: Data Quality

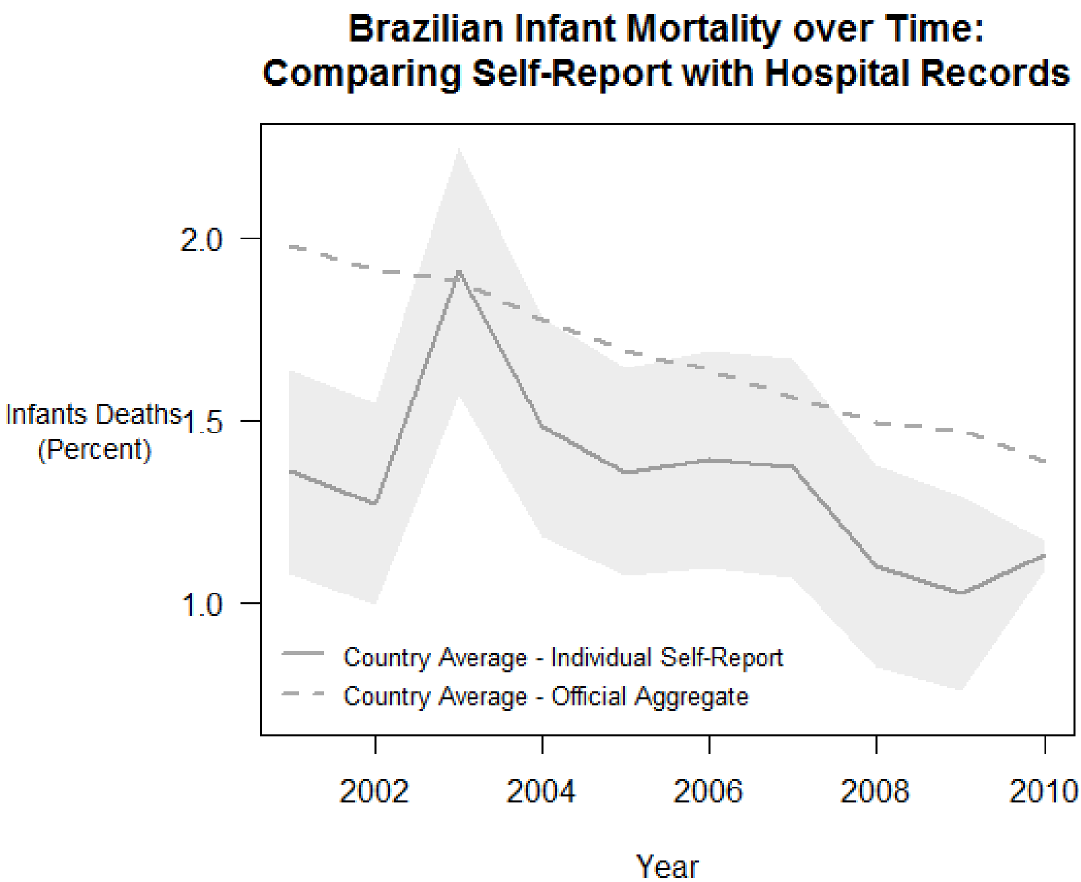

implications of development programs like the BFP by adding context to the ambiguous "mixed" health results that studies which marginalize over recipients and non-recipients tend to find. From the standpoint of causal inference, if we look only at the average treatment effect of the Program on the treated (i.e., on recipients), thus ignoring any spillover effects, we may conclude, as many have, that the BFP is associated with improved overall health outcomes.

However, if we return to our objective of evaluating the development impacts of the Program with respect to overall health outcomes, then we ought to take into account spillover effects of the Program onto the non-recipient poor. Once we allow for this possible SUTVA violation and attempt to estimate heterogenous effects of the Program we confront an inevitable reality—it becomes impossible to resolve the overall development question without first specifying some objective function that takes as inputs not just average health outcomes, but also the distribution among them.

Policy Implications

Since not all the poor receive *Bolsa Família* transfers and non-recipient status seems associated with some negative externalities, certain policy implications follow. Perhaps the most important implication is that the practice of using BFP status as the default preferential assignment mechanism for most other social programs operated under the Ministry of Social Development could potentially exacerbate negative development externalities. Currently BFP recipients receive preferential access or assignment to programs such as *Minha Casa, Minha Vida* (government housing), the *Programa Nacional de Acesso ao Ensino Técnico e Emprego* and *Próximo Passo* (technical training and employment), *Brasil Alfabetizado* (literacy program), *Olhar Brasil* (eyeglasses for Brazil), *Aqui Tem Farmácia Popular* (free distribution of medications), *Rede Cegonha* (specialized treatment for pregnant mothers and infant children), and *Saúde nas Escolas* (health consultations in schools).[15]

This approach takes what may well be random noise in BFP assignment and makes it non-random by correlating treatment status across various development programs. Exclusion from the BFP, even given true eligibility then implies systematic exclusion from many social programs which could contribute to individual health and welfare. Ironically, the *Cadastro Único*, or the Single Unified Registry, which serves as the basis for the assignment mechanism of the BFP, is oft taken as the most developmentally-progressive component of the Program. Santos, et al., capture the optimism regarding the generalization of the use of the Registry beyond BFP assignment when they write, "systems integration looks like the next step for building a better and more egalitarian society" (Santos et al., 2011). In fact, the single

[15]See also `http://www.mds.gov.br/falemds/perguntas-frequentes/superacao-da-extrema-pobreza%20/plano-brasil-sem-miseria-1/plano-brasil-sem-miseria`.

assignment mechanism, though logically-consistent and administratively-efficient, may be statistically-flawed and this fact should be included in the evaluation of its institutional implications.

6 Conclusion and Further Research

This paper presents evidence that the *Bolsa Família* Program in Brazil seems to have a net harmful effect on infant mortality rates and that the program is associated with both gains for recipients and negative externalities for the non-recipient poor. Several lessons may be taken from this research. First, it is important to match the subpopulation under evaluation with the question at hand. When the success of a policy is evaluated from a development perspective, it may well be that the relevant population is not just the group assigned to the policy or the target group, but the whole of the poor population or an even broader subset, depending on the plausibility of the SUTVA assumption and one's definition of "development."

With the recent emphasis in the literature on identifying causal relationships, there has often been a cautious narrowing of the generality of claims of findings. This cautiousness is appropriate; however, there is no reason why the quest to identify causal relationships should be treated as mutually exclusive with the answering of broad practical questions. A failure to identify all the appropriate subpopulations of interest and to parse out the relevant heterogeneous effects attributable to spillover effects may undergird the considerable variability of causal claims regarding the relationship between the BFP and health outcomes.

Second, this paper highlights the fact that logically-consistent policy is not always statistically reasonable policy. For example, the use of a single assignment mechanism to identify the group of poor voters eligible for a wide range of pro-poor policies is logically consistent. However, if the assignment mechanism is noisy (i.e., some people that should be enrolled in these programs are not and others that should not be are), then the errors and spillover effects will be correlated across policies. That is, the single-but-noisy assignment mechanism implies policy externalities will pile-up on the same group of poor voters who are accidentally excluded from recipient-status. What might be the minor consequence of essentially random exclusion from one program becomes a major consequence due to subsequent systematic exclusion from a host of pro-poor programs. Overlooking this fact can have problematic development implications. If we acknowledge noise in assignment to treatment we may be forced to acknowledge that unobserved assignment noise is best left uncorrelated by assignment design—in particular, when the spillover effects of the policy have yet to be identified or have been identified as negative. This implies that, at least in some situations, the efficiency gain of simplifying assignment to one mechanism across programs can have a significant development cost.

The mechanisms through which the spillover effects identified in this paper operate are not completely clear. Several relevant factors are at play. Many who have found ambiguous health results point to supply-end provision issues with the local government (Rocha, 2009; Soares, Ribas and Osório, 2010). However, since the local government is democratically elected, provision should be, at least in part, endogenous to demands and institutions of preference aggregation. If the BFP shifts preferences or demands through one or more of its income-shifting or behavior-incentivizing mechanisms, we may see changes in the provision or supply of health services. This may include non-excludable health services, such as primary care, used by much of the population including the non-recipient poor. These avenues are to be explored in later research.

References

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

Barham, Tania. 2011. "A healthier start: the effect of conditional cash transfers on neonatal and infant mortality in rural Mexico." *Journal of Development Economics* 94(1):74–85.

Brandão, A, S Dalt, VH Gouvêa, J Vaitsman and R Paes-Sousa. 2007. "Segurança alimentar e nutricional entre os beneficiários do Programa Bolsa Família." *Vaitsman J, Paes-Sousa R. Avaliação de políticas e programas do MDS: resultados: Bolsa Família e assistência social. Brasília: Ministério do Desenvolvimento Social e Combate à Fome* pp. 99–115.

Cataife, Guido and Charles Courtemanche. 2014. "Income-based disparities in health care utilisation under universal health coverage in Brazil, 2002–2003." *Global public health* 9(4):394–410.

Chioda, Laura, João De Mello and Rodrigo R Soares. 2012. "Spillovers from conditional cash transfer programs: Bolsa Família and crime in urban Brazil.".

de Bem Lignani, Juliana, Rosely Sichieri, Luciene Burlandy and Rosana Salles-Costa. 2011. "Changes in food consumption among the Programa Bolsa Família participant families in Brazil." *Public health nutrition* 14(05):785–792.

Díaz-Cayeros, Alberto, Federico Estévez and Beatriz Magaloni. 2012. "Strategies of vote buying: Democracy, clientelism and poverty relief in Mexico." Book Manuscript.

Imai, Kosuke and David A Van Dyk. 2004. "Causal inference with general treatment regimes." *Journal of the American Statistical Association* 99(467).

Lagarde, Mylene, Andy Haines and Natasha Palmer. 2009. "The impact of conditional cash transfers on health outcomes and use of health services in low and middle income countries." *Cochrane Database Syst Rev* 4(4).

Lindert, Kathy, Anja Linder, Jason Hobbs and Bénédicte De la Brière. 2007. "The nuts and bolts of Brazils Bolsa Família Program: implementing conditional cash transfers in a decentralized context.".

Menezes, Francisco, Edmar Gadelha, Mariana Santarelli, Rozi Billo, Delaine Martins Costa, Jacy Corrêa Curado, Luciene Burlandy, Rosana Magalhães, Rosana Salles da Costa, Ismênia Blavatsky de Magalhães et al. 2008. "Repercussões do Programa Bolsa Família na segurança alimentar e nutricional das famílias beneficiadas." *Rio de Janeiro: IBASE* .

Nerlove, Marc. 1971. "Further evidence on the estimation of dynamic economic relations from a time series of cross sections." *Econometrica: Journal of the Econometric Society* pp. 359–382.

Paes-Sousa, Rômulo, Leonor Maria Pacheco Santos and Édina Shisue Miazaki. 2011. "Effects of a conditional cash transfer programme on child nutrition in Brazil." *Bulletin of the World Health Organization* 89(7):496–503.

Rasella, Davide, Rosana Aquino, Carlos AT Santos, Rômulo Paes-Sousa and Mauricio L Barreto. 2013. "Effect of a conditional cash transfer programme on childhood mortality: a nationwide analysis of Brazilian municipalities." *The lancet* 382(9886):57–64.

Rocha, Cecilia. 2009. "Developments in national policies for food and nutrition security in Brazil." *Development Policy Review* 27(1):51–66.

Ross, Michael. 2006. "Is democracy good for the poor?" *American Journal of Political Science* 50: 860–874.

Santos, L. 2010. Bolsa Familia Programme: Economic and Social Impacts under the Perpective of the Capabilities Approach. In *13th BIEN Congress, Sao Paolo, Brazil.*

Santos, Leonor Maria Pacheco, Romulo Paes-Sousa, Edina Miazagi, Tiago Falcão Silva and Ana Maria Medeiros da Fonseca. 2011. "The Brazilian experience with conditional cash transfers: A successful way to reduce inequity and to improve health." Paper in the World Conference on Social Determinants of Health in Rio de Janeiro on October 19–21.

Shei, Amie, Federico Costa, Mitermayer G Reis and Albert I Ko. 2014. "The impact of Brazil's Bolsa Familia conditional cash transfer program on children's health care utilization and health outcomes." *BMC international health and human rights* 14(1):10.

Soares, Fábio Veras, Rafael Perez Ribas and Rafael Guerreiro Osório. 2010. "Evaluating the impact of Brazil's Bolsa Familia: cash transfer programs in comparative perspective." *Latin American Research Review* 45(2):173–190.

Soares, Sergei, Rafael Perez Ribas and Fabio Veras Soares. 2010. "Targeting and coverage of the Bolsa Família Programme: Why knowing what you measure is important in choosing the numbers". Technical report Working Paper, International Policy Centre for Inclusive Growth.

Szwarcwald, Célia Landmann, Maria do Carmo Leal, Carla Lourenço Tavares de Andrade and Paulo Roberto Borges de Souza Jr. 2002. "Infant mortality estimation in Brazil: what do Ministry of Health data on deaths and live births say?" *Cadernos de Saúde Pública* 18(6):1725–1736.

Zucco, Cesar. 2013. "When payouts pay off: Conditional cash transfers and voting behavior in Brazil 2002–10." *American Journal of Political Science* 57(4):810–822.