

Working paper

S-37120-PAK-1

Using Preference Parameter Estimates to Optimise Public Sector Wage Contracts

A Field Study in
Pakistan

Michael Callen
Karrar Hussain
Yasir Khan
Charles Sprenger

August 2015

IGC

International
Growth Centre



DIRECTED BY



FUNDED BY



Using Preference Parameter Estimates to Optimise Public Sector Wage Contracts: A Field Study in Pakistan

Michael Callen, Karrar Hussain, Yasir Khan & Charles Sprenger

Final Project Update

IGC Award No. 1-VCS-VPAK-VXXXX-37120

Overview

This document provides a final report for the project “Using Preference Parameter Estimates to Optimize Public Sector Wage Contracts.” We introduced a set of technological and incentive contract innovations aimed at making door-to-door polio vaccination efforts more effective. We implemented the project in Lahore District with the help and support of Rashid Langrial (former Commissioner of Lahore), Capt. Retired Muhammad Usman (District Coordinating Officer Lahore), and Zulfiqar Ali (Executive District Officer Health Lahore).

Eliminating polio in Pakistan is a central global public health challenge; Pakistan is one of three polio endemic countries remaining in the world. In 2014, Pakistan saw 306 new polio cases, amounting to more than 85% of new global cases and constituting, according to the World Health Organization, a ‘global public health emergency’. While the disease is preventable through inexpensive vaccinations, public health bodies in Pakistan face severe challenges to achieving universal coverage

In response to the contemporary resurgence of polio in Pakistan, the Punjab Health Department has initiated door-to-door vaccination campaigns. Prior to our intervention, the performance of polio vaccinators was extremely costly to measure. The status quo procedure was to provide vaccinators with a map of potential households and ask them to self-report the vaccinations completed on each day of a multi-day drive. In collaboration with the Department of Health, we designed a smartphone-based monitoring system that allowed us to track the door-to-door vaccination activities of polio vaccinators by geo-stamping and time-stamping their vaccination attempts. Data from the smartphone system are aggregated in real-time on a dashboard available to senior health administrators so that they can observe the extent of vaccination coverage.

The smartphone system also allows the implementation of pay-for-performance contracts, because vaccinators’ efforts can be observed precisely. For a two-day vaccination drive, the Health Department required workers to complete a given number of vaccinations (300) in exchange for a fixed bonus (10 USD). Vaccinators could allocate vaccinations either to the first day or to the second day of the drive, according to their preferences. In a between-subject design, we assigned vaccinators different ‘interest rates’, or the rate at which one could trade vaccinations on day 1 of the drive for vaccinations on day 2 of the drive. Additionally, between subjects, some vaccinators made their allocations three days prior to the vaccination drive, while some made their allocations on the first day of the vaccination drive. This design permits the identification of time preferences and dynamic inconsistency for the sample of vaccinators.¹ We found that 55 percent of vaccinators prefer to delay tasks more when they are making an immediate choice than when they are making an advance choice based on their decisions from consecutive drives. Our between-subject tests also provide substantial evidence of dynamic inconsistency.²

¹ This approach to measuring time preferences closely follows the method used in Andreoni and Sprenger (2012), Gine et al (2014), and Augenblick, Niederle, and Sprenger (2015).

² Vaccinators allocating tasks three days in advance of the drive, on average, allocate 6.4 percent fewer tasks to day one than to day two. Vaccinators allocating tasks on the morning of the drive, on average, allocate 19.1 percent fewer tasks to day one than to day two, corresponding to a treatment effect of 12.7 percentage points. The treatment effect is roughly equal to the effect moving from a scheme where tasks on day one and day two count equally toward the bonus (i.e., and interest rate of 1), to a scheme where tasks on day one count as one task and tasks on day two count as 1.1 tasks (i.e., and interest rate of 1.1) toward the total requirement.

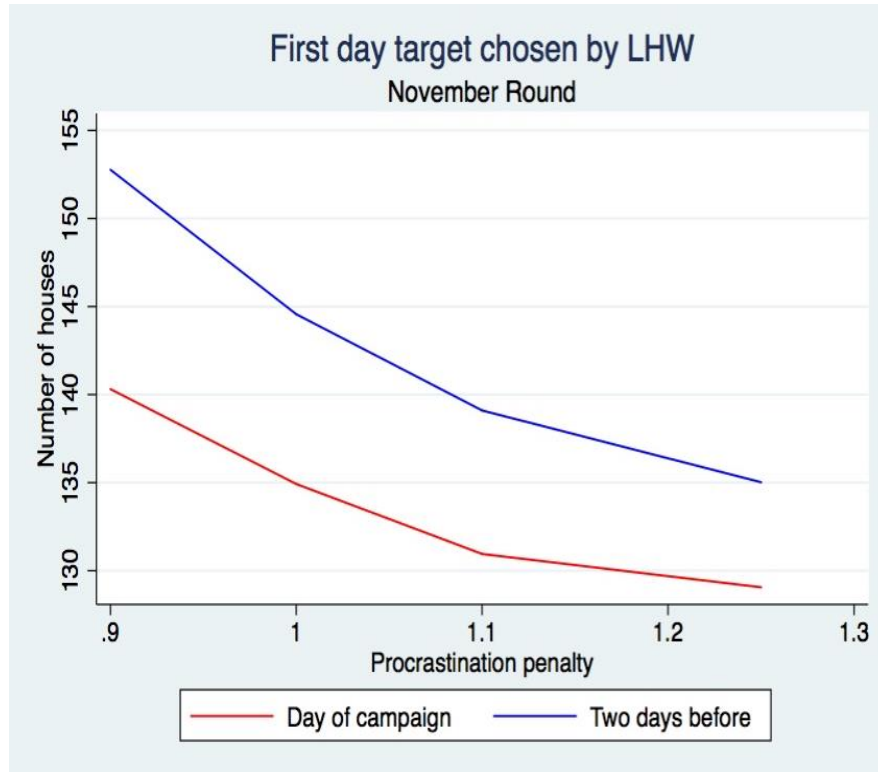


Figure 1: Dynamic Inconsistency in target setting

Building from this finding, our experiment turned to the policy potential of individually tailoring ‘interest rates’ based on the discounting pattern of individual vaccinators. Behavioral and experimental economics researchers have introduced a range of methods for measuring time, risk, and social preferences. Our innovation is to use these measures to customize incentives. We believe this idea might have a broad range of applications.

In a second drive, we tailored a subset of vaccinators’ interest rates to their intertemporal substitution patterns measured in the first drive. With an objective of achieving equal provision of vaccination across drive days, vaccinators whose contracts are tailored to do so achieve the objective substantially more frequently than a randomized control group. A simple way of summarizing these results is by measuring the reduction of unevenness in task allocations. Figure 2 plots vaccinations on day 1 of the drive against vaccinations on day 2 of the drive separately for subjects assigned a random interest rate (left panel) and subjects assigned an individually tailored interest rate (right panel). The effect of tailoring is to reduce the average Euclidean distance between vaccinator allocations and the 45-degree line from 12.2 tasks in the untailored group to 7.7 tasks in the tailored group, a reduction of 4.5 tasks (SE=2.11). This result provides the primary motivation for further investigation into the potential of using experimental measures of preference to improve incentives.

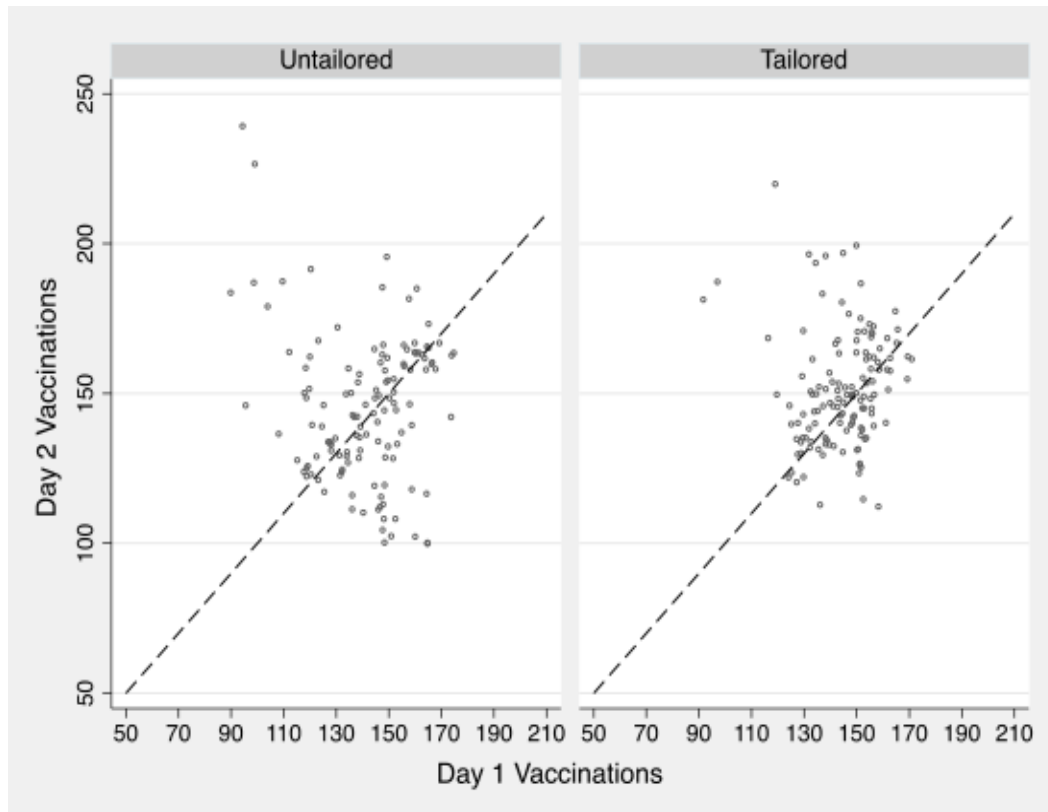


Figure 2: Comparing Tailored and Untailored Vaccinator Allocations

Background: Polio Vaccination Campaigns

Effective service delivery requires a mix of monitoring and monetary incentives. This is especially true in developing countries where monitoring systems and performance incentives are limited. Similar issues affect the system tasked with eradicating polio from Pakistan through door-to-door eradication vaccination campaigns.

The door-to-door campaigns are designed as follows: Teams of two vaccinators, typically a government worker and a volunteer from the community or from another department of the government go door to door providing oral polio drops. The makeup of the team fills two important needs: Helping government outreach workers engage with the community and creating ownership from within the community through the engagement of volunteers.

The project helped the City District Government of Lahore pilot the use of smartphones for monitoring and enabling the use of performance contracts for polio workers. The system created a paper-free monitoring system, which can be used to check the performance of each worker in a particular locality in real-time. We used inexpensive smartphones (about 30 USD each) that are readily available in Pakistan as monitoring devices. Each phone was preloaded with the application we designed as a monitoring system, allowing us to track the door-to-door vaccination activities of polio vaccinators. The smartphone system aggregates data in real-time on an online dashboard available to senior health administrators. The data from the monitoring system also provided the basis for pay-for-performance contracts.

In the status quo vaccination system, four or five roving teams of health workers work to vaccinate a neighborhood. They use paper maps like the one in (Figure 3), record their efforts on paper forms like the one in (Figure 4), and employ a protocol of writing in chalk on the walls of houses (Figure 5) to indicate to other teams that they have been there. For an independent monitor to check up on a team's efforts, she would have to physically follow their footpath through the neighborhood, reading the chalk marks. As a result, health workers often fail to achieve their targets, but report that they have.

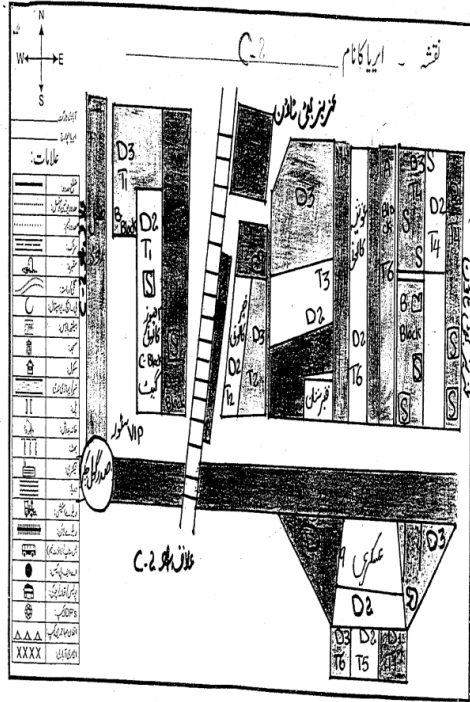


Figure 3: Example paper map that was previously used by health workers.

ایریا گنیر دیان شیٹ - ستمبر 1 - پوے سے چنوا کی گھر گھر مہریم تاریخ: 20 نومبر 2013

گھر نمبر	گھر کا نام	گھر کا پتہ	گھر کا مالک	گھر کا پتہ	گھر کا مالک	گھر کا پتہ	گھر کا مالک	گھر کا پتہ	گھر کا مالک
10	235	12	200	182	18	18	18	18	18
13	196	12	204	185	19	19	19	19	19
15	237	10	186	170	16	16	16	16	16
18	255	11	179	156	23	23	23	23	23
15	260	12	195	174	21	21	21	21	21
24	256	11	174	158	16	16	16	16	16

Figure 4: Example paper form that was previously used by health workers.

EPI 12/07/10 H6 بھئی والدین اپنے پانچ سال سے کم عمر بچوں کو قطرے دینے سے انکاری ہیں۔	R			EPI 12/07/10 H6 اس گھر میں 5 سال سے کم عمر تین بچوں اور دو مہمان بچوں کو ویکسین پلائی گئی۔	3/3+2			EPI 12/07/10 H6 تو سبھی پروگرام برائے حفاظتی ٹیکہ جات اور تاریخ		
EPI 12/07/10 H6 نمبر 3۔	T3			EPI 12/07/10 H6 پانچ سال سے کم عمر تین بچوں میں سے ایک بچے کو قطرے نہیں پلائے جاسکے۔	2/3			EPI 12/07/10 H6 گھر کا نمبر 6 ہے۔		
EPI 12/07/10 H6 نمبر 3 طرف گئی ہے۔	T3	→		EPI 12/07/10 H6 اس گھر میں 5 سال سے کم عمر کے بچے نہیں ہیں۔	0			EPI 12/07/10 H6 گھر میں 5 سال سے کم عمر بچوں کی تعداد 3 ہے اور 3 ہی کو ویکسین پلائی گئی۔	3/3	
EPI 12/07/10 H6				EPI 12/07/10 H6 گھر کو تالا لگا ہوا ہے۔	L					

Figure 5: Example of the protocol used for health workers to mark houses they've visited through writing in chalk on the house walls.

Smartphone System and Incentive Scheme

In this study, we tested a system in which, in order to get a financial bonus, the health workers used a smartphone to set their daily targets for the first two days of the vaccination campaign. The incentive structure expected workers to complete roughly 300 vaccinations over the campaign, and permitted them to move tasks from day 1 to day 2 in setting their target. The workers recorded their goals, along

with all attempted vaccinations, on a custom-designed app. Information about actual vaccinations was aggregated in a central server and sent to an online portal that visualized the data for policymakers' use, identifying whether or not the targets were being met (Figure 6). If a worker met the targets she set in advance for the first two days of the campaign, she got her bonus.

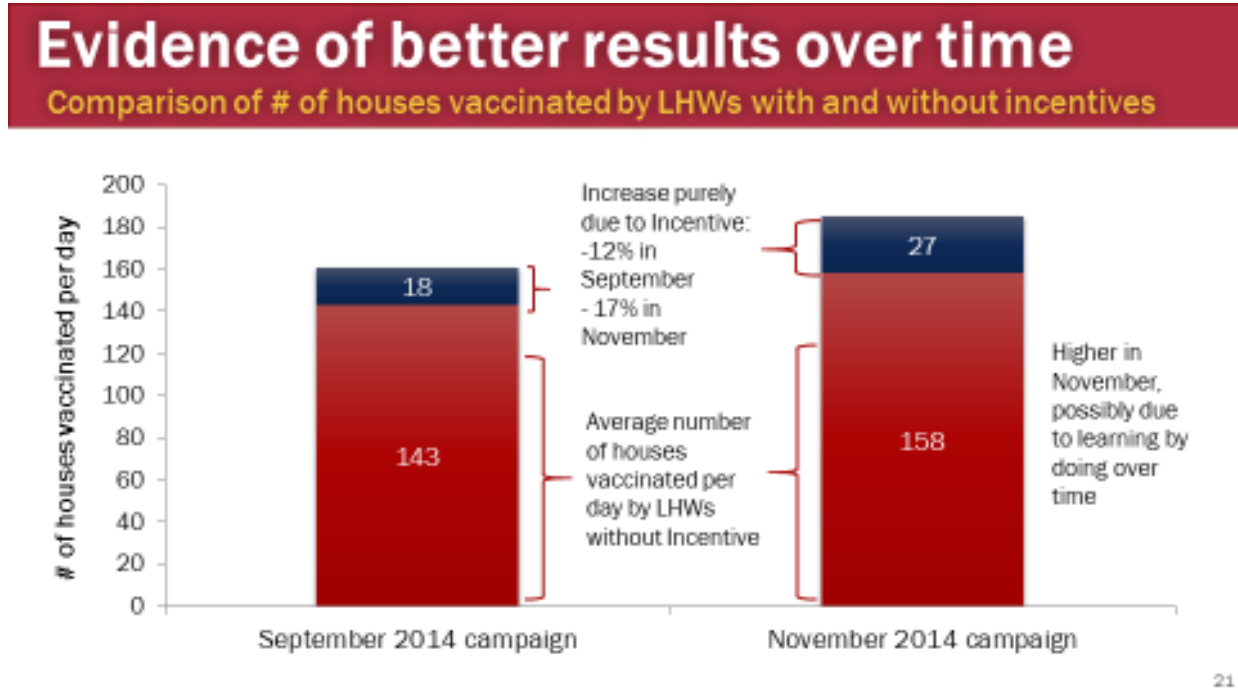


Figure 6: The main page of the online dashboard available to policymakers.

The system was tested in Allama Iqbal town and Nishtar town in Lahore over a course of three polio vaccination campaigns. Before every campaign the workers gathered for training sessions in centrally located health facilities to learn how to use the smartphone application. During the course of the campaign two help lines were always active to help workers if they faced any issues or errors. A total of 505 workers were selected to be part of the program and were further assigned to different arms of the experiment. The main bulk of workers, 336 out of 505, were assigned to a smartphone plus bonus scheme. Further 85 workers were provided with a smartphone but no bonus and the remaining 85 were treated as pure control group. All these assignments were made randomly. We estimate that providing a bonus (which was substantial) increased vaccinations by 16 percent per day, adding about 25 additional vaccination attempts per day over a base of 150 vaccination attempts. This difference is statistically significant at the 99 percent confidence level.

Our incentive system also provided a means of measuring time preference. In making her bid for how she prefers to work over the two days of the drive that we incentivized, the health worker let the app know how she prefers to allocate work over time. This allowed us to design optimal preference-specific incentives that we applied in the second round of the campaign. During this round, we had a control group that received an incentive scheme that was picked at random and not individualized to her

preferences. We find that providing individualized incentives substantially improves the intertemporal profile of effort. Given smartphone technology, this was not only easy and cheap to implement, but also potentially a game-changing way to incentivize public servants. The system gives a worker a degree of control over her activity (not to mention the chance for a bonus), and increases and evens out her effort in a way specific to her alone. All the while, it gives the manager the ability to track her and her coworkers' progress toward vaccination goals in real time. And it does this in a way that, we project, will cost around the same as existing systems, once implemented at scale.



21

Figure 7: Effect of Incentives

Data from the mobile phones show that a performance bonus encourages workers to complete 15 to 17 percent more vaccinations than workers that are just being monitored and not positively incentivized – this is in addition to the advantage of giving management better monitoring. Perhaps more importantly, the self-tailoring incentive structure succeeded in evening out efforts: the rise in vaccinations was achieved by spreading the effort over the two days of the drive. Furthermore, in the few cases where parents didn't allow their children to be vaccinated, the online dashboard allowed policymakers to identify those households for follow-up.

Conclusion

Using smartphones to collect data on the performance of polio vaccination campaign workers is a cheap and reliable way to improve the existing monitoring activities. The use of geo stamped and time stamped data provides a good picture of the extent of vaccination coverage and can easily identify blind spots in paper-based administrative divisions. It has also been useful in identifying vaccinator performance in an objectively verifiable manner. Indeed, observing the intertemporal allocation of vaccinations allowed us to infer the preferences of vaccinators and to design more effective individualized incentives. This is the first project, to our knowledge, that makes use of preference

measures to customize and to improve worker incentives. The monitoring of each individual provides an incentive to the worker to increase effort in itself. However the data can be easily used to administer a performance bonus scheme which we have shown to increase the performance by fifteen percent over and above what monitoring can achieve.

Right now a number of government programs in Pakistan, and especially in Punjab, are using mobile phone based monitoring systems. These systems are collecting data for monitoring by senior government officials but they can be easily adapted to construct incentive based pay structure for government employees. Through this experiment we have shown that such a scheme is possible for front line public health workers and can lead to significant improvement in performance at minimum cost.

The International Growth Centre (IGC) aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research.

Find out more about our work on our website www.theigc.org

For media or communications enquiries, please contact mail@theigc.org

Subscribe to our newsletter and topic updates www.theigc.org/newsletter

Follow us on Twitter [@the_igc](https://twitter.com/the_igc)

Contact us
International Growth Centre,
London School of Economic and Political Science,
Houghton Street,
London WC2A 2AE

When citing this paper, please use the title and the reference number listed on the front page.

IGC

**International
Growth Centre**

DIRECTED BY



FUNDED BY



Designed by soapbox.co.uk