

**Pitfalls of Participatory Programs:
Evidence from a randomized evaluation in education in India**

Abhijit V. Banerjee (MIT)

Rukmini Banerji (Pratham)

Esther Duflo (MIT)

Rachel Glennerster (MIT)

Stuti Khemani (The World Bank)

Abstract

Participation of beneficiaries in the monitoring of public services is increasingly seen as a key to improving their efficiency. In India, the current government flagship program on universal primary education organizes both locally elected leaders and parents of children enrolled in public schools into committees and gives these groups powers over resource allocation, and monitoring and management of school performance. However, in a baseline survey we found that people were not aware of the existence of these committees and their potential for improving education. This paper evaluates three different interventions to encourage beneficiaries' participation through these committees: providing information, training community members in a new testing tool, and training and organizing volunteers to hold remedial reading camps for illiterate children. We find that these interventions had no impact on community involvement in public schools, and no impact on teacher effort or learning outcomes in those schools. However, we do find that the intervention that trained volunteers to teach children to read had a large impact on activity *outside* public schools—local youths volunteered to be trained to teach, and children who attended these camps substantially improved their reading skills. These results suggest that citizens face substantial constraints in participating to improve the public education system, even when they care about education and are willing to do something to improve it.

*This study was undertaken as a collaboration among Pratham Mumbai Education Initiative (Pratham), MIT Abdul Lateef Jameel Poverty Action Lab, and the World Bank, and thus involves the participation of many people from these institutions. We would especially like to thank Dan Keniston, Sanjib Kundu, and Mukesh Prajapati for implementation, and Marc Shotland for managing the project. We thank Dhruva Kothari and Jacqueline Merriam for excellent research assistance. We thank Amit Dar, Jishnu Das, Dipak Dasgupta, Shanta Devarajan, Venita Kaul, Reema Nayar, Priyanka Pandey, Lant Pritchett, Michelle Riboud, V. J. Ravishankar, Shekhar Shah, Venki Sundararaman, and Salman Zaidi of the World Bank for numerous discussions and comments on the project. We thank the Trust Fund for Environmentally and Socially Sustainable Development (TFESSD) and the DIME initiative at the World Bank for the funding to undertake the surveys and analysis. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors, and do not necessarily represent the views of the World Bank, its Executive Directors, or the countries they represent.

1. Introduction

The deplorable state of publicly provided social services in many developing countries has attracted considerable attention in recent years, from academics and policy makers alike. The 2004 World Development Report "Making Services Work for Poor People" details the dismal quality of education and health services offered to the poor in developing countries, and concludes that "social services fail the poor" (World Bank, 2004). Both sectors are plagued by high provider absenteeism, little on-the-job effort from those who do come to work, and overall poor performance. For example, teachers in primary schools and medical staff at primary health centers in India have absence rates of 25 percent and 40 percent, respectively (Chaudhury et al., 2006).¹ In India's public education system, these problems are manifest in the failure of the public schools to impart even the most elementary skills. A 2005 nationwide rural survey on educational attainment of 7-14 year olds found that, even though most children are enrolled in school, 35 percent of primary school age children could not read a simple paragraph, and 41 percent could not do simple subtraction, competencies that should, according to government guidelines, be achieved by grade 2 (Pratham, 2005).

Inadequate funding does not appear to be the only reason for the systems' poor performance. In fact, in the case of education, rigorous, randomized evaluations have found little evidence that more resources on their own, with no changes to the way education is delivered, can improve test scores.² In contrast, studies have found improvements in outcomes when modest incentives have been given to teachers (see e.g., Duflo, Hanna and Ryan, 2007, and Muralidharan and Sundararaman, 2007³). However, in both these cases the incentives were implemented by the non-governmental organization (NGO) that ran the program. When government bureaucrats have to implement them, the incentives were ineffectual. In Kenya, for example, head teachers were given grants to reward teachers who attend regularly. Even though the head teacher could use the money to get other things for the school if they did not use it to reward the teachers, they almost universally chose to give the money to the teachers, irrespective of whether the teacher had done anything to deserve it, with the result that the incentive had no impact (Kremer and Chen, 2002). In India a reform that was meant to link government nurses' pay to their attendance was initially very effective, but failed to have any impact after the local bureaucracy started providing official excuses for most of the nurses' absences (Banerjee, Duflo, and Glennerster, 2008).

The poor incentives of bureaucrats and public providers to deliver quality services contrasts with evidence of strong incentives of governments to deliver quantity, especially in education, by building schools and recruiting teachers (Glewwe and Kremer, 2006, Keefer and Khemani, 2004, and Banerjee and Somanathan, 2007). Keefer and Khemani (2005), building on a large political economy literature⁴, argue that this can be explained by political incentives of governments that are skewed toward the provision of private and verifiable benefits to some citizens, at the expense of quality improvements that are more diffuse and harder to verify. Providing (secure) jobs through teacher recruitment, or scholarships and other transfers, is more effective in winning elections than getting teachers to teach.

¹ For health, Das and Hammer (2005) show that the average time spent by a public doctor to examine a patients is 2 minutes, and that the average visit does not include any physical exam. In education, Chaudhury, et al. (2006) find that half of the teachers present in school in India were not teaching at the time of the investigator's visit.

² For example, Glewwe, Kremer, and Moulin (2002) find no impact from textbooks, Glewwe, Kremer, Moulin and Zitzewitz (2004) find no impact from flip charts, and Banerjee, Kremer and Jacob, (2004) find no impact from additional teachers in India. For discussion and more references, see Glewwe and Kremer (2006).

³ See, on the other hand, Glewwe, Ilias, and Kremer (2003) for a cautionary tale on multitasking by teachers even when incentives are properly implemented.

⁴ See, e.g. Acemoglu and Robinson..(2001); Banerjee and Iyer (2005), Acemoglu and Robinson (2006); Rajan (2007); Hoff and Stiglitz (forthcoming).

Such political economy arguments suggest that government bureaucracies may be ill equipped to improve the quality of public services. This has led development practitioners to believe that the involvement of beneficiaries is essential to make services work for poor people. Development projects funded by international organizations are increasingly required to include “beneficiary participation” components, such as the constitution of users’ committees, parents-teachers associations, etc. or, where such institutions already exists, interventions aimed at mobilizing people to make them more effective. The World Bank’s World Development Report “Making Services work for Poor People” describes a range of alternative institutional designs that enable beneficiaries to exercise better control over the quality of services that they are receiving (World Bank, 2004).

Despite this enthusiasm, and some encouraging experiments we describe below, whether external interventions actually affect the beneficiaries, and whether this participation can in turn improve public services, remain largely open empirical questions. It is challenging to answer these questions definitively because there is no single leading candidate for what encourages beneficiary participation and/or makes it more effective and a wide range of possible models.

At a very minimum, there are at least two dimensions along which interventions designed to promote participation can (and do) vary.

- First, community members might either be able to:
 - intervene directly to improve service delivery (by contributing additional inputs, by punishing or rewarding the agents, etc.) or
 - they may only be able to intervene indirectly (say by contributing to a fund that is administered by some bureaucrat or complaining to the agent or the agent’s supervisor).⁵
- Second, community members might be able to exercise control (direct or otherwise)
 - individually or as a part of a small group (for example they may ask to be paid for the work they did on the local road project), or
 - it may require joint action by some larger group.⁶

Both the direct and indirect dimensions allow for continuous variation, creating a continuous plane of possible models of popular participation.

The different models have their advantages and disadvantages: for example, direct control obviously makes it easier to exercise control but opens the door for capture of the entire process by a locally dominant group. This is what happened in the sugar cooperatives in Maharashtra, India (Banerjee et al., 2001). Indirect controls, such as complaint systems, have the advantage that it is possible to bypass the local elites, but at the cost of potential slippage (the complaint may get lost or ignored).⁷ Small groups require less coordination to exercise control (and hence are less subject to coordination failures), but large groups make sense when the good is not excludable, and complaints are less likely to be ignored when they come from a large enough group.

This raises the possibility that promoting some forms of beneficiary control might be more rewarding than focusing on others. At the same time, given a specific model of beneficiary control, there are many different ways to

⁵ The World Development Report (2004) calls these respectively the “short route” and the “long route” to control.

⁶ There is actually a natural third dimension: there may be specific, pre-specified mechanisms for the exercise of control, or the precise mechanics may be left to the community to determine. This is potentially very important, but impinges less on the issues emphasized in this paper.

⁷ See Bardhan and Mookherjee (2006) for an illuminating discussion of the trade-offs here.

encourage the beneficiaries to get more involved in exercising control. At one extreme are interventions where people are simply told their rights within the system and then left alone to figure out how they want to claim those rights. At the other extreme, beneficiaries may be organized and trained to participate in demonstrations, or other forms of public action, including community interventions that substitute for the publicly provided services.

This paper reports on a randomized experiment in the state of Uttar Pradesh (UP) in India that compares the effects of three different interventions intended to promote the greater exercise of beneficiary control on two very different forms of beneficiary control within the same setting. One of the forms of beneficiary control that we study involves the exercise of indirect control by a large group; the other is all about direct control by individuals. The treatments vary between the relatively hands-off (tell people what their rights are with respect to their children's education and alert them to their responsibilities) and the highly interventionist (training village volunteers in how to identify the children in the village who are lagging behind, as well as in techniques for teaching these children more effectively).

The possibility of affecting large group indirect control comes from the structure of the Indian federal government's flagship current program on elementary education, the Sarva Shiksha Abhiyan (SSA). This program, under which the federal government will contribute 65 percent of the educational budget in UP in 2007-08 (22,374 million rupees or 559 million dollars), gives a prominent role to the Village Education Committees (VECs), though VECs had existed prior to the SSA. In Uttar Pradesh, the VEC is a body consisting of three parents, the head-teacher of the village school and the head of the village government, which is supposed to be the key intermediary between the village and the district educational authorities and is supposed to exercise both indirect control (by monitoring the performance of the schools, reporting problems to higher authorities, and requesting additional resources), as well as some direct control (by deciding whether the contracts of the existing community-based teachers should be renewed, recruiting new hires, and allocating the additional resources for improvement that the school receives under the SSA).⁸ While the VEC itself is a small group, we argue that the three parent members of the VEC could not have real control without the active backing of a large number of other parents and that this larger group of parents also have the ability to exercise control without going through the VEC. In this sense the model of beneficiary control in education in UP is one of direct and indirect control exercised by a large group.

The opportunity to do these experiments arose because by 2005, more than four years after SSA was launched, the prima facie evidence suggested that at least in Jaunpur district of the state of UP, beneficiary control was not being exercised particularly effectively. A survey of children, parents, and teachers in 280 villages in that district found that while most villages did have a VEC, very few parents knew of its existence, sometimes even when they were supposed to be members of it. VEC members were also unaware of even the most important responsibilities that they have been assigned under the SSA: hiring of additional teachers, allocation of school resources, and monitoring of performance. At the same time the state of education in these villages bordered on the disastrous: 15 percent of children age 7 to 14 could not recognize a letter; only 39 percent could read and understand a simple story (of grade 1 level); 38 percent could not recognize numbers. Yet parents, teachers, and VEC members were often unaware of the scale of the problem and tended to overestimate what children in the village really knew (Banerjee et al. 2006).

⁸ It is important to emphasize that even though SSA is national program, its implementation, as well as the structure and role of the VECs varies a lot across states in India and UP is not necessarily representative.

Given that the institutional frame for large group control (the VEC etc.) were extant everywhere, but almost entirely non-functional, one goal of the interventions was to make large group control work better by mobilizing public sentiments and providing better information. A second goal was to sensitize people about the importance of education for their children and the state of education in the village, with the expectation that this would encourage them to try to do something about it, either in their private capacity or as a part of a small group--in other words a small-group direct control strategy. Some training that would make them more effective in this capacity was part of the intervention.

The three interventions were designed and implemented by Pratham, India's most prominent educational NGO. The interventions that Pratham implemented were designed after several months of piloting in the area. There was a conscious attempt to be sensitive to the specificities of the local environment (in particular the constraints and opportunities resulting from the existing policies of the state government), to follow the available best practice guidelines for community engagement (World Bank, 1996, 2001), and to streamline the design for replicability and external validity.

One of the assumptions underlying the design of the intervention was that lack of relevant knowledge was a factor holding back participatory action. It was clear from the baseline survey that community members did not know what they were entitled to, what they were actually getting, and how they could put pressure on the providers—all of which are critical for participation to be effective (Jenkins and Goetz, 1999; Goetz and Jenkins, 2001; Paul, 1987). A second element in the Pratham's program design was the need (if participation was to be a reality) to involve more people with the state of education in the village, particularly on learning outcomes, and to turn that interest into some form of coordinated action.

All three interventions shared a basic structure: Pratham activists spent several days in a village facilitating small group discussions in each of the village's neighborhoods or hamlets around issues of elementary education. In each of these small group meetings, people were invited to a large village-wide meeting. These large village-wide meetings were the culmination of the whole process and were attended by teachers, parents, community members and members of the village administration.

In the simplest of the three interventions using this format, the Pratham team facilitated the small and big group meetings in the village. In the meetings, school staff and village local government representatives were encouraged to share information about the structure and organization of local service delivery, including the role and activities of the VECs. In the design of this intervention, the focus was on ensuring that the relevant information about norms and provisions (usually on inputs) for schools were shared in the village. The meetings were followed by the distribution of pamphlets that described the various roles and responsibilities of VEC members. The pamphlet was also explained to as many individual VEC members as possible. In this intervention, there were no direct discussions on the learning levels of children, unless the topic was brought up by community members.

The second treatment included the activities described in the first intervention. In addition, in the small neighborhood or hamlet level meetings, Pratham team members demonstrated the process of creating "learning report cards" by conducting simple assessments of reading and arithmetic with the local children. In this process, Pratham team members actively invited and immediately oriented local community members to participate in the making of the report

cards. The most villages of Uttar Pradesh, parents are not very literate and thus unable to really understand the issues related to learning levels of their children. One of the objectives behind the design of the second intervention was that community members and parents become sensitized to the status of their children's learning. By participating or observing how children were being assessed in reading or arithmetic, community members and parents would get a first hand view of what was meant by basic learning and how their children were performing on simple tasks. The various neighborhood report cards were then also discussed at the village-wide meeting. The aggregation of the hamlet or neighbourhood report cards led to the village report card which was discussed at the big meeting at the village level.

The first two interventions were therefore aimed at strengthening the SSA model of community participation by a combination of informing village members of the provisions and norms and giving them a good idea of what processes had been envisaged in SSA for facilitating collective action.

The third intervention added a "direct-control small group" component or a component that demonstrated that local "solutions" are possible. In addition to all the activities undertaken under the first and second interventions, in the third case a "demonstration" class was started in the village. This demonstration class showed how simple activities with children helped to improve their reading skills. The demonstration class ran in the village for a period of 3-4 days. Local community members usually youth attended and watched the activities. Those who had shown keen interest were then given several days training and materials to work with children of their village. This was an entirely voluntary effort. No money was paid to volunteers for their work with children. The pedagogical technique, training and materials that were used in Jaunpur were the one that Pratham had developed and used throughout India for teaching basic reading skills. After training, volunteers then held reading classes in the villages. The typical "reading course" lasted two to three months, with classes held every day outside of school. This intervention thus offered committed individuals the opportunity and the competence needed to directly improve learning among children.

I

The results from the evaluation show that none of the three intervention methods managed to effect large group indirect or direct control over public schools, in terms of participation by any of the players (the parents, the VEC, the teacher), nor did they improve school performance. This is not because the mobilization entirely failed: the meetings organized by Pratham were well attended (on average more than 100 people attended in a village of 360 households). Nor is it because villagers are unwilling to get involved or do not care about education, or because they are pessimistic about the possibility of improvement. This becomes evident when we look at the impact of the small group component of the third intervention, which was a clear success. It succeeded in mobilizing a large number of volunteers from the villages, who signed up for the Pratham training and then set up reading classes in their village: More than 400 reading camps were held across 55 villages. Almost 7,500 children enrolled (more than 130 children per village) suggesting that there were enough parents and children who were keen to improve their children's education.

The results also demonstrate that teaching these children how to read is not an impossibly difficult task. In contrast to the failure of the first two interventions, the reading camps had very large effects on learning: after a year, we

see evidence of very substantial progress for the children who attended the camps. Our instrumental variables estimate suggest that the average child who could not read anything at baseline and who attended the camp was 60 percentage points more likely to decipher letters after a year than a comparable child in a control village. The average child who attended the camp and who could decipher letters, but not words, in the baseline was 26 percentage points more likely to be able to read and understand a story than a comparable child in the control villages. Combined with the natural progress of children over the course of a year, these results imply that, after a year, all the children who could not read at all and attended the camp were able to decipher letters, and 35 percent of the children who could read letters and attended a camp were able to read and understand a story.

However it is also noteworthy that neither of the first two interventions generated the kind of educational volunteering that the third intervention so successfully generated. Villagers got the same information about the state of education in the village in treatments two and three. What the third intervention did differently was to train the volunteers, which presumably made volunteering more salient and gave the volunteers some confidence in what they were doing. But in principle, there was nothing to stop educated villagers in intervention 2 villages (or for that matter, in intervention 1 villages) from reacting to the campaigns by starting some classes on their own. Indeed they could have asked the Pratham facilitators for help with starting the classes. This did not happen: there was no increase in the availability of tutoring classes in intervention 1 and 2 villages.

Thus, none of the intervention had any effect on large-group strategies from any of the three interventions, and the first two interventions failed to generate small group interventions. The contrast between these five negative results and the enormous impact of the last intervention on coaching and child outcomes, is striking. The difference between what happened to the large group and small group behaviors may lie in the fact that large group mechanisms make very different demands on the community than small group mechanisms. The fact that the third intervention got so many people to volunteer for read camps but almost entirely failed to influence involvement with the VECs, might then reflect the community's expectations about the efficacy of the large group mechanisms—if you do not believe that these mechanisms work, there is no reason to invest in them. We go on to argue that thinking about the different models of participatory action can also help us understand the relation between these results and the (in some instances, much more encouraging) results from the other recent evaluations of participatory programs (Kremer and Vermeersch, 2002; Olken, 2005; Banerjee and Duflo, 2006; Bjorkman-Svensson, 2006; and Duflo, Dupas, Kremer, 2007).

The contrast between the impact of the first and particularly the second intervention on volunteering and the impact of the third is perhaps less surprising, but it reinforces the evidence presented in Banerjee and Duflo (2006) showing that merely giving villagers information about the state of public goods in the village, without facilitating the use of that information, may not always be very useful. While people in the village were clearly prepared to volunteer, they did not feel comfortable coming forward to do so without the training and encouragement from the facilitators (although we do not know whether it was the encouragement or the training that did it).

In the remaining sections of this paper we describe the institutional context of participatory action, emphasizing the nature of the SSA model (section 2); compare the interventions that are evaluated (section 3); and discuss the data and

the empirical strategy (section 4). We present the results in section 5, and conclude with a section that links these results to the related literature and discusses what this evidence tells us about how and when community participation can be effective.

2. The Sarva Shiksha Abhiyan Model of Participatory Action in Education

2.1 Policy

The VECs are conceived of as the primary channel of participatory action under the SSA, which is the central initiative of the Indian central government toward achieving universal primary education in India. States in India differ in the design of VECs and the roles and responsibilities assigned to them. In UP, the VECs consist of the elected head of the village government, the head teacher of the local government school, and three parents of children enrolled in the village public schools. The parent members of the VECs are nominated by block-level public officials (the block is the first administrative level above the village). VECs are responsible for monitoring performance of public schools by visiting the schools, scrutinizing records, meeting teachers, parents, and children, and devising solutions to address any problem encountered. They are entitled to claim specified public monies and powers for this purpose—such as public grants for school development, the power to mobilize community contributions toward school improvement, and the power to identify and hire community members if the school needs additional teachers⁹. These community teachers are called Shiksha Mitras: they are usually paid by the state government (in some cases, the community pays for them), but the community has the responsibility for overseeing them.

Even as conceived, the VECs' control over the school is primarily indirect. It can petition for resources that the school is entitled to ask for—for example for hiring a Shiksha Mitra—but whether the funds show up or not depends on how much pressure they can put on the bureaucracy at the district level. There is no official guarantee that any village is entitled to them. It can also complain about the teachers or the level of education in the schools, but, once again, the ultimate decision on whether anything will be done about it is not in their hands. Nevertheless the VEC can intervene directly in some areas: It can put direct pressure on the Shiksha Mitra to teach better and come to school more often, and it has the right not to rehire a Shiksha Mitra if his/her performance is deemed unsatisfactory. It gets a small amount of money each year (Rs. 7000, about \$170) from the SSA to spend on school maintenance and ways to improve teaching in the school. It can also raise and spend resources from the village, but, given that almost all the resources in public schools have traditionally come from the government and the poverty of the average villager in this part of UP, this is probably a relatively limited option.

In the case of Uttar Pradesh, the VEC by itself is a relatively small group. Two members – the head of the local government and the head teacher already have powerful roles that they play in the village and in the school respectively. The remaining three members are nominated not elected. Since these committee members are nominated, rather than elected, it is likely that it would need the support of a larger group of parents to take any action. It is true that the village head (who is part of the VEC) might be able to exercise control on his own, but given the limited political salience of education, this is probably not a priority for him. In order to get him to act or to bypass him the parents in the VEC would need the mandate of a large group of villagers. Indeed, if, for example, they were the only three parents who were

⁹ This description is taken from Uttar Pradesh state government published documents on SSA and VECs.

complaining about the quality of the teaching, while the rest of the parents seemed content, it is not clear that any action against the teacher would even be warranted. In other words, it seems likely that the VEC can only be effective when it reflects the collective will of a large group of engaged parents.

On the other hand, it is also worth emphasizing that India has an elaborate system of local self-governance with each village sending an elected representative to a local governing council (the panchayat) with a range of statutory powers and a direct link to the higher level of government. In principle the villagers did not need to go through the VEC in order to put pressure on the schools. They could directly lobby the panchayat, or even the MLA, who is their representative in the UP state assembly. We will therefore look at all forms of protest and lobbying by villagers and not just what they do vis a vis the VEC.

2.2 Practice

In practice, although the VECs had been (statutorily) constituted in 2001 in every village, well before our study began, there is no evidence of parental involvement in the running of the public schools, either through the VEC or through other community or village mechanisms, based on a pre-intervention survey we conducted in 280 villages.

The survey included an interview of 1,029 VEC members (their names and addresses were provided by each school's headmaster). The salient results from this survey are presented in table 1 (column 1). A striking fact is that only 38 percent of the VEC members interviewed mention the VEC spontaneously when asked the names of organizations they belong to, and 25 percent still say they are not members after being asked specifically if they are a member of the VEC. Only 26 percent of them have heard of the SSA. Only 21 percent know that their committee is entitled to receive resources from the SSA (panel A). Most startlingly, only 3.6 percent mention the ability to request government funds to hire a Shiksha Mitra (an additional para-teacher) when the school is over-crowded as one of the VEC's prerogatives (and responsibilities) (Panel D). Yet, this is probably the most important power they have, since this gives them not only an extra teacher, but an extra teacher they directly control (at least on paper).

Given the ignorance of the VEC members, themselves, it is not surprising that parents know even less about the VEC and its responsibilities. Column 1 in table 2 reports some summary statistics from parents' responses to this survey. For example, household respondents were asked whether there was any committee in the village that dealt with issues relating to education. A startling 92 percent responded that they did not know of any such committee. Only 3 percent could name actual members of the VEC.¹⁰ It is worth noting that ignorance and lack of participation in the institutions of local governance, is not just a problem for education. Only 14.2 percent of respondents report that a household member ever attended a Gram Sabha, village meetings that were institutionalized as part of a country-wide decentralization initiative in 1993. However, as panel C of table 2 shows, for the households that do participate in local governance, education seems to be a low priority. Of those who have attended any Gram Sabha, only 3.3 percent mention education when asked about which issues were covered in the last meeting. More generally, when parents were asked what

¹⁰ Moreover, the proportion of people without any knowledge of VECs remains as high even when we look only at parents whose children are enrolled in government schools (Banerjee et. al. 2006, Figure 9).

they consider the most pressing issues in the village, education ranks fifth on the list of village problems, with just 13.1 percent of respondents mentioning it at all.

The baseline survey also found evidence that community members—parents, head teachers, and village leaders—do not know how bad things are in the village as far as education is concerned. The actual learning levels of all children between the ages of 7 and 14 from over 8,000 randomly selected households across our study villages were measured using tools developed by Pratham to rapidly assess whether a child can read, write, and do simple arithmetic. The results are presented in figure 1 (see also panel A of table 6), and shows how low the learning levels were: 15 percent of children age 7 to 14 could not recognize letter; Only 39 percent could read and understand a simple story (of grade 1 level); 38 percent could not recognize numbers.

The survey asked parents what they knew about learning levels of children in the village, as well as their own children. Twenty percent and 21 percent of parents, when asked, said that they had no idea of the village children's ability to read a paragraph or write a sentence, respectively. The average parent overestimated by 12 percent the proportion of children in the village who could read a paragraph (table 2, panel B, column 1). Furthermore, 42 percent of parents were too optimistic about their own children's ability to read, and 25 percent overestimated their ability to write. For example, 66 percent of the parents of the children who could read nothing thought they could at least read letters, and 36 percent of the parents of the children who could barely decipher letters thought their children could read and understand a story (see figure 2). The picture is even more distorted in math, where a full 82 percent of the parents of the children who could only recognize numbers, but could neither subtract or divide, believed that their kids could subtract (see figure 3).

3. The interventions

The fact that there were large gaps in what the average villager knows about the state of education in his village or what he can do to about it, suggested the possibility of getting them more involved by sharing this information with them. Moreover the palpable lack of any urgency associated with the problems in education suggested that motivating the villagers and helping them coordinate on doing something about education might also make the VEC more effective. To achieve this, Pratham, a very large non governmental organization long involved in trying to improve the quality of education in India, developed and experimented with three different methods of mobilizing the communities and sharing information with them.

3.1 Intervention description

All three interventions adopted the same basic structure to share information on education and on the resources available to villagers to improve the quality of education. The intervention started with small-group discussions carried out in each hamlet over at least two days. In all these meetings, Pratham staff members acted as facilitators, encouraging discussion by raising questions, rather than providing facts: Do you know about the status of education in your village? Do you think children are learning? What aspects of education concern you the most? The intervention culminated in a general village meeting, typically attended by the school headmaster, teachers, parents, VEC members and others from the community. The intervention teams tried to facilitate the discussion in this meeting so that local key actors of the village (the school teachers or Gram Pradhans) provided both general information about the provisions and resources available at village

level, as well as village-specific information on resources available to the village, schemes that were in place, about the existence of VECs, its membership, what resources it receives, and the different roles it can play. However, Pratham facilitators were provided a fact sheet covering information about the public education system and VECs, and checked whether all these facts were shared at the village meeting. If something was missing, they would raise it themselves. In the following weeks, facilitators visited each VEC member and gave him or her a written pamphlet on the roles and responsibilities of the VEC, which they also discussed with the VEC member.

This formed the basic structure of all three interventions. The first intervention stopped at this point. The key additional element in the second intervention was that villagers also generated their own information about children's learning outcomes. As already mentioned, it was clear from the baseline survey that a large fraction of children currently enrolled in school were unable to read a simple text or do basic arithmetic, but many parents overestimated their children's learning levels. In addition, during the piloting the field staff noted that even when people talked about education, it was rarely about learning—if anything got them excited it was the state government's scholarship program, or the new mid-day school-meals program. The second intervention aimed at sharing information about the status of learning in the villages with parents, teachers, village leaders, and VEC members, to help parents focus on the issue of learning in their discussions about education. To this end, the Pratham staff taught interested villagers how to evaluate a child using the simple testing instrument used by Pratham everywhere (including for our own data collection), which we describe in more details below. In each neighborhood, a number of citizens tested all the children, and in just a few days the villagers generated their own "reading report card," which was then discussed at the village meeting. Villagers who had participated in creating the report card for their locality were encouraged to speak out at the village meetings and present their findings and experiences. This had the impact of generating the necessary information, actively engaging the community, and shifting the conversation in the general meeting toward learning issues. In addition, this intervention also transferred to the community a specific monitoring tool, which could make it easier for villagers to monitor progress.

The third intervention supplemented interventions 1 and 2 by providing a way for a motivated citizen to directly act to improve education levels. Pratham introduced the villagers to a simple technique for teaching children how to read used in its own flagship "Read India" program. It invited local volunteers to get trained in the technique and start reading classes to be held after school. The training lasted for four days. Volunteers were then supported by periodic visits from Pratham staffers, who checked that the classes were held and provided some in-service training. Each village in treatment 3 received about 7 visits from staffers. The third intervention therefore offered the community a direct and individual way to improve learning levels. Such direct interventions by the village community is specifically mentioned in the UP state government pamphlet on VECs as one of the things that a community member can and should do to improve education in their village. However, it is worth emphasizing one key difference between this piece of the third intervention and everything else that was implemented in this experiment. This was the one thing that sections of the community could do *without engaging at all with the school system* or even the majority of the village.

Each of these interventions was implemented in 65 villages, randomly selected out of the 280 villages in the baseline between September and December 2005. A fourth group of 85 villages formed the control group. Monitoring data suggests that the interventions were well implemented: All treated villages had at least one meeting and some had more than one. A total of 215 village-level meetings were held in the 195 villages. The meetings were well attended (the general meeting had on average 108 participants, 95 percent of whom were present during the entire meeting), with good

representation from different hamlets and castes of the village (37 percent of the meeting attendees were women). In terms of who spoke, 72 percent of the meetings had participation from a wide range of groups and castes, and in 55 percent of the meetings, men and women were equally likely to be talking (in 84 percent of the remaining meetings, males did most of the talking). In 55 of the 65 treatment 3 villages (i.e., 84 percent of the total), there were volunteers who started reading classes. On average, there were 7.4 reading camps per treatment 3 village, each led by a different volunteer—but with the considerable variation of between 0 and 16 groups per village—and 7,453 children attended camps. Therefore, in each of the 55 villages with reading camps, about 135 children were enrolled in these camps.

3.2 Comparison with best practice—why these interventions?

While these interventions were based on Pratham’s extensive knowledge of the situation on the ground, and on extensive piloting of ways to conduct effective meetings, it is reasonable to ask whether the interventions were actually designed to work well. There is of course no conclusive way to answer this question—there is always the possibility that something else would have worked better in any particular context. On the other hand, the intervention needed to be simple enough to be replicable and the lessons to be generalizable. Recognizing this, the community of practitioners has developed a set of “best practices” for external interventions aimed at inducing greater popular participation in the monitoring and improvement of public services. The fact that the interventions followed these guidelines should go a long way toward assuring that what we evaluate in this paper corresponded *ex-ante* to what policymakers would consider to be an effective participation intervention. Thus, in this section we summarize the best practice guidelines available from the World Bank’s Participation Sourcebook (and other notes on best practice from this same source)¹¹ and argue that the intervention that Pratham designed and implemented met most of these criteria. Moreover, we will suggest that Pratham was a natural candidate for being the implementing organization. In this sense, the intervention we study should provide us a “best case” scenario for the effectiveness of these kinds of interventions, at least in a context similar to UP.

The first guideline is that the intervention should be inclusive: there needs to be an attempt to include all sections of the village and make them feel that they are a part of the intervention. Specifically, it is not enough to have one big meeting where the outsiders come and deliver their message and leave. The Pratham volunteers were in the field at least for two days (for four days in the cases of interventions two and three). Facilitating teams went from hamlet to hamlet within a village, making sure to cover “low-caste” hamlets, carrying out conversations about education in small and large groups (which enabled women to participate, for example), and inviting local people to take the lead.

Second, the mobilization should not create unrealistic expectations. Here, the objective was to raise the learning levels of the children, in particular with respect to reading. Pratham’s experience shows that it is indeed possible for a child to considerably improve his or her reading level in a few months, provided that some targeted attention is paid to the child (Banerjee et al., 2006). The tools transferred to the community through the mobilization efforts were appropriate for the direct and indirect controls available to the village community to improve education. They could use these tools for indirect action—such as monitoring children’s learning, complaining about poor performance to higher authorities, and demanding additional resources or training for their teachers. They could use the tools for direct action—for example,

¹¹ The Sourcebook can be found at the following web site, with links to other notes on designated best practice: <http://www.worldbank.org/wbi/sourcebook/sbhome.htm>

giving greater incentives to Shiksha Mitras by threatening to withhold contract renewal unless performance improves, or allocating the resources they receive under SSA or through other village discretionary budget to improve school conditions. In intervention 3, they could use the teaching tool directly to hold classes to teach children.

Third, the intervention should not bypass or undermine existing institutions. Interventions 1, and 2, which build on the role of VECs in facilitating change, clearly satisfy this criterion. Intervention 3 was the only one where an alternative to the existing institutions was proposed, though as a complement to the system, rather than a substitute.

Fourth, practitioners emphasize the value of community-owned “score-cards” in mobilizing communities to take action. In interventions 2 and 3 the community created its own report card by testing children in math and reading. Both the results and the tools were transferred to community members.

Pratham is the largest educational NGO in India. It has demonstrated success in several randomized evaluations of its programs and reaches millions of children throughout India. Pratham designed and runs the Annual State of Education Report (ASER), which tests children in all of India’s nearly 600 districts every year, and is extremely prominent in the discourse on education in India. The organization takes community participation in education very seriously (it is the backbone of its flagship “Read India” program) and devotes considerable resources to make sure that the program is implemented as well as possible. Pratham’s motivation and expertise thus made it an obvious candidate for implementing these interventions.

Taken together this suggests that we can be reasonably sure that what we are evaluating was well-designed and effectively implemented programs. It is also clear that they reached their intermediate goals: encouraging participation, holding meetings focused on education, and generating discussion, interest, and willingness on the part of at least some people to act (as evidenced by the fact that Pratham was effective at recruiting volunteers for the reading camps).

3.3 Why might we have expected these interventions to work?

These interventions involved a combination of more information (about learning levels and levers for change such as the VEC) and more coordination on doing something about education. The predicted impact of more coordination seems unambiguous: More action and better results is what we would have expected.

Giving people more information, on the other hand, need not always promote greater collective action. For example, if the only people who were acting initially were those who believed that the returns to collective action were very large, information that gave everyone a more realistic sense of the effectiveness of collective action might very well reduce the total effort put in. However, in our particular context this effect seems unlikely because for most people the effect of the information should have been to make them, *ceteris paribus*, more willing to act: First because they were being told that they had more powers than they had previously imagined, and second because they discovered that the problem was more serious than most of them had known and therefore there was a greater need to act.

However, this presumes that they registered the information that they were being given.¹² If they were sufficiently pessimistic about being able to use these powers to start with—possibly based on their previous experience with trying to

¹² At least in the case of information about learning levels, treatment 2 was designed to help ensure community members understood the information by including them in its generation.

put pressure on the education system—it would not be surprising if they did not pay very much attention when the powers were being detailed. In particular, the decision to pay attention probably only pays off if others are also paying attention, so it is easy to see that there may be a coordination failure at this preliminary stage. In that case we would see no change in their willingness to participate and no change in outcomes.

4 Evaluation: Data collection, baseline results, and empirical approach.

4.1 Data Collection

The evaluation took place in 280 villages in Jaunpur district in the state of Uttar Pradesh (UP), India. The state of UP is the most populous state in India with a larger population than Mexico (166 million inhabitants according to the 2001 census). It is also among the five worst performing states in terms of basic literacy. Jaunpur district was chosen in part because it is close to the state average in terms of literacy, and in part because it was one of the districts where Pratham was not already working. Districts in India are divided into administrative blocks. Four of these blocks were randomly selected to participate in the study, and the study villages were then randomly selected within each block. The survey and the study are thus representative of Jaunpur district (and its 3.9 million population) as a whole.

In both treatment and comparison villages, a baseline survey was conducted in March and April 2005, and the endline survey took place in March and April 2006. The evaluation instruments included a detailed data set on 10 randomly selected households per village; a dataset on reading and math outcomes for all children aged 7 to 14 (at the time of the baseline) in 30 randomly selected households in the villages; a dataset of headmasters of government primary schools and all VEC members; data on school infrastructure and functioning; and an average of 6.7 observations for each school on teacher presence, obtained during random, unannounced visits during school hours. Every survey was conducted at least twice; hence, each dataset has a panel structure. The data on learning outcomes were collected using a simple tool developed by Pratham, and used in the ASER survey. For reading, children were presented with a simple paragraph and asked to read it. If they could not read it the tester quickly switched to a panel with a few single words on it, and if they could read it, the tester moved up to a longer and more complex story. If the child struggled to read the words on the panel, the tester showed a panel with several letters and asked the child to identify them. A similar approach is used for math—with the levels in the math case ranging from easy to hard in addition, subtraction, multiplication, and division.

The final sample for the baseline survey consists of 2,800 households, 316 schools, 17,533 children (ages 7-14) tested in reading and math, and 1,029 VEC member interviews from the 280 villages. In the endline 17,419 children were tested. There was very little attrition from the baseline survey: The endline sample includes all but 716 of the children in the baseline (the attrition is evenly spread across the various treatment and control groups). An additional group of 2,131 newly-turned seven-year-olds were added to the sample at endline (although these are not included in most of the analysis that follows since we have no baseline for them).

4.2 Baseline Results

Baseline data broken down by treatment group is shown in an appendix (available from the authors). It confirms that there are no systematic differences between treatment and comparison groups or between the different treatment groups: the allocation to groups was indeed random.¹³ The baseline data for the combined sample is presented in column 1 of each regression table. Several of the results from the baseline survey have already been discussed above: We saw that the VEC was not active (table 1), that villagers knew very little about the local structures for impacting school policy (table 2), that the learning results were dismal (figure 1), and that the community's perceptions about learning did not match the unfortunate reality (table 2).

The baseline provides some additional relevant data on the situation in the schools. From the school surveys (table 3) we found nearly all the schools had at least some textbooks, covered classrooms, a water source, and material, such as maps on the wall. As public schools they had all received some money from the government but only ten percent had received money from the panchayat. About half the schools had toilets, and less than a third had a separate toilet for girls. Sixty-nine percent provided midday meals (a legal requirement). Only 2 percent had electricity. Involvement of parents in school is limited—with only about four fifths of schools even holding parents' meetings. Very few parents make donations—although about a third of schools have had some parental volunteer input.

Table 4 presents data on teaching resources, and teaching and children absence. On average, schools in the baseline had 3.2 teachers, of which 0.62 were Siksha Mitras. Headmasters and teachers were absent roughly a quarter of the time and in only 44 percent of the cases was a teacher actually teaching at the time of an unannounced visit to the school by the enumerators (these results are virtually identical to the nationwide results reported in Chaudhury et al., 2006¹⁴). Children are absent a lot as well: only 53 percent of the boys and 50 percent of the girls enrolled were present at the time of the unannounced visit (Table 5, panel B). This is confirmed by parents: in the last 14 days, parents report that their children were in schools only 7.3 days (out of 12 school days). Panel A in Table 5 shows that most children are enrolled in school: only 7 percent of children aged 7 to 14 were not in school at the baseline. Given the poverty in the village, it is striking that a large minority of children (37 percent) are enrolled in private or NGO schools (mainly private schools).

There was thus considerable margin for improvement: more resources could have been provided (including Shiksha Mitras), and teachers and students could have attended more.

¹³ Of 88 characteristics presented in the appendix tables, only 3 were found to have statistically significant differences at the 5 percent level between the groups—i.e. somewhat less than would be predicted by chance. The three were the number of seats in the classroom (where treatment 1 and 3 had slightly more than control), the number of parents who saw education as a problem (where families in treatment villages were slightly less likely to say that education was a problem), and the distance between parents beliefs about what their child could do and reality. As a family, school facilities were not better in treatment than control.

¹⁴ Unfortunately, at the baseline, the random check did not collect data on regular teachers and Shiksha Mitras separately. However, this data was collected at the endline, and, in the comparison group, Shiksha Mitras were both more likely to be present and more likely to be teaching than regular teachers.

4.3 Empirical Strategy

- **Intermediate outcomes**

Given the randomization, the basic empirical strategy is fairly straightforward. We first group each outcome into “families” of related variables (each family corresponds to a panel in our regression tables). We then regress the endline measure of each outcome in that family on an indicator for each treatment group (the control group is the omitted category) and baseline measures for each of the outcomes in the family, i.e.¹⁵

$$y_{ijk} = \alpha + \beta_{1k}T_1 + \beta_{2k}T_2 + \beta_{3k}T_3 + X\gamma_k + \epsilon_{ijk} \quad (1)$$

where i indexes the households, j indexes the village, k indexes the outcome, and X are the baseline values for all the outcomes in the family. The standard errors are clustered at the village level (using White standard errors) to reflect the fact that the treatments were implemented at that level. We also run a specification where we pool all the treatments together.

The only empirical difficulty is that there are a large number of outcomes that could have been affected by the interventions. This embarrassment of richness means that there is a danger of “cherry picking”: We could choose to emphasize the results that show large effects. To solve this problem, we present results on all the outcomes on which we collected data, and, for each family of outcomes, we follow Katz, Kling, and Liebman (2007), and calculate the average standardized effect over the family of outcomes. For a family with K different outcomes, each indexed by k , the average effect of treatment 1, $\widehat{\beta}_1$, is for example calculated as:

$$\widehat{\beta}_1 = \frac{1}{k} \sum_{k=1}^K \frac{\widehat{\beta}_{1k}}{\widehat{\sigma}_{1k}},$$

where $\widehat{\sigma}_{1k}$ is the standard deviation of the control group for outcome k . The other average effects are calculated in a similar way.

The families of intermediate outcomes we consider are: what VEC members know about their role (Table 1, panel A); VEC activism (Table 1, panels B and D); what VEC members know about the education situation in the village (Table 1, panel C); parental involvement with the school (Table 2, panel A; Table 3, panel B); parental knowledge about the education situation in the village (Table 2, panel B); the priority given to education in village discussions (Table 2, panel C); school resources (Table 3, panel A); teacher and student presence in school (Table 4, panel A; Table 5, panels A, B, and C).

¹⁵ In the appendix, we also present simple differences, without any control variables.

- **Learning**

Learning is obviously the main outcome of interest. Pratham classifies reading level in 5 groups: cannot read at all, can read letters, can read words, can read a short paragraph, and can read (and understand) a story. Reading letters is a simple decoding exercise. Reading a story requires some understanding of the story and some fluency. Reading a word or paragraph is in between (ability to combine letters, but no understanding). We therefore group paragraphs and words into a single category. Since the intervention was designed to bring children from a level of reading to a superior level, we then look at how the interventions affect the proportion of children who, having started at a given level, end up reading at least at a given level. For example, we restrict the sample to children who could not read anything at baseline, and look at the effect of each intervention on the proportion of these children who could read at least letters, at least words or paragraph, or could read a story by the endline. Thus a child who can read a story at the endline gets a 1 for the “letter,” “word or paragraph” and “story” levels. We then estimate equation (1) in for the three outcomes, for all the subsamples.

5 Intervention Results

5.1 Knowledge of the participatory institutions

Reassuringly, the intervention did affect what VEC members know. Panel A in Table 1 shows the impact of the program on what the VEC members know about their role. Looking at the average effect of treatments 1 through 3 (in column 6), the average impact on VEC knowledge about their role is large: 0.35 standard deviations. The effect of the treatment is positive on all the variables in this family. It is significant for their knowledge of the SSA (there is an increase of 7.5 percentage points in the fraction who have heard about the SSA, and 7.8 percentage point in the number of VEC members who know they can access funds through the SSA), and the probability that they have been trained (13 percentage points). However, these improvements are counterbalanced by a worsening of what VEC members know in the control villages, so that the level of knowledge of the VEC members did not actually increase in the treatment villages between the baseline and the endline surveys.

Panel C of Table 1 shows that VEC members also know more about the village’s state of education in the intervention villages: their knowledge of what children know has improved on average by 0.21 standard deviations (significant at 10 percent level) on average across all three treatments (column 6). Curiously, they seem to have learnt the most in intervention 1.

We also find a significant difference in parents’ knowledge of the VEC between treatment and control villages (Table 2, panel A): They are 2.6 percentage points more likely to know there is a VEC (compared to only 4 percent in the control at endline). These effects are however strikingly small. The proportion of people who know about the VEC even after the interventions remains very small, barely 7 percent. The 360 or so households in an average village sent 108 adults to the meeting. Even under the extreme assumption that 2 adults came from every household that was represented, this means that 1 in 6 households were at the meeting. If everyone who was at the meeting registered the fact that there was a VEC in the village, we would have expected the fraction of those who knew about the VEC by the endline to be at least 15 percent (and plausibly much more since 62 percent of the participants were male, and therefore probably represented different families). It seems that either many of the participants in the meetings did not register the information about the

VEC or they promptly forgot (the fact that the share of parents who know about the VEC went down from above 8 percent to 4 percent in the comparison villages, suggests that people do forget).

Parents are also more aware of the village's status of education in villages that came under treatment three, which is the treatment where this information was generated and discussed, and reading camps were held (Table 2, panel B). The average effect on the knowledge of education status in the village is 0.10 standard deviations in treatment 3 villages, and 0.07 standard deviations across all three treatments on average (column 6). These are sizeable effects and the effect for treatment 3 villages is significant at the 10% level. All the treatment effects on the variables in this family indicate improvement for these two groups of villages although different variables are individually significant in the two cases.

5.3 Engagement

Despite the real, if modest, difference in awareness, we see very little difference between the VEC's performance in treatment and control villages. They are no more likely to report that they have complained to anybody, or tried to raise resources in any of the treatment groups (Table 1 panel B). Panel D in the same table shows that they neither showed more awareness of the VEC's responsibilities for hiring Shiksha Mitras nor were they planning to do anything more about hiring Shiksha Mitras.

The intervention also did nothing to increase the level of engagement of parents with the schools (Table 2, panel C). Parents are no more likely to have visited the school or to have volunteered time and/or money in the treatment villages than in the control villages—and this holds for each of the three treatments (Table 2, panel A). The parents' reports are confirmed by the head teachers (Table 3, panel B), who do not report any more visits from parents, or having received any more inputs of time or money from parents, or having exercised any more effort to involve parents

As reported before, the one place where we do see a difference is in the intervention 3 villages where volunteers ran more than 400 reading courses in the 55 villages. We did not collect systematic direct information on volunteer classes in the other villages, but there was no report of such classes being started. An indirect indication that the supply of tutoring classes did not increase is that we did not see an increase in the fraction of children attending tutoring classes in intervention 1 and 2 villages. In contrast, where the volunteers provided a readily available outside option (the read class) parents did take advantage of it. Eight percent of the children in intervention 3 villages have attended a read class (Table 6). Reassuringly, the impact was concentrated among children who did not already read fluently. Thirteen percent of the children who could not read anything or could read only letters in the baseline attended the class, while 7.4 percent of the children who could read words or a short paragraph at baseline attended the class. Also, 3 percent of the children who could read a story at baseline attended the class. While it is not nearly the case that all the children attended, this still represents 135 children per village on average in the 55 villages that held the camps.

Another way parents could have reacted, is to choose a purely individual course of action: exit. The private school market is very active in UP, and almost every village has a private school. Interestingly, however, there is no evidence of parents reacting to the information about the learning of the children by moving their children to a private school. In intervention 3 villages we actually see that the fraction of out of school children go up by 1.3 percentage points (significant at the 5 percent level), a not insignificant 16 percent increase over the comparison group figure, which turns out to be

entirely due to children dropping out of private or NGO school (results omitted to save space). It may be that parents consider the reading classes to be an adequate alternative to a private school.

5.4 School resources, teachers, and students

With headmasters doing no more to mobilize parents, and neither parents nor VEC members complaining more about school resources, it is not surprising that there is no evidence that any of the treatments generated additional non-teaching resources for the schools (Table 3, panel A). The effect on that family of outcomes is actually negative (and even significantly so, for treatment 2).

However, the most expensive and important resource in a school is the teaching staff. Results on teaching resources are displayed in Table 4, panel A. Thanks to funds available under the SSA, the community has the right to a new, local teacher to relieve the burden on the teaching staff. The VEC can play a role in getting approval for the position, and in the choice of whom to hire. This possibility was discussed during all the interventions. The number of such teachers (Shiksha Mitra) increased between the baseline and the endlines (it almost doubled in the control villages). In treatment 2 villages, there is a significant increase in the number of Shiksha Mitra. This is encouraging, as this is one immediate action that was in the power of the villagers in this group, after they learnt about the State of Education in the village.¹⁶ In intervention 3, they could instead focus on starting reading classes, which may explain why the effect was lower there. However, there was no impact on the teacher presence (not even Shiksha Mitra)¹⁷ so that the average effect on the family of outcome indicating “teaching resources” available to the school is insignificant.

Another margin of action at the individual level for parents, to improve their children learning, could have been to push them to attend school more regularly. Children are in school only about half the time, leaving a large margin for improvement. Panel B in Table 4 suggests that the interventions did not affect child effort either (other than through their participation in the reading camps). As we saw previously, they are no more likely to be enrolled in government school. They are also no more likely to be present, according to random checks (panel B) or parental reports (panel C). It is possible that parents did not see the value of attending more given what schools are delivering.

5.5 Learning: Mobilization and Information

Given that there appear to have little or no action at the school level (except for the hiring of 20% more Shiksha Mitras in intervention 2), no initiative by small groups outside the school system, and no individual action by parents to take their children out of the public school system, seek extra help, or force them to attend school more regularly, we probably should not expect interventions 1 and 2 to have an impact on learning levels. And indeed, as Tables 6 shows, neither intervention has an effect on either reading levels or math levels. Both reading and math levels increased as much in the comparison group (presumably because the children became older) as in either Treatment 1 or Treatment 2. None of the treatment effects estimated for various reading levels and subgroups is significant at the 5 percent level or better.

¹⁶ Note that this contrasts with the report of the VEC that they had not done anything about Shiksha Mitras. The Shiksha Mitra may have been hired by the Pradhan without informing the rest of the committee.

¹⁷ Given that the number of Shiksha Mitras went up (though mostly not by enough to have a significant effect) this measure of attendance does not separate between changes in the attendance of Shiksha Mitras who were already working and the effect of adding a new generation of Shiksha Mitras (who may be different from the previous generation).

5.6 Learning: Impact of the “teaching” intervention

In contrast, the third intervention (which had the reading camps) had a very large impact. Column 5 in panel A of Table 6 displays the learning results of intervention 3. Overall, children in the villages that received intervention 3 are 1.7 percent more likely to read at least letters (significant at the 5 percent level). They are 1.8 percent more likely to read words or paragraphs (significant at the 5 percent level), and 1.7 percent more likely to read stories (significant at the 10 percent level).

This average masks considerable heterogeneity, however: when we run the same regression separately for children at different reading levels, the result reveals a clear pattern. Children who could not read anything at the baseline are 7.9 percent more likely to be able to read at least letters at the endline in intervention 3 villages. But their improvement stops at the letter recognition stage: they are no more likely to be able to read paragraphs or stories (Table 6, panel B). Those who could read only letters at baseline are 3.5 percent more likely to read at least paragraphs or words, and 3.3 percent more likely to read stories if they were in intervention 3 villages. Those who could read either at the word or paragraph level are 4 percent more likely to still read at last paragraphs or words (note that 19 percent of the children who started at that level have regressed in the comparison group).

These increases may not seem that large at first, but we have to remember only a small fraction of the village’s children attended the classes: on average, only 8 percent of children (including 13 percent of those who could not recognize letters) in our sample attended the reading class in intervention 3 villages. Since none of the interventions (including intervention 3) seems to have affected any other dimension of the child’s learning experience, it is reasonable to assume that the effect of intervention 3 on reading outcomes came entirely from attending the read classes. In this case, being in an intervention 3 village is a valid instrument for attending the read class. We thus run a two-stage least squares regression, where the variable “attended a reading class” is instrumented using “intervention 3”, i.e. we run

$$y_{ij} = \alpha + \lambda R + \beta_1 T_1 + \beta_2 T_2 + X_{ij} \gamma + \epsilon_{ij} \quad (2)$$

where R is a dummy for whether the child attended the reading program, and T_3 is the instrument for R (other variables instrument for themselves). X is a set of child-specific controls including age, gender, school status at baseline, and reading level at baseline.

The results are presented in column 7 in Table 6. Provided the exclusion restriction is valid (i.e., that the effect of intervention 3 is entirely channeled through attendance to the reading classes), this instrumental variable estimate tells us the effect of attending a read class on the ability to read at various levels, for the sample of kids who chose to attend. On average the results in panel A suggests that attending a read class makes these children 22 percent more likely to be able to read at least letters, 23 percent more likely to read at least a word or paragraph, and 22 percent more likely to read a story (though this last effect is not significant). We should not forget, however, that this is a test of the impact of the read class on those who participated, and it is possible that if the program had been successful at enrolling more children the extra children would not have benefited as much—for example if they were somewhat less motivated than those who did enroll.

Again, to understand the impact of the program, it is important to disaggregate by initial learning level. A child who could not read at all is 60 percentage points more likely to read letters after attending the reading class. A child who could read only letters at baseline and who attended the camp is 26 percent more likely to read stories. Thus, the effect on those who could already read letters is mainly to allow many of them to directly switch to the story level. Children who were able to read at the word or paragraph level are 61 percent more likely to still read at the word or paragraph level, and 46 percent more likely to read at the story level (though this last effect is not significant). These are very large effects. In fact, they had not much scope to be larger. Adding the endline reading level in the comparison group (column 2) to the point estimate of the treatment effects (column 7) implies that *every child* who could not read anything at baseline but attended a read class could read letters at the endline, and almost every child (98 percent) who could read at the word or paragraph level can now read at the story level. Almost 35 percent of those who could read letters at baseline and who attended the read class are able to read at the story level.

In summary, the reading program, which offered the villagers an alternative form of participatory action where a single individual could directly affect learning outcomes, did lead to dramatic increase in reading ability for those who attended. These results confirm Pratham's intuition that, combined with the traditional school system, a two to three month camp attended two hours a day is sufficient to get many children who could not read fluently (the letters readers) to read fairly fluently. Children who start with nothing can be taught letters. This suggests that either a second camp would be needed for them, or that they are harder to teach. Not surprisingly, there is not much evidence of an effect of the reading camps on math (although there is a surprising 2 percentage point increase in the number of children who can divide in treatment 3 villages (Table 6, panel E)). These results are important in and of themselves (they do suggest the Pratham intervention, which is carried out in many parts of India, is an extremely cost-effective way to improve test scores), and also because they suggest that teaching a child to read is not particularly difficult. If a village level volunteer with 10 or 12th grade education can achieve this goal after four days of training, it suggests that the failure of the government school to achieve it owes more to a lack of incentives or motivation, than to the innate difficulty of the task.

However, not everyone who should have attended did so, and as a result the effects of the program appear more muted when aggregated at the village level. The next step for Pratham therefore seems to be to find ways to effectively increase the outreach of the program, while maintaining its effectiveness. In ongoing work we are evaluating the impact of training the teacher in the reading method.

6 Interpretation and Conclusion

“Citizen Participation” is often touted as a general purpose solution to the many deficiencies of publicly provided services. Aid agencies, such as the World Bank, now recommend that some institutionalized community participation should be part of all the government programs they fund. Governments also are increasingly beginning to count on “people’s power” (and contributions) as a way to revitalize their struggling education and health sectors. Whether participation can indeed be achieved through external interventions or top-down policy design, and further, whether it can be effective in improving service delivery or development outcomes, remains an open question on which rigorous evidence is beginning to emerge. This paper contributes to the evidence and begins to distil lessons on the elements of intervention design that are likely to be successful.

The results from the study show sharply divergent results. Participation in all forms of large group control, both direct and indirect, did not change under any of the interventions. The first two interventions also had no effect on any form of small group control including any form of volunteering by villagers. As one may have expected, in these two groups, there was no impact on learning. In contrast, in treatment 3 villages there was quite remarkable community participation in response to the offer of being trained in a teaching tool, something that gave the volunteers (presumably with the support of some parents) a large measure of direct control over educational outcomes for a group of village children. More than 400 community members volunteered to take up the tool, and held reading camps in which almost 7,500 children enrolled. These camps were remarkably effective in teaching illiterate children to begin to read.

In retrospect, this contrast between the lack of effect of intervention 2 and the success of intervention 3, is clearly a result of the fact that a small group action was facilitated in intervention 3, while intervention 2 emphasized large group action and did not directly facilitate anything that could be done in a small group. (though there was nothing to prevent villagers from independently taking up a small group action). The contrast between a small group action and a large group action is presumably in part a result of the challenges of sustaining collective action in a large group. If parents and the VEC expected that the action would eventually collapse, one can see why they would choose to close their mind to the whole thing from the beginning and not even register the information that was being given in the meetings. The small group action, by contrast, required little coordination, especially since the classes could be held pretty much anywhere in the village by anyone who was interested and motivated. In part it may also reflect general skepticism among parents and other community members that the school system can deliver much more unless there is a major systemic reform (this may well be a misjudgment—there is a tendency to judge schools by their appearance---but it is certainly an opinion that one encounters quite often). If so, they may not feel that it’s worth their while to fight to make it work better. They may be more optimistic about initiatives taken by individual citizens, such as the reading classes.

The power of individual actions that can generate collective benefits can also explain the impacts of the participatory interventions studied by Olken (2005). He carried out a randomized evaluation of a program in Indonesia. In some villages, non-elite villagers were invited to a meeting where the financial details of road construction in the village were discussed—the idea was that they had the option of speaking up if they did not agree with what was being claimed. In another, randomly chosen, set of villages, anonymous complaint forms were distributed to villagers to expose any corruption in road construction that they knew about. While the first program had no effect on overall corruption, it substantially reduced the amount of leakage that happened through inflated claims about the number of labor hours,

which is the only part of the spending that directly impacted the lives of the average villager (since they supply the labor). The second program reduced overall corruption, but only in villages where the comment form was distributed via school students to take to their parents, completely bypassing the village government. Both gave individuals the scope to act on their own, essentially by complaining, though neither offered any direct control. The second program also protected the identity of the complainant, which might explain its greater success.

There is also other evidence that it is hard to make large group indirect control work. In Kenya, for example, encouraging school committees to report on teacher performance to the district administration (Kremer and Vermeersch, 2002) did little to improve absence rates. Banerjee and Duflo (2006) describe an intervention where the village community was alerted to the fact that the health worker in the government-run village sub-center was absent. This had absolutely no effect even though, in principle, the community had the right to complain through the village institutions.

The one potentially important exception to this rule is an intervention in the health sector in Uganda that seems in many ways to be very similar to ours (Bjorkman and Svensson, 2006). Communities in Uganda were provided with baseline information on the status of health services in public facilities. Over the course of several days facilitators used the information to generate conversations about the quality of care amongst the community and between the community and the providers. The project was implemented in 25 communities, randomly selected out of 50. Bjorkman and Svensson (2006) find that, in treatment communities, action was taken by communities to reform and rejuvenate the beneficiary-control institutions (called HUMCs) which, like the VECs in Uttar Pradesh, were found to be inactive at the baseline. The action included sacking and replacing many members of the existing community health committees. Systematic monitoring by the community of health clinics was often organized. Both provider attendance and quality of service measures (including wait time and quality of care) improved. The final result was an increase in the immunization rate, and a substantial drop in child mortality.

It is clear that this model was very much like the SSA model in the sense that the HUMCs had only indirect influence.¹⁸ They could only complain about non-performing employees and/or dispensaries: Any action against them would have to be taken by a senior bureaucrat. However there are two very salient differences between the two interventions. First, the dispensaries in Uganda are significantly larger and more visible than the schools in UP: Each dispensary serves on average about 2,200 households (Bjorkman and Svensson, 2006) or more than 13,000 people assuming a household size of six. That means that the 50 dispensaries in the study cover nearly 3.3 percent of Uganda's population. Being on the HUMC is therefore potentially more prestigious than being on the VEC and might therefore attract a more elite group of villagers, which would make it easier for the HUMC to act on its own. The fact that the interventions led to the replacement of the HUMC members may be an indicator that HUMC membership was perceived, at least post-intervention, as being something to aspire to, whereas no one seemed to value VEC membership. Second, the Community-Based Organizations (CBOs) that facilitated the intervention in Uganda seem to have played a much more active role in pressuring public providers to improve performance than Pratham chose to play. For example, facilitators from these organizations in Uganda directly negotiated with the dispensary staff before involving the villagers, and the villagers who got involved were hand picked by the CBO. Pratham felt that any intervention that would require them to be any more involved could not have been scaleable; given that health centers are much larger than the schools, and that CBOs were themselves largely local organizations, sustainability may have been less of a concern in the Uganda case. In

¹⁸ Arguably the HUMCs had less direct influence than the SSA model, in which the VEC has some direct power over the Shiksha Mitras.

any event, the involvement of the CBO meant that the local (or non-local) elites were much more involved in the collective action in Uganda than they were in UP. This may be important because most villagers in Jaunpur had no experience of complaining about education to anybody (who do you complain to?) let alone getting a response: Indeed some of them commented to the Pratham team that even the idea of having a village wide-discussion about education was novel. In the circumstances it is not implausible that some more “leadership” would have helped, though its worth remembering that the head of the village was invited to the meeting and was present in a substantial fraction of them (on the other hand, at least in India, the village head does not command a lot of influence beyond the village).

Another related possibility, along the lines argued by Banerjee and Duflo (2006), is that the Ugandan intervention’s “action plan,” drawn by the public providers and the community, was more specific in its goal to improve health services and this might have been quite important given the relative novelty of what the villagers were being asked to do. On the other hand, the outcome of the meetings in the villages in Uttar Pradesh could easily have been an action plan to improve public schools. In fact the meetings were structured to lead to such a plan. The fact that the education meetings in Uttar Pradesh did not yield an action plan, while the health meetings in Uganda did, is thus probably part of what we need to explain.

Of course these are only some of the many potential differences between the two models. Another possibility is that health is different, either because people care more about it (although people cared enough about education to volunteer the time to teach classes) or because, the non-delivery of health services is much more easily and reliably observed by an individual (“I was entitled to that service but I did not get it”), than the non-delivery of effective teaching (what constitutes effective teaching?). Indeed, this might introduce an element of individual action, since people could start going to the clinics and demand the services they did not know they were entitled to before. This kind of individual action is less straightforward in the case of education since monitoring non-delivery requires regular visits to the school (lack of teacher effort is the main problem). Finally, the differential response by public providers in UP and Uganda may reflect the fact that Indian teachers are less sensitive to social opprobrium than Ugandan nurses. It appears that Ugandan nurses are often from the community itself, whereas teachers in Uttar Pradesh generally come from outside the village, belong to upper social strata compared to parents, and are powerfully unionized. In the field-work during this study, Pratham facilitators often noted that teachers tended to blame *parents* for their lack of interest in education as the most relevant source of poor learning outcomes. It is therefore possible that citizens are unable to exert pressure on these teachers, while social pressure seemed to have played a role in Uganda.¹⁹

The results of this intervention also contrast with those of an educational intervention in Kenya. As noted above, in UP the large group failed to have an effect even where it had direct control—over the Shiksha Mitras, for example. It is true that taking any punitive action against the Shiksha Mitras is still not entirely straightforward: The Shiksha Mitra is typically from the village and presumably has his/her protectors within the village community; to overcome their resistance would require pressure from a large enough group of parents. Nevertheless, a very similar intervention in Kenya where parent-teacher committees in primary schools were given the right (and the resources) to hire extra teachers on short-term contracts and to decide whether the teachers ought to be renewed (Duflo, Dupas, Kremer, 2007) did work. And in the (randomly chosen) half of the schools where the committee was provided a day-long training that affirmed and

¹⁹ This is unlikely to be the whole story since the Shiksha Mitras in India do come from the community (through often the upper strata of the community) and they did not respond any more to the intervention than the regular teachers.

explained their exact rights in this regard, the program was more effective in improving children's learning experience, mainly because regular teachers were less likely to bully the extra teachers into teaching their class for them.

Why did the VEC not exercise its direct control on the Shiksha Mitras more effectively? One possibility is that in UP, unlike in Kenya, quality of education is not seen by poor villagers as something worth fighting for. This seems consistent with the fact that even after parents found out how little their children had learnt in the government school (in interventions 2 and 3), there was no detectable movement of children to private schools. On the other hand, we do see other forms of private response. Almost 400 read camps were started and more than 135 children attended camps in each village, suggesting that at least some people cared about education. An alternative view is that collective action needs to be learnt over time. In Kenya, school committees have historically raised their own funds and played an important part in running the local schools. When school fees were abolished in Kenya, these committees stopped having to raise money (the government started giving the committees a grant to replace the money they used to raise), but a long tradition of activism remains, and VEC members are respected members of the community. In India there is no comparable history of local collective action with respect to education. Finally, another important difference is that in the Kenya program, the VECs were given a very specific task: to hire and monitor an extra-teacher for the school. Resources were clearly available and earmarked for this purpose, and it was easy to access them (an NGO with a long standing presence in these schools was administering them). The extra teacher was hired by the VEC following a transparent process, and it was common knowledge that it was accountable to its members (especially in the schools where the role of the VEC in this regards had been reinforced). These last two elements made this intervention much more of a "small group-direct control" intervention.

Whatever the explanation, it seems clear that the current faith in participation as a panacea for the problems of service delivery is unwarranted. It is possible that it can be made to work on a more systematic basis, but it would take a lot of patience and experimentation to get there.