Expanding Governance as Development: Evidence on Child Nutrition in the Philippines

Eli Berman
elib@ucsd.edu
UCSD, NBER

Mitch Downey
pmdowney@ucsd.edu
UCSD

Joseph Felter
jfelter@stanford.edu
Stanford

July 15, 2015

Abstract

Worldwide, extreme poverty is often concentrated in violent, unstable, ungoverned spaces. Researchers and practitioners struggle to effectively reach these areas with traditional development assistance. Expanding governance by controlling territory through coercion may have both benefits and costs for local residents, especially if the transition to government control is violent. We estimate for the first time whether a large counterinsurgency program improves development outcomes, exploiting the staggered roll-out of the Philippine Army program ("Peace and Development Teams"), which treated 12 percent of the population between 2002 and 2010. Though treatment increased violence, the program progressively reduced child malnutrition, by 32% after two years and 49% after six. That reduction compares favorably with conventional child health interventions, which might anyways be infeasible in the weakly governed spaces of the rural Philippines. These findings invite an evidence-based discussion of expansions of governance, an extensive margin of development.

Keywords: Conflict and development; Counterinsurgency; Malnutrition; Philippines

JEL Classification Numbers: F51, I15, O22.

"Thousands of children are killed every year as a direct result of fighting - from knife wounds, bullets, bombs and landmines, but many more die from malnutrition and disease caused or increased by armed conflicts... Any disease that caused as much large-scale damage to children would long ago have attracted the urgent attention of public health specialists."

— Graça Machal, Expert of the Secretary-General of the United Nations (1996)

1 Introduction

Conflict and instability are central challenges in the implementation of development assistance. As Table 1 illustrates, countries receiving Western development assistance tend to be not only poor, but also conflict-cursed. The table lists the top 15 recipients of Official Development Assistance (ODA) per capita from the World Bank, the United States, and the United Kingdom. Along with average annual ODA, it also lists measures of poverty, political instability, and conflict. These 15 countries include four of the top five and five of the top ten most unstable nations. Seven of these countries have been involved in conflict for at least 15 of the last 25 years, with Sudan and Colombia involved in conflict for all 25.

[Table 1 about here.]

The top recipients of foreign aid are typically violent, unstable places for two reasons. First, poor governance in general and violence in particular undermines investment of all types, including in human capital and in institutions. This has occurred in Somalia, Afghanistan and Sudan, for instance, where average GDP per capita is less than $2 per day.[1] Second, ODA is often directed at countries such as Iraq and Colombia with the objective of stabilizing their ungoverned spaces, even when they are not among the poorest. Foreign assistance is a major segment of these economies. From 2003-2012, ODA from these three sources made up more than 5% of GDP in nine of these countries and more than 25% in the case of Afghanistan.

Despite the centrality of conflict in development, policymakers continue to struggle to pair assistance with security. There are two key challenges. First, many development programs require the transfer or use of valuable resources like cash, food, or equipment. These resources invite predation and targeting by rebel groups.[2] Second, if development programs win popular support for government (Berman, Shapiro, and Felter, 2011b), rebel groups may violently disrupt project implementation or deter the government from continuing.

[1] The World Development Indicators had no GDP per capita on record for Somalia.

[2] See Grossman (1999), Collier (2000), and Nunn and Qian (2014).

Recent empirical work underscores this challenge. Using a regression discontinuity design Crost, Felter, and Johnston (2014) show sharp increases in rebel-initiated attacks in Philippine villages following the award of small-scale infrastructure grants. Using a multi-country panel, Besley and Persson (2011) show that in the absence of strong institutions, increases in foreign aid significantly increase the onset of large-scale political violence. Nunn and Qian (2014) show that US food aid to conflict-prone countries increases the likelihood and duration of civil conflict. These findings suggest that aid *can* intensify conflict.

How might agencies pair development assistance with security? Policymakers have struggled to answer this question. One option implemented during the wars in Afghanistan and Iraq was development programs explicitly selected and protected by the military. Discussing development as part of a counterinsurgency strategy in Afghanistan, General McChrystal, commander of NATO forces, said: “We view it as a process, and not an event, which enables Afghan ownership and reinforces Afghan sovereignty... In some areas, it will be security assistance. In some areas, it will be less military, and it will be more based on help with governance and development...” (McChrystal, 2010), and that “We will not win simply by killing insurgents. We will help the Afghan people win by securing them, by protecting them from intimidation, violence and abuse (Hall and McChrystal, 2009).” Berman et al. (2011b) provide empirical evidence that small-scale military development projects reduced violence in Iraq. Moreover, Berman et al. (2013) show that both military and nonmilitary development projects are more violence-reducing the greater the security presence.

Does enhancing security and governance improve development outcomes? The answer is theoretically ambiguous. We tend to think that a state will provide institutions that are more welfare enhancing than will rebels, but states may also neglect populations in the periphery, and even if they did not the transition to state governance is often a coercive and perhaps destructive process, sometimes accompanied by abuses of human rights. To date, no research has addressed this question, despite the volume of assistance targeted towards unstable countries. In this paper, we attempt to fill that void in the context of a large counterinsurgency program operated by the Armed Forces of the Philippines: “Peace and Development Teams” (PDT). Exploiting the program’s staggered roll-out over nine years, we estimate effects on one of the few development outcomes available annually for Philippines municipalities: child malnutrition.

Treatment is clearly associated with increased violence –much of which is initiated by government. Yet the program progressively reduced child malnutrition, by about 32% after two years and 49% after six. Those figures are subject to caveats about possible selection bias and about feasible scale, yet they are quite large –as large as the most successful child health interventions in the development literature.

In the next section, we discuss the Philippine conflict and the design of the PDT program. Section 3 describes our data. In Section 4 we consider non-random selection of where PDT is implemented. Since we lack a strictly exogenous source of variation in implementation, in principle Section 4 provides guidance in selecting control variables to estimate the treatment effect of PDT free of selection bias. In practice, Section 5 reveals that estimated treatment effects seems robust to the inclusion (or absence) of variables predicting selection, beyond local trends and location x year fixed effects. Section 6 concludes with policy recommendations, comparing these estimated effects to those of other malnutrition interventions and speculating on the larger question of secure governance in development.

2 Context

The Philippines has suffered low-grade civil conflict for decades.[3] While there are a number of rebel groups, they can be broadly classified into two categories. The first are Islamic separatists, primarily the Moro Islamic Liberation Front (MILF) and the Moro National Liberation Front (MNLF), but also including the smaller and more radicalized Abu Sayyaf Group (ASG). These are primarily active in the country's South in the southwestern areas of Mindanao and the Sulu Sea, ostensibly fighting for an independent Islamist state. Past compromises between Islamist militants (primarily the MNLF) and the federal government have significantly expanded the scope of local authority, with the establishment of the Autonomous Region of Muslim Mindinao (ARMM) in 1996, one of the Philippines' seventeen regions.[4]

The New People's Army (NPA), the armed wing of the Communist Party of the Philippines (CPP), form a second category. Over the last decade, the NPA has been the most active rebel group, accounting for nearly two-thirds of violent incidents (Crost et al., 2014). Unlike the MILF and MNLF, the NPA is a Maoist revolutionary group seeking to overthrow and replace the established government and is broadly active in the Philippines as a whole.

Importantly, the formal government faces little risk of being overthrown. The asymmetry of power is heavily tilted in its favor. The conflict is characterized by small-scale insurgent attacks rather than frequent full-scale battles. In situations like these, governments struggle for enough information to help them battle insurgents (Berman and Matanock, 2015). Government authority is relatively uncontested in urban areas, where rebel groups have limited popular support and opportunities for cover. Figure 1 displays violent incidents and mal-

[3]Here, we provide only a brief review of these conflicts. The interested reader should see Crost and Johnston (2010), Felter (2005), Hernandez (2014), Quimpo (2012), or Schiavo-Campo and Judd (2005).

[4]Regions are the largest subnational division of government in the Philippines.

nutrition, our key development outcome, by population density, showing that both conflict, and poverty are concentrated in rural areas.[5]

[Figure 1 about here.]

As part of its efforts to gain popular support for government, the Philippine Army launched the PDT program in 2002. According to the program manual special Army units were designated to enter selected villages (*barangay*),[6] clear out entrenched rebels, assess community needs, and connect the village to government programs and services. This might include building schools or clinics, protecting local business or markets from rebel extortion, securing roads to nearby villages or cities, or simply providing sufficient security for other government or international agencies to do their work.[7] Each PDT implementation is relatively short, averaging less than three months, but is meant to establish a basis for continuing government involvement.

The PDT program is quite significant in scale. Table 2 reports PDT implementations per year, as well as cumulative implementations through 2010. In any given year, 500-1,000 of the Philippines' 42,000 villages,[8] accounting for 1-2% of the population, received PDT.[9] Over the full period, over 5,000 villages (accounting for 12% of the population) received PDT. The geographic unit immediately larger than villages is the municipality, which have considerable political authority.[10] The Philippines has about 1,600 municipalities, nearly half of which (47% of the population) countained a village treated with PDT during the sample period.

[Table 2 about here.]

Figure 2 presents a map of implementations. Note that PDT treatment is spread throughout the Philippines, though disproportionately in peripheral locations with low population density.

[5]Figure 1 displays the annual count of violent incidents. Violent incident *rates*, one of our key variables, are even more skewed as population appears in the denominator.

[6]Officially, the program is implemented within barangays, subnational political units smaller than municipalities. Barangays are mutually exclusive and exhaustive and are located entirely within municipalities. The Philippines has approximately 42,000 barangays, with an average population of about 3,000. Throughout the paper, we refer to barangays as "villages."

[7]Unfortunately available data does not include details about the specific set of services offered in particular PDT operations.

[8]The exact number of villages and municipalities in the Philippines changes every year, as these units merge and split. For our analyses, we use a consistent set of village and municipality definitions that closely corresponds to the official 2009 definitions used in Felter (2005).

[9]For details on the calculation of village-level populations, see the appendix.

[10]The Philippines draws a distinction between "municipalities" and "cities." Technically speaking, there are about 1,500 municipalities and 140 or so cities, which have larger populations. The political distinction between these units is small so for simplicity, we use the term "municipalities" to collectively refer to cities and municipalities.

[Figure 2 about here.]

In conclusion, the PDT program represents a concerted effort by the government to expand control into areas that are relatively poor, violent and rural. It is a large program that includes both security and development elements, at least nominally. Implementation occured slowly over a period of nine years. We exploit this staggered rollout in our empirical strategy.

3 Data

We have three primary variables: PDT implementations, malnutrition rates, and violent incidents. Additional variables (e.g., population, geography, etc.) are discussed in an appendix.

3.1 PDT data

Data on PDT implementations comes from the Armed Forces of the Philippines (AFP). For the universe of PDT implementations, it includes start dates,[11] an identifier the implementing unit, and a geographic code for the village.[12]

Our data includes all PDT implementations.[13] The program began in 2002 and concluded in 2010. While the AFP had been expanding governance and providing development services long before 2002, we lack systematic data on efforts preceding the PDT campaign. A limitation of our analysis is that we can only study the effects of this particular wave of interventions, without being able to account for how previous programs might influence estimated treatment effects.

3.2 Violent incident data

Our incident data includes the full universe of violent incidents reported by the AFP. They were first compiled and analyzed in Felter (2005)[14] and subsequently updated as part of the Empirical Studies of Conflict (ESOC) Project.[15] The data are based on underlying

[11] For 75% of implementations, we also observe end dates.

[12] For a small fraction (0.6%) of implementations, the geographic code corresponded to a municipality. We coded these as having occurred in each village within that municipality.

[13] Because of changes in the definition of villages over time, we were unable to merge two of the 6,819 implementations with the rest of our data.

[14] See also Berman et al. (2011a), Berman et al. (2012), Crost et al. (2014), and Crost et al. (2013).

[15] esoc.princeton.edu

Armed Forces of the Philippines and Philippine Army incident reports.[16] For each incident, the data includes rebel, civilian, and government casualities; an indicator for whether the incident was rebel- or government-initiated; the number of rebels captured or surrendered; and a geographic code corresponding to the village where the incident occurred.

3.3 Malnutrition data

Malnutrition data is from the Philippines National Nutrition Council's (NNC) Operation Timbang (OPT) project. Operation Timbang is the NNC's largest program, seeking to annually weigh every child in the country aged 0-71 months. In the late 1970's, the Philippine government established village-based health care provision as a national strategy (Phillips, 1986). Since that time, it has conducted a number of large-scale programs to systematically establish permanent health care experts in local villages. Currently operating programs include the Barangay (village) Nutrition Scholars (BNS) program (established in 1978), Day Care Centers (established in 1990), the Barangay Health Worker (BHW) program (established in 1995), and the Rural Health Midwives Placement Program (RHMPP) through which placements began in 2008.

Weighing is conducted by an OPT Plus team, which includes village health and day care workers, members of the Barangay Council, and sometimes other local community leaders and mothers. This team seeks to compose an exhaustive list of all children in the village age 5 and under.[17] The team designates an accessible location where weighing can occur. The NNC specifies that this "may be held in a barangay hall, day care center, barangay health station, health and nutrition post, home or any place easily accessible to the target population." Beginning in January, the OPT Plus Team is provided with instructions and materials from the federal government for the weighing procedure. Weighing occured between January and March each year, with results reported to the federal government. The OPT program receives significant attention in the local media and results are widely publicized and discussed. They inform government resource allocation decisions.

The details of the Operation Timbang process are important for two reasons. First, the village-centric measurement process probably increases data reliability. Particularly because we are interested in unstable places where the government has limited authority, we might be concerned if federal agencies were directly responsible for weighing children.[18] Since OPT

[16] Because the data originally come from the Army's incident reports, they likely undercount attacks in which the Army was not involved. This complicates the variable's interpretation, but does not bias our results.

[17] Recall that the average village is 3,000 people.

[18] Although Operation Timbang results are interpreted as indicators of poverty and are connected to resouce allocation decisions, the Philippine government has sought to de-politicize the program. This also

is implemented by local staff from the same village they are likely to know of and have access to all children.[19] Second, the systematic nature of the program (e.g., the standards for the establishment of the OPT Plus Team, the provision of a consistent set of instructions and materials, etc.) gives us some confidence in comparability over time and across locations.

We use the official estimated malnutrition rate, based on weight-for-age z-scores. This definition of malnutrition,[20] and of the reference population used, follow the recommendations of the World Health Organization (WHO).[21] The data the NNC makes publicly available have two limitations, relative to the underlying data collected through OPT. First, it is aggregated over villages up to the municipality level. This still provides quite detailed data for analysis, as the Philippines has approximately 1,600 municipalities, but does not match the village-level precision of the PDT data. Second, rather than report the distribution of weight for age, the data report only the malnutrition rate (defined as the percent of children who are two standard deviations below the age-specific mean of an internationally-recognized reference population).[22]

Table 3 displays descriptive statistics for our three key malnutrition variables: malnutrition rates, severe malnutrition rates, and the OPT estimate of coverage – the percent of children in the municipality who were weighed. The first panel presents all available observations. It suggests significant heterogeneity in malnutrition: the 75^{th} percentile has two and a half times the malnutrition rate of the 25^{th} percentile. The second panel weights municipalities by their population to obtain estimates more representative of the country.

[Table 3 about here.]

Of particular concern is the coverage measure, which is often quite low and sometimes quite high.[23] We might be concerned that the malnutrition esitmates from these municipalities are not reliable. Thus, for our main analyses, we exclude municipalities for which *a*)

increases the reliabilty of the data.

[19]Similarly, the involvement of local midwives and day care workers increases the likelihood that weighed children are actually aged 0-71 months, as the program designates.

[20]This definition of malnutrition, weight-for-age z-scores (WAZ) dates to Gomez et al. (1956). Since then, Seoane and Latham (1971) have proposed splitting (WAZ) scores into height-for-age and weight-for-height z-scores (Cole et al., 2007). Height-for-age (HAZ) is considered a measure of long-term malnutrition ("stunting") and weight-for-height (WHZ) is considered a short-term acute measure ("wasting"). Unfortunately the OPT data include only malnutrition defined according to WAZ scores.

[21]In 2010, following the recommendation of the World Health Organization, the Philippines switched from the International Reference Standard (IRS) to the WHO Child Growth Standard (CGS), which defines malnutrition in the same way, but uses a different reference population (Group, 2006). This was done to maintain consistency with international standards.

[22]It also includes severe malnutrition, defined as being three standard deviations below the mean.

[23]Coverage is the number of children weighed as a percentage of the *estimated* number of children age 0-71 months. Thus, it can exceed 100 for a number of reasons, including children being weighed multiple times or, more likely, inaccurate population estimates.

coverage is less than 66% or greater than 110%,[24] and *b*) population greater than twice the mean.[25] The third panel displays malnutrition characteristics for this sample, which trims about 13% of observations.

Finally, the fourth panel presents descriptive statistics from the Autonomous Region of Muslim Mindanao (ARMM), where Islamist rebel groups are most active. We have data on only 27 municipality-years from the region (which includes 116 municipalities). This is likely because that autonomous regional government has overwhelming legal authority and is not required to cooperate with federally sponsored data collection. As such, our malnutrition results are primarily driven by PDT implementations outside the ARMM, under-representing regions affected by Islamist rebels and over-representing those affected by NPA. Recall that the NPA is responsible for nearly two-thirds of violent incidents (in data which *are* representative).

Importantly, malnutrition rates decline throughout the sample period, as illustrated in Figure 3, which reports the mean (weighted and unweighted) and interquartile range over time. To account for this national trend our analyses below will include year fixed effects. Because this rate of decline may well vary across municipalities, we will focus on results that also allow for municipality-specific linear trends in malnutrition.

[Figure 3 about here.]

3.4 Summary statistics

Our main analyses will be conducted at the municipality-year level. Table 4 presents summary statistics, including separate figures for municipalities that received PDT and those that did not.[26]

[Table 4 about here.]

Municipalities receiving PDT tend to be relatively disadvantaged: They have higher malnutrition rates (10.5% compared to 8.7%), are more likely to experience violent incidents (51% of municipality-years compared to 19%) and, when experiencing violence, tend to have more intense violence (7 incidents per year per 100,000 residents compared to 4.1). These descriptive statistics foreshadow the formal results regarding non-random PDT implementation discussed in the next section.

[24] The NNC recommends caution with measurements outside 80 and 110. We felt 80 was too restrictive.

[25] Because our analyses weight by population, we were primarily concerned about inaccurate measurement among particularly large municipalities.

[26] The determination of key variables is discussed in Section 4. The Table reports statistics for our main estimation sample, and uses the population weights used in the final analysis.

Table 4 also reports several other facts relevant to estimation. First, municipality fixed effects alone account for over 81% of variation in malnutrition, our dependent variable, leaving less than 19% with which to estimate treatment effects conditional on those effects. Second, conditional on a PDT implementation within a given municipality-year, only 13.6% of the population live in a treated village. Thus, our estimated municipality-level treatment effects are based on treatment experienced directly by, on average, only about a seventh of the measured municipality population. Finally, on average, we observe malnutrition 2.4 years after the most recent PDT implementation, allowing us to estimate delayed effects.

4 Non-random selection

We seek to estimate the effect of PDT implementation on malnutrition, but are sensitive to the possibility of selection bias –namely that villages may be selected for PDT treatment on criteria that themselves predict malnutrition. In this section we investigate the selection mechanism in order to avoid that selection bias.

As a favor to the impatient reader, we state up front that, though this discussion of selection is interesting in its' own right, it turns out to have very little bearing on our estimated selection effects, which are described in the next section (to which said reader should feel free to skip).

Formally, we aspire to estimate the coefficents of this equation,

$$ln(MalnutritionRate)_{it} = \alpha_i + \beta_1 FractionPDT_{it-1} + \varepsilon_{it} \tag{1}$$

where i indexes municipalities and t years, the malnutrition rate (weight for age) is measured as explained above, and the fraction of the municipality's population in PDT-treated villages is measured by $FractionPDT$. A municipality-specific fixed effect accounts for fixed factors that might predispose municipalities to have high malnutrition absent PDT.

Of course villages were not randomly assigned to receive PDT, which complicates estimating its causal effect on malnutrition. In this section, we explore the selection decision for PDT implementation. In Section 4.1 we summarize basic facts that inform our specification; we present a model of selection in Section 4.2.

4.1 Facts informing selection

Three basic facts are important in understanding PDT implementation: PDT implementations are geographically clustered; repeat implementations are common; and violence rises during PDT.

Geographic clustering of PDT makes operational sense. Villages are quite small, with an average population of about 3,000, so a collection of five neighboring villages remains a relatively small area and clustering implementations would simplify the logistics of organizing military units locally. Moreover, PDT implementations might simply push rebels to the next village, which would then require attention.

We investigate spatial clustering by estimating

$$PDT_{it} = \alpha_i + \eta_{m(t)} + \phi_{y(t)} + \beta_3 NeighPDT_{i,t-1,t-3} + \beta_6 NeighPDT_{i,t-1,t-6} + \varepsilon_{it}\,. \quad (2)$$

Here PDT_{it} is an indicator that PDT was implemented in village i during month t, $\eta_{m(t)}$ and $\phi_{y(t)}$ are month of year and year effects, respectively, and $PDT_{i,t-1,t-k}$ is the fraction of the four villages nearest to i in which PDT was implemented between 1 and k months ago. As shown in Table 5 below, these neighboring implementations are highly predictive, a result that is robust to including month and village fixed effects. The final column, including month and year effects, suggests that having all four neighboring villages receive PDT in the last six months increases the likelihood that a village will receive PDT in a given month by about .0011, or about 85% of the unconditional mean. Having all four neighbors receive PDT in the last three months increases this probability by .0019, more than double the unconditional mean.[27]

Given this evidence of geographical clustering, and the possibility that it is due to relocation of rebels across space, we will control for spillovers between neighbors in estimating the effects of PDT on malnutrition in the analysis below.

[Table 5 about here.]

The second fact underlying our selection specification is that repeat implementations of PDT are common, suggesting that the expansion of governance is not a monotonic, universally successful process. Table 6 reports on the distribution of implementations across the 5,188 villages that received PDT at least once between 2002 and 2010. About one quarter of villages received multiple implementations.

[Table 6 about here.]

To better understand repeat implementations, Figure 4 plots the distribution of start dates for the first PDT implementation a village receives. The gray bars indicate the distribution for villages that received only one implementation, while the clear bars report that of villages which received multiple implementations. Clearly, the date of first implementation

[27] $.0019 = .0011 + .0008$

was much earlier for villages that would eventually receive multiple PDT's. This suggests that many of the single-PDT villages may not have experienced persistent success either, they simply received PDT too late to experience a repeat by the end of 2010 when our sampling period ends.

Taken together, the evidence of repeated PDT interventions suggests that, like job-training, counselling, or many other programs evaluated by social scientists, PDT does not always work. We will return to this point below because a history of treatemnt has implications for selection and for the size of treatment effects.

[Figure 4 about here.]

We turn now to violence. High levels of rebel violence might predispose a village to be selected for PDT (Berman et al., 2011b) so we are interested in violence as a predictor of selection. Violence might also rise (or fall) on implementation of PDT, either because of reporting bias or because coercive force is required to implement.[28]

To assess rates of violence before and after PDT, we use an event study specification. We estimate the following specification using monthly (indexed t) village (indexed v) data.

$$Violence_{vt} = \alpha_v + \delta_t + \sum_{\tau=-18}^{18} \beta_\tau PDT_{vt} + \varepsilon_{vt} \tag{3}$$

Figure 5 study plots the estimated β_τ coefficients for the 37 months around a PDT implementation (month 19 in the figure corresponds to the month of implementation). For reference, mean violence for the ommitted category (months more than a year and a half before or after PDT) is .0024 incidents per month per 100,000 population.

[Figure 5 about here.]

Figure 5 shows that violence rises about 9 months before a PDT implementation. This might imply that PDT responds to spikes in violence. However, results not included here show that this rise is entirely due to an increase in government-initiated incidents. This implies that the increase in violence during the 9 months before PDT (which roughly corresponds to a 50% increase in the incident rate) may be part of the program itself. The military may deploy to an area in order to clear out rebels and prepare for PDT or to assess whether PDT would be appropriate or feasible.

During the months immediately surrounding the implementation, violence is a factor of 2-4 times higher than normal, which is likely a combination of a true effect and a reporting

[28] Increases in troop strength in Iraq correlated with increased reported violence (Berman et al., 2013).

effect. Violence remains high during the 6 months after the implementation, before returning to its normal level.

The results plotted in Figure 5 suggest both that PDT is predicted by violence (and particularly AFP military action preceding PDT) and that it has a short-term effect of increasing violence.

4.2 Selection estimates

Informed by the three facts just presented (in Section 4.1), we estimate a formal model of PDT selection with the following features. First, because PDT is implemented in geographic clusters, we control for PDT implementations in the nearest four villages in the last three and six months. Second, to account for complicated dynamics of PDT targeting and violence, we include three and six month lags of incident rates in the village and in the village's municipality. Finally, to account for complicated patterns of repeat implementations, we include an indicator for whether the village has previously received PDT and a linear (column 4) and quadratic (column 5) time trend in months since last PDT implementation.[29]

$$\begin{aligned} PDT_{it} = {} & \gamma_1 NeighborPDT_{i,t-1,t-3} + \gamma_2 NeighborPDT_{i,t-1,t-6} \\ & + \gamma_3 AvgViolence_{i,t-1,t-3} + \gamma_4 AvgViolence_{i,t-4,t-6} \\ & + \gamma_5 PreviousPDT_{it} + \gamma_6 MonthsSincePDT_{it} + \gamma_7 MonthsSincePDT_{it}^2 \\ & + \alpha_i + \eta_{m(t)} + \phi_{y(t)} + \delta_i t + \varepsilon_{it} \end{aligned} \tag{4}$$

Table 7 reports these results, estimated at the village-month level. The three major findings of the previous section again express themselves robustly in predicting PDT implementation. PDT in neighboring villages during the previous 3 and 6 months predict new PDT starts, as we saw in Table 5, and even in the presence of fixed effects. Lagged incidents predict PDT, over the previous 3 and 6 months, again even when allowing for fixed effects. Finally, recent PDT in the same village predicts a lower probability of repeat treatment, though that effect fades over time. Those three findings are robust across specifications, including column (6) which allows for village-specific linear time trends.

[Table 7 about here.]

To summarize, the AFP seems to select villages for treatment by violence, proximity to treated neighbors and no very recent treatment. That selection has the potential to bias our

[29]To avoid imputing PDT end dates, we use a time trend in months since the last PDT *began*. Note that since we include an indicator for whether the village had previously received PDT, the value of the time trend for villages never receiving PDT does not matter.

estimates of PDT on malnutrition (equation (1)), which we now turn to.

Unfortunately, our key outcome measure (child malnutrition) is not observed at the village-level. Instead, we must aggregate to the municipality-year level. Appendix B shows the model presented in Equation (4) aggregated first to the municipality-month level, then to the municipality-year level. The key results in Table 7 hold at these higher levels of aggregation.

With a set of predictors of PDT implementation in hand, we have a selection model that should help us avoid selection bias in estimating the treatment effects of PDT on malnutrition. Short of an instrument for selection, or the ethically disturbing idea that coercive force would be randomly assigned across municipalities, this is the best we can hope for in minimizing possible selection bias.

5 Estimated Treatment Effects

5.1 Results

We can now estimate the effect of PDT on malnutrition while checking for possible selection effects. We use a combination of tools to deal with selection bias: including a set of lead coefficients in the "event study" tradition (Hoynes, Page, and Stevens, 2011; Sandler and Sandler, 2013), including municipality specific trends, including province×year effects as controls, and including controls suggested by the analysis of selection in the previous section.

Our base estimating equation is

$$\ln(MalnRate)_{mt} = \alpha_m + \delta_{p(m)t} + \gamma_m t + \sum_{\tau=-k}^{k} 1\{\tau \neq -1\}\beta_\tau PDT_{mt-\tau} + \varepsilon_{mt} \tag{5}$$

where m indexes municipalities and t indexes years, $\ln(MalnRate)$ is the natural logarithm of the malnutrition rate, $PDT_{mt-\tau}$ is the fraction of municipality m's population living in a village receiving PDT in year $t-\tau$. We use log malnutrition because we expect a proportional response to treatment.[30] For a robustness check measure malnutrition in levels:

[30]Note that the malnutrition rate is the percent of children whose weight is below some exogenously set threshold value (as recommended by the World Health Organization on the basis of a standard reference population); i.e., it is the value of the CDF of the distribution of children's weights, evaluated at a particular value. (More precisely this value differs by age, which we do not observe.) In a model where PDT has a linear effect on children's weights the effect will be larger where the mass near the threshold is larger. Intuitively, a linear effect on weights moves more children over the threshold when there are more children near the threshold (at higher malnutrition rates). Thus we expect heterogeneous effects where the effects are larger (i.e., in villages where the mass of children near the threshold is higher). Unfortunately we do not observe the mass near the threshold. (As explained below, the effect needs only to be locally linear,

the results are similar, though less precise.

Note that the β_τ coefficients for $\tau < 0$ are the coefficients on PDT leads, which we include in order to check for pre-existing trends in malnutrition for municipalities that will be treated. β_τ coefficients for $\tau > 0$ estimate lagged impulse-response as the result of PDT treatment. This may occur because governance is persistent, or because the economic and health effects of governance take some time to set in.

Our specification drops the 1-year lag of PDT to form a reference point (the regression constant) in checking for pre-existing trends. We allow for estimated contemporaneous effects though they are likely to be very small: Since children are weighed January-March of each year, PDT in year t could affect those malnutrition rates only if it occurred in the first quarter. On the other hand, patterns of government-initiated violence in the year preceding PDT treatment indicate that contemporaneous and transition effects are possible up to one year before PDT treatment. Given the reference point, treatment effects should safely be reflected in the lag coefficients β_τ for $\tau >0$, with some ambiguity about how to interpret the contemporaneous and reference coefficient.

Estimated effects of PDT on malnutrition are reported Table 8. In order to preserve precision, but allow for a flexible lag and lead structure, we start by restricting coefficients to be equal for PDT in adjoining years for lag combinations (3,4), (5,6) and 7+, and for lead combinations (2,3) and (4,5). The first column reports a restricted specification which allows only for municipality and year fixed effects, but no municipality specific trends. It shows no consistent treatment or selection effect. The second column adds to this specification municipality specific trends, which we believe are necessary based on the analysis in the previous section and the observation that malnutrition is trending downward. It shows no evidence of selection in the lead coefficients, but perhaps a hint in the contemporaneous coefficient (-0.093) and strong, statistically significant evidence of a growing treatment effect at lags of 1, 2, 3-4, 5-6 and 7+, which peaks at 58 logarithmic points (approximately 44 percent), seven years out. An F test on the combined significance of all the lags is appropriate,

---

increasing weights of all children near the threshold.) Under the following two assumptions, the mass of children near the threshold is increasing in the malnutrition rate. First, suppose that children's weights are distributed according to some (possibly asymmetric) single-peaked distribution. (In calculating the z-scores that are used in the WHO Child Growth Standards to define malnutrition –the definition used in our data for later years, children's weights are assumed to be distributed according to the Box-Cox power exponential (BCPE), which is a bell curve (Group, 2006).) Second, suppose the threshold value below which children are classified as malnourished (the value at which the CDF if evaluated) is below the peak of this distribution. (As shown in Table 3, the average malnutrition rate in our estimation sample is about 10% and the 75^{th} percentile is about 13%. Thus, this assumption seems reasonable.) Because we believe these two assumptions are reasonable, we expect the effect of PDT on the malnutrition rate to proportional. (Note from Table 3 that the malnutrition rate is never zero, so no observations are lost when taking logs. In a robustness check, we consider severe malnutrition rates, which does not have this property.)

since the coefficients are correlated (and share the same sign): it is strongly significant, with a p-value of 0.007.

[Table 8 about here.]

Column (3) reports an even less restricted specification, which replaces year effects with province-year effects. The estimated treatment effects are essentially the same, with treatment effects growing from 21 log point at the first lag to 69 for lags 7 or more, and a p-value of joint significance of 0.01. The specification in column (3) also shows some sign of selection in the lead coefficients. The malnutrition rate appears to be lower during the year before PDT than the 4 prior years, suggesting that there may have been a secular decline preceding the intervention.

To investigate the possibility of trend reduction in malnutrition in villages selected into PDT Column (4) relaxes assumptions by allowing 2- and 3-year leads separate coefficients. The results suggest a 7.3 log point decline between three and two years before PDT, followed by a 16.6 log point decline between two years and one year before PDT. This again suggests some selection on villages with a trend reduction already, though none of the coefficients are statistically different from zero, using individual or joint tests (p=0.414). Column (5) reports the result of continuing this exercise by estimating separate coefficients for the 4-year and 5-year leads. The estimated coefficients suggest that relative to comparison villages, malnutrition rose steadily (by 22 log points) between five and three years pre-PDT, before falling between 3 and 1 year before, back to the about the level it was five years before PDT.[31] One interpretation might be that the "pre-trend" in malnutrition during the 3-years before PDT was simply mean reversion, in which case it poses little risk of biasing our estimated treatment effects. This interpretation is speculative since, again, all pre-PDT coefficients are individual and jointly statistically insignificant.

Figure 6 graphs the results from column (5) of Table 8, providing a visual version of the "event study" test. The figure plots the estimated coefficients for all lags and leads, with the pre-treatment period normalized to zero. Mimicking our discussion of the estimated coefficients, there is no statistical evidence of a pre-trend. Yet the optics arouse some suspicion that a village with a temporary increase in malnutrition in t-3 might be at high risk of treatment, so that mean reversion could be confounded with a treatment effect.

[Figure 6 about here.]

[31] Recall that PDT_{t+1} is ommitted so that each coefficient can be interpreted as the level of malnutrition, relative to the level during the year before PDT.

Evidence against a simple mean reversion explanation is the accumulated reduction in malnutrition following PDT, compared to comparison villages (actually, comparison municipalities with a higher population share treated). In Table 8 the coefficient corresponding to the 7-year lag of PDT is large and marginally statistically significant, ranging from 58 to 80 log points, though imprecisely estimated. To understand the lack of precision, table 9 displays the distribution of the number of years since the most recent PDT, across municipality-years in the estimation sample. About 75% of observations are within four years of a PDT implementation and 90% are within six years. We believe that our analysis is most able to predict the effects of PDT over this time scale. We are therefore more comfortable reporting estimates 5-6 years after treatment, which range from 49 to 58 log points, depending on specification. These are hard to reconcile with a mean reversion explanation for the treatment effect, even in the presence of a 23 point decline in the two years preceding treatment.

[Table 9 about here.]

A more flexible visual version of that same test is reported in Figure 7, which frees up all the lagged coefficients year by year. Despite the loss of precision that results from estimating more coefficients, statistically it contains the same information, rejecting the hypothesis of no treatment effect (p=0.008). Visually, it shows a clear trend in the estimated coefficients of reduction in malnutrition, which monotonically decline from treatment year onwards. That trend would again be hard to reconcile with a mean-reversion explanation for the estimated treatment effect.

[Figure 7 about here.]

An alternative approach we can take to the possibility of pre-existing trends is to see if estimated treatment effects are robust to including predictors of PDT, which we investigated in Section 4. The results in Table 8 include municipality and province-year fixed effects, as well as municipality-specific linear trends. However, they do not control for lagged violence and treatment of neighboring villages, which we found to be predictors of PDT at both the village and municipality level. As Figure 5 made clear, PDT is preceded by violence, often government intitiated.

Table 10 shows that these controls have little effect of the estimated effects of PDT. Column (1) replicates our preferred specification, column (5) from Table 8. Adding violent incidents comes at a cost: because we don't observe violent incidents beyond 2010, controlling for lagged incidents requires dropping 2012 from the estimation sample. Column (2) replicates column (1) on the sample from 2005-2011. These results are less precisely estimated, but are very similar to our preferred specification. After dropping nearly 15% of the sample (specifically that which is most likely to be informative about post-PDT effects), the F-test now fails to reject the null that coefficients on the PDT lags are jointly zero ($p = .126$). This is apparently due to lost precision, as the coefficient estimates themselves are actually slightly larger (more negative). One observation worth noting is that the 3-year PDT lead is now marginally significant, consistent with the discussion above, though again the leads are not jointly significant (p=0.478).

Column (3) then adds the controls discussed in Section 4, estimating the following specification:

$$\begin{aligned}\ln(MalnRate)_{mt} =& \alpha_m + \delta_{p(m)t} + \gamma_m t + \sum_{\tau=-k}^{k} 1\{\tau \neq -1\}\beta_\tau PDT_{mt-\tau} \\ &+ \theta_1 Violence_{mt-1} + \theta_2 NeighborPDT_{mt-1} + \varepsilon_{mt}\end{aligned}$$

where $Violence_{mt-1}$ is the lagged rate of violent incidents in the municipality and $NeighborPDT_{mt-1}$ is the population-weighted fraction of the four nearest municipalities which received PDT last year. Including these controls has almost no effect on either the pre- or post-PDT coefficients, perhaps suggesting that the municipality and province-year effects and municipality trends are sufficient to capture selection of PDT.

[Table 10 about here.]

A final approach to possible pre-existing trends is to include repeat PDT interventions in the analysis. As discussed above, villages seem to suffer a type of recidivism and are frequently selected for repeat treatment. Our preferred approach has been to estimate the effect of the initial PDT treatment in a village, for two reasons. First, we suspect that follow-up implementations are not always as substantial as initial implementations. Sometimes PDT units seem to return to a previously visited village to simply check that rebels have not taken back control. Second, and related, the decision to implement a follow-up PDT may indicate that the first PDT was unsuccessful. Thus, we think of the PDT intervention itself as being an initial implementation followed, where appropriate, by subsequent interventions which are likely less intense and are carried out on an even more selected sample.

Column (4) replicates the specification from column (1), but is based on all PDT implementations, rather than simply the first one each village receives. The estimated coefficients are similar in both size and precision. Given that the selection mechanisms behind repeat implementation likely differ significantly from those behind initial implementations, we are encouraged that our estimated impacts are similar for both. Taken together, we

In summary, PDT provides a statistically significant medium term malnutrition that accumulates to almost 50 percent after six years, and probably continues afterward. Selection might reduce that estimate by up to half, though its hard to think of the mechanism by which that would occur.[32]

Returning to Figure 6 the interpretation of our estimated coefficients requires a little interpretation. The simplest interpretation is that we've estimated the effects on malnutrition of treating a single village in isolation, in which case our estimate (49% after six years) would be valid at the village level (subject perhaps to caveats about selection). That simple interpretation assumes that none of the estimated effects were due to spillover from treatment to control villages within the same municipality –recalling that we aggregated up to the municipality level to estimate. Spillovers are likely, though. Malnutrition might decline because PDT improved security, infrastructure and services in treatment villages –access to markets and regional clinics for instance–, which could easily benefit children in neighboring villages as well. On the other hand, PDT might relocate rebels from treatment villages to neighboring villages, to the detriment of nutrition among neighboring children. Allowing for spillovers, we can interpet the treatment effects as applying to a municipality, but we should be careful about extrapolation to predict the results of scaling up treatment to an entire municipality: the average treated municipality-year had only 13.6% of population treated at a time, so that 100% treatment is well outside the estimation sample.

6 Conclusions

A major challenge confronting modern development is how to effectively provide assistance in the violent, unstable places, where many of the world's poorest citizens live. Many programs require a minimally functional governance, without which implementing agents face unacceptable risks or entire populations are left out of reach. Even when interventions can reach these target populations, implementation is often compromised by leakage, and the implied insertion of capturable rents into an insecure environment may actually increase violence.

One prospective way to deliver development assistance is through military-centric development programs. To be sure, these programs are inherently coercive, and run the risk of violence, human rights abuses and degradation of social welfare for residents. We take a first

[32]An grislier interpretation of the findings would be that PDT increased child mortality, and the malnourished children were more likely to suffer mortality. In this scenario, we would expect to observe PDT effects on the estimated total number of children in the village, which can be calculated using the number of children weighed and the estimated coverage from the OPT data. Using our specification, we find no effect of PDT on the number of children in the village.

step towards measuring those risks, as well as the possible benefits of military intervention in ungoverned spaces.

The Philippines' Peace and Development Teams, operated by the Armed Forces of the Philippines, directly reached 12% of the population over nine years. The average implementation (which reached only one seventh of the population in treated municipalities) reduced that municipality's malnutrition rate by 7 percent over six years, net of any increase due to the sometimes violent transition to governance. Assuming that the treatment effect is entirely concentrated in treated villages, that corresponds to a 49 percent reduction in malnutrition for a very vulnerable population.

An important caveat regards interpretation. Even compared to other multifaceted interventions, there are huge gaps in our knowledge of the actual content of the PDT intervention, and even larger gaps in our understanding of the mechanism by which health improvements occurred. Research on those topics, which is likely to be qualitative, would be welcome.

How important is this finding? To put these effect sizes in context, Table 11 summarizes several evaluations of child malnutrition treatments which use weight-for-age z-scores (WAZ) as an outcome.[33] For each study, we summarize the intervention, the measurement of the dependent variable, the estimated effect in its original form, and the implied reduction in the malnutrition rate.

[Table 11 about here.]

As the Table makes clear, PDT induced improvements in child nutrition are comparable to the leading examples in the literature.[34] Each intervention estimates a 35-45% reduction in the malnutrition rate, which is a little larger than the effect of PDT after 2 years (32%) but smaller than the six year reduction associated with PDT (49%).[35]

Moreover, the interventions described in Table 11 would likely be infeasible in PDT villages, given their lack of governance. In the absence of a persistent security presence, it is

[33] Many studies separate WAZ into weight-for-height and height-for-age (Duflo, 2003; Graff-Zivin, Thirumurthy, and Goldstein, 2006; Lavy et al., 1996), something we cannot do with Operation Timbang data.

[34] Not included in the table are a large number of interventions for which no effects have been found. See Masset et al. (2011) for a review of mostly unsuccessful agricultural interventions, Morris et al. (2004) for a study estimating that a Brazilian conditional cash transfer program increased malnutrition, Lind et al. (2008) for an RCT showing that iron supplements can increase malnutrition for non-deficient infants, and Singh, Park, and Dercon (2013) for a study estimating that India's school lunch program did not reduce malnutrition among participants, although it did prevent the deleterious effects of unexpected drought.

[35] That estimated effect should not be assumed to scale in linear proportion if an entire (rebel-controlled) municipality were treated. Empirically, that would require an extrapolation from 1.9% of population treated in the average municipality-year to 100%. Furthermore, a proportional treatment effect assumes no spillovers to neighboring villages in the same municipality. PDT may have positive spillovers on neighboring villages (by securing local markets and roads) or negative ones (by pushing rebels into nearby villages and increasing local violence).

difficult to imagine programs distributing medicine or cash operating safely or successfully.[36] How can one reach these citizens? While we do not believe that military-centric counterinsurgency programs are the final answer to this question, we do think that these findings invite an evidence-based discussion of the role of expansions of governance, as an extensive margin of development.

[36] Crost et al. (2014) is a striking example.

References

Eli Berman and Aila Matanock. The empiricists' insurgency. *Annual Review of Political Science*, 18:443–464, 2015.

Eli Berman, Michael Callen, Joseph H Felter, and Jacob N Shapiro. Do working men rebel? Insurgency and unemployment in Afghanistan, Iraq, and the Philippines. *Journal of Conflict Resolution*, 55(4):496–528, 2011a.

Eli Berman, Jacob N Shapiro, and Joseph H Felter. Can hearts and minds be bought? The economics of counterinsurgency in Iraq. *Journal of Political Economy*, 119(4):766–819, 2011b.

Eli Berman, Joseph Felter, Ethan Kapstein, and Erin Troland. Predation, Economic Activity and Violence: Evidence from the Philippines. *NBER Working Paper*, 2012.

Eli Berman, Joseph H Felter, Jacob N Shapiro, and Erin Troland. Modest, secure, and informed: Successful development in conflict zones. *American Economic Review*, 103(3): 512–17, 2013.

Timothy Besley and Torsten Persson. The Logic of Political Violence. *The Quarterly Journal of Economics*, 126(3):1411–1445, 2011.

Kenneth H Brown, Janet M Peerson, Juan Rivera, and Lindsay H Allen. Effect of supplemental zinc on the growth and serum zinc concentrations of prepubertal children: a meta-analysis of randomized controlled trials. *The American journal of clinical nutrition*, 75(6):1062–1071, 2002.

Tim J Cole, Katherine M Flegal, Dasha Nicholls, and Alan A Jackson. Body mass index cut offs to define thinness in children and adolescents: International survey. *British Medical Journal*, 335(7612):194, 2007.

Paul Collier. Rebellion as a quasi-criminal activity. *Journal of Conflict Resolution*, 44(6): 839–853, 2000.

Benjamin Crost and Patrick B Johnston. *Aid Under Fire: Development projects and Civil Conflict*. Harvard Kennedy School, Belfer Center for Science and International Affairs, 2010.

Benjamin Crost, Joseph H Felter, Hani Mansour, and Daniel I Rees. Election Fraud and Post-Election Conflict: Evidence from the Philippines. *IZA Discussion Paper*, 2013.

Benjamin Crost, Joseph Felter, and Patrick Johnston. Aid Under Fire: Development Projects and Civil Conflict. *The American Economic Review*, 2014.

Esther Duflo. Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in south africa. *The World Bank Economic Review*, 17(1):1–25, 2003.

Joseph H Felter. *Taking guns to a knife fight: a case for empirical study of counterinsurgency.* Stanford University, 2005.

Federico Gomez, Rafael Ramos Galvan, Silvestre Frenk, Joaqíun Cravioto Munoz, Raquel Chávez, and Judith Vázquez. Mortality in second and third degree malnutrition. *Journal of Tropical Pediatrics*, 2(2):77–83, 1956.

Joshua S Graff-Zivin, Harsha Thirumurthy, and Markus Goldstein. Aids treatment and intrahousehold resource allocations: Children's nutrition and schooling in kenya. *NBER Working Paper*, 2006.

Herschel I Grossman. Kleptocracy and revolutions. *Oxford Economic Papers*, 51(2):267–283, 1999.

WHO Multicentre Growth Reference Study Group. WHO Child Growth Standards based on length/height, weight and age. *Acta Paediatrica*, 450:76, 2006.

Michael T. Hall and Gen. Stanley A. McChrystal. Isaf commander's counterinsurgency guidance. *Internation Security Assistance Force Technical Report*, 2009.

Ariel Hernandez. Managing intractable identity conflicts–a concert of measures in the southern philippines. In Ariel Hernández, editor, *Nation-building and Identity Conflicts*, pages 149–185. Springer, 2014.

Hilary Hoynes, Marianne Page, and Ann Huff Stevens. Can targeted transfers improve birth outcomes? evidence from the introduction of the wic program. *Journal of Public Economics*, 95:813–827, 2011.

Marta M Jankowska, David Lopez-Carr, Chris Funk, Gregory J Husak, and Zoë A Chafe. Climate change and human health: Spatial modeling of water availability, malnutrition, and livelihoods in mali, africa. *Applied Geography*, 33:4–15, 2012.

Victor Lavy, John Strauss, Duncan Thomas, and Philippe De Vreyer. Quality of health care, survival and health outcomes in ghana. *Journal of Health Economics*, 15(3):333–357, 1996.

Torbjörn Lind, Rosadi Seswandhana, Lars-Åke Persson, and Bo Lönnerdal. Iron supplementation of iron-replete indonesian infants is associated with reduced weight-for-age. *Acta Paediatrica*, 97(6):770–775, 2008.

John Maluccio and Rafael Flores. *Impact evaluation of a conditional cash transfer program: The Nicaraguan Red de Protección Social*. International Foor Policy Research Institute, 2005.

Edoardo Masset, Lawrence Haddad, Alex Cornelius, and Jairo Isaza-Castro. *A systematic review of agricultural interventions that aim to improve nutritional status of children.* 2011.

Gen. Stanley A. McChrystal. Dod media roundtable with gen. mcchrystal nato headquarters in brussels. *DoD Transcript (June 10, 2010)*, 2010.

Saul S Morris, Pedro Olinto, Rafael Flores, Eduardo AF Nilson, and Ana C Figueiro. Conditional cash transfers are associated with a small reduction in the rate of weight gain of preschool children in northeast brazil. *The Journal of Nutrition*, 134(9):2336–2341, 2004.

Nathan Nunn and Nancy Qian. Us food aid and civil conflict. *American Economic Review*, 104(6):1630–66, 2014.

David R Phillips. Primary health care in the philippines: Banking on the barangays? *Social Science & Medicine*, 23(10):1105–1117, 1986.

Nathan Gilbert Quimpo. Mindanao, southern philippines: The pitfalls of working for peace in a time of political decay. In Michelle Ann Miller, editor, *Autonomy and Armed Separatism in South and Southeast Asia*. Institute of Southeast Asian Studies, Singapore, 2012.

Marie T Ruel, Purnima Menon, Jean-Pierre Habicht, Cornelia Loechl, Gilles Bergeron, Gretel Pelto, Mary Arimond, John Maluccio, Lesly Michaud, and Bekele Hankebo. Age-based preventive targeting of food assistance and behaviour change and communication for reduction of childhood undernutrition in haiti: a cluster randomised trial. *The Lancet*, 371 (9612):588–595, 2008.

Danielle H. Sandler and Ryan Sandler. Multiple event studies in public finance and labor economics: A simulation study with applications. *Working Paper*, 2013.

Salvatore Schiavo-Campo and Mary P Judd. *The Mindanao conflict in the Philippines: Roots, costs, and potential peace dividend*, volume 24. Conflict Prevention & Reconstruction, Environmentally and Socially Sustainable Development Network, World Bank, 2005.

Nicole Seoane and Michael C Latham. Nutritional anthropometry in the identification of malnutrition in childhood. *Journal of Tropical Pediatrics*, 17(3):98–104, 1971.

Abhijeet Singh, Albert Park, and Stefan Dercon. School meals as a safety net: an evaluation of the midday meal scheme in india. *Young Lives Working Paper*, 75, 2013.

A Additional data details

Three important variables not discussed in the main text are: municipality populations, village populations, and geography (used to identify the four nearest neighbors).

We seek annual estimates of population variables. We obtain these from three data sources:

1. Annual province-level population estimates from the National Statistical Coordination Board (NCSB) from 2000-2011
2. Municipality-level population estimates from 2000, 2003, 2007, and 2010
   - Estimates from 2000, 2007, and 2010 are from the Philippines Census
   - Estimates from 2003 are from the NCSB small area poverty estimates
3. Village-level population estimates from the 2007 Census

With these data sources, we estimate annual populations in three steps. First, we linearly interpolate and extrapolate municipality populations for missing years. Second, we proportionally adjust these municipality populations so that they sum to the annual province-level estimates. Finally, we divide a municipality's population among its villages using the 2007 population distribution.

As a basis for the geography variables, we rely on a dataset from the NCSB with the longitude and latitude of each villages' centroid. A similar dataset was available for municipality longitude and latitdue, but inspection revealed it was fraught with inconsistencies. The village-level longitudes and latitudes were much more reliable.

Thus, we estimated each municipality's longitude and latitude using the village data. To do so, for each municipality, we took the midpoint of the longitude and the latitude from the various villages within the municipality. This identifies a point which is in the center of the smallest rectangle that could be drawn to include each village's centroid. However, there is no guarantee that this point is actually within the municipality.[37] Thus, we define the longitude and latitude of the municipality to be the village centroid with the smallest Euclidean distance from this rectangle's center. This does not guarantee that the municipality's "location" is its centroid, but provides a reasonable approximation which guarantees that the location will actually be within the municipality.

Finally, we calculate a distance matrix containing the Euclidean distance between each municipality. The four nearest neighbors were selected to be the four municipalities for which the Euclidean distance was the smallest.

B Additional selection models

In Section 4.1, we presented evidence for three important considerations regarding selection in the PDT implementation decision. These were formalized in Equation (4) in a model

[37] Consider, for instance, a u-shaped municipality. This point would lie outside of its borders, despite being the center of the smallest inclusive rectangle.

estimated in Table 7 at the village-month level. Unfortunately, our key outcome measure (child malnutrition) is not observed at the village-month level, but is only available at the municipality-year level, and so we must aggregate this selection model. The next two Tables do that in stages, first aggregating to municipality-months, and then to municipality-years. The main findings of 4.1 turn out to be robust to both those aggregations –over periods and over space.

Table B.1 presents the same selection model at this higher level of aggregation. Instead of a binary indicator of a PDT implementation beginning in month t, the dependent variable is the fraction of the municipality's population living in a village with a PDT implementation beginning in month t.[38] The municipality-level results (Table B.1) are broadly consistent with the village-level results (Table 7).[39]

[Table B.1 about here.]

Again, however, our outcome data is not available at the monthly level. Therefore, all analyses are run at the municipality-year level. Table B.2 aggregates the data to this level, modeling the fraction of a municipality's population living in a village where PDT was implemented at all during the year, as a function of lagged neighboring PDT implementations, lagged incidents, and aggregated measures of PDT history.[40] The results are broadly consistent with those in Table 7, although the higher level of time aggregation makes the results less consistent across columns. In most specifications, neighboring PDT and violent incidents continue to positively predict PDT implementation.

[Table B.2 about here.]

[38] Most independent variables are analogous to their village-level counterparts. The fraction of four nearest neighbors is now the fraction of population in the four nearest municipalities living in a village where PDT began during the previous three and six months. We now include an indicator of whether any village has previously received PDT and the time trend counts the number of months since any village in the municipality received PDT.

[39] The fact that neighboring areas' PDT is predictive with a longer lag for municipalities makes sense because they are larger than villages.

[40] As we aggregate to higher levels, it becomes increasingly misleading to report past PDT receipt as binary. Thus, for past PDT receipt, we separately include a (continuous) variable capturing the fraction of the population living in a village that has previous received PDT and a binary indicator that this continuous variable is equal to zero. This allows for a discontinuous (at zero) effect of past PDT receipt.

Figure 1: Violent incidents per year and malnutrition by population density quintile

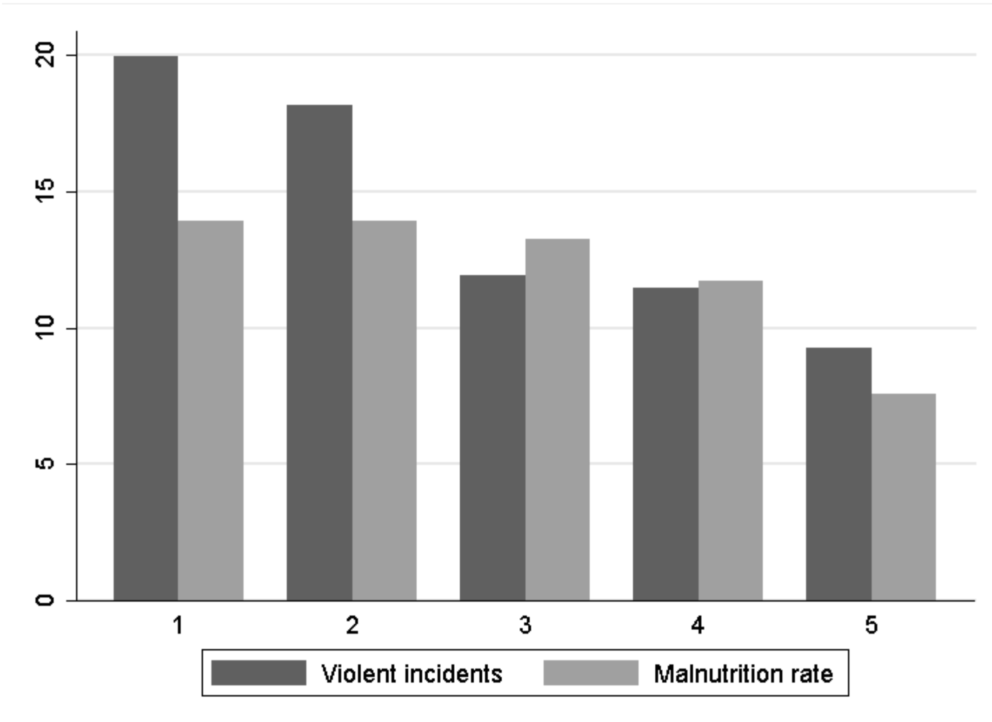

Figure 2: Villages receiving PDT

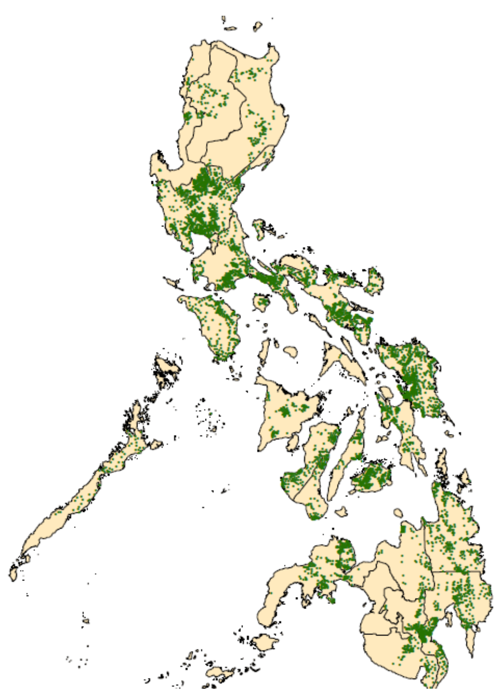

Figure 3: Malnutrition rates over time

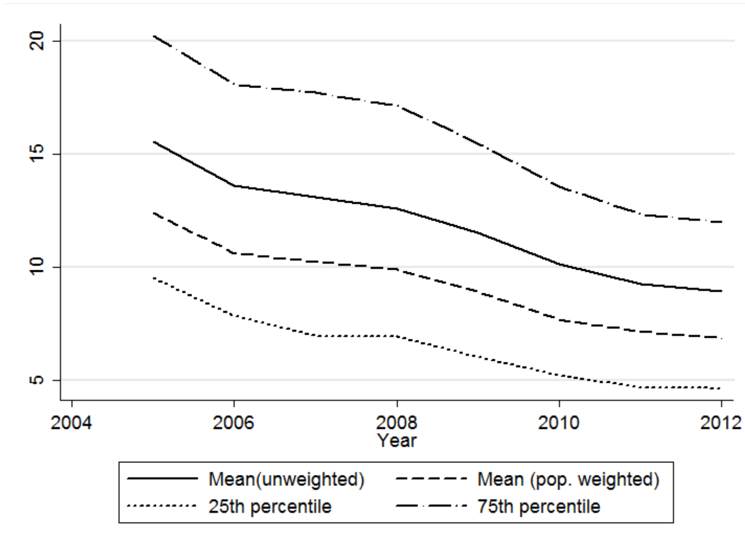

Figure 4: Start dates of first implementation

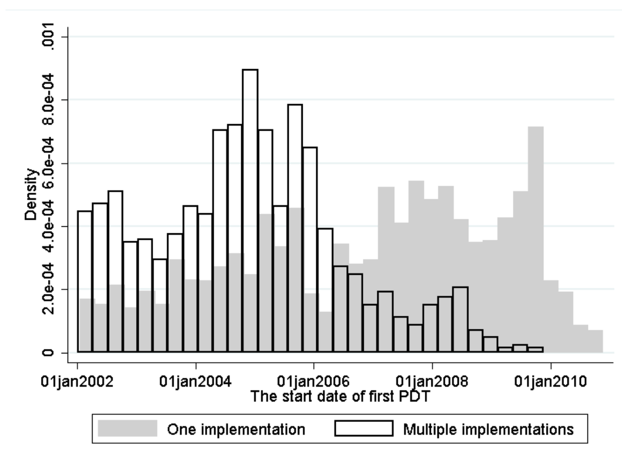

Figure 5: Village-level violence surround PDT

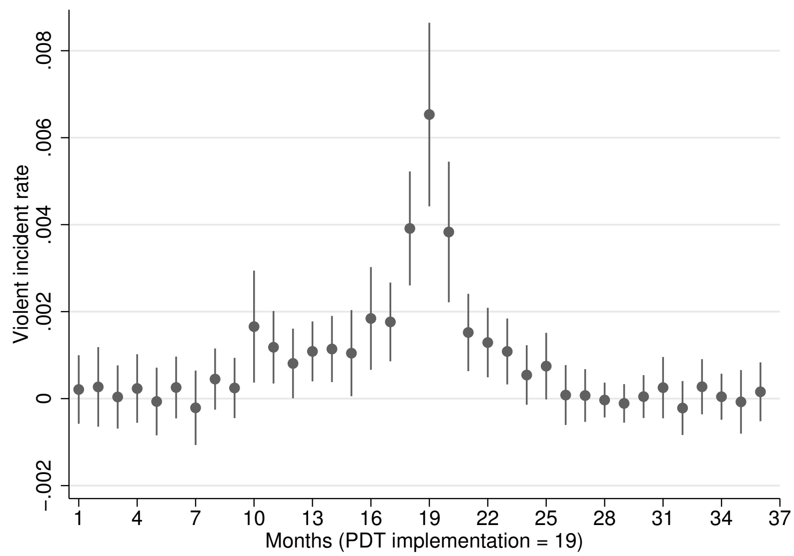

Figure plots coefficients, and 95% confidence intervals, corresponding leads and lags of PDT in equation (3), which includes month and village fixed effects. Month 19 corresponds to the month of PDT implementation. Mean violence for the ommitted category (months more than a year and a half before or after PDT) is .0024 incidents per month per 100,000 population.

Figure 6: Estimated effects of PDT

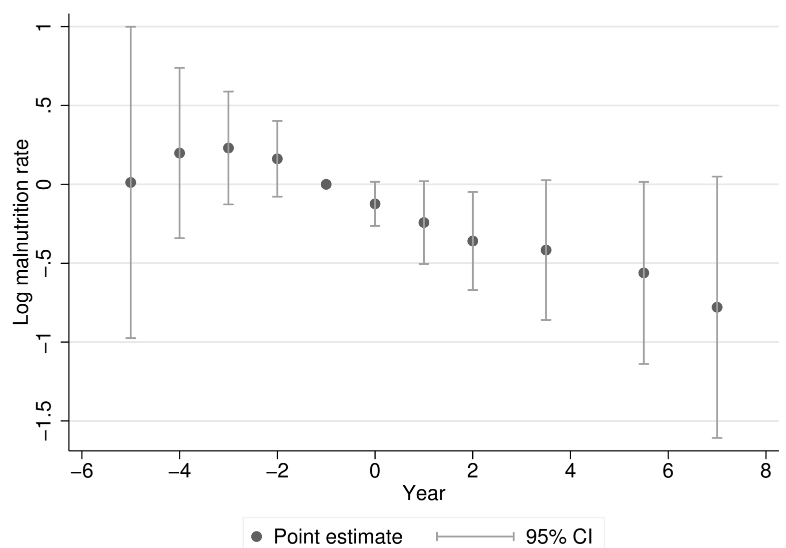

Figure plots coefficients, and 95% confidence intervals, corresponding to the effects of an average PDT (in which 12% of the population is treated) implemented at time 0. Coefficients are based on Column (5) of Table 8. Log malnutrition rates are normalized to zero during the year before the intervention.

Figure 7: Estimated effects of PDT

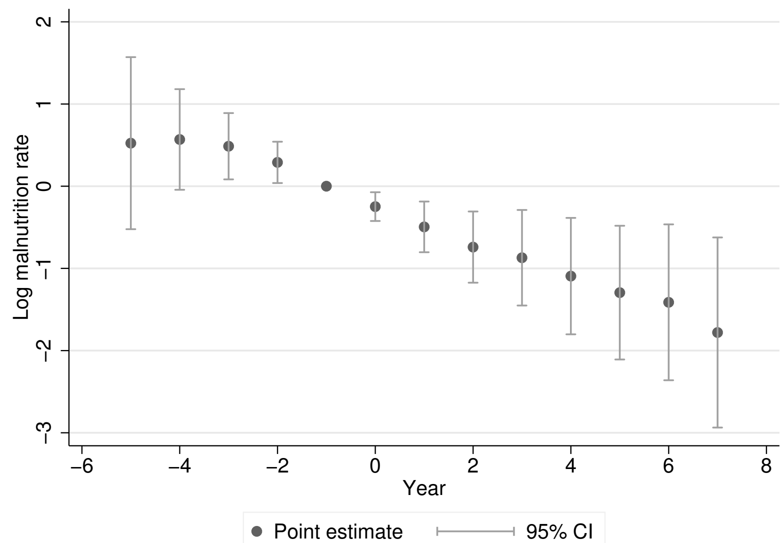

Figure plots coefficients, and 95% confidence intervals, corresponding to the effects of an average PDT (in which 12% of the population is treated) implemented at time 0. Log malnutrition rates are normalized to zero during the year before the intervention. An F-test fails to reject the null that the pre-PDT coefficients are jointly zero ($F = 1.74$, $p = .149$), but does reject the null that the post-PDT coefficients are jointly zero ($F = 2.96$, $p = .008$).

Table 1: Conflict and instability among top foreign aid recipients

	Average from 2003-2012						Conflict
	Annual ODA per capita				GDP	Rank of Pol.	Years
Country	World Bank	USA	UK	Total	per capita	Instability	1988-2012
Iraq	$0	$136.9	$11.9	$148.8	$2,279.5	2	19
Afghanistan	1.9	76.8	11.4	90.1	351.8	3	24
Haiti	0	52.4	1.0	53.5	524.5	23	3
Somalia	0	22.1	14.3	36.3	–	1	18
Zambia	6.0	17.4	10.4	33.8	789.2	130	0
Rwanda	4.3	14.1	9.1	27.6	362.3	65	16
Ghana	13	6.3	7.8	27.1	655.5	103	0
Colombia	13.4	13.3	.2	26.8	4,284.3	9	25
Mozambique	9.8	10.3	5.3	25.5	399.2	123	5
Tanzania	11.5	6.4	6.4	24.3	457.6	78	0
Sudan	0	18.3	5.4	23.6	858.1	5	25
Uganda	8	10.6	4.7	23.3	410.3	32	23
Malawi	2.8	7.5	10.8	21.2	255.4	101	0
Senegal	12.9	6.3	.5	19.7	896.7	83	10
Mali	9.1	10	.1	19.1	540.7	84	6
Philippines	-.9	1.2	-.1	.3	1,499.9	18	25

The table displays the top 15 recipients, among countries with populations of 10 million or more, of total official development assistance (ODA) from the World Bank, the United States, and the United Kingdom. World Bank ODA refers to the International Bank for Reconstruction and Development (IBRD) and the International Development Association (IDA) only. USA and UK ODA refer to all sources of bilateral flows. The "Total" column refers to the total of these three sources. GDP per capita is adjusted for Purchasing Power Parity. ODA and GDP per capita are in 2011 US dollars and are obtained from the World Bank's World Development Indicators. "Pol. Instability" refers to the World Bank's Political Instability measure in the World Governance Indicators (WGI), which ranks 214 countries according to "perceptions of the likelihood of political instability and/or politically-motivated violence, including terrorism." Conflict refers to the definition from the Peace Research Institute Oslo (PRIO), which defines a conflict as involving 25 or more battle deaths in a year.

Table 2: Scale of PDT

	Cumulative				Annual			
	Villages		Municipalities		Villages		Municipalities	
Year	*N* of 42,013	% of Pop.	*N* of 1,648	% of Pop.	*N* of 42,013	% of Pop.	*N* of 1,648	% of Pop.
2002	457	1.21%	156	9.7%	457	1.21%	156	9.7%
2003	907	2.10	247	14.7	459	.89	143	8.9
2004	1,606	3.32	361	20.6	743	1.35	197	11.7
2005	2,445	4.72	448	27.4	928	1.61	202	13.7
2006	2,938	6.51	502	34.1	647	2.16	172	14.4
2007	3,600	8.10	604	38.4	929	2.21	254	18.1
2008	4,286	9.86	652	40.7	1,008	2.81	237	16.5
2009	4,996	11.5	693	46.2	1,027	2.49	213	18.7
2010	5,176	11.9	712	47.1	426	1.00	128	10.6

Source: Authors' calculation based on data provided by the Armed Forces of the Philippines (AFP). Note: Treatment is at the village level. Municipalities are considered treated if they include a treated village.

Table 3: Malnutrition summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max	P25	P75
All, unweighted							
Malnutrition rate	11506	11.805	7.333	.057	64.989	6.113	16.107
Severe malnutrition rate	11506	1.641	1.841	0	36.292	.523	2.094
Coverage	10824	87.222	18.645	6.38	877.45	76.92	100
All, population weighted							
Malnutrition rate	11506	9.163	6.527	.057	64.989	4.201	12.461
Severe malnutrition rate	11506	1.306	1.444	0	36.292	.503	1.62
Coverage	10824	89.223	15.732	6.38	877.45	81.3	100
Main estimation sample, population weighted							
Malnutrition rate	9957	9.565	6.687	.057	64.989	4.396	13.169
Severe malnutrition rate	9957	1.294	1.472	0	36.292	.493	1.594
Coverage	9289	90.388	15.064	6.38	877.45	82.77	100
ARMM, population weighted							
Malnutrition rate	27	10.84	5.187	.727	25.181	7.693	11.973
Severe malnutrition rate	27	2.689	2.248	0	10.228	1.206	3.311
Coverage	26	62.705	22.731	6.76	124.76	50.54	71.51

Source: Authors' calculations based on Operation Timbang (OPT) data. Note: Malnutrition is measured as a weight-for-age Z-score (WAZ). Full sample is based on an imbalanced panel of 1,537 municipalities with at least one malnutrition measurement from 2005-2012. Estimation sample is based on an imbalanced panel of 1,516 municipalities with at least one malnutrition measurement from 2005-2011, in which that measurement meets some minimal data quality standards, and in which other key variables are available. ARMM refers to the Autonomous Region of Muslim Mindinao, one of the Philippines' 17 regions, known to have particular data reliability issues. See text for additional details.

Table 4: Summary statistics

X	All municipalities Mean (s.d.)	R^2 from mun. fixed effects	PDT municipalities $P(X=0)$	Mean (s.d.) given $X>0$	Non-PDT municipalities $P(X=0)$	Mean (s.d.) given $X>0$
Malnutrition rate	9.56 (6.69)	.813	0	10.46 (6.98)	0	8.68 (6.26)
Violence	.022 (.067)	.466	.49	.070 (.111)	.81	.041 (.074)
PDT pop. frac.	.019 (.081)	.217	.72	.136 (.179)	–	–
Neighboring PDT	.018 (.054)	.264	.59	.077 (.093)	.89	.041 (.059)
Years since PDT	.90 (1.52)	.676	.23	2.38 (1.62)	–	–
N	9957		4697		5260	

"Malnutrition rate" is the percent of children age 5 and under who are two standard deviations or more below the mean of an age-specific reference population. "Violence" is the rate of violent incidents reported to the Armed Forces of the Philippines, per 1,000 population. "PDT pop. frac." is the fraction of the municipality living in a village that received PDT in the given year. "Neighboring PDT" is the PDT population fraction, averaged over the four nearest municipalities. "Years since PDT" is the number of years since at least one village in the municipality received PDT. See text for further details.

Table 5: Evidence of geographic clustering of PDT implementation

DV: PDT implementation	(1)	(2)	(3)	(4)
Fraction of 4 nearest neighbors where	0.0008***	0.0008***	0.0008**	0.0008**
PDT began in last 3 months	(0.0003)	(0.0003)	(0.0004)	(0.0004)
Fraction of 4 nearest neighbors where	-0.0002	0.0023***	0.0026***	0.0011***
PDT began in last 6 months	(0.0002)	(0.0004)	(0.0005)	(0.0004)
Constant	0.0013***	0.0010***	0.0005	–
	(0.0002)	(0.0004)	(0.0003)	
N (village-months)	5,044,920	5,044,920	5,044,920	5,044,920
R^2	0.0000	0.0005	0.0006	0.0129
Year effect	No	Yes	Yes	Yes
Month of year effect	No	No	Yes	Yes
Village effect	No	No	No	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$

Standard errors clustered at the province ($n = 86$) level in parentheses.

Sample includes a balanced panel of 49,460 villages from 2002-2010. Because of lags, the first 6 months of 2002 are excluded.

Table 6: Repeat PDT implementations

Implementations	Villages	Percent
1	3,873	74.0%
2	1,106	21.3
3	215	4.1
4	25	0.5
5	5	0.1
6	1	<0.1
Total	5,188	100%

Source: Authors' calculations based on data provided by the Armed Forces of the Philippines (AFP). Implementations refers to PDT implementations.

Table 7: Predicting PDT at the village-month level

DV: PDT began in month t	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of four nearest neighbors where PDT began during last 3 months	0.015***	0.015***	0.015***	0.014***	0.014***	0.008**
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.003)
Fraction of four nearest neighbors where PDT began during last 6 months	0.004	0.004	0.003	0.016***	0.018***	0.014***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
Number of incidents, average rate (t-1 through t-3)		0.006***	0.003***	0.005***	0.005***	0.002
		(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Number of incidents, average rate (t-4 through t-6)		0.003***	0.002**	0.004***	0.004***	0.003**
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Number of incidents, municipality average rate (t-1 through t-3)			0.076***	0.083***	0.084***	0.101***
			(0.024)	(0.025)	(0.026)	(0.028)
Number of incidents, municipality average rate (t-4 through t-6)			0.014	0.027**	0.028**	0.049**
			(0.010)	(0.012)	(0.012)	(0.020)
Indicator that the municipality/barangay received PDT at least once in the past				-0.027***	-0.030***	-1.042***
				(0.003)	(0.002)	(0.003)
Number of months since last PDT implementation began				0.000***	0.000***	0.038***
				(0.000)	(0.000)	(0.002)
Number of months since last PDT implementation began (squared)					-0.000**	-0.000***
					(0.000)	(0.000)
N (village-months)	4,540,104	4,540,104	4,540,104	4,540,104	4,540,104	4,540,104
R^2	0.015	0.015	0.015	0.022	0.022	0.471
Year effect	Yes	Yes	Yes	Yes	Yes	Yes
Month of year effect	Yes	Yes	Yes	Yes	Yes	Yes
Village effect	Yes	Yes	Yes	Yes	Yes	Yes
Vil.-specific trend	No	No	No	No	No	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$

Standard errors clustered at the province ($n = 86$) level in parentheses.

Village-specific trends are estimated by taking first differences and including a village fixed effect.

Table 8: Estimated effects of PDT

DV: $\ln(MalnRate)$	(1)	(2)	(3)	(4)	(5)
PDT_{t+5}					0.012
					(0.503)
PDT_{t+4}					0.198
					(0.275)
$\sum_{\tau=4}^{5} PDT_{t+\tau}$	-0.117	-0.060	0.132	0.195	
	(0.145)	(0.185)	(0.233)	(0.277)	
PDT_{t+3}				0.239	0.230
				(0.178)	(0.183)
PDT_{t+2}				0.166	0.162
				(0.119)	(0.122)
$\sum_{\tau=2}^{3} PDT_{t+\tau}$	-0.024	0.095	0.170		
	(0.079)	(0.122)	(0.122)		
PDT_{t}	-0.031	-0.093	-0.106	-0.126*	-0.124*
	(0.064)	(0.061)	(0.067)	(0.072)	(0.071)
PDT_{t-1}	-0.018	-0.183**	-0.211*	-0.248*	-0.242*
	(0.091)	(0.091)	(0.119)	(0.133)	(0.134)
PDT_{t-2}	-0.006	-0.260**	-0.316**	-0.367**	-0.359**
	(0.088)	(0.110)	(0.128)	(0.157)	(0.158)
$\sum_{\tau=3}^{4} PDT_{t-\tau}$	0.112	-0.306*	-0.356*	-0.427*	-0.417*
	(0.118)	(0.157)	(0.179)	(0.225)	(0.226)
$\sum_{\tau=5}^{6} PDT_{t-\tau}$	0.169	-0.388*	-0.485**	-0.575*	-0.561*
	(0.149)	(0.215)	(0.236)	(0.293)	(0.294)
$\sum_{\tau=7}^{\infty} PDT_{t-\tau}$	0.070	-0.580**	-0.687*	-0.795*	-0.779*
	(0.189)	(0.267)	(0.360)	(0.422)	(0.423)
N	11409	11409	11383	11383	11383
R^2	0.879	0.936	0.943	0.943	0.943
F (pre)	0.327	2.074	1.443	0.964	0.967
(p val.)	0.722	0.132	0.243	0.414	0.431
F (post)	2.939	3.463	3.269	3.026	2.825
(p val.)	0.017	0.007	0.010	0.015	0.021
Mun. FE	Yes	Yes	Yes	Yes	Yes
Mun. trends	No	Yes	Yes	Yes	Yes
Prov.×year FE	No	No	Yes	Yes	Yes
Years	2005-2012	2005-2012	2005-2012	2005-2011	2005-2011

* $p < .10$, ** $p < .05$, *** $p < .01$. Table shows estimated effects of PDT based on analyses at the municipality-year level. Effects are estimated using leads and lags of the fraction of the municipality living in a Barangay receiving PDT. To improve precision and save degrees of freedom, we impose many coefficients are equal (e.g., the 5-year effect is the same as the 6-year effect). The "pre" F-statistic corresponds to the null hypothesis that the coefficients on all leads are jointly zero. The "post" F-statistic corresponds to the null hypothesis that the coefficients on all lags are jointly zero. Standard errors, in parentheses, are clustered at the province-level (n=78).

Table 9: Distribution of years since PDT

Years since PDT	Municipality-years	Percent	Cumulative
1	1,004	23.6%	23.6%
2	876	19.8	43.3
3	788	17.8	61.1
4	588	13.3	74.4
5	413	9.3	83.7
6	284	6.4	90.1
7	210	4.7	94.8
8	131	3.0	97.8
9	69	1.6	99.3
10	30	0.7	100
Total	4,433	100%	

Source: Authors' calculations based on data provided by the Armed Forces of the Philippines (AFP). Table is based on the main estimation sample (2005-2012). Because PDT end dates are often missing, "Years since PDT" refers to the number of years since the most recent PDT implementation began, not ended. Because most PDT implementations for which end dates are available are less than 3 months, we do not consider this a major problem.

Table 10: Robustness to Selection

DV: $\ln(MalnRate)$	(1)	(2)	(3)	(4)
PDT_{t+5}	0.012	0.407	0.424	0.444
	(0.503)	(0.554)	(0.557)	(0.309)
PDT_{t+4}	0.198	0.369	0.376	0.218
	(0.275)	(0.294)	(0.293)	(0.202)
PDT_{t+3}	0.230	0.305*	0.310*	0.251*
	(0.183)	(0.182)	(0.181)	(0.140)
PDT_{t+2}	0.162	0.194	0.196	0.145*
	(0.122)	(0.120)	(0.119)	(0.085)
PDT_{t}	-0.124*	-0.166	-0.169	-0.097*
	(0.071)	(0.107)	(0.106)	(0.054)
PDT_{t-1}	-0.242*	-0.301	-0.308*	-0.243**
	(0.134)	(0.188)	(0.183)	(0.112)
PDT_{t-2}	-0.359**	-0.427*	-0.433**	-0.336**
	(0.158)	(0.219)	(0.216)	(0.133)
$\sum_{\tau=3}^{4} PDT_{t-\tau}$	-0.417*	-0.475*	-0.482*	-0.381*
	(0.226)	(0.262)	(0.259)	(0.201)
$\sum_{\tau=5}^{6} PDT_{t-\tau}$	-0.561*	-0.546*	-0.554*	-0.487*
	(0.294)	(0.314)	(0.311)	(0.275)
$\sum_{\tau=7}^{\infty} PDT_{t-\tau}$	-0.779*	-0.824*	-0.832*	-0.690*
	(0.423)	(0.451)	(0.447)	(0.387)
$Violence_{t-1}$			0.048	
			(0.047)	
$NeighborPDT_{t-1}$			-0.025	
			(0.065)	
N	11383	9933	9933	11383
R^2	0.943	0.948	0.948	0.943
F (pre)	0.967	0.883	0.897	1.099
(p val.)	0.431	0.478	0.470	0.363
F (post)	2.825	1.807	1.783	3.640
(p val.)	0.021	0.121	0.126	0.005
PDT	First	First	First	Any
Years	2005-2012	2005-2011	2005-2011	2005-2012

* $p < .10$, ** $p < .05$, *** $p < .01$. Table shows estimated effects of PDT based on analyses at the municipality-year level. Because incident data for 2011 is unavailable, including lagged incidents requires restricting to 2005-2011. All columns include municipality fixed effects, municipality-specific trends, and province-year fixed effects. Columns (1)-(3) are based on only a Barangay's first PDT implementation. Column (4) is based on any PDT implementation. The "pre" F-statistic corresponds to the null hypothesis that the coefficients on all leads are jointly zero. The "post" F-statistic corresponds to the null hypothesis that the coefficients on all lags are jointly zero. Standard errors, in parentheses, are clustered at the province-level (n=78).

Table 11: Malnutrition impacts for various interventions

Study	Study context	Dep. Var.	Estimated effect	Reduction
Brown et al. (2002)	Meta-analysis of 31 RCTs of zinc supplements	Most studies used individual-level WAZ scores	Zinc supplements increase WAZ by 0.309 SD (95% CI: [.178,.439])	43.9%
Jankowska et al. (2012)	Estimated effects of expanding drought and Sahel dessert in Mali, extrapolated from past droughts' effects	Village-level average WAZ scores	One standard deviation decrease in "drought" increases village average WAZ by .159 (off a base of -1.22, with standard deviation 0.49)	16.3%
Maluccio and Flores (2005)	RCT evaluating 2-year effects of conditional cash transfers (CCT) in Nicaragua	Community-level WAZ scores	Program reduced malnutrition rate by 6.2 percentage points (off a base of 16.6%)	37.3%
Ruel et al. (2008)	RCT evaluating 3-year effects of World Vision maternal and child health programs in Haiti	Community-level WAZ scores for children under 5	Program reduced malnutrition rate by 6 percentage points (off a base of 17.8%)	33.7%

Brown et al. (2002): To convert the impact on individual z-scores to the impact on the malnutrition rate, we used the following procedure. First, we assumed that the child weight distribution in the Philippines has the same variance as the reference population (this assumption is surely wrong). Then a malnutrition rate of 9.59% (our sample mean) implies that the average child's z-score is -.7, approximately. The zinc supplements, then, would move children's weights by .309, which implies a change in the fraction of children with z-scores below -2 (the fraction who are malnourished) from .0959 to .0538, or a 43.9% reduction in the malnutrition rate.

Jankowska et al. (2012): Formally, "drought" is defined as rainfall minus potential evapotranspiration (PET). To convert the impact on individual z-scores to the impact on the malnutrition rate, we used the following procedure. First, we assumed that the child weight distribution has the same variance as the reference population. We also assume that village size is orthogonal to malnutrition rates. Then we can directly convert the summary statistics of the village-level average WAZ score (mean: -1.22, standard deviation: 0.49) into a distribution of malnutrition rates (mean: 23.9%, standard deviation: 13.9%). A reduction in drought would increase the average village-level WAZ score by .159, implying the new average malnutrition rate is 20%, or 16.3% lower.

Table B.1: Predicting PDT at the municipality-month level

DV: Fraction of population receiving PDT	(1)	(2)	(3)	(4)	(5)
Fraction of 4 nearest neighbors' pop. where PDT began during last 3 months	0.015	0.014	0.012	0.012	-0.029
	(0.015)	(0.015)	(0.015)	(0.015)	(0.029)
Fraction of 4 nearest neighbors' pop. where PDT began during last 6 months	0.036***	0.036***	0.041***	0.040***	0.005
	(0.011)	(0.011)	(0.013)	(0.013)	(0.013)
Number of incidents, average rate (t-1 through t-3)		0.049***	0.050***	0.049***	0.053***
		(0.017)	(0.017)	(0.017)	(0.018)
Number of incidents, average rate (t-4 through t-6)		0.011	0.015*	0.014*	0.026**
		(0.008)	(0.008)	(0.008)	(0.013)
Fraction of municipality that received PDT at least once in the past			-0.021***	-0.021***	-1.017***
			(0.003)	(0.003)	(0.017)
Number of months since any barangay had a PDT implementation			-0.000***	-0.000***	0.001***
			(0.000)	(0.000)	(0.000)
Number of months since any barangay had a PDT implementation (squared)				0.000*	-0.000***
				(0.000)	(0.000)
N	1.78e+05	1.78e+05	1.78e+05	1.78e+05	1.78e+05
R-squared	0.027	0.028	0.034	0.034	0.401
Fixed effects	Yes	Yes	Yes	Yes	Yes
First difference	No	No	No	No	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$

Standard errors clustered at the province ($n = 86$) level in parentheses.

Municipality-specific trends are estimated by taking first differences and including a village fixed effect.

Table B.2: Predicting PDT at the municipality-year level

DV: Frac. of pop. with PDT	(1)	(2)	(3)	(4)
Fraction of 4 nearest neighbors' pop. receiving PDT (lagged)	-0.031	0.122**	-0.187**	0.063*
	(0.033)	(0.047)	(0.074)	(0.036)
Incident rate (lagged)	0.047	0.068**	-0.073*	0.073***
	(0.029)	(0.033)	(0.043)	(0.027)
Fraction of municipality that received PDT in the past		-0.459***	-0.929***	-1.250***
		(0.067)	(0.035)	(0.045)
Indicator that no barangay in municipality has received PDT		-0.006	0.112***	-0.028**
		(0.014)	(0.016)	(0.013)
Years since at least one barangay received PDT		0.004	0.101***	-0.005
		(0.003)	(0.013)	(0.011)
Years since PDT (squared)			-0.011***	0.002
			(0.002)	(0.002)
N	6,798	6,798	6,798	6,798
R-squared	0.283	0.413	0.464	0.514
Year effect	Yes	Yes	Yes	Yes
Municipality effect	Yes	Yes	Yes	Yes
Mun.-specific trend	No	No	No	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$

Standard errors clustered at the province ($n = 86$) level in parentheses.

Municipality-specific trends are estimated by taking first differences and including a village fixed effect.