

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

Eli Berman
Mitch Downey
Joseph Felter

December 21, 2015

Abstract

Worldwide, extreme poverty is often concentrated in spaces where people and property are not safe enough to sustain effective markets, and where development assistance is dangerous –and might even induce violence. Expanding governance by coercively taking control of territory may enable markets and development programs, but costs to local residents may exceed benefits, especially if that expansion is violent. We estimate for the first time whether a large counterinsurgency program improves welfare. We exploit the staggered roll-out of the Philippine “Peace and Development Teams” counterinsurgency program, which treated 12% of the population between 2002 and 2010. Though treatment temporarily increased violence, the program progressively reduced child malnutrition: by 10% in the first year, and by 30% from year three onwards. Improved nutritional status was not due to increased health and welfare expenditures, but instead to improved governance. Treatment effects are comparable to those of conventional child health interventions, though conventional programs are likely infeasible in this setting. Rebels apparently react to treatment by shifting to neighboring municipalities, as malnutrition worsens there –with statistically significant ‘treatment’ effects of similar size. Thus overall program effects are close to zero. These findings invite an evidence-based discussion of governance expansion, an extensive margin of development.

*Corresponding author: Eli Berman, 9500 Gilman Dr. #0508, La Jolla, CA 92093. Email: elib@ucsd.edu. Telephone: (858) 534-2858. The authors thank discussants Christopher Blattman and Matthew Webb, and participants in seminars and sessions at the NBER Economics of National Security, UC San Diego, University of British Columbia, the Empirical Studies of Conflict annual meetings at the US Institute of Peace, the American Economic Association, the Pacific Conference on Development Economics at UCSD, the Seminar for the Study of Development Strategies at Columbia University – for which Kolby Hanson and Anselm Rink replicated our results, the “Political Economy of Social Conflict” conference at Yale University, and the “Defence and Security Economics Workshop” at Carleton University. This research is supported by the Office of Naval Research (ONR) through Award N000141110735 at the National Bureau of Economic Research and the Development and Conflict Research Program at the UC Institute on Conflict and Cooperation. Any opinions, findings, conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of ONR. All mistakes are ours. Word count (including appendices): 15,306.

Keywords: Development and Conflict; Counterinsurgency; Malnutrition; Philippines

JEL Classification Numbers: F51, I15, O22.

Contents

1	Introduction	1
2	Context	4
3	Data	5
3.1	PDT	6
3.2	Violent incidents	6
3.3	Malnutrition	6
3.4	Summary statistics	8
3.5	PDT implementation	9
4	Selection into PDT	11
5	Estimated Treatment Effects	12
5.1	Estimating equation	12
5.2	Treatment effects	14
5.3	Spillovers	16
5.4	Robustness checks	18
5.5	Scaling	20
5.6	Net program effect	20
6	Interpretation	21
7	Conclusions	22
A	Additional data details	24
B	Additional results	25

“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 (1996), Expert Report to the Secretary-General of the United Nations

1 Introduction

Conflict and unstable governments are currently central challenges to development. Consider development assistance: as Table I illustrates, countries receiving Western Official Development Assistance (ODA) tend to be not only poor, but also conflict-cursed. The table lists the top 15 recipients of ODA per capita from the World Bank, the United States, and the United Kingdom, along with measures of poverty, political instability, and conflict. These 15 include four of the five least stable countries on the globe, and five of the bottom ten. Seven have suffered conflict for at least 15 of the last 25 years, with Sudan and Colombia involved in conflict for all 25. 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.

[Table I about here.]

Current recipients of development assistance are typically violent and unstable for two reasons. First, poor governance in general and violence in particular undermines investment of all types, including in human capital and in institutions. So while countries such as China, with politically stable governments that protect property and people (at least for elites) have experienced sufficient growth to graduate out of a need for assistance, the conflict-cursed countries in the table have not. Somalia, Afghanistan and Sudan, for instance, have average GDP per capita of less than \$2 per day.¹

Second, ODA is often directed at countries such as Iraq and Colombia (who are not among the poorest) with the objective of helping control their ungoverned spaces as part of a separate agenda of countering terrorism, narcotics, human trafficking and other global public bads. Economists might object that this agenda distorts allocation of assistance to the very poorest. Yet, if development assistance actually helped improve governance, that allocation might be justifiable on efficiency grounds. That extensive margin of development, bringing governance to unsafe spaces, is the focus of this paper.

The centrality of conflict in modern development raises both practical and conceptual challenges. Practically, development programs insert lootable and extortable resources like cash, food, equipment and personnel into insecure spaces, inviting predation and targeting by rebel groups.² Conceptually, the usual logic by which development programs improve welfare

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

²See Collier (2000), Nunn and Qian (2014).

requires that people and property are safe enough to allow contracting. If so, an injection of resources implies a welfare improvement if the other conditions of the first welfare theorem are also met. In the absence of those conditions, development assistance of any form may reduce welfare of residents by shifting rents to extortionary actors or by inducing welfare loss through violence and human rights abuses.³

Recent empirical results underscore 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 intensity and duration of civil conflict. Aid *can* intensify conflict, perhaps undermining the possibility of economic development by reducing security of people and property.⁴

Rather than implementing development programs in poorly governed spaces, an alternative might be to first expand governance into that space. This is often an option; many countries, like the Philippines, have pockets of ungoverned or poorly controlled space bordering fairly well controlled areas. Expansion of governance is a new topic for economists, though recently Callen et al. (2015) provide an economic explanation for the expansion of Pakistan into formerly ungoverned space, and both Vanden Eynde (2015) and Berman et al. (2012) have looked at whether potential tax revenue induces government to contest rebel control, in India and the Philippines respectively. While as economists we may have normative concerns about recommending a coercive policy, it would nevertheless be a valuable positive exercise to investigate what the net welfare effects of such an expansion of governance are.

Does expanding governance improve development outcomes? The answer is theoretically ambiguous. We tend to think that a state will provide institutions that are more conducive to functioning markets (and therefore welfare) than will rebels or informal governance. Yet states sometimes neglect populations in the periphery, and rebels often provide services.⁵ Moreover, even if state governance is preferable to that by rebels, the transition is sometimes coercive and destructive (Machal, 1996), and accompanied by abuses of human rights, as it has been in the Philippines (Sales, 2009). So the net effect of an expansion of governance may make it undesirable for residents, even if government is preferred to rebels in a steady

³Grossman (1999).

⁴One option implemented in Afghanistan and Iraq was development programs selected and protected by the military. Berman et al. (2013) show that development projects, when accompanied by adequate force levels in Iraq, were security enhancing (in the sense of reducing violent incidents). Discussing development and counterinsurgency in Afghanistan, NATO force commander McChrystal said: “We view it as a process... which enables Afghan ownership and reinforces Afghan sovereignty... In some areas, it will be security assistance. In some areas... it will be more based on help with governance and development...” (McChrystal, 2010).

⁵Berman (2009) describes provision of services by Hamas, Hezbollah, the Mahdi Army and the Taliban; Heger (2010) documents community services provided by the Irish Republican Army; Keister (2011) describes services provided by the Moro Islamic Liberation Front and the Moro National Liberation Front in the Southern Philippines. These sources and anecdotal evidence suggest that when rebels control territory they typically provide at least some form of security and dispute adjudication services to noncombatants, apparently at low cost to themselves.

state.

To date, no empirical research has directly assessed the net benefits of expanding governance, despite the volume of assistance now directed at unstable spaces. In this paper, we attempt to fill that void by evaluating a large counterinsurgency program operated by the Armed Forces of the Philippines: “Peace and Development Teams” (PDT). Exploiting unprecedented access to precise program data on location and timing, and 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.

The program progressively reduced child malnutrition, by about 10% after the first year and 30% after three years, an effect that is persistent for the seven year duration that we can estimate.⁶ These improvements in child welfare occurred despite increased violence associated with treatment –much of which is initiated by government forces. These improvements are slightly smaller than the treatment effects of other successful child health interventions, yet those alternatives would be difficult to implement in such weakly governed spaces. While we can provide little detail about the mechanism by which PDT reduced child malnutrition, we can rule out the possibility that it proxies for a general expansion of social services, including health services. We therefore conclude that improvements in security, in access to markets, or in access to existing government services must be the causal mechanisms.

How do rebels react to PDT implementation? One possibility is that they are deterred, and retreat into inaction. This is often the working assumption of one-location models, even when estimated with data from multiple locations (e.g., Berman et al. (2013)). Our evidence indicates the opposite. Rebels apparently respond to PDT by relocating to neighboring municipalities, as evidenced by worsening malnutrition among neighbors following PDT. Thus, despite the improvements in child nutrition in treated municipalities, the aggregate effect of the program seems to be nullified by a displacement of malnutrition to neighboring municipalities which experience an approximately equal and opposite increase in malnutrition. This is the first evidence, to our knowledge, of a localized strategic reaction by opponents to a program of this type.

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, Section 4 is useful for understanding threats to causal inference. In Section 5 we use an “event study” design to estimate the malnutrition effects of PDT. We show that PDT was not preceded by trends in malnutrition and that our results are unaffected by the selection controls identified in Section 4. Section 6 presents additional results that are useful for interpreting the estimated effect before Section 7 concludes with a discussion of policy implications, comparing these estimated effects to those of other malnutrition interventions and speculating on the larger question of secure governance in development.

⁶Estimates are subject to caveats about scaling, which we discuss below.

2 Context

The Philippines has suffered low-grade civil conflict for many decades,⁷ which continues despite a history of fairly functional multiparty democracy since 1986 (polity score of eight, since 1989).⁸ In that sense it is similar to Colombia, Mexico, Indonesia, or Turkey.

While a number of rebel groups are active, they can be broadly classified into two categories. The New People’s Army (NPA), the armed wing of the Communist Party of the Philippines (CPP), form the first category. The NPA has a following among the very poor, which is exacerbated by severe income inequality and the very gradual pace of land reform, even in the period of democracy (Montinola, 2013). During the first decade of this century, the NPA has accounted for nearly two-thirds of violent incidents (Croston et al., 2014). It is active throughout the Philippines, though more so in rural areas.

Islamic separatists form a second category, which includes the Moro Islamic Liberation Front (MILF), the Moro National Liberation Front (MNLF), and the smaller and more extreme Abu Sayyaf Group (ASG). Separatists are primarily active in the country’s South, in Mindanao and the Sulu Sea, and are 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 Mindanao (ARMM) in 1996, one of the Philippines’ seventeen regions.⁹

Importantly, the formal government faces little risk of being overthrown. The asymmetry of power is heavily tilted in its favor.¹⁰ The conflict with the NPA is characterized by small-scale insurgent attacks in poor rural areas, rather than frequent full-scale battles. Government authority is relatively uncontested in urban areas, where rebel groups have limited popular support and opportunities to organize. Figure I displays this pattern, showing that both violent incidents and malnutrition –our key development outcome, are associated with low population density. Both conflict and poverty are concentrated in rural areas.

[Figure I about here.]

As part of an effort to expand governance, the Philippine Army launched the PDT program in 2002, which continued earlier “stabilization operations.” Our sample period roughly coincides with the Presidency of Gloria Macapagal-Arroyo¹¹ (January 2001 until June 2010), though the program has continued under her successor.¹² According to the

⁷Here, we provide only a brief review of these conflicts. The interested reader should see Croston and Johnston (2010), Felter (2005), Hernandez (2014), Quimpo (2012), or Schiavo-Campo and Judd (2005).

⁸Since 1989, the Philippines has maintained a polity score of eight. This score, originated by Ted Gurr, ranges from -10 (least democratic) to 10 (most democratic) and is based on executive recruitment, executive constraints, and political competition (Marshall and Cole, 2011).

⁹Regions are the largest subnational division of government.

¹⁰In asymmetric conflicts like these, governments struggle to obtain information that helps them locate insurgents, rather than struggle to defend territory (Berman and Matanock, 2015).

¹¹A professor of economics from an elite Philippine family.

¹²The sample period also includes a Presidential election in May 2004 and three small suppressed military

program manual, special Army units were designated to enter selected villages (*barangay*),¹³ clear out entrenched rebels, assess community needs, and connect the village to government programs and services. In principle, this might include building schools or clinics (though we report evidence below to the contrary).¹⁴ Treatment more likely included 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. 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 II 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, accounting for 1-2% of the population, were treated. Over the full period, over 5,000 villages (fully 12% of the population) received treatment. The Philippines has about 1,600 municipalities, nearly half of which (47% of the population) contained a village treated with PDT during the sample period.¹⁵

[Table II about here.]

Figure II maps implementations. Note that PDT treatment is spread throughout the Philippines, though disproportionately in peripheral locations with low population density.

[Figure II about here.]

In conclusion, the PDT program represents a concerted effort by the government to systematically expand control into areas that are relatively poor, violent and rural. It includes security and, in principle, development elements. To estimate treatment effects we will exploit staggered roll-out over nine sample years.

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

uprisings: a mutiny in July 2003, a state of emergency in February 2006 in reaction to an alleged coup attempt and a small rebellion in November 2007. None of these seem to affect program rollout (analysis available upon request).

¹³Officially, 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.” The exact number of villages and municipalities in the Philippines changes every year, as these units merge and split. For analysis we use a consistent set of village and municipality definitions that closely corresponds to official 2009 definitions used in Felter (2005).

¹⁴Unfortunately, we lack data about the specific set of services offered in PDT operations.

¹⁵The geographic unit immediately larger than villages is the municipality, which have considerable political authority. 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.

appendix.

3.1 PDT

This research is enabled by access to a dataset unique in the study of counterinsurgencies. The Armed Forces of the Philippines (AFP) have shared with us the location and timing of all PDT implementations from 2002 through 2010.¹⁶ While the AFP had been expanding governance and providing development services long before 2002, we lack systematic data on those efforts. (So we can only study the effects of this particular wave of interventions, without being able to account for how previous programs might influence current treatment effects.)

3.2 Violent incidents

Our incident data include the full universe of violent incidents reported by the AFP. They were first compiled and analyzed in Felter (2005)¹⁷ and subsequently updated as part of the Empirical Studies of Conflict project. The data are based on underlying AFP incident reports.¹⁸ For each incident, the data includes rebel, civilian, and government casualties; the number of rebels captured or surrendered; and a geographic code corresponding to the village where the incident occurred.¹⁹ A unique characteristic of these data is an indicator for whether the incident was rebel- or government-initiated.

3.3 Malnutrition

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 (RHMP) through which placements began in 2008.

¹⁶For 75% of implementations, we also observe end dates. 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. 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.

¹⁷See also Crost et al. (2014) and Crost et al. (2013).

¹⁸Because the data originally come from AFP incident reports, they likely undercount attacks in which the AFP was not involved. This complicates the variable's interpretation, but does not bias our results.

¹⁹For about 7% of incidents, the geographic code corresponds to a municipality. These incidents are excluded from village-level analyses. Since our main analyses are conducted at the municipality-level, this has little bearing on our empirical results.

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.²⁰ 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 occurred 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. Since OPT is implemented by local staff from the same village they are likely to know of and have access to all children.²¹ 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,²² and of the reference population used, follow the recommendations of the World Health Organization (WHO). The publicly available NNC data has two limitations relative to the underlying data collected. 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).

Table III 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 75th percentile has two and a half times the malnutrition rate of the 25th percentile. The second panel weights

²⁰Recall that the average village is 3,000 people.

²¹Similarly, 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.

²²This 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. 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 (Growth Reference Study Group, 2006). This was done to maintain consistency with international standards.

municipalities by their population to obtain estimates more representative of the country.

[Table III about here.]

Of particular concern is the coverage measure, which is often quite low and sometimes quite high. We might be concerned that the malnutrition estimates from these municipalities are not reliable. Thus, for our main analyses, we exclude municipalities for which *a*) coverage is less than 66% or greater than 110%, and *b*) population is greater than twice the mean.²³ The third panel displays malnutrition characteristics for this sample, which trims about 13% of observations.

Finally, OPT contains no data from the Autonomous Region of Muslim Mindanao (ARMM), where Islamist rebel groups are most active. This is because the autonomous regional government has overwhelming legal authority and is not required to cooperate with federally sponsored data collection. As such, our malnutrition results are 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 III, which reports the weighted mean and 90/10 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 III about here.]

3.4 Summary statistics

Our main analyses will be conducted at the municipality-year level. Table IV presents summary statistics, including separate figures for municipalities that received PDT and those that did not.²⁴

[Table IV about here.]

Municipalities receiving PDT tend to be relatively disadvantaged: They have higher malnutrition rates (10.1% compared to 8.6%), are more likely to experience violent incidents

²³Coverage 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. The NNC recommends caution with measurements outside 80 and 110. We felt 80 was too restrictive. We only exclude municipality-years with poor coverage when they are large because we weight by population and were primarily concerned about inaccurate measurement among particularly influential observations.

²⁴The 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.

(70% of municipality-years compared to 41%) and, when experiencing violence, tend to have more intense violence (7 incidents per year per 1,000 residents compared to 4.4). These descriptive statistics foreshadow the formal results regarding non-random PDT implementation discussed in the next section.

Table IV also reports several other facts relevant to estimation. First, municipality fixed effects alone account for over 79% of variation in malnutrition, our dependent variable, leaving less than 21% with which to estimate treatment effects conditional on municipality. Second, conditional on a PDT implementation within a given municipality-year, only 12.5% 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, we observe malnutrition for years after PDT implementation (2.6 years after the most recent, on average) allowing us to estimate delayed effects.

3.5 PDT implementation

Three aspects of the program are important in understanding PDT implementation: PDT implementations are geographically clustered; repeat implementations are common; and nominal implementation is preceded by a period of government-initiated incidents.

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 might simply push rebels to the next village, which would then require attention, so treating clusters would be efficient.

We investigate spatial clustering by estimating

$$PDT_{it} = \alpha_i + \delta_t + \beta_3 NeighborPDT_{i,t-1,t-3} + \beta_6 NeighborPDT_{i,t-4,t-6} + \varepsilon_{it} . \quad (1)$$

Here PDT_{it} is an indicator that PDT was implemented in village i during month t , δ_t is a month effect, $NeighborPDT_{i,t-1,t-3}$ is the fraction of the four villages nearest to i in which PDT was implemented between 1 and 3 months ago, and $NeighborPDT_{i,t-4,t-6}$ is the fraction of these villages receiving PDT between 4 and 6 months ago. As shown in Table V below, these neighboring implementations are highly predictive. The first column shows that the effect of having all four nearest neighbors receive PDT in the last three months increases the probability of PDT receipt by over 3 percentage points, from a base of only .12%. Columns (2)-(4) show that village fixed effects reduce this effect to about 2 percentage points, but it remains highly statistically significant ($p < .01$).

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 selection into PDT treatment.

[Table V 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 VI 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. A possible predictor of PDT treatment will therefore be time since last implementation in the same village.

[Table VI about here.]

To better understand repeat implementations, Figure IV plots the probability that a repeat PDT is observed against the year of first PDT implementation (in our sample period). The probability of a repeat PDT implementation declines in program year, from between 40 and 50% in the first three program years to zero in the last. One interpretation of these data is that the program improved over time in reducing recidivism. Another is that the most difficult villages were selected first. A third interpretation is simply that the data are right-truncated, so 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. That last interpretation would suggest that about half of villages will eventually experience multiple treatments.

[Figure IV about here.]

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 estimate the effects of the average PDT implementation, rather than restricting to a municipality’s first implementation.

We turn now to violence. PDT implementation in principle follows a “clear, hold, build” logic: offensive force first is used to reduce rebel presence, then a defensive posture is adopted to retain control, followed by implementation of governance and development improvements. That sequencing is consistent with complementarity between security and development (Berman et al., 2013).

To assess whether PDT follows that sequencing we use an event study specification which regresses government-initiated violent incidents on indicators for the months preceding and following PDT implementation, estimating the following specification using monthly (indexed t) village (indexed v) data,

$$Violence_{vt} = \alpha_v + \delta_t + \sum_{\tau=-18}^{18} \beta_{\tau} PDT_{vt} + \varepsilon_{vt} \quad (2)$$

Figure V plots the estimated event study β_{τ} coefficients for the 37 months around PDT implementation (month 0 in the figure being the month of implementation).²⁵ For reference,

²⁵For regression results, see Table BI of the appendix.

mean violence for the omitted category (months more than a year and a half before or after PDT) is .0024 incidents per month per 1,000 population.

[Figure V about here.]

Figure V shows that government-initiated violence increases about fivefold nine months before a nominal PDT implementation date, increases again the month before implementation, remains high for about four months, and then drops to the pre-implementation plateau and declines gradually back to the year-before mean over the next half-year. For our purposes the key is that the “clear” stage can be dated to nine months before the nominal implementation date, a fact reflected clearly in the precisely estimated spike reported in Table BI.²⁶

4 Selection into PDT

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 so that estimated treatment effects are confounded with selection. In this section we investigate the selection mechanism in order to help design estimates which avoid that bias. In the next section we will take two approaches to selection bias: using an event-study specification that examines estimated pre-implementation trends directly, and including predictors of selection as regressors.

Formally, we aspire to estimate the coefficients of an equation like this one,

$$\ln(\text{MalnutritionRate})_{it} = \alpha_i + \beta_1 \text{FractionPDT}_{it-1} + \varepsilon_{it} \quad (3)$$

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.

We estimate a PDT selection equation with four features. First, to account for the “clear” stage, we date PDT implementations as beginning nine months before their nominal start date. Second, because PDT is implemented in geographic clusters, we control for PDT implementations in the nearest four villages over the previous three and six months (column 1). Third, to allow for previous violence to predict implementation, we include three and six month lags of incident rates in the village (column 2) and in the village’s municipality (column 3). 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

²⁶One might be concerned that violence actually causes reactive PDT at a nine month lag, rather than the opposite. To check, we created the same figure using rebel-initiated violence (rather than government-initiated), and found no similar lead pattern, leaving us to conclude that causality more likely flows from planned PDT to pre-PDT, government-initiated violence.

quadratic (column 5) time trend in months since the last PDT implementation.²⁷

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

Table VII reports these results, estimated at the village-month level.²⁸ The three major findings of the previous section again express themselves in predicting PDT implementation. PDT in neighboring villages during the previous three and six months predict new PDT starts, as we saw in Table V. Lagged incidents predict PDT, over the previous three and six months. Finally, recent PDT in the same village predicts an increased probability of repeat treatment, though that effect fades over time. While all of these findings hold after including village fixed effects (column 7), all coefficients drop substantially. This suggests that a large part of the non-random selection is accounted for by time-invariant unobserved heterogeneity.

[Table VII about here.]

To summarize, the AFP seems to select villages for treatment by violence, proximity to treated neighbors and recent treatment. That selection has the potential to bias our estimates of PDT on malnutrition (equation (3)), which we now turn to.

Unfortunately, our outcome measure (child malnutrition) is available for municipalities rather than villages and only annually. Appendix B reports the model presented in Equation (4) aggregated to the municipality-year level to match the units of observation we will use to estimate treatment effects. The key results in Table VII hold at these higher levels of aggregation (see Table BII).

With a set of predictors of PDT implementation in hand, we can check for robustness to 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.

5 Estimated Treatment Effects

5.1 Estimating equation

We can now estimate the effect of PDT on malnutrition while checking for possible selection effects. We use two methods to deal with selection bias: first, we include a set of

²⁷To avoid imputing PDT end dates, we use a time trend in months since the last PDT *began*. Since we include an indicator of whether the village had previously received PDT, the value of that time trend for villages never receiving PDT does not apply to those villages.

²⁸Table BII in the appendix shows the same specification estimated at the municipality-year level.

lead coefficients in an “event study” specification (Jacobson, LaLonde, and Sullivan, 1993; Hoynes, Page, and Stevens, 2011; Sandler and Sandler, 2013), that also allows municipality specific trends, and province-year effects as controls; second, we include predictors of PDT suggested by the analysis of selection in the previous section.

Our base estimating equation is

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

where m indexes municipalities, p indexes provinces, and t years. $\ln(\text{MalnRate})$ is the natural logarithm of the malnutrition rate, $\text{PDT}_{mt-\tau}$ is the fraction of municipality m ’s population living in a village receiving PDT in year $t - \tau$. As in the previous section, we date PDT at nine months before the nominal start date to account for the “clear” evident in Figure V of Section 3.5 above. We use log malnutrition because we expect a proportional response to treatment.²⁹

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 before treatment. The coefficients β_τ for $\tau > 0$ estimate lagged impulse-response as the result of PDT treatment. This may occur because governance is persistent, or because the economic effects of governance take some time to set in. In particular, we might expect investment in physical and human capital to accumulate over years as a result of improvements in property rights, safety, and access to markets. The health effects of improved health and nutrition would also take time to accumulate in weight for age.

Our specification drops the 1-year lead of PDT to form a reference point (the regression constant) in checking for pre-existing trends. Given the reference point and the timing of measurement (January through March of every year), treatment effects should primarily be reflected in the lag coefficients β_τ for $\tau > 0$.

Rebels may respond to PDT treatment by displacing their activity to untreated areas. To the extent that this happens within the same municipality it will be captured in the estimated treatment effect, since our unit of observation is a municipality. Should rebels relocate their activity to neighboring municipalities, we can estimate the effect of possible spillovers by including as a regressor the average level of PDT in the four nearest neighbor municipalities (measured by geometric distance from midpoint).

²⁹The malnutrition rate is the percent of children whose weight falls below an age-specific threshold recommended by the World Health Organization (on the basis of a standard reference population). In a model where PDT has a linear effect on children’s weights near the threshold, the effect will be larger where the mass near the threshold is larger. Unfortunately we do not observe the mass near the threshold. However, because children’s weights tend to have bell-curve type distribution (The WHO assumes Box-Cox Power Exponential) and because the threshold is likely to the left of this distribution’s peak (since the average malnutrition rate is only about 10%), we expect the mass to be larger when the malnutrition is higher. The log transformation generates larger effects when the malnutrition rate is higher. The malnutrition rate is never zero (Table III), so no observations are lost by taking logarithms.

Our estimating equation including spillovers is then

$$\begin{aligned} \ln(\text{MalnRate})_{mt} = & \alpha_m + \delta_{p(m)t} + \gamma_{mt} + \sum_{\tau=-k}^k 1\{\tau \neq -1\} \beta_{\tau} \text{PDT}_{mt-\tau} \\ & + \sum_{\tau=0}^k \theta_{\tau} \text{NeighborPDT}_{mt-\tau} + \varepsilon_{mt} \end{aligned} \tag{6}$$

where $\text{NeighborPDT}_{mt-\tau}$ is the population-weighted average of PDT in the four nearest municipalities so that the lag coefficients θ_{τ} estimate spillovers from neighboring municipalities. Following the discussion above, those spillovers cannot be signed in theory. They might increase malnutrition due to transition effects and degradation in governance between government and rebel control, or they might improve nutrition if rebel governance is actually an improvement.

5.2 Treatment effects

Our main results are illustrated in Figure VI which plots the estimated coefficients of equation (6). Visually, there is little evidence of a pre-existing trend in PDT before treatment, in years $t=-5$ through $t=-2$. The year $t-1$ coefficient is normalized to zero. From the year that treatment begin ($t=0$), we see a persistent reduction in malnutrition, which levels off five or six years after treatment. Standard error bands indicate that none of the estimated coefficients are statistically different from zero (at the 5% level).

[Figure VI about here.]

Table VIII reports estimated coefficients. In order to gain precision coefficients are restricted to be equal for lag combinations (3,4), (5,6) and 7 or more, as reported in column 1. (The full set of lags and leads illustrated in the figure is reported in Appendix table BIV.) Specifications reported in this table all include a full set of municipality fixed effects, municipality specific trends, and province x year fixed effects. (Appendix table BIV reports that results are sensitive to excluding municipal trends, and that province x year fixed effects improve precision.)

[Table VIII about here.]

In this specification (and those that follow) the coefficients on PDT leads are generally small. We test for selection or some other source of pre-existing trends with an F-test of the joint significance of the four lead coefficients (which fortunately share the same sign). The F-test has a p value of 0.56, which indicates no evidence of a pre-existing trend.

Treatment effects of PDT on malnutrition are reflected in the contemporaneous and lagged coefficients: a reduction of 6.6 log points in the treatment year, 18.1 log points the

year after, 25.9 log points after two years... through a 33.6 log point reduction after seven years. These are large reductions. Nevertheless, the F-statistic on the joint significance of the six contemporaneous and lagged coefficients has a p-value of 0.06, not quite significant at the 5% level.

More precision can be gained for testing by further restricting the set of lagged coefficients, but before we pursue that, the true size of our F-tests requires some attention. Our standard errors are clustered at the province level (for the 78 Philippine provinces). Recent results by MacKinnon and Webb (2014) indicate that when clustered entities are unbalanced (i.e., come in very different sizes), or when the proportion of clusters treated is small, sampling distributions of the cluster robust variance estimator diverge from asymptotic approximations, even for fairly large sample sizes, resulting in test size distortions and frequent overrejection of the null hypothesis.

As a check, Column 2 reports the same set of coefficient estimates with standard errors estimated by the wild cluster bootstrap method of Cameron et al. (2008), which are consistent. Note that these bootstrapped standard errors are about ten percent larger than the analytic in Column 1, indicating a possible size distortion in testing. More importantly, a wild cluster bootstrap of the F-test indicates that our F-statistic rejects a true null hypothesis fully 14 percent, more than twice the 6 percent estimated of the asymptotic p (labeled “analytic” in the Table).³⁰

To illustrate this point, Figure VII compares the empirical distribution of the bootstrapped F-statistic, the histogram, to the analytical F distribution, which is drawn as a curve. The histogram clearly has a lot of density to the right of the curve. The vertical red line indicates the realized F-statistic of 2.228, with six percent of the curve to the right of it, but 14 percent of the histogram. This discrepancy suggests skepticism about the analytical p-value, and about our asymptotic approximations in general, when clustering by province.

[Figure VII about here.]

Faced with a lack of precision, despite our wealth of observations, we must make some assumptions. One approach would be to impose uncorrelated errors across municipalities (as well as across provinces). This assumption would be violated if rebels responded to PDT by relocating their activity to neighboring municipalities in a way that affected child nutrition. Column 3 reports estimates of the same specification with standard errors clustered by 1513 municipalities. Standard errors of estimated treatment effects (lags) are generally larger under this assumption, leading to an F-statistic of 1.035 with a p-value of 0.346 (bootstrapped), even further from rejection. Larger standard errors under an assumption of uncorrelated residuals indicate that ε_{mt} may be negatively correlated across municipalities in equation (5), which would be consistent with spillover effects of PDT across municipalities through rebel relocation. Henceforth, we stick with clustering standard errors at the province level.

Before turning to a test for spillovers across municipalities, we try an alternative route to precision by restricting lead and lag coefficients. This should yield precision gains since

³⁰We thank our discussant Matthew Webb for alerting us to this possible problem, and for generously assisting with Stata code for testing.

PDT is highly serially correlated within municipality, so that estimated coefficients have high (negative) sample covariance. The event study plot in Figure VI indicates many options for constraining coefficients, since lead coefficients look statistically identical and lagged coefficients change very little from three years after treatment onwards. Column 4 reports the effect of restricting the coefficients on leads two through four to be equal, while also restricting coefficients on lags three onward to equality. Estimated treatment effects are similar through lag three, with a long run estimated reduction in malnutrition of 27.8 log points. In this specification the F-test on leads indicates joint insignificance ($p=0.47$), while the F-test on lagged and contemporaneous coefficients rejects at $p=0.026$ (wild clustered bootstrap) and 0.017 (analytical).

Further precision can be obtained by restricting the contemporaneous and lagged coefficients to a spline function that restricts coefficients in years zero, lag one and lag two to be linearly related with a constant slope $(\tau+1)*\beta$. That specification is reported in Column 5, yielding a slope of 0.084 for treatment year, the next year and the following year (reductions in malnutrition of .084, .168 and .252 log points) and a long run effect of -0.286 after year three. With that restriction, the joint F test rejects a null of no treatment effect at $p=0.01$ (wild clustered bootstrap).

Figure VIII illustrates the splined specification of Column 5 with a solid line in a shaded confidence region. Compared to the unconstrained results, the estimated treatment effects are indeed smaller (i.e., less malnutrition reduction). Nevertheless, from one year after treatment onwards the treatment effect is precisely enough estimated to be statistically different from zero, with a long run, precisely estimated, reduction in malnutrition of 28.6 log points, or 25 percent.³¹

[Figure VIII about here.]

5.3 Spillovers

We return now to the possibility of spillover effects of PDT on malnutrition in neighboring municipalities. Geographical relocation of rebels is common in asymmetric civil wars, and is anecdotally present in the Philippine context in particular. Having established that PDT has a treatment effect, it would not be surprising to find that relocation of rebels has an opposite effect, through reduced security of people and property or reduced access to markets. That type of spillover was also suggested by the discovery of negatively correlated residuals across municipalities in Table VIII.

Table IX reports the results of estimating equation (6), which augments estimated treatment effects with spillover coefficients reflecting the effects of PDT in neighboring municipalities. Column 1 reports again the estimated treatment effects from Table VIII Column 5, using the spline specification to maximize precision.

³¹The estimated five year lead coefficient has very low precision for lack of data. Table BIII (below) 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.

[Table IX about here.]

Column 2 adds analogous coefficients for neighbor PDT, again imposing the spline restriction on lagged coefficients, as in Column 1, but this time on both own and Neighbor PDT. The coefficients on PDT in neighboring municipalities (averaged over the four nearest neighbors) are positive, indicating that PDT in neighboring municipalities increases malnutrition. They are about the same size as the own-PDT coefficients. The long run (3+ years) coefficient on neighboring PDT is statistically different from zero, indicating a 34.4 log point increase in malnutrition, as opposed to the long term own PDT coefficient estimate indicating a 35.5 long point reduction. The neighboring PDT coefficients are jointly just barely statistically significant at the ten percent level (with an analytical p-value of 0.106 reported for an F test at the bottom of the table, and a bootstrapped p-value of 0.099) but provide highly suggestive evidence that PDT in neighboring municipalities worsens malnutrition, especially in the long term.

Those effects are illustrated by Figure IX, which plots spline coefficients and confidence intervals for both own PDT and neighbors' PDT. The effects are approximately symmetric and of opposite sign, both over the first three years, and thereafter.

[Figure IX about here.]

One concern with this evidence of spillover is that it may be a statistical “reflection” due to the correlation of neighbor PDT with own PDT, causing a bias to be transmitted from the latter coefficients to the former. In column 3 we explore that possibility by estimating the spillover coefficients (the effects of PDT in neighboring municipalities) for 6001 municipalities that do not receive PDT during the sample period.³² We find strong evidence of spillover even into municipalities that do not receive any PDT: the F-test for joint significance of the two spline coefficients has a p-value of 0.01. We infer that PDT in neighboring municipalities truly does increase malnutrition through spillover.

To complete the split-sample analysis, Column 4 reports estimates of the full spillover specification for the complementary sample of municipalities that ever experience PDT. Own-PDT and neighboring PDT coefficients are very close to those reported for the full sample (Column 2) and statistically indistinguishable. Compared to estimates for municipalities that never experience PDT, the spillover effects are much smaller in municipalities that experience at least some PDT. One possible explanation is that a small amount of PDT is protective against spillovers. Another is that treatment effects differ with the level of PDT, so that a log-linear specification provides a poor approximation. Municipalities that never experience PDT average much lower rates of PDT among neighbors (4.1% vs. 7.7%) as well as lower malnutrition rates (8.6% vs. 10.1%), as reported in Table IV.

Returning to the estimated effects in Column 2, the estimated spillover effects invite interpretation. PDT induces negative spillovers among neighbors, in addition to positive improvements in child nutrition. Including those coefficients in the longer regression corrects

³²We thank David Green and Francesco Trebbi of the University of British Columbia for this suggestion.

a small omitted variable bias (since own PDT is apparently correlated with neighbors' PDT), resulting in slightly larger own-treatment effects. The effect for 3+ years is stronger, a 35.5 log point reduction. That effect might also have been slightly overstated in the short regression due to contamination of comparisons from treatment, but that contamination would have been diluted among the 1,512 comparison municipalities.

We have no direct evidence on the mechanism by which spillover takes place, yet anecdotal evidence suggests that PDT displaces rebels from treated municipalities to neighboring municipalities, reversing the mechanism by which PDT reduces malnutrition. The only other plausible alternative we can think of by which PDT would lower malnutrition for children in treated municipalities while increasing it for children in neighboring municipalities is that PDT redirects other government resources from the neighbors to the treated. We rule that out with direct evidence on spending below, leaving us to conclude that rebels react to PDT not by being deterred into inaction but by relocating to neighboring municipalities, which in turn worsens malnutrition. To our knowledge this is the first empirical evidence of strategic reaction by rebels to a large scale development program of this type.

5.4 Robustness checks

The results in Table VIII include municipality and province-year fixed effects and municipality-specific linear trends. The lead coefficients show no evidence of pre-existing trends before treatment. However, we can additionally test for selection bias by checking if estimated treatment effects are robust to including predictors of PDT, which we investigated in Section 4. There we found that treatment of neighbors was a predictor of PDT, which we have covered in our discussion of spillover. We also saw that lagged violence predicted PDT, which we turn to here. In addition, we will check for robustness of results to using only the first instance of PDT (in a village) and to the calculated timing of treatment.

Table X shows that the estimated treatment effects are virtually unaffected by adding violence as a predictor of PDT. Column 1 replicates our preferred specification, column 2 of Table IX. 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 very similar to those in column 1 except that the estimated treatments effects are less precisely estimated.³³

Column 3 adds violence as a predictor, estimating

$$\begin{aligned} \ln(\text{MalnRate})_{mt} = & \alpha_m + \delta_{p(m)t} + \gamma_m t + \sum_{\tau=-k}^k 1\{\tau \neq -1\} \beta_\tau \text{PDT}_{mt-\tau} \\ & + \sum_{\tau=0}^k \theta_\tau \text{NeighborPDT}_{mt-\tau} + \eta \text{Violence}_{mt-1} + \varepsilon_{mt}. \end{aligned} \quad (7)$$

³³The loss of precision occurs because we've lost nearly 15% of the sample by dropping 2012, including the part most likely to be informative about post-PDT effects.

where $Violence_{mt-1}$ is the lagged rate of violent incidents in the municipality and $NeighborPDT_{mt-1}$ is, again, the population-weighted average of PDT in the four nearest municipalities. Comparing columns 2 and 3, including a measure of violence has almost no effect on the estimated treatment effect of PDT, except for a small lead effect suggesting selection on high malnutrition when lead violence is included as a regressor. So there is no evidence of omitted variable bias due to omission of violence which predicts PDT. We conclude that the event study specification (including municipality and province-year effects and municipality-specific trends) is fairly robust to any selection on malnutrition that we can predict, though the specification in column 3 suggests that the trend reduction might be slightly overstated.

[Table X about here.]

As an alternative robustness check we exclude repeat PDT interventions from the analysis, using only first instances. As discussed above, treated villages sometimes suffer a type of recidivism and are selected for repeat treatment. Our preferred approach has been to estimate the effect all PDT treatments in a village in order to report an average treatment effect, which includes the possibility that repeat PDT indicates villages predisposed to malnutrition or that a second PDT is less (or more) intense than the first. Column 4 replicates the specification from column 1, but uses only first PDT implementations received by a village (and aggregated to the municipality level). Estimated coefficients are similar, though they show some sign of pre-treatment selection. Given that the selection criteria of villages for repeat implementation might differ from that of initial implementations, the fact that estimated effects are similar for both again demonstrates robustness, supporting a causal interpretation. The estimated effects for first PDT are a little smaller in absolute value, but statistically the same as those for any PDT. They are less precisely estimated.

As an additional robustness test, we check whether results are robust to the estimated start date of PDT, which we coded as nine months before the nominal start date, as indicated by the evidence of “clear” activity in Figure V. The lead coefficients indeed show evidence that malnutrition rates were higher beyond 1-year before PDT (the omitted category), though the t tests are statistically unimpressed (as is the joint F, which is not reported). Importantly, the lead coefficients should now be interpreted as malnutrition relative to the year before the “official” PDT start date, and, as Figure V shows, government forces were active in the village for most of that year. Thus, we feel that the positive lead coefficients in column 5 are evidence of short-term treatment effects, rather than selection bias.

In summary, PDT provides a statistically significant medium term malnutrition reduction that accumulates to about 30 percent after three years (35 log points) and is constant until at least seven years out.³⁴

³⁴A disturbing possible source of bias is that PDT increased child mortality among the malnourished, reducing the proportion of malnourished children. We checked by estimating whether PDT reduced the number of children age 5 and under (reported in the OPT data) and found no evidence of such an effect. We also checked whether PDT increased OPT coverage (the share of children weighed), which was already quite high, and again found no effects.

5.5 Scaling

The scaling of our estimated treatment effects requires some explanation. To this point, we have described PDT’s effects on child malnutrition by directly interpreting the coefficient on the lagged PDT variables. This is the estimated reduction that would occur if the share of the population receiving PDT went from 0 to 1, which is realistic for a village. However our unit of observation for regression analysis is a municipality. The average municipality has 26 villages and (as Table IV shows) the average implementation only treats 12.5% of the municipality’s population. Our preferred specification (column 2 of Table IX) implies that PDT reduces the malnutrition rate by about 30% in the long run. So our estimated treatment effect is based on linear extrapolation from a program that on average treats only about one eighth of the municipality’s population to inference about what would happen if the full population were treated.

If the effects of PDT accrued only to the population in the treated villages (the 12.5% of the municipality’s population, on average), then linear extrapolation of our estimates would imply that a village which receives PDT sees malnutrition fall by 30%. Beyond extrapolation of a possibly nonlinear effect, there is also the potential for externalities, since PDT may well affect non-recipient villages within municipalities, just as we have seen that it affects neighboring municipalities. Other villages may benefit from a nearby PDT implementation if PDT improves security, infrastructure, or services in the area –access to markets and regional clinics, for instance, could benefit children in neighboring villages. In that case, some of the malnutrition reduction we estimate is realized by non-recipient villages in the municipality and recipient villages must experience a less than 30% decline in malnutrition. On the other hand, PDT might relocate rebel activity from treatment villages to neighboring villages within the municipality, to the detriment of child nutrition, (as we have seen across municipalities in Table IX). In this case, the malnutrition reduction we estimate includes improvements for recipient villages, offset by worsening malnutrition in neighboring villages, and recipient villages actually experience more than a 30% decline in malnutrition. With those caveats, our estimated treatment effects should be interpreted as a linear approximation to the effect of a small change in treated population in a municipality, including possible within-municipality spillovers, possibly positive or negative.

5.6 Net program effect

Given the evidence of spillovers, what’s the aggregate effect of the PDT program? One way to answer that question is to test whether the malnutrition increase due to neighbors’ PDT (the θ s in equation (6)) negate the decrease due to own PDT (the β s). Specifically, equation (6) with a spline imposed has two coefficients on lagged own PDT, a linear and a long term level –call them β_{lin} and β_{lt} , and two coefficients on lagged neighbors’ PDT, $-\theta_{lin}$ and θ_{lt} . We test whether $\beta_{lin} = -\theta_{lin}$, and whether $\beta_{lt} = -\theta_{lt}$. Those test statistics are reported in the bottom four rows of Table IX. In neither case are they rejected, indicating that the own effects of PDT are statistically negated by those of neighbors’ PDT. That would imply a net treatment effect of about zero if population and malnutrition rates were balanced across treated municipalities and neighbors experiencing spillover.

Estimating the aggregate effect is a little more complicated because population and malnutrition rates differ across municipalities (they tend to be higher in treated municipalities than in their neighbors). Predicting the effects of the specification in column (2) of Table IX by converting log changes to levels and aggregating yields a small aggregate increase in malnutrition due to PDT of one hundredth of one percent.

An alternative, more flexible, approach would be to predict an aggregate effect using the split sample estimates reported in Columns 3 and 4 of Table IX, which allow the spillover effects to differ between ever-treated and never-treated municipalities. The never-treated have higher spillover effects but lower baseline malnutrition rates (as we saw in Table III). That approach yields an estimated aggregate program effect of a one tenth of one percent improvement (i.e., reduction) in child malnutrition. Our conclusion is that across specifications, spillovers either largely or completely negate the reductions in malnutrition due to the PDT program.

6 Interpretation

While PDT reduces malnutrition quite substantially, we know very little about how. One possibility is that security and removal of extortionary rebels allow markets to function, enabling private expenditure on food and medical care for children. Alternatively, given the evidence of “clear” activity and the program rhetoric about enabling government services, one might suspect that PDT is only the visible tip of a vast iceberg of associated security and governance programs, including medical care.

To check the hypothesis that improved government services are an omitted variable correlated with PDT we estimate the effects of PDT on municipal government spending on health and nutrition as well as total spending (again coding start dates with a nine month lead, to reflect “clear” activity). Figure X reports coefficients from an event study regression of the logarithm of health and nutrition spending (Panel (a)) and total spending (Panel (b)) at the municipality level on leads and lags of PDT. There is no evidence that PDT had any effect on either type of spending, or that PDT was preceded by trends in these spending variables.

[Figure X about here.]

A related interpretation would be that a PDT implementation coincides with a broader increase in federal interest in a municipality. In this case, we would expect to see PDT accompanied by an increase in federal-municipal transfers, an important revenue source for municipal governments.³⁵ Figure XI shows that, to the contrary, PDT typically occurs in the middle of a sustained decrease in federal-municipal transfers.

³⁵While transfer amounts are legally determined by a fixed formula, federal authorities have some discretion. See Troland (2014) for a detailed discussion of Federal-Municipal transfers and their effects on local public finances.

[Figure XI about here.]

In summary, the evidence weighs against inferring that PDT improves nutrition through effects on government spending. That leaves two possibilities: either enhanced security improved the quality of government service provision at constant spending rates, in ways that favored children's nutrition; or enhanced government control enabled the functioning of markets in a way that advantaged nutrition of children.

7 Conclusions

A major challenge confronting development economics is how to effectively assist people who unfortunately live in violent, unstable places. Many programs require a minimally functional government, 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.

An alternative, or precursor, to effective development assistance may be expansion of governance, hopefully enabling provision of basic services, development assistance, and the functioning of markets. Coercive expansion of governance runs the risk of violence and human rights abuses, so the net effect on welfare of residents is unclear. We take a first step towards measuring those net welfare effects, in the context of the Philippines, a middle income economy with a military subject to weak democratic oversight, experiencing a festering, low intensity insurgency.

The Philippines' Peace and Development Teams, operated by the Armed Forces of the Philippines, directly reached 12% of the population over nine years. We estimate that the average implementation reduced child malnutrition by about 30% in treated villages after three years, and persisted afterwards until at least year seven. As to mechanism, PDT is not associated with expansions in government expenditure on welfare or health, so we conclude that improvements in security, access to markets or access to existing government services must have enabled improved infant health.

How important is this finding? To put these effect sizes in context, Table XI summarizes several evaluations of child malnutrition treatments which use weight-for-age z-scores (WAZ) as an outcome.³⁶ 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.

³⁶Many 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. Also 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 protect against the deleterious effects of unexpected drought).

[Table XI about here.]

As the Table makes clear, PDT induced improvements in child nutrition comparable to the leading examples in the literature. Each intervention estimates a 35-45% reduction in the malnutrition rate, which is a little more than the estimated long term effects from our preferred specification (about 30%).

Moreover, the interventions described in Table XI would likely be infeasible in PDT villages, given their lack of governance. In the absence of a persistent security presence, it is difficult to imagine programs distributing medicine or cash to those vulnerable populations operating safely or successfully. A precursor program such as PDT can enable those interventions, and has substantial welfare benefits for local residents, even if the rebels it displaces are relatively benign Marxists conducting a fairly low casualty insurgency.

Our conclusions come with three important caveats. First, even compared to other multifaceted interventions, there are deficits in our knowledge of the actual content of the PDT intervention, and in our understanding of the mechanism by which health improvements occurred. Research on those topics, which is likely to be qualitative, would be welcome. Second, the health effects of PDT come from an intervention that connects peripheral villages to a middle income economy governed by a fairly stable multiparty democracy, making it more like Turkey or Indonesia than like a less functional setting such as the undergoverned spaces of the Middle East or Africa.

Third, and most importantly for policy, the aggregate effects of the intervention on child wellbeing were approximately zero, since the negative spillover effects of PDT on neighboring municipalities negated the positive treatment effects. Those negative spillovers suggest that not only is government good for child development, rebels are bad for it –even in a relatively low-intensity conflict like that in the Philippines. Nevertheless, though the intervention benefited children locally, it was not conducted at sufficient scale to accomplish more than displacing malnutrition associated with rebel activity, unfortunately evoking a “whack-a-mole” analogy.

We do not believe that military-centric counterinsurgency programs are a panacea for human development in poorly governed spaces. We do think that these findings invite an evidence-based discussion of the role of expansions of governance, as an extensive margin of development.

Affiliations:

Department of Economics, University of California, San Diego and NBER

Department of Economics, University of California, San Diego

Center for International Security and Cooperation and Hoover Institution, Stanford University

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 Philippine 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 latitude, 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.

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

B Additional results

Table BI reports the coefficient estimates illustrated in Figure V.

[Table BI about here.]

In Section 3.5, we presented evidence for three important considerations regarding selection in the PDT implementation decision. These were formalized in Equation (4) in a model estimated in Table VII 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 main findings of 3.5 turn out to be robust to both those aggregations – over periods and over space.

Table BII 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 year t .³⁸ As in Table VII, Table BII codes PDT as starting nine months before the nominal start date to account for the “clear” period. The municipality-level results (Table BII) are broadly consistent with the village-level results (Table VII).

[Table BII about here.]

To fully represent the data and help the reader understand precision issues in estimating long-term effects, Table BIII shows the distribution of years since most recent PDT for the municipality-year observations included in the main estimation sample.

[Table BIII about here.]

[Table BIV about here.]

³⁸Most 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.

References

- Berman, E. (2009). *Radical, Religious and Violent: The New Economics of Terrorism*. MIT Press.
- Berman, E., J. H. Felter, E. Kapstein, and E. Troland (2012). Predation, Taxation, Investment and Violence: Evidence from the Philippines. *NBER Working Paper* (18375).
- Berman, E., J. H. Felter, J. N. Shapiro, and E. Troland (2013). Modest, Secure, and Informed: Successful Development in Conflict Zones. *American Economic Review* 103(3), 512–17.
- Berman, E. and A. Matanock (2015). The Empiricists’ Insurgency. *Annual Review of Political Science* 18, 443–464.
- Besley, T. and T. Persson (2011). The Logic of Political Violence. *The Quarterly Journal of Economics* 126(3), 1411–1445.
- Brown, K. H., J. M. Person, J. Rivera, and L. H. Allen (2002). 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.
- Callen, M., S. Gulzar, A. Rezaee, and J. N. Shapiro (2015). Choosing Ungoverned Space: Pakistan’s Frontier Crimes Regulation. *UCSD mimeo*.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics* 90(3), 414–427.
- Cole, T. J., K. M. Flegal, D. Nicholls, and A. A. Jackson (2007). Body mass index cut offs to define thinness in children and adolescents: International survey. *British Medical Journal* 335(7612), 194.
- Collier, P. (2000). Rebellion as a quasi-criminal activity. *Journal of Conflict Resolution* 44(6), 839–853.
- Crost, B., J. Felter, and P. Johnston (2014). Aid Under Fire: Development Projects and Civil Conflict. *The American Economic Review*.
- Crost, B., J. H. Felter, H. Mansour, and D. I. Rees (2013). Election Fraud and Post-Election Conflict: Evidence from the Philippines. *IZA Discussion Paper*.
- Crost, B. and P. B. Johnston (2010). *Aid Under Fire: Development projects and Civil Conflict*. Harvard Kennedy School, Belfer Center for Science and International Affairs.
- Duflo, E. (2003). Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. *The World Bank Economic Review* 17(1), 1–25.
- Felter, J. H. (2005). *Taking guns to a knife fight: a case for empirical study of counterinsurgency*. Stanford University.

- Gomez, F., R. R. Galvan, S. Frenk, J. C. Munoz, R. Chávez, and J. Vázquez (1956). Mortality in second and third degree malnutrition. *Journal of Tropical Pediatrics* 2(2), 77–83.
- Graff-Zivin, J. S., H. Thirumurthy, and M. Goldstein (2006). AIDS Treatment and Intra-household Resource Allocations: Children’s Nutrition and Schooling in Kenya. *NBER Working Paper*.
- Grossman, H. I. (1999). Kleptocracy and revolutions. *Oxford Economic Papers* 51(2), 267–283.
- Growth Reference Study Group, W. M. (2006). WHO Child Growth Standards based on length/height, weight and age. *Acta Paediatrica* 450, 76.
- Heger, L. (2010). In the Cross-Hairs: Examining the Factors that make Terrorists Likely to Attack Civilians. *PhD Thesis, UC San Diego*.
- Hernandez, A. (2014). Managing Intractable Identity Conflicts—A Concert of Measures in the Southern Philippines. In A. Hernández (Ed.), *Nation-building and Identity Conflicts*, pp. 149–185. Springer.
- Hoynes, H., M. Page, and A. H. Stevens (2011). Can targeted transfers improve birth outcomes? Evidence from the introduction of the WIC program. *Journal of Public Economics* 95, 813–827.
- Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993). Earnings Losses of Displaced Workers. *The American Economic Review*, 685–709.
- Jankowska, M. M., D. Lopez-Carr, C. Funk, G. J. Husak, and Z. A. Chafe (2012). Climate change and human health: Spatial modeling of water availability, malnutrition, and livelihoods in Mali, Africa. *Applied Geography* 33, 4–15.
- Keister, J. (2011). States within States: How Rebels Rule. *PhD Thesis, UC San Diego*.
- Lavy, V., J. Strauss, D. Thomas, and P. De Vreyer (1996). Quality of health care, survival and health outcomes in Ghana. *Journal of Health Economics* 15(3), 333–357.
- Lind, T., R. Seswandhana, L.-Å. Persson, and B. Lönnerdal (2008). Iron supplementation of iron-replete Indonesian infants is associated with reduced weight-for-age. *Acta Paediatrica* 97(6), 770–775.
- Machal, G. (1996). Impact of Armed Conflict on Children. *Report of the expert of the UN Secretary-General, submitted pursuant to General Assembly resolution 48/157*.
- MacKinnon, J. G. and M. D. Webb (2014). Wild Bootstrap Inference for Wildly Different Cluster Sizes. *Queen’s Economics Department Working Paper, No. 1314*.
- Maluccio, J. and R. Flores (2005). *Impact evaluation of a conditional cash transfer program: The Nicaraguan Red de Protección Social*. International Food Policy Research Institute.

- Marshall, M. G. and B. R. Cole (2011). Global Report 2011: Conflict, Governance, and State Fragility. *Working Paper, Center for Systemic Peace*.
- Masset, E., L. Haddad, A. Cornelius, and J. Isaza-Castro (2011). *A systematic review of agricultural interventions that aim to improve nutritional status of children*.
- McChrystal, G. S. A. (2010). DoD Media Roundtable with Gen. McChrystal NATO Headquarters in Brussels. *DoD Transcript (June 10, 2010)*.
- Montinola, G. R. (2013). Change and Continuity in a Limited Access Order: The Philippines. In D. C. North, J. J. Wallis, S. B. Webb, and B. R. Weingast (Eds.), *In the Shadow of Violence: Politics, Economics, and the Problems of Development*, pp. 149–197. Cambridge.
- Morris, S. S., P. Olinto, R. Flores, E. A. Nilson, and A. C. Figueiro (2004). 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.
- Nunn, N. and N. Qian (2014). US Food Aid and Civil Conflict. *American Economic Review* 104(6), 1630–66.
- Phillips, D. R. (1986). Primary health care in the Philippines: Banking on the Barangays? *Social Science & Medicine* 23(10), 1105–1117.
- Quimpo, N. G. (2012). Mindanao, Southern Philippines: The Pitfalls of Working for Peace in a Time of Political Decay. In M. A. Miller (Ed.), *Autonomy and Armed Separatism in South and Southeast Asia*. Singapore: Institute of Southeast Asian Studies.
- Ruel, M. T., P. Menon, J.-P. Habicht, C. Loechl, G. Bergeron, G. Pelto, M. Arimond, J. Maluccio, L. Michaud, and B. Hankebo (2008). 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.
- Sales, P. M. (2009). State terror in the Philippines: the Alston Report, human rights and counter-insurgency under the Arroyo administration. *Contemporary Politics* 15(3), 321–336.
- Sandler, D. H. and R. Sandler (2013). Multiple Event Studies in Public Finance and Labor Economics: A Simulation Study with Applications. *Working Paper*.
- Schiavo-Campo, S. and M. P. Judd (2005). *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.
- Seoane, N. and M. C. Latham (1971). Nutritional anthropometry in the identification of malnutrition in childhood. *Journal of Tropical Pediatrics* 17(3), 98–104.
- Singh, A., A. Park, and S. Dercon (2013). School meals as a safety net: An evaluation of the Midday Meal scheme in India. *Young Lives Working Paper* 75.

Troland, E. (2014). Can Fiscal Transfers Increase Local Revenue Collection? Evidence From The Philippines. *Working Paper*.

Vanden Eynde, O. (2015). Mining Royalties and Incentives for Security Operations: Evidence from India's Red Corridor. *Paris School of Economics mimeo*.

Table I: Conflict and instability among top foreign aid recipients

Country	Average from 2003-2012				GDP per capita	Rank of Pol. Instability	Conflict Years 1988-2012
	Annual ODA per capita						
	World Bank	USA	UK	Total			
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 II: Scale of PDT

Year	Cumulative				Annual			
	Villages		Municipalities		Villages		Municipalities	
	<i>N</i> of 42,013	% of Pop.	<i>N</i> of 1,648	% of Pop.	<i>N</i> of 42,013	% of Pop.	<i>N</i> 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 III: 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	11383	9.327	6.589	.057	64.989	4.304	12.805
Severe malnutrition rate	11383	1.317	1.470	0	36.292	.491	1.638
Coverage	10715	89.771	15.047	6.38	877.45	82.08	100

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 1,513 municipalities with at least three malnutrition measurements from 2005-2012, in which that measurement meets some minimal data quality standards, and in which other key variables are available.

Table IV: Summary statistics

X	All municipalities		PDT municipalities		Non-PDT municipalities	
	Mean (s.d.)	R^2 from mun. fixed effects	$P(X = 0)$	Mean (s.d.) given $X > 0$	$P(X = 0)$	Mean (s.d.) given $X > 0$
Malnutrition rate	9.33 (6.59)	.786	0	10.06 (6.86)	0	8.57 (6.21)
Incidents	.026 (.073)	.542	.30	.070 (.111)	.59	.044 (.076)
PDT pop. frac.	.012 (.064)	.149	.81	.125 (.169)	–	–
Neighboring PDT	.016 (.051)	.223	.64	.077 (.093)	.90	.041 (.059)
Years since PDT	1.07 (1.72)	.645	.21	2.64 (1.78)	–	–
N	11,383		5,374		6,009	

“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. “Incidents” is the rate of incidents (mostly violent) reported by the Armed Forces of the Philippines, per 1,000 population. It is unavailable in 2012, and therefore based on a smaller sample. “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 V: Evidence of geographical clustering of PDT implementation

$DV : 100 \times PDTStart$	(1)	(2)
$NeighborPDT_{t-1,t-3}$	3.295*** (0.386)	1.944*** (0.419)
$NeighborPDT_{t-4,t-6}$	1.701*** (0.308)	0.353 (0.271)
Constant	0.128*** (0.019)	
N	4.25e+06	4.25e+06
R^2	0.002	0.015
Barangay effect	No	Yes
Month effect	Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$. Dependent variable is 100 times a binary indicator that PDT started in a given Barangay (village) during a given month. $NeighborPDT_{t-1,t-3}$ is the share of the four nearest Barangays where PDT began in the previous one to three months. $NeighborPDT_{t-4,t-6}$ is the share of the four nearest Barangays where PDT began in the previous four to six months, but not the last three months. Sample includes a balanced panel of 42,038 Barangays from 2002-2010. Because of lags, the first six months of 2002 are excluded. Standard errors clustered at the province ($n = 78$) level in parentheses.

Table VI: 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 VII: Predicting PDT at the village-month level

DV: $100 \times PDT_t$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>NeighborPDT</i> _{<i>t-1,t-3</i>}	0.783*** (0.097)					0.744*** (0.101)	0.424*** (0.101)
<i>NeighborPDT</i> _{<i>t-4,t-6</i>}	0.260*** (0.066)					0.219*** (0.066)	-0.090 (0.067)
<i>IncidentRate</i> _{<i>t-1,t-3</i>}		0.954*** (0.187)				0.551*** (0.126)	0.271*** (0.096)
<i>IncidentRate</i> _{<i>t-4,t-6</i>}		0.648*** (0.148)				0.234** (0.095)	-0.017 (0.081)
<i>MunIncidentRate</i> _{<i>t-1,t-3</i>}			3.768*** (1.356)			2.474** (1.144)	0.739 (1.021)
<i>MunIncidentRate</i> _{<i>t-4,t-6</i>}			3.319** (1.312)			2.603** (1.241)	1.353 (1.036)
<i>AnyPastPDT</i>				0.337*** (0.082)	0.209** (0.103)	0.127 (0.093)	-2.261*** (0.147)
<i>MonthsSincePDT</i>				0.012** (0.005)	0.024*** (0.008)	0.024*** (0.007)	0.013** (0.006)
<i>MonthsSincePDT</i> ² /10					-0.002 (0.001)	-0.002 (0.001)	0.000 (0.001)
<i>Constant</i>	0.134*** (0.020)	0.108*** (0.017)	0.099*** (0.016)	0.085*** (0.014)	0.085*** (0.014)	0.095*** (0.016)	
N	3951572	3995652	3995658	3995658	3995658	3951572	3951572
<i>R</i> ²	0.002	0.001	0.001	0.002	0.002	0.004	0.021
Barangay effect	No	No	No	No	No	No	Yes
Month effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the province ($n = 78$) level in parentheses.

Table VIII: Estimated effects of PDT

DV: $\ln(\text{MalnRate})$	(1)	(2)	(3)	(4)	(5)
PDT_{t+5}	0.156 (0.416)	0.156 (0.378)	0.156 (0.416)	0.174 (0.695)	0.166 (0.594)
PDT_{t+4}	0.018 (0.186)	0.018 (0.201)	0.018 (0.174)		
PDT_{t+3}	0.082 (0.090)	0.082 (0.094)	0.082 (0.106)		
PDT_{t+2}	0.057 (0.045)	0.057 (0.045)	0.057 (0.064)		
$\sum_{\tau=2}^4 PDT_{t+\tau}$				0.055 (0.059)	0.047 (0.090)
PDT_t	-0.066 (0.079)	-0.066 (0.086)	-0.066 (0.064)	-0.061 (0.097)	
PDT_{t-1}	-0.181** (0.087)	-0.181* (0.098)	-0.181* (0.094)	-0.170* (0.091)	
PDT_{t-2}	-0.259*** (0.098)	-0.259** (0.113)	-0.259** (0.122)	-0.242** (0.103)	
$\sum_{\tau=0}^2 (\tau + 1)PDT_{t-\tau}$					-0.084** (0.035)
$\sum_{\tau=3}^4 PDT_{t-\tau}$	-0.303** (0.115)	-0.303** (0.127)	-0.303* (0.170)		
$\sum_{\tau=3}^{11} PDT_{t-\tau}$				-0.278** (0.119)	-0.286*** (0.110)
$\sum_{\tau=5}^6 PDT_{t-\tau}$	-0.322** (0.138)	-0.322** (0.147)	-0.322 (0.200)		
$\sum_{\tau=7}^{11} PDT_{t-\tau}$	-0.336* (0.183)	-0.336* (0.190)	-0.336 (0.253)		
N	11383	11383	11383	11383	11383
R^2	0.943	0.943	0.943	0.943	0.943
F (pre)	0.750	0.750	0.340	0.759	0.366
p (analytic)	0.561	0.561	0.851	0.472	0.695
p (bootstrap)		0.707	0.846		
F (post)	2.228	2.228	1.035	3.611	6.590
p (analytic)	0.060	0.060	0.396	0.017	0.002
p (bootstrap)		0.140	0.346	0.026	0.010
Clustering	Prov. (78)	Prov	Mun. (1513)	Prov	Prov
Standard errors	Analytic	Bootstrap	Bootstrap	Bootstrap	Bootstrap

* $p < .10$, ** $p < .05$, *** $p < .01$. Effects of PDT estimated at the municipality-year level. PDT is the fraction of the municipality's population treated. To improve precision we impose equality of coefficients (e.g., the 5-year effect is the same as the 6-year effect) in Columns 1-3 and linear splines in columns 4 and 5. The "pre" F-statistic tests the null that coefficients on all leads are jointly zero. The "post" tests the null that the contemporaneous and all lagged coefficients are jointly zero. Standard errors, in parentheses, clustered by province ($n=78$), except in Column 3. Wild clustered bootstrapped standard errors use 999 replications.

Table IX: Spillover effects

DV: $\ln(MalnRate)$	(1)	(2)	(3)	(4)
PDT_{t+5}	0.166 (0.560)	0.218 (0.607)		0.156 (0.600)
$\sum_{\tau=2}^4 PDT_{t+\tau}$	0.047 (0.087)	0.046 (0.079)		0.034 (0.088)
$\sum_{\tau=0}^2 (\tau + 1)PDT_{t-\tau}$	-0.084** (0.034)	-0.100** (0.039)		-0.080** (0.038)
$\sum_{\tau=3}^{11} PDT_{t-\tau}$	-0.286** (0.114)	-0.355*** (0.129)		-0.295** (0.125)
$\sum_{\tau=0}^2 (\tau + 1)NeighPDT_{t-\tau}$		0.071 (0.050)	0.403*** (0.149)	0.022 (0.052)
$\sum_{\tau=3}^{11} NeighPDT_{t-\tau}$		0.344** (0.172)	0.920** (0.466)	0.294 (0.210)
N	11383	11383	6001	5350
R^2	0.943	0.943	0.937	0.950
F (own)	6.590	7.174		3.316
p (analytic)	0.002	0.001		0.043
p (bootstrap)	0.013	0.008		0.078
F (neigh.)		2.309	6.378	1.733
p (analytic)		0.106	0.003	0.185
p (bootstrap)		0.114	0.010	0.194
F ($\beta_{lin} + \theta_{lin} = 0$)		0.654		1.211
p (analytic)		0.421		0.275
F ($\beta_{long} + \theta_{long} = 0$)		0.006		0.000
p (analytic)		0.936		0.995
Sample	All	All	non-PDT mun.	PDT mun.
N_{prov}	78	78	75	64

* $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors are based on wild cluster bootstrap of Cameron et al. (2008), clustered on province. All bootstrapped results are based on 999 replications. All columns include municipality fixed effects and trends, and province \times year effects. See note to table VIII for details. The number of municipalities in the split sample (columns 3 and 4) does not add to the total because province \times year fixed effects absorb some observations.

Table X: Estimated Effects of PDT: Robustness checks

DV: $\ln(\text{MalnRate})$	(1)	(2)	(3)	(4)	(5)
PDT_{t+5}	0.218 (0.666)	0.291 (0.632)	0.115 (0.632)	-0.294 (0.376)	0.236 (0.267)
$\sum_{\tau=2}^4 PDT_{t+\tau}$	0.046 (0.077)	0.074 (0.082)	0.088 (0.080)	0.034 (0.108)	0.104 (0.097)
$\sum_{\tau=0}^2 (\tau + 1)PDT_{t-\tau}$	-0.100*** (0.036)	-0.118*** (0.040)	-0.117*** (0.039)	-0.089 (0.059)	-0.075** (0.032)
$\sum_{\tau=3}^{11} PDT_{t-\tau}$	-0.355*** (0.134)	-0.424*** (0.115)	-0.413*** (0.133)	-0.320* (0.181)	-0.174* (0.103)
$\sum_{\tau=0}^2 (\tau + 1)NeighPDT_{t-\tau}$	0.071 (0.050)	0.076 (0.050)	0.094* (0.050)	0.016 (0.078)	0.053 (0.043)
$\sum_{\tau=3}^{11} NeighPDT_{t-\tau}$	0.344** (0.166)	0.341* (0.174)	0.405** (0.177)	0.142 (0.218)	0.094 (0.168)
Lagged violence rate			0.046 (0.043)		
N	11383	9933	9912	11383	11383
R^2	0.943	0.948	0.948	0.943	0.943
F (own)	7.174	6.643	5.814	2.633	5.344
p (analytic)	0.001	0.002	0.004	0.078	0.007
p (bootstrap)	.009	.006	.004	.097	.024
F (neigh.)	2.309	1.299	2.034	0.868	1.173
p (analytic)	0.106	0.279	0.138	0.424	0.315
p (bootstrap)	.099	.120	.171	.425	.264
F ($\beta_{lin} + \theta_{lin} = 0$)	0.654	0.428	0.124	2.065	0.262
p (analytic)	0.421	0.515	0.726	0.155	0.610
F ($\beta_{long} + \theta_{long} = 0$)	0.006	0.130	0.001	1.166	0.203
p (analytic)	0.936	0.719	0.973	0.284	0.653
PDT	Any	Any	Any	First	Any
Start date	9 mon. lead	9 mon. lead	9 mon. lead	9 mon. lead	Actual
Years	2005-2012	2005-2011	2005-2011	2005-2012	2005- 2012

* $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors are based on wild cluster bootstrap of Cameron et al. (2008). All bootstrapped results are based on 999 replications. See note to Table VIII for details. Lagged violence rate is incidents per 1,000 population during the year before malnutrition is measured. Incident data for 2011 is unavailable so including lagged incidents requires restricting to 2005-2011 (columns 2 and 3). Columns (1)-(3) and (5) use all of a Barangay's PDT implementations, while column (4) uses only the first. Column (5) uses nominal program start date.

Table XI: 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 desert 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 BI: Government-initiated violence: Event study coefficients

Distant pre-PDT		Shortly before PDT		PDT and immediate post		Distant post-PDT	
PDT_{t+18}	0.00009 (0.00034)	PDT_{t+9}	0.00147** (0.00057)	PDT_t	0.00429*** (0.0008)	PDT_{t-8}	0.00021 (0.00033)
PDT_{t+17}	0.00031 (0.00038)	PDT_{t+8}	0.00085** (0.00034)	PDT_{t-1}	0.00618*** (0.00105)	PDT_{t-9}	0.00011 (0.00027)
PDT_{t+16}	0.00013 (0.00033)	PDT_{t+7}	0.00069** (0.00032)	PDT_{t-2}	0.00346*** (0.00079)	PDT_{t-10}	0.00005 (0.0002)
PDT_{t+15}	0.00027 (0.00037)	PDT_{t+6}	0.00106*** (0.00033)	PDT_{t-3}	0.00139*** (0.00037)	PDT_{t-11}	-0.00001 (0.00018)
PDT_{t+14}	-0.00012 (0.00032)	PDT_{t+5}	0.00111*** (0.00031)	PDT_{t-4}	0.00146*** (0.0004)	PDT_{t-12}	-0.0002 (0.00019)
PDT_{t+13}	0.00001 (0.00023)	PDT_{t+4}	0.00078** (0.00037)	PDT_{t-5}	0.0009*** (0.00032)	PDT_{t-13}	0.00067 (0.00034)
PDT_{t+12}	-0.00025 (0.00031)	PDT_{t+3}	0.00189*** (0.00052)	PDT_{t-6}	0.00043 (0.0003)	PDT_{t-14}	-0.00006 (0.0003)
PDT_{t+11}	0.00016 (0.0003)	PDT_{t+2}	0.00151*** (0.00041)	PDT_{t-7}	0.00075** (0.00032)	PDT_{t-15}	0.00014 (0.00023)
PDT_{t+10}	0.00014 (0.00028)	PDT_{t+1}	0.00367*** (0.00056)			PDT_{t-16}	-0.00001 (0.00021)
						PDT_{t-17}	0.00008 (0.00034)
						PDT_{t-18}	0.00015 (0.00028)

* $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors, clustered at the province level, shown in parentheses. The Table reports coefficients corresponding to monthly leads and lags of PDT in equation (2), which includes month and village fixed effects. See Figure V to see the coefficients plotted over time.

Table BII: Predicting PDT at the municipality-year level

DV: $PDTFraction_t$	(1)	(2)	(3)	(4)	(5)	(6)
$NeighPDT_{t-1}$	0.125*** (0.024)				0.033 (0.026)	0.003 (0.028)
$IncidentRate_{t-1}$		0.076*** (0.027)			0.046** (0.018)	0.010 (0.014)
$AnyPastPDT$			0.022*** (0.006)	0.018** (0.008)	0.014* (0.008)	0.010 (0.009)
$FracPastPDT$			0.041** (0.016)	0.041** (0.016)	0.035** (0.017)	-0.268*** (0.017)
$YearsSincePDT$			-0.001 (0.001)	0.003 (0.005)	0.005 (0.004)	0.007* (0.004)
$YearsSincePDT^2$				-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.000)
Constant	0.015*** (0.002)	0.015*** (0.003)	0.009*** (0.002)	0.009*** (0.002)	0.008*** (0.002)	
N	13184	13184	13184	13184	13184	13184
R^2	0.017	0.016	0.039	0.039	0.042	0.273
Municipality effect	No	No	No	No	No	Yes
Year effect	Yes	Yes	Yes	Yes	Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the province ($n = 78$) level in parentheses. The table uses a PDT start date that is nine months ahead of the nominal, official start date to account for the pre-implementation “clear” stage documented in Figure V and Table BI.

Table BIII: 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 three months, we do not consider this a major problem.

Table BIV: Appendix table with FE's and unconstrained coefficients

DV: $\ln(MalnRate)$	(1)	(2)	(3)	(4)
PDT_{t+5}	0.036 (0.877)	0.175 (0.585)	0.156 (0.416)	0.209 (0.447)
PDT_{t+4}	-0.053 (0.200)	-0.159 (0.160)	0.018 (0.186)	0.063 (0.215)
PDT_{t+3}	-0.046 (0.084)	-0.019 (0.073)	0.082 (0.090)	0.110 (0.124)
PDT_{t+2}	-0.004 (0.046)	0.008 (0.055)	0.057 (0.045)	0.066 (0.063)
PDT_t	0.025 (0.065)	-0.039 (0.081)	-0.066 (0.079)	-0.082 (0.091)
PDT_{t-1}	-0.030 (0.067)	-0.122 (0.096)	-0.181** (0.087)	-0.211* (0.110)
PDT_{t-2}	-0.026 (0.084)	-0.180 (0.112)	-0.259*** (0.098)	-0.304* (0.154)
PDT_{t-3}				-0.373* (0.194)
PDT_{t-4}				-0.374 (0.249)
$\sum_{\tau=3}^4 PDT_{t-\tau}$	0.035 (0.087)	-0.206 (0.153)	-0.303** (0.115)	
PDT_{t-5}				-0.394 (0.268)
PDT_{t-6}				-0.476 (0.333)
$\sum_{\tau=5}^6 PDT_{t-\tau}$	0.132 (0.119)	-0.190 (0.189)	-0.322** (0.138)	
$\sum_{\tau=7}^{\infty} PDT_{t-\tau}$	0.072 (0.149)	-0.215 (0.221)	-0.336* (0.183)	-0.483 (0.374)
N	11409	11409	11383	11383
R^2	0.879	0.936	0.943	0.943
F (pre)	0.128	0.428	0.750	0.452
p (analytic)	0.972	0.788	0.561	0.771
F (post)	2.672	1.060	2.228	2.285
p (analytic)	0.028	0.389	0.060	0.036
Mun. FE	Yes	Yes	Yes	Yes
Mun. trends	No	Yes	Yes	Yes
Prov. \times year FE	No	No	Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$. Observations are municipality-years. PDT is the fraction of the municipality's population living in a village receiving PDT. The "pre" F-statistic corresponds to the null hypothesis that coefficients on all leads are jointly zero. The "post" F-statistic corresponds to the null hypothesis that coefficients on all lags are jointly zero. Observations are dropped between columns (2) and (3), when province \times year effects are added because there is only one municipality in the province-year. Standard errors are clustered by province ($n=78$). Analytic rejection probabilities assume an asymptotic F distribution.

Figure I: Violent incident and malnutrition rates by population density quintile

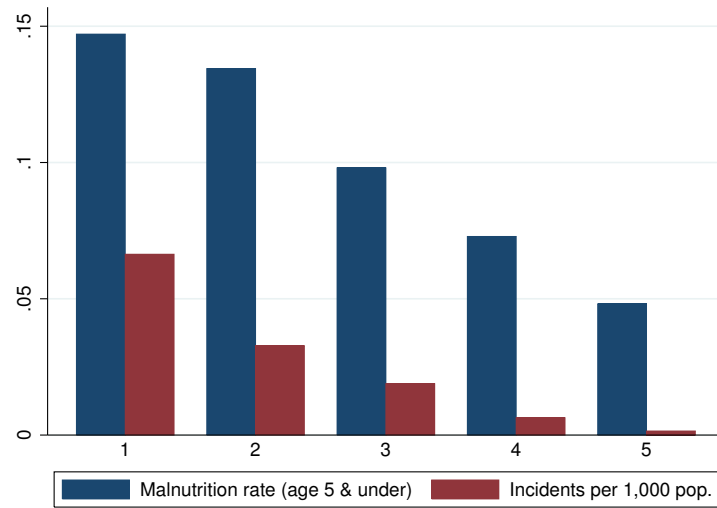


Figure displays population-weighted average violent incident and malnutrition rates for 1,513 municipalities by population density quintile (1 = least dense, 5 = most dense). Quintiles are constructed to include equal population shares. Results are similar based on violent incident counts, rather than rates.

Figure II: Villages receiving PDT





Figure displays locations of all 6,819 PDT implementations. See Section 3 for further description of the data.

Figure III: Malnutrition rates over time



Source: Authors' calculations based on Operation Timbang (OPT) data. Note: Malnutrition is based on weight-for-age Z-scores (WAZ) for children age 0-5. All three rates based on population-weighted estimates across municipalities. See Section 3 for further discussion of the data.

Figure IV: Probability repeat PDT is observed



Figure displays the probability that a Barangay receiving its first PDT in year x receives another implementation by the end of the program (in 2010).

Figure V: Village-level violence around PDT




Figure plots coefficients, and 95% confidence intervals, corresponding to leads and lags of PDT in equation (2), which includes month and village fixed effects. Month 0 corresponds to the month of PDT implementation. Mean violence for the omitted category (months more than a year and a half before or after PDT) is .0024 incidents per month per 1,000 population.

Figure VI: Estimated effects of PDT






Figure plots coefficients, and 95% confidence intervals, corresponding to a PDT implementation at time 0. Coefficients are based on Column (4) of Appendix table BIV. Log malnutrition rates are normalized to zero during the year preceding treatment. An F-test fails to reject the null that the pre-PDT coefficients are jointly zero ($F = .452, p = .771$), but does reject the null that the post-PDT coefficients are jointly zero ($F = 2.285, p = .036$).

Figure VII: Bootstrapped F-statistic



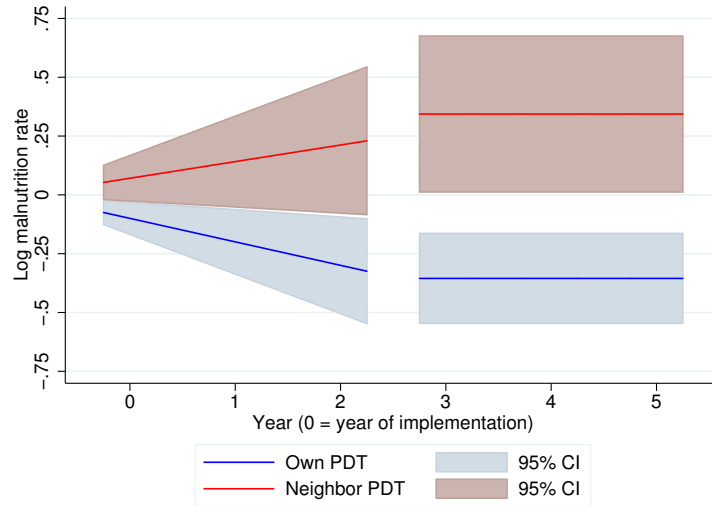
Bootstrapped F-distribution is for the test of lagged coefficients, Table VIII Column 2, 78 (province) clusters. The wild clustered bootstrap of Cameron et al. (2008) is implemented in Stata with 999 replications. The vertical line marks $F=2.23$, which has an asymptotically correct p of 0.06 but sits to the left of 14% of the empirical density.

Figure VIII: Spline results alongside unconstrained results




Green dots and vertical bars correspond to coefficients (and 95% confidence intervals) for the unconstrained event study model displayed in Figure VI and Column (4) of Appendix table BIV. Blue lines and shaded areas correspond to estimated effects (and 95% confidence intervals) for the spline model in Column (5) of Table VIII. Log malnutrition rate is normalized to zero for the year preceding PDT implementation.

Figure IX: Direct and spillover effects of PDT on malnutrition



Lower (blue) lines and shaded areas correspond to estimated effects (and 95% confidence intervals – asymptotic rather than bootstrapped) of PDT in the same municipality. Upper (red) lines illustrate estimated effects of PDT averaged over four nearest neighboring municipalities. Estimates are from Column (2) of Table IX. Log malnutrition rate is normalized to zero for the year preceding PDT implementation.

Figure X: Spending effects of PDT




(a) Health and Nutrition Spending

(b) Total Municipal Spending

The figure plots coefficients, and 95% confidence intervals, corresponding to the effects of an average PDT implemented at time 0. Underlying regressions, available upon request, include municipality and province-year effects and municipality-specific trends. Log spending rates are normalized to zero during the year before the intervention.

Figure XI: Federal Transfers and PDT



The figure plots coefficients, and 95% confidence intervals, corresponding to the effects of an average PDT implemented at time 0. Underlying regressions, available upon request, include municipality and province-year effects and municipality-specific trends. Log spending rates are normalized to zero during the year before the intervention.