

Crowdsourcing government accountability: Experimental evidence from Pakistan*

Arman Rezaee[†] Ali Hasanain[‡] Yasir Khan[§]

December 12, 2015

Abstract

We develop and implement a novel, mobile phone-based information clearinghouse, and experimentally evaluate its ability to overcome information asymmetries and improve public service delivery to farmers in Punjab, Pakistan. Like many crowdsourcing websites, our clearinghouse collects and disseminates ratings—here, on the success of government veterinarians in inseminating livestock. We find that, compared to control, farmers receiving ratings enjoy 27 percent higher insemination success. This effect is entirely due to increased veterinarian effort, rather than farmers switching veterinarians. Treatment farmers are also 33 percent more likely to return to a government veterinarian rather than seeking a private provider. These results suggest large welfare benefits from a low-cost information intervention, which holds out hope for improved government accountability for the poor using basic mobile technology.

* *Authors' Note:* We thank Eli Berman, Michael Callen, Julie Cullen, Clark Gibson, Craig McIntosh, Edward Miguel, Karthik Muralidharan, and faculty at UC San Diego for their support at all stages of this project. We also thank Saad Gulzar, the International Growth Centre Pakistan office, the Punjab Livestock and Dairy Development Department, and the World Bank Pakistan office for help designing and implementing the project. Excellent research assistance was provided by Amanullah Haneef, Umair Khawaja, Zia Mehmood, and Zarak Sohail. We thank Sarojini Hirshleifer and Janna Rezaee for their excellent feedback. This research was supported by the University of California Office of the President UC Lab Fees Research Program Grant ID No. 23855, by funding from the Abdul Latif Jameel Poverty Action Lab and the Center for Effective Global Action through the Agricultural Technology Adoption Initiative, and by the International Growth Centre. Support for Rezaee's time was provided by AFOSR # FA9550-09-1-0314 and ONR # N00014-14-1-0843.

[†]University of California, San Diego. email: arezaee@ucsd.edu

[‡]Lahore University of Management Sciences. email: ali.hasanain@gmail.com

[§]University of California, Berkeley. email: yasir.khan@berkeley.edu

1 Introduction

Asymmetric information between citizen principals and service-providing agents often leads to sub-optimal outcomes for the rural poor across the developing world (World Bank, 2004; Wild et al., 2012). In the case of *government agents*, asymmetric information has led to corruption in elected officials (Ferraz and Finan, 2011), waste in government processes (Bandiera et al., 2009), leakage between public service allocations and expenditures (Reinikka and Svensson, 2004), and more generally poor public service delivery across sectors, countries, and even continents (Chaudhury et al., 2006). In the case of *private agents*, asymmetric information has led to inefficient market allocations and rent capture at the expense of consumers (Jensen, 2007; Svensson and Yanagizawa, 2009; Aker, 2010).

Monitoring can decrease asymmetric information, but it is particularly costly to implement monitoring schemes in rural developing settings. This is because poor infrastructure makes information collection and transmission expensive in these contexts. In addition, research shows monitoring may not be effective without complimentary financial incentives (Duflo et al., 2012) and its effects attenuate as agents find alternative strategies to pursue rents (Olken and Pande, 2012).

Information clearinghouses, such as yelp.com, angieslist.com, and amazon.com, decrease asymmetries inexpensively. These crowdsourcing websites collect, aggregate, and disseminate masses of ratings at costs much lower than traditional reviewers such as the New York Times, though to date, their application has been limited to commercial settings. Furthermore, such sites have yet to take hold in the rural developing world, characterized by thin markets, low literacy rates, and 2G wireless networks.

We design and implement an information clearinghouse to reduce government agent shirking in a context fraught with asymmetric information: agricultural service provision in the developing world. Our clearinghouse provides citizens in rural Punjab, Pakistan with government veterinarians' success rates at artificially inseminating livestock, an objective measure of veterinarian effort. It gathers and disseminates locally relevant information from a large

base of farmers automatically, in real time, using a call center.

Our clearinghouse model stands in contrast to government monitoring schemes that provide information to agents' superiors, relying on the "long route" of accountability in which citizens must influence policymakers to improve service provision (Callen et al., 2015). It approaches the problem more directly; it strengthens the "short route" of accountability by increasing citizens' direct power over government agents (World Bank, 2004).

And our clearinghouse strengthens government agent accountability in providing a service that is important for the livelihood of people across the developing world—renewing livestock through artificial insemination (AI). Livestock agriculture accounts for 12 percent of GDP in Pakistan, and is a key growth sector for the rural poor (Pakistan Economic Survey 2013-14). AI is crucial to renewing livestock. Most households only keep female cows because of the dual advantage of producing milk and calves, both of which require cows be pregnant. But government veterinarian shirking leads to AI success rates lower than what is possible given the technology, costing farmers potential income.

We evaluate this clearinghouse using a randomized controlled trial. Using data generated by the clearinghouse, we find that farmers treated with information on local government veterinarians' AI success rates have a 27 percent higher AI success rate than controls when they subsequently return for government services. In addition, treatment farmers are 33 percent more likely to return to a government veterinarian for AI rather than to seek a private provider.

Multiple mechanisms could explain this treatment effect on AI success rates, including treatment farmers selecting better veterinarians and/or veterinarians exerting more effort for treatment farmers. Several of our results suggest the latter—that government agents work harder when the ratings system is in place. First and foremost, treatment farmers are no more likely than control farmers to switch veterinarians after treatment. Thus the effect cannot be driven by farmers simply switching to the 'best vet' in terms of AI success and/or

price. Second, treatment farmers pay lower prices after treatment.¹ While farmers may be able to improve AI success rates through their behavior alone, a change in prices requires a change in veterinarian behavior.²

Our estimated treatment effects on AI success are potentially subject to both selection and reporting biases since they use data from the clearinghouse. In this data, we only observe farmers who return for government AI after treatment and not those who switch to private providers, as these are not part of our clearinghouse. Returning farmers must then also choose to answer the phone and to report AI success to the clearinghouse. Importantly, we find analogous results using a representative in-person survey not subject to selection or reporting biases but with lower precision. We find an overall 26 percent treatment effect in this representative sample, which averages a treatment effect of 83 percent for farmers that select back into government AI after treatment and a treatment effect of 4 percent for attritors.³

Our results fit the context—artificial insemination requires unobserved effort in at least two ways. First, veterinarians must keep semen straws properly frozen in liquid nitrogen canisters from the time when they are delivered to AI centers until right before insemination. Second, veterinarians must then precisely insert these straws during insemination. At the same time, farmers cannot infer a veterinarian’s effort from outcomes alone. Even when executed properly, AI will not be successful 100 percent of the time, and success rates may vary based on animal health and nutrition.

In addition, while government veterinarians collect a salary and are protected from punishment for poor performance, they are legally allowed to charge a ‘show-up’ fee to farmers

¹Note the estimated treatment effect on log AI price has a p-value of 0.12 in our primary specification.

²It is also possible that learning something about AI success rates in general causes farmers to take better care of their livestock and that this in turn increases AI success rates. However, we find that treatment farmers who subsequently switch to private providers do not have increased AI success rates. If our treatment effects were driven by changes in livestock care, we would expect to see effects regardless of which provider farmers subsequently choose.

³Note the estimated overall treatment effect has a p-value of 0.12 in our primary specification. The treatment effect for farmers that select back into government AI, analogous to the AI success rate result using clearinghouse data, is significant at 5 percent.

for their services on top of the fixed cost of AI. Therefore, in response to their low unobserved effort being revealed to farmers, government veterinarians may prefer to exert more effort and continue to collect a fee than to lose a customer. In other words, they may internalize the benefits of their marginal effort, a characteristic more common to private than public markets.

In a standard agency model with a stochastic outcome and inability to contract on this outcome, either unobserved agent effort (moral hazard) or unobserved inherent agent ability (adverse selection) a priori predicts both sub-optimal outcomes at baseline and that outcomes will improve as unobserved effort is revealed. We find both of these predictions to be true. However, because treatment farmers see increased AI success rates without switching veterinarians, our results rule out a pure adverse selection model and support one of moral hazard.

Several additional results from our representative in-person survey support a standard agency model. First, we find that farmers' baseline expectations about the average AI success rate of their own government veterinarians do not correlate with actual average AI success rates. This suggests the existence of asymmetric information ex ante. Second, treatment causes farmers' endline expectations about their veterinarian to become strongly correlated with the truth. This suggests that farmers indeed update their beliefs. Third, farmers who received more negative information relative to their expectations saw larger treatment effects. This suggests that the amount of information farmers receive determines their benefit.

More generally, the market for AI in rural Punjab is one in which informationally disadvantaged consumers pay more than the marginal cost of AI provision through two channels—prices and veterinarian effort. In this market, treatment-induced veterinarian effort implies consumer welfare gains so long as there are no compensating price increases or negative spillovers onto control farmers, which we do not find. Furthermore, this implies overall social welfare gains so long as the cost to veterinarians' increased effort is not too great.⁴

⁴We do not believe the marginal cost to veterinarians' increased effort induced by treatment to be very large in this setting, as travel costs are paid either way. Government veterinarians also do not spend any

Our study differs from previous evaluations of the effect of information on markets with only a price channel, where changes in prices are pure transfers and any social welfare gains must come from increased market efficiency (Jensen, 2007; Svensson and Yanagizawa, 2009; Aker, 2010). Many other markets have multiple channels for rents and thus expect similar social welfare gains, including education (Andrabi et al., 2014), elections (Ferraz and Finan, 2011), and markets for private restaurants (Jin and Leslie, 2003).

In such related studies, with the exception of previous clearinghouses evaluated in Fafchamps and Minten (2012) and Mitra et al. (2014) (in both cases, the authors find no treatment effects), interventions to reduce asymmetric information are costly, static, and/or do not lead to clear social welfare gains. Our clearinghouse, on the other hand, relies on crowdsourcing technology that is cost-effective, self-sustaining, and scalable. Conservative estimates suggest a 27 percent higher AI success rate translates into nearly an additional half of one month’s median income per AI provided, a 300 percent return on the cost of the intervention. These effects hold out hope for improved government accountability as cellular technology improves and becomes cheaper.

The paper proceeds as follows: Section 2 provides background on our study district and government AI service provision there, Section 3 outlines our research design, including providing more information on the clearinghouse and the randomized controlled trial embedded within it, Section 4 provides results, Section 5 discusses the interpretation and social welfare implications of these results, and Section 6 concludes.

2 Background

2.1 The market for AI in Sahiwal, Punjab, Pakistan

We implemented our clearinghouse in the Sahiwal district of Punjab province, Pakistan.

While we selected Sahiwal based on several logistical constraints, we view it as representative

more time visiting treatment farmers. Any costs must be in terms of concentration, etc.

of the whole of Punjab, and of similar agricultural districts across the country, though with a slightly higher prevalence of livestock.⁵

Sahiwal has a vibrant market for artificial insemination for at least two reasons. First, almost all livestock in the district are female. Second, artificial insemination decreases the costs of selectively breeding to increase milk yields, as only the semen from high-yielding bulls needs to be transported and not the bulls themselves.⁶

The government is the largest supplier in this market, offering low-cost AI services by veterinarians who have required AI training. The official cost of government AI is 50 PKR per insemination (approximately 0.5 USD), but government veterinarians are legally allowed to charge a ‘show-up’ fee to cover the cost of their gasoline, as well as any other costs or risks. This results in average costs of approximately 200 PKR per visit. The government has 92 one-room artificial insemination centers or veterinary offices spread throughout the district, staffed by roughly 70 active veterinarians.⁷ These veterinarians’ sole job is to provide artificial insemination.⁸

The only other organized supplier in this market is Nestle, but they have far fewer active veterinarians providing AI services in Sahiwal. Most private veterinarians are self-employed, buying semen from large private suppliers and providing AI services without any training. At baseline, these private veterinarians collectively provide approximately 57 percent of AI services across Sahiwal, with government veterinarians making up the remainder.

⁵According to the 2010 Punjab’s Multiple Indicator Cluster Survey, households in Sahiwal on average have 1.4 fewer acres of agricultural land and .24 more cattle than households in other districts in Punjab. Sahiwal’s average wealth, labor force participation rates, and child mortality rates are representative of Punjab.

⁶The provincial government selectively breeds livestock in two main centers in Punjab. It then distributes the semen produced to government veterinarians across the province, including in Sahiwal.

⁷Throughout our study period, a total of 77 veterinarians were active in Sahiwal for any amount of time. Only a handful of veterinarians transferred in or out of Sahiwal.

⁸In some cases they may provide vaccinations during AI service provision, but this occurs very rarely. A smaller, distinct group of veterinarians care for sick animals.

2.2 Asymmetric information in the market for AI

On a single visit, a farmer can never fully observe veterinarian effort. However, even before our intervention, farmers could have decreased asymmetries by aggregating information about their veterinarians' success rates across visits and across households. Our data suggests that they do not. At baseline, farmers' estimates of their current government veterinarian's AI success rate are uncorrelated with the truth. This can be seen in Figure 6, Panel A.

This asymmetric information contributes to AI success rates that are lower than what veterinarians can achieve. At baseline, AI success rates average approximately 70 percent, while success rates of 85-90 percent are possible with the training and equipment in Sahiwal.

3 Research design

3.1 The clearinghouse

To measure veterinarian prices and effort and to subsequently disseminate that information to consumers, we developed a novel cellular-based information clearinghouse. Figure 1 diagrams the four components of our intervention.

Pre-treatment: During the study, government veterinarians in Sahiwal were required to collect real time information on all AI service provisions using an Android smartphone equipped with an Open Data Kit-based application.⁹ The data was immediately sent to the clearinghouse. We denote this data collection as $t = 0$ in Figure 1.

Data collection and aggregation: Each service provision generated two subsequent phone calls. First, one day later (denoted $t = +1$ day in Figure 1), a representative from the clearinghouse call center called the farmer to verify that the veterinarian had provided service and to ask what price he had charged. Then, sixty days later ($t = +60$ days), they called again to ask if the artificially inseminated livestock were pregnant. The clearinghouse continuously aggregated this price and AI success rate data for each veterinarian.

⁹In practice, veterinarians did not always comply. See Section 4.3 for discussion.

Treatment: The clearinghouse collected and aggregated information from January to September, 2014. Treatment began in October 2014, once we had sufficient data on veterinarians to have meaningful measures of price and AI success rates. Treatment took place during the second call (at $t = +60$). Only this time a randomized group of farmers was provided information on local veterinarians' prices and AI success rates. The uninformed farmers became the control group.

Post-treatment: The clearinghouse allowed us to link farmers over time, so we observe post-treatment government AI provision for both treatment and control farmers (if they return; Figure 1 depicts the return of a treatment farmer but not a control farmer). These post-treatment observations also generate two follow-up phone calls.¹⁰

3.2 Information provision

In the treatment group, the clearinghouse representative presented farmers with information on the top three veterinarians within three kilometers of their household in terms of AI success rates for cows, and the top three veterinarians in terms of AI success rates for buffalo.¹¹

We gave treatment farmers AI success rates for these three to six veterinarians, and the average price of the service, during the second follow-up call.¹² The clearinghouse then sent a follow-up SMS with the same information. If farmers requested it, we also gave them veterinarians' phone numbers, information on average farmer-reported satisfaction with veterinarians on a 1-5 scale, and information on any other veterinarian in our system.

The clearinghouse administered treatment at the farmer level through a coin-flip stratified on the nearest government veterinary clinic to a farmer's household. Farmers who returned for service provision after treatment assignment retained their initial assignment. Note that

¹⁰Note, however, that treatment selection is carried forward in time. See Section 3.2.

¹¹When we had fewer than 25 observations for a veterinarian, we weighted success by $\sqrt{n}/5$, where n was the number of observations. By design, almost every veterinarian had more than 25 observations each for cows and buffalo once the treatment began. The exceptions were two veterinarians hired after our treatment began in October 2014.

¹²There can be overlap in the most successful veterinarians in terms of cows and buffalo.

treatment occurred at a different time for each farmer, 60 days after they first entered our clearinghouse. This means that the post-treatment period differs for each farmer.¹³

3.3 Representative survey

In addition to the clearinghouse data, we independently surveyed a representative sample of farmers from across Sahiwal. We did so because the clearinghouse sample is not representative: to enter the clearinghouse, farmers first selected government AI over private, then their government veterinarian complied to record their service provision, then we were able to reach them on the phone to collect price and AI success information; and then we only observed post-treatment outcomes for clearinghouse farmers who subsequently returned to a government veterinarian for AI (as opposed to a private provider).

For these surveys, we sampled 90 of Sahiwal’s approximately 500 villages from a district village census.¹⁴ Within each village, we selected ten households using the Expanded Program on Immunization (EPI) cluster sampling method (Henderson and Sundaresan, 1982). We selected households that reported owning at least two livestock (cows and/or buffalo) and having regular access to a cellular phone.

We manually entered survey farmers’ phone numbers into our clearinghouse to generate treatment or control follow-up calls. These calls were near identical to those to farmers that entered our clearinghouse on their own, and the treatment information provision component was identical.¹⁵

Sample villages can be seen in Figure A.1. Figure 2 presents a timeline of the clearinghouse and survey data collection. The baseline survey occurred prior to our clearinghouse

¹³Unfortunately, the coin used for randomization was shaved, due to a glitch in the clearinghouse algorithm. This resulted in 52 percent of farmers being treated. However, the probability of treatment remained fixed across farmers across time.

¹⁴We stratified the sample by whether or not a government veterinarian center was in each village and on whether each village bordered an irrigation canal. The sample is representative of Sahiwal in terms of: area, settled area, cultivated area, area of wheat, rice, cotton, sugar cane, pulses, orchards, and vegetables, having a river, distance to the nearest veterinarian center, number of livestock in the village, literacy rates, religion, age, and standard wealth index characteristics. Results available upon request.

¹⁵The only difference was that instead of asking questions about a specific recorded service provision from 60 days ago as is the case with clearinghouse calls, we asked about farmers’ last AI service.

implementation, and the endline survey occurred immediately prior to the clearinghouse being shut down.¹⁶

Tables 1, 2, and A.1, report the balance of our clearinghouse and representative survey samples between treatment and control farmers.

3.4 Empirical specifications

We use the following specification for our primary analysis:

$$outcome_{ft} = \alpha + \beta T_f + \Gamma_{ft} + \epsilon_{ft} \quad (1)$$

where $outcome_{ft}$ is an outcome for farmer f from post-treatment AI visit t . T_f is a treatment indicator, Γ_{ft} are treatment strata and other baseline controls to improve precision, and ϵ_{ft} is an idiosyncratic error term. While we administered treatment at the farmer level, treatment information provision was localized at the village-cluster level. We cluster standard errors at this village-cluster level to allow for correlation in outcomes between farmers in the same village-cluster. Village-clusters are groups of villages that share the same government veterinarians within a three kilometer radius. There are roughly two villages per village-cluster.

We define post-treatment for control farmers as all observations after the phone call in which they were selected into control rather than treatment. This ensures balance in the length of the post period between treatment and control farmers.

We have four primary outcomes:

Switched veterinarians_{ft}: a dummy variable equal to one if a farmer’s veterinarian at visit t differed from the farmer’s veterinarian at visit $t - 1$.

Log price_{ft}: the log price paid for AI at visit t , as reported by the farmer when called

¹⁶We conducted a purely technical survey at midline to collect new phone numbers for those households that changed numbers between the baseline and the first round of treatment phone calls. This allowed us to treat as many independently surveyed farmers as possible.

the next day.

AI success rate $_{ft}$: a dummy for the success of the AI provided at visit t , as reported by the farmer when called 60 days later.

Returned $_f$: a dummy variable equal to one if a farmer returned for government AI after treatment by the end of the project.¹⁷

4 Results

In this section, we present results. First, we present treatment effects using our representative sample (Section 4.1) and our clearinghouse sample (Section 4.2). Second, we show that treatment does not induce veterinarian reporting bias (Section 4.3) or farmer reporting or selection biases (Section 4.4) in the clearinghouse sample. Third, we explore the primary mechanism for our treatment effects, decreased moral hazard or increased effort by veterinarians for the treated, through heterogeneity analyses (Sections 4.5 and 4.6).

4.1 Treatment effects—representative sample

Table 3 presents treatment effects using our representative sample. We report first effects on price. Column (3) shows a statistically insignificant price reduction for the entire sample, which remains insignificant if we disaggregate into the subsamples of farmers who either returned to government AI (1) or attritted to private providers (2) after treatment. In column (4), we find that treatment farmers who return to government AI have a 47 percentage point, or 83 percent, higher AI success rate. In contrast, column (5) reports an insignificant treatment effect on AI success for farmers who attritted, indicating that treatment does not induce farmers to seek out a better private provider. In column (6) we find that, while it is not quite significant, overall AI success rates are large and positive even when including

¹⁷We pre-specified our empirical specification in our pre-analysis plan, registered in the AEA RCT registry. We did not pre-specify *Returned* $_f$. We did pre-specify *Switched veterinarians* $_{ft}$, *Log price* $_{ft}$, and *AI success rate* $_{ft}$. We pre-specified the latter two outcomes conditional on veterinarian switching, but we have made them unconditional since we do not observe veterinarian switching.

those farmers that attrited: treatment farmers have a 17 percentage point, or 26 percent, higher AI success rates after treatment.¹⁸

While these results are not subject to reporting or selection biases, the size of our representative sample allows for less precision than with our clearinghouse sample, which we will now turn to.

4.2 Treatment effects—clearinghouse sample

Table 4 presents treatment effects of information provision on our primary outcomes using the clearinghouse sample. In column (1), treatment farmers are 3.2 percentage points, or 33 percent, more likely than control farmers to return for government AI after treatment.¹⁹ As a visualization, we present an added-variable plot of this result in Figure 3.

In columns (2) through (4), we present effects on those farmers that return after treatment selection. In columns (2) and (3) we find that there are no statistically significant treatment effects on veterinarian switching or on log prices, though the coefficient on log price is nearly significant with a p-value of 0.12. In column (4), we find that treatment farmers have a 17 percentage point, or 27 percent, higher AI success rate after treatment.

This treatment effect on AI success rates is substantially smaller in magnitude than the analogous 47 percentage point treatment effect we report in Table 3, column (4) in the representative sample. However, we cannot reject that the effect in the representative sample is equal to that in the clearinghouse sample.

In Figure 4, we present the treatment effect on AI success rates in real time (as opposed to in pre/post time, where post begins at a different time for each farmer). The top panel illustrates that treatment farmers have higher AI success rates consistently over time, while the bottom panel traces the size and significance of this treatment effect over the post period.

¹⁸The p-value of this estimate is 0.12.

¹⁹The low overall return rate is likely because the average time for farmers between treatment and the end of our study period is five months and AI is only required roughly once a year per animal. As we see in Table 5 as well, only 30 percent of return visits were recorded by veterinarians, so even in five months the true return rate is likely 40 to 50 percent.

These results suggest that any information spillovers between treatment and control farmers are either small or fixed throughout time. The latter is unlikely given the rolling nature of treatment. If anything, there is a small bump up in AI success rates for control farmers in the first month of the treatment, which suggests positive information spillovers. This would attenuate our results. The figure also suggests that there are no negative spillovers onto control farmers from veterinarian effort constraints.

The most likely cause of the across-the-board downward trend in AI success rates beginning in March 2014 is changes in leadership of the Punjab Livestock and Dairy Development Department at both the provincial and Sahiwal district levels—the new regime was less focused on veterinarian performance than the last had been.

In Figure 5, we present the treatment effect on log AI prices in real time. We find that the same visual trends hold for prices, and that when we bootstrap standard errors, the treatment effect is significant in six of eight months.

We reproduce our primary treatment effects on our representative survey sample, selecting on returning for government AI after treatment, in Table A.2. The point estimates are of a similar magnitude.²⁰

4.3 Treatment does not induce a veterinarian reporting bias

In order to believe the internal validity of our clearinghouse sample, it is important to note in Table 5 that treatment does not induce a reporting bias among government veterinarians. We measure reporting bias by comparing farmer reports of service provision from our representative survey with entries in the clearinghouse. While government veterinarians only comply by reporting AI approximately 30 percent of the time, they are equally likely to report for treatment and control farmers.

²⁰Note that the mean return rate of control farmers is higher in this sample, but not three times that of the clearinghouse sample. This is consistent with the fact that we do not rely on veterinarian reporting for this data. Also, these farmers had less time after treatment to return to our sample on average.

4.4 Ruling out farmer selection and reporting biases in the clearinghouse sample

For the same reason, we must also rule out farmer selection and reporting biases. Our estimates would include a farmer selection bias if farmers that would otherwise see higher success rates are those that select back into government AI after treatment. Our estimates would include a farmer reporting bias if treatment farmers are more or less likely to answer the phone when we call to ask about AI success.

We have already presented evidence against both farmer selection and response biases in Table 3, column (6). Accounting for attriters removes possible selection bias. In addition, the representative survey had a successful follow-up rate of 96 percent with no differential attrition, which removes possible response bias.

As an additional check for farmer selection bias, in Table 6 we show balance on all measured pre-treatment outcomes, including AI success rates, between returning treatment and control farmers in the clearinghouse data. While this does not rule out selection on unobservables, we believe that it does rule out the most likely type of selection that could drive such a large increase in AI success rates in our post-treatment sample—selection back into government AI by farmers who have younger, healthier livestock more likely to get pregnant. If this selection were occurring, such younger and healthier animals should have then been more likely to get pregnant in the pre-periods as well, yet we do not see this. We also do not see any differences in past prices paid, past veterinarian switching, or other administrative variables.

4.5 Treatment effects by government veterinarian rank

In order to explore the mechanism for our treatment effects, we present a series of heterogeneous treatment results that support a standard moral hazard model.

First, in Table 7, we present treatment effects for two important sub-populations, sepa-

rated according to the ranking of the last government veterinarian who served them—those for whom this veterinarian was ranked in the top three in their village-cluster, and those for whom he was not. This aligns with those veterinarians on whom treatment farmers received information regarding AI success rate and price. We separate control farmers based on what they would have been told, had they been treated.²¹

We find suggestive evidence that our main results are localized to farmers whose past veterinarian was not ranked in the top three in their area at the time of treatment.²² Again, this is in line with a standard moral hazard model. The more a farmer learns a veterinarian can increase unobserved effort, the more s/he is able to then bargain away rents from the veterinarian.²³

Perhaps the most surprising result in Table 7 is that farmers whose past veterinarian was not ranked in the top three are more likely to return. To investigate this, we show in Table A.3 that farmers in Table 7 Panel B tend to live almost twice as far away from their closest veterinary center.²⁴ This is consistent with farmers living in more remote areas settling for lower effort veterinarians because of higher switching costs. And it is exactly these farmers with higher switching costs that receive the largest benefits from treatment.

²¹Note that at the beginning of our treatment phone calls we verify farmers' villages as they were automatically generated by GPS. This verification is not done with control farmers. To avoid measurement error correlated with treatment, we separate treatment farmers based on what they would have been told had we not verified their village. This hypothetical information set correlates with the truth at over 90 percent.

²²These results are suggestive because, while the point estimates are qualitatively different, we cannot reject this difference with significance.

²³We should also expect heterogeneous treatment effects based on whether or not a farmer's past government veterinarian was ranked top in their village-cluster versus second best, or second best versus third best, etc. We do not have power to accurately detect these differences, but results are consistent with the same simple model. Results available upon request.

²⁴In addition, these farmers have more buffalo. We control for baseline means of both of these variables in Table 7.

4.6 Results using farmer expectations from the representative survey sample

If we are to believe that our results are in line with a standard moral hazard model, we should expect the level of asymmetric information between farmers and veterinarians at baseline to be important. We present three results in this vein, in this case using farmers' stated expectations. These expectations come from our representative survey sample, in which we asked farmers what they expect the average AI success rate of their past veterinarians to be.

In Figure 6, we compare farmers' expected average AI success rate for their veterinarian prior to treatment with the actual average AI success rate of that veterinarian. Actual average AI success rates are drawn from our clearinghouse data prior to October 2014 when treatment calls began.

Our first result is in Panel A of the figure—at baseline there is no correlation between farmer expectations and the true AI success rate of their veterinarian. This suggests there is room to improve service delivery by relieving asymmetric information.

Our second result is in Panel B of the figure—at endline there is a strong correlation between expectations and the truth for treatment farmers. In other words, treatment changes expectations. This is a crucial test that information was passed on through our treatment. Panel C presents the endline correlation for control farmers—while much smaller than with treatment farmers and insignificant, there is a positive correlation. This suggests potential information spillovers between treatment and control farmers, which would attenuate our treatment results above.

Point estimates for these two results are reported in Table 8. The null hypothesis that the coefficients in columns (2) and (3) are equal is almost rejected, with a p-value of 0.115.

Third, using farmer expectations we can also separate treatment effects by the level of asymmetric information between farmers and veterinarians at baseline. To do so, we difference farmers' expected average AI success rate with the truth. We then split our sample according to whether farmers had above or below the median in this difference. Positive values

in this difference occur when farmers are told that their veterinarian is better than they expected; negative values occur when farmers are told their veterinarian is worse than they expected. The median is .012.

Table 9 presents results from this heterogeneity analysis. We find that, as with treatment effects by government veterinarian rank, the more unexpectedly negative the information a farmer receives about their veterinarian, the more s/he is able to then bargain away rents from the veterinarian.

5 Discussion

5.1 Interpretation: Unobserved effort or inherent ability?

Several results suggest that the treatment effect on AI success rates is entirely due to increased veterinarian effort for the treated. To illustrate this, we can walk through the process by which farmers select a veterinarian and negotiate prices and effort. First, farmers decide whether to get AI at all when a cow is in heat. Next, they decide whether to stick with their previous veterinarian. If farmers switch, they then decide whether to call a government or private veterinarian. Finally, they decide how to engage with this veterinarian in pre-visit negotiations over the phone as well as during the AI visit (and veterinarians have to decide how to respond).

In our setting, farmers almost always choose to inseminate their livestock in heat, so we would not expect any changes in this decision. Next, we show in Table 4 that treatment farmers are no more likely than control farmers to switch veterinarians after treatment. Thus the treatment effect cannot be driven by farmers simply switching to the ‘best vet’.

We do see changes in whether farmers call a government or private veterinarian, however. Importantly, we show in Table 3 that treatment farmers who subsequently switch to private providers do not have increased AI success rates. If our treatment effect is driven by changes in farmer behavior towards their livestock, we would expect effects regardless of

which veterinarian the farmer selects after treatment. The same argument can be applied to the results from Section 4.5. If our treatment effect is driven by changes in farmer behavior, farmers' past veterinarian ranking should not matter.

Thus, we can turn to the final part of the decision process as the likely mechanism—farmers' engagement with veterinarians. Our results are consistent with farmers using the information we provide to them to negotiate reductions in government veterinarians' informational rents through higher effort and lower prices. And while farmers may be able to improve AI success rates through their behavior alone, the decrease in prices that we find requires a change in veterinarian behavior.

If we are to view increased veterinarian effort as the driver of our results, then that effort must be easily varied across visits. Anecdotes suggest that this is true. One commonly cited example of low veterinarian effort is the way in which veterinarians treat semen straws. As mentioned above, the provincial government delivers these straws to veterinary centers in liquid nitrogen canisters, and they must be kept frozen until just before use. Veterinarians sometimes take straws out before leaving on a visit rather than transporting the canister to the farm. This likely results in the semen spoiling, though the veterinarian still performs AI and charges the farmer. And because farmers call veterinarians before AI to negotiate a time and price, treatment farmers could pressure them to take better care transporting semen. Veterinarians would have to exert more effort but farmers would likely still pay them positive rents rather than having to pay the cost to find a new veterinarian.

5.2 Social welfare implications

To understand the social welfare implications of this intervention, we consider benefits and costs to farmers and to veterinarians as well as the cost of the intervention itself.²⁵

Benefit to farmers: if the treatment effect of 27 percent on AI success rates translates into just three percent more calves born per year per farmer (i.e., if farmers with a failed

²⁵We do not consider changes in price as such is a transfer with no net social welfare implications.

AI attempt are able to successfully impregnate their animal two months later), and the expected value of a calf is roughly 107,500 PKR (approximately 1075 USD) at the market, then treatment farmers would earn an additional 3,225 PKR (32 USD) per year, equal to nearly half of one month's median income.²⁶ This is a conservative estimate. It does not count the additional net value of two months of milk nor the cumulative net present value effect of an increased future stream of livestock.

Cost to farmers: we showed that farmer treatment effects are not due to changes in farmer behavior, we do not consider there to be costs to farmers of this intervention.

Benefit to veterinarians: farmers do not switch veterinarians more as a result of treatment, which suggests no change in veterinarian market shares that could impact social welfare. However, treatment farmers are more likely to return for government AI. Thus, if anything, government veterinarians benefit from this intervention. This would be at the cost of private veterinarians, however, so we will not consider it.

Cost to veterinarians: we do not believe the marginal cost to veterinarians' increased effort induced by treatment to be very large in this setting, as travel costs are paid either way. Government veterinarians also do not spend any more time visiting treatment farmers. Any costs must be in terms of concentration, etc.

Cost of the intervention: including one-time fixed costs to develop our clearinghouse technology, this intervention cost approximately 50,000 USD to reach over 6,000 farmers for treatment or control calls, or approximately 8 USD per farmer.

Adding it up, we find benefits of 32 USD per farmer from an intervention that cost 8 USD per farmer. This suggests a large, 300 percent return.

²⁶This calf value is the average of male and female calf prices reported at <http://www.pakdairyinfo.com/feasibility.htm>, accessed 10/8/2015. The monthly median income of households in Paksitan, according to the World Bank, is 73.26 USD per month, accessed 10/8/2015.

6 Conclusion

In this paper, we present results from the randomized controlled trial of a novel solution to a common government accountability failure: shirking by government agents in a setting of asymmetric information. Our solution is novel not only in that it leverages the cost-effective, self-sustaining nature of crowdsourcing to help the poorest, but also in that it does so in a tough setting. In rural Punjab, the market for artificial insemination is thin, literacy rates are low, and cellular networks are very limited—yet we were able to employ an information clearinghouse with success.

The very fact that our clearinghouse was successful purely through providing information confirms the existence of asymmetric information in this setting. And the fact that veterinarians respond with increased effort confirms that this asymmetric information is about unobserved effort. While these confirmations are neither novel nor heartening in and of themselves, they allow us to fit the livestock sector in Punjab into a context that is much more general. Moral hazard has been documented in numerous sectors, public and private, across the developing world. We might expect our clearinghouse to help citizens in any of these sectors, so long as they answer the phone.

And given the low cost of our clearinghouse, we might expect similarly large returns in other sectors. Conservative estimates suggest a 300 percent return to farmers on the cost of the intervention. This is driven by a 27 percent increase in AI success rates for treatment farmers. In other words, thousands of poor, rural Pakistanis who were treated are now more likely to have milk to drink and calves to raise or to sell for substantial income. This is heartening.

As a testament to the scalability of our clearinghouse, we have already begun conversations within the Livestock Department about expanding the program to all of Punjab. This would require no additional fixed costs and less than proportional marginal costs. Across contexts, we are already experimenting with scaling our information clearinghouse to relieve asymmetric information between citizens and pollution regulators in Punjab. We hope to

learn how crowdsourcing can work in a regulatory rather than a market environment, and for public rather than private goods.

Some of these effects of improving the flow of information depend on how connected the population is, and the price of connectivity. In another project, we are experimenting with the placing of cellular towers to understand how economic, social, and political outcomes are impacted by across-the-board decreases in the cost of transmitting information.

We hope this paper and other new studies will improve our understanding of how technology can be leveraged to improve the feasibility and impact of already tried-and-true interventions, such as monitoring to reduce asymmetric information. As cellular networks improve and as technology to collect, aggregate, and disseminate information advances, our results suggest we may see improved outcomes for citizens across the rural developing world.

References

- Aker, Jenny C**, “Information from markets near and far: Mobile phones and agricultural markets in Niger,” *American Economic Journal: Applied Economics*, 2010, 2 (3), 46–59.
- Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja**, “Report cards: The impact of providing school and child test scores on educational markets,” 2014.
- Bandiera, Oriana, Andrea Prat, and Tommaso Valletti**, “Active and Passive Waste in Government Spending: Evidence from a Policy Experiment,” *American Economic Review*, 2009, 99 (4), 1278–1308.
- Callen, Michael, Saad Gulzar, Ali Hasanain, Yasir Khan, and Arman Rezaee**, “Personalities and Public Sector Performance: Evidence from a Health Experiment in Pakistan,” Technical Report, National Bureau of Economic Research 2015.
- Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, and F Halsey Rogers**, “Missing in action: teacher and health worker absence in developing countries,” *The Journal of Economic Perspectives*, 2006, 20 (1), 91–116.
- Duflo, Esther, Rema Hanna, and Stephen P Ryan**, “Incentives work: Getting teachers to come to school,” *The American Economic Review*, 2012, pp. 1241–1278.
- Fafchamps, Marcel and Bart Minten**, “Impact of sms-based agricultural information on indian farmers,” *The World Bank Economic Review*, 2012, 26 (3), 383–414.
- Ferraz, Claudio and Frederico Finan**, “Electoral accountability and corruption: Evidence from the audits of local governments,” *American Economic Review*, June 2011, 101 (4), 1274–1311.
- Henderson, Ralph H and T Sundaresan**, “Cluster sampling to assess immunization coverage: a review of experience with a simplified sampling method.,” *Bulletin of the World Health Organization*, 1982, 60 (2), 253.
- Jensen, Robert**, “The digital divide: Information (technology), market performance, and welfare in the South Indian fisheries sector,” *The Quarterly Journal of Economics*, 2007, pp. 879–924.
- Jin, Ginger Zhe and Phillip Leslie**, “The effect of information on product quality: Evidence from restaurant hygiene grade cards,” *The Quarterly Journal of Economics*, May 2003.
- Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria**, “Asymmetric information and middleman margins: An experiment with west bengal potato farmers,” 2014. Working paper.
- Olken, Benjamin A. and Rohini Pande**, “Corruption in Developing Countries,” *Annual Review of Economics*, 2012, 4, 479–509.

Reinikka, Ritva and Jakob Svensson, “Local Capture: Evidence from a Central Government Transfer Program in Uganda,” *The Quarterly Journal of Economics*, 2004, 119 (2), 679–705.

Svensson, Jakob and David Yanagizawa, “Getting prices right: the impact of the market information service in Uganda,” *Journal of the European Economic Association*, 2009, 7 (2-3), 435–445.

Wild, Lena, Vikki Chambers, Maia King, and Daniel Harris, “Common Constraints and Incentive Problems in Service Delivery,” Technical Report, Overseas Development Institute 2012.

World Bank, *World Development Report 2004: Making Services Work for the Poor*, World Bank, 2004.

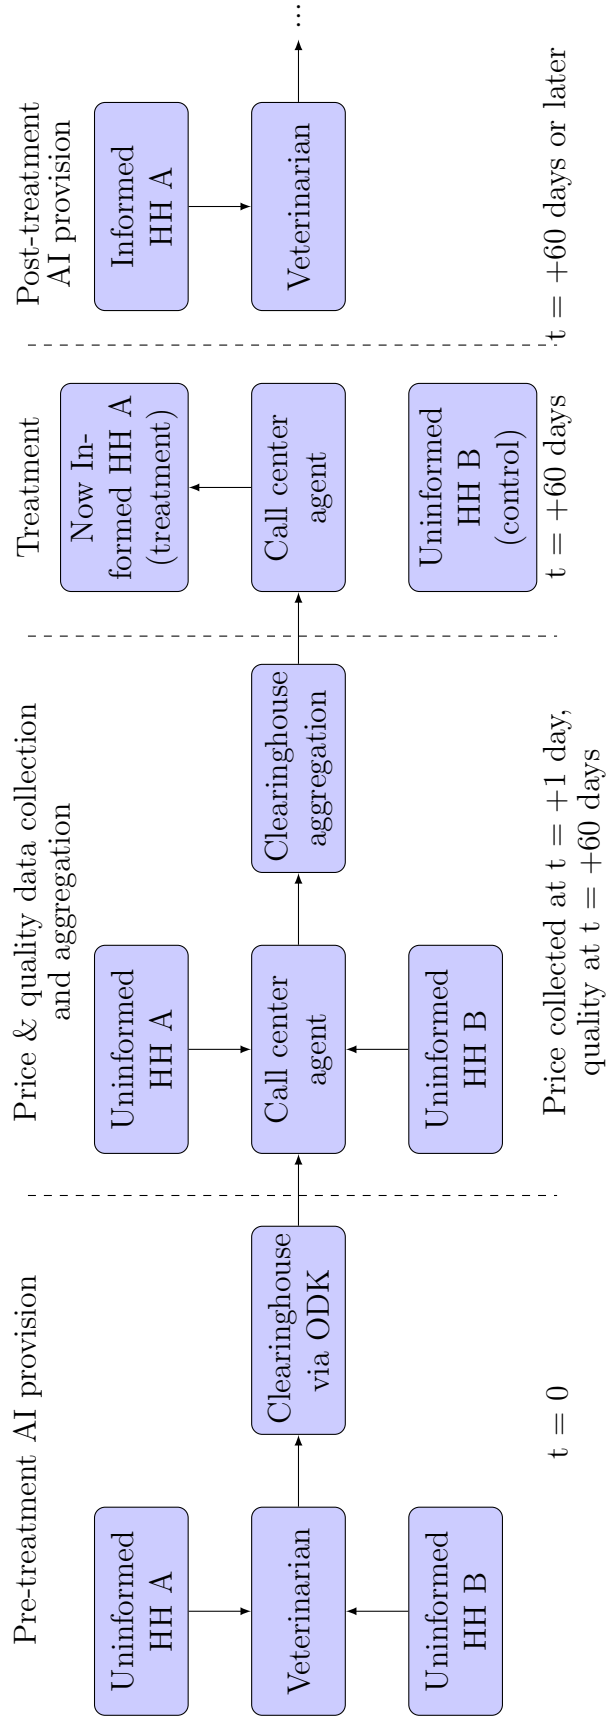
7 Tables and figures

Table 1: Treatment balance—clearinghouse data

	Treatment	Control	Difference	P-value
Satisfaction with AI service provision (1-5)	4.185 [0.736]	4.136 [0.760]	0.049 (0.029)	0.123
Farmer switched vets since last AI visit	0.052 [0.222]	0.047 [0.213]	0.005 (0.0100)	0.133
AI visit charges (PKR)	196 [180]	203 [250]	-7 (9)	0.479
AI visit success rate (pregnancy / AI attempts)	0.686 [0.458]	0.687 [0.457]	-0.002 (0.016)	0.432
No of cows owned by farmer	2.544 [3.439]	2.447 [3.053]	0.097 (0.155)	0.312
No of buffalo owned by farmer	3.121 [3.777]	3.315 [6.347]	-0.195 (0.366)	0.771
Distance to closest AI center (km)	2.170 [2.254]	2.277 [2.259]	-0.107 (0.114)	0.825

Notes: Standard deviations reported in brackets. Standard errors reported in parentheses. Means and differences are unconditional. P-values are from OLS regressions with randomization strata fixed effects and standard errors clustered at the village-cluster level. The sample consists of 6,473 pre-treatment farmer-visit-level observations from 3,094 unique farmers across 202 village-clusters. Some regressions have fewer observations due to missing data. Beginning in October 2014, treatment farmers received information about the AI success rates of their local government veterinarians. Satisfaction, AI visit charges, and numbers of cows and buffalo are reported by farmers on the phone one day after AI service provision. AI visit success rate is reported by farmers on the phone 60 days after AI service provision. Farmer switched vets and distance to closest AI center are automatically generated administrative data.

Figure 1: Clearinghouse flowchart



Notes: Arrows indicate the flow of information. The collection of quality data and treatment occur during the same follow-up phonecall 60 days after service provision. Beginning in October 2014, treatment farmers received information about the AI success rates of their local government veterinarians.

Figure 2: Clearinghouse and representative survey timelines

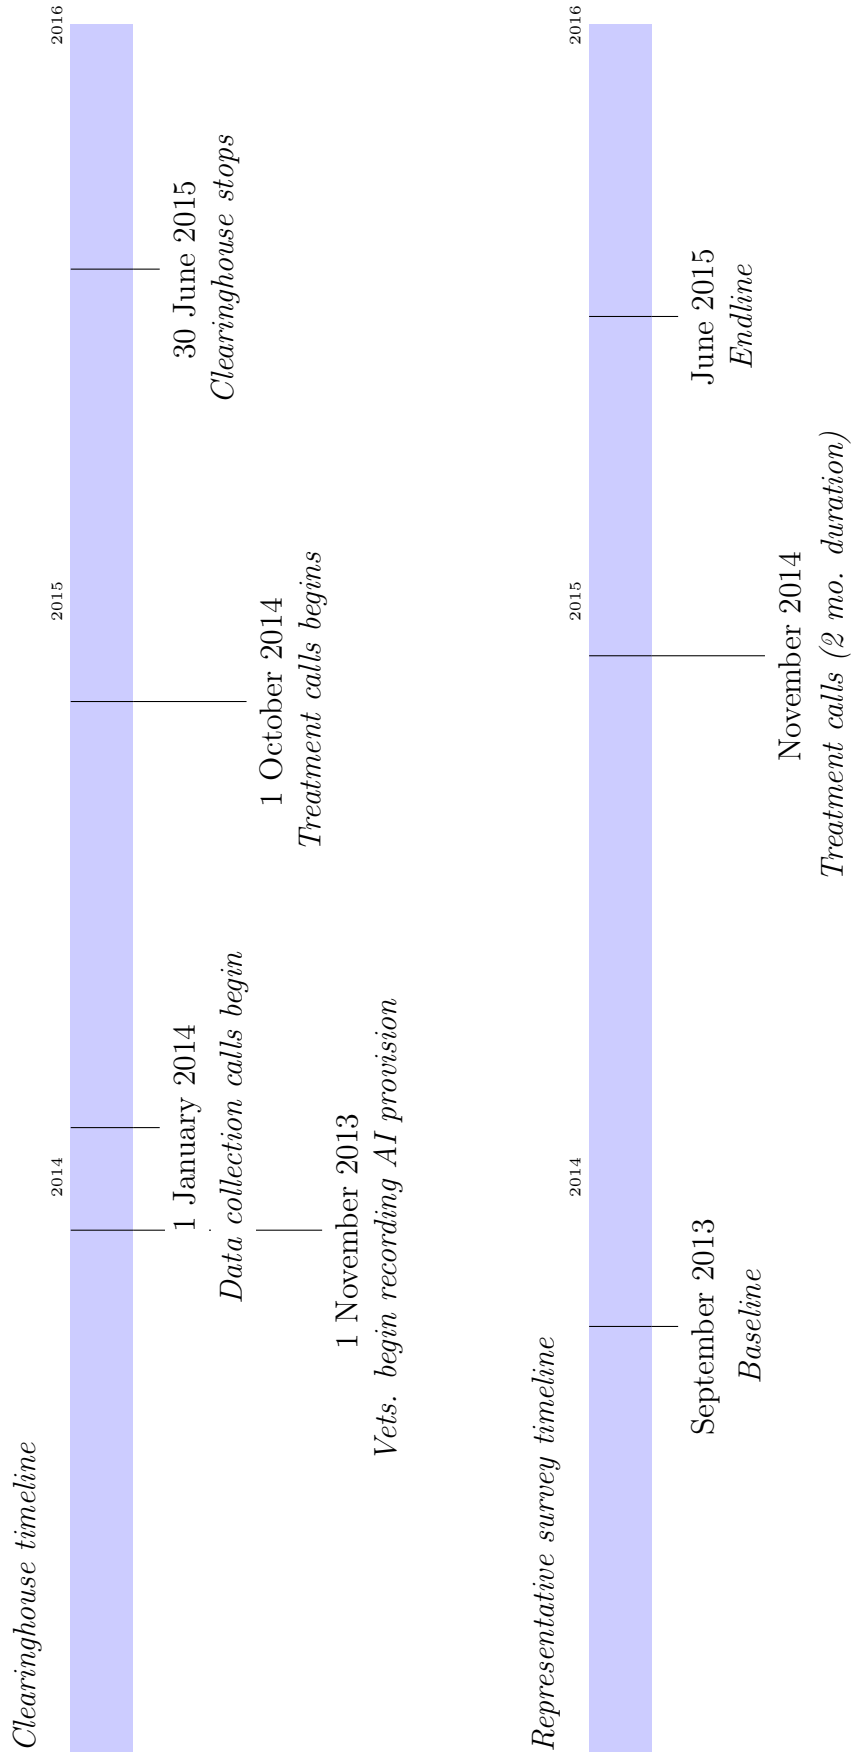


Table 2: Treatment balance—representative survey sample

	Treatment	Control	Difference	P-value
Farmer-level baseline variables—190 observations across 61 village-clusters				
Livestock is primary source of HH’s income (=1)	0.085	0.097	-0.012	0.748
	[0.281]	[0.297]	(0.042)	
1-10 effort household puts into selecting veterinarian	6.200	5.575	0.625	0.491
	[2.361]	[2.049]	(0.537)	
Farmer attrited from in-person endline	0.021	0.011	0.011	0.812
	[0.145]	[0.104]	(0.018)	
Farmer-visit-level variables—356 pre-treatment observations from 190 farmers across 61 village-clusters				
Farmer switched vets since last recorded AI visit (=1)	0.179	0.190	-0.011	0.879
	[0.385]	[0.393]	(0.055)	
AI visit charges	367	356	10	0.771
	[373]	[361]	(48)	
AI visit success rate	0.703	0.750	-0.047	0.159
	[0.447]	[0.431]	(0.049)	
1-10 AI visit farmer satisfaction	7.694	9.302	-1.608	0.290
	[2.184]	[22.333]	(1.754)	
1-10 farmer estimated AI visit veterinarian success rate	6.636	6.315	0.321	0.606
	[1.739]	[1.981]	(0.276)	

Notes: Standard deviations reported in brackets. Standard errors reported in parentheses. Means and differences are unconditional. P-values are from OLS regressions with randomization strata fixed effects and standard errors clustered at the village-cluster level. Some regressions have fewer observations due to missing data. All data come from baseline surveys fielded in August and September 2013, with the exception of “Farmer attrited from endline survey”. This variable is a dummy equal to one if a farmer was present during our baseline survey and not our endline survey. The sample of farmers was selected to be geographically representative of Sahiwal and is drawn from 90 different villages. The sample is limited to farmers that report receiving services from a government veterinarian at baseline. Treatment farmers received information about the AI success rates of their local government veterinarians. Treatment calls were conducted in November 2014 and January 2015.

Table 3: Treatment effects—representative survey sample

Outcome:	Log price			AI success rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment farmer (=1)	0.027	-0.146	-0.062	0.470**	0.028	0.172
	(0.405)	(0.216)	(0.164)	(0.186)	(0.187)	(0.109)
Mean of dependent variable	5.856	5.888	5.874	0.567	0.765	0.672
# Observations	69	87	156	63	79	142
# Village-clusters	27	39	53	29	35	51
R-Squared	0.633	0.655	0.540	0.498	0.281	0.271
Sample	Returned	Attrited	Both	Returned	Attrited	Both

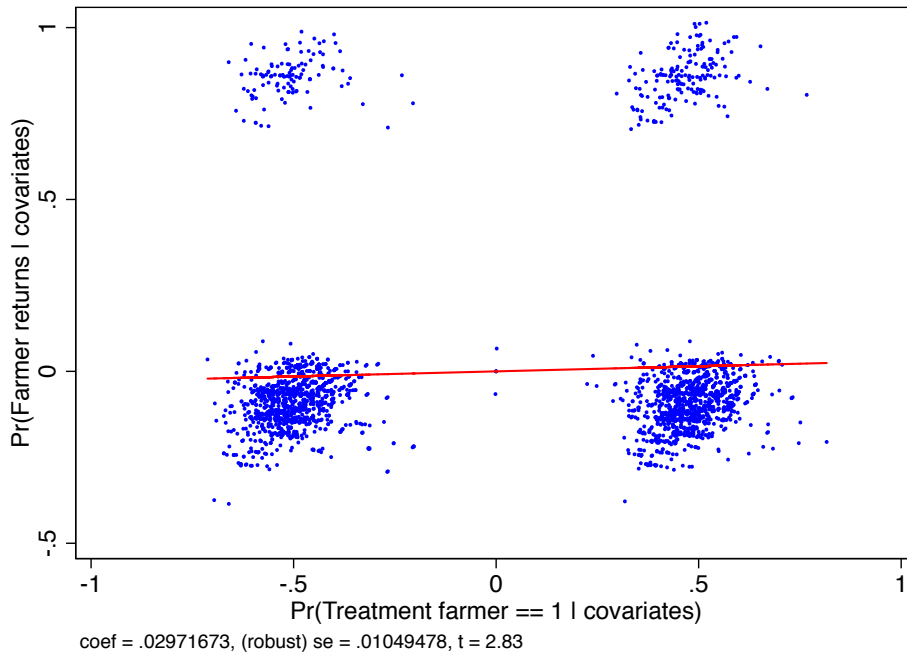
Notes : * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the village-cluster level reported in parentheses. All regressions include randomization strata fixed effects, survey wave fixed effects, and controls for baseline mean outcomes (switched veterinarians, log price, AI success rate) as well as dummies for whether the given observation is missing each baseline mean outcome. The sample is limited to post treatment reports of AI service provision from farmers during our endline survey, conducted in June 2015. Treatment farmers received information about the AI success rates of their local government veterinarians. Treatment calls were conducted in November 2014 and January 2015. Returned indicates farmers that received government AI before treatment and subsequently returned for government AI after treatment by the end of the project. Attrited indicates farmers who received government AI before treatment and instead subsequently received private AI by the end of the project. Log price and AI success rates are recalled by farmers from service provisions two to seven months ago.

Table 4: Treatment effects—clearinghouse data

Outcome:	Returned	Switched veterinarians	Log price	AI success rate
	(1)	(2)	(3)	(4)
Treatment farmer (=1)	0.032*** (0.011)	0.007 (0.028)	-0.270 (0.170)	0.168** (0.083)
Mean of dependent variable	0.098	0.084	5.248	0.623
# Observations	3184	629	312	240
# Village-clusters	205	111	103	98
R-Squared	0.192	0.305	0.596	0.489
Sample	Pre	Post	Post	Post

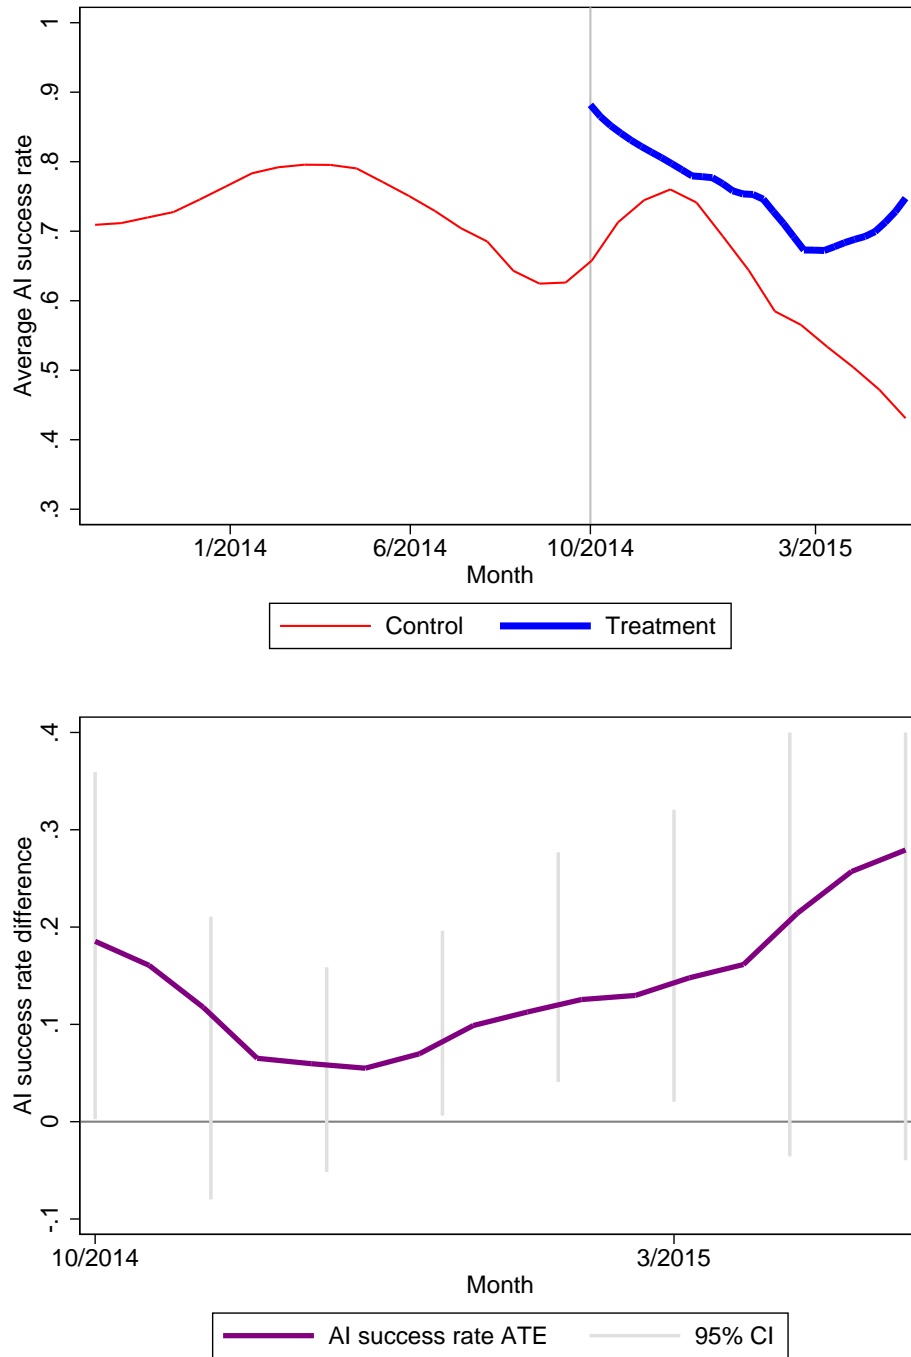
Notes : * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the village-cluster level reported in parentheses. All regressions include randomization strata fixed effects and controls for baseline mean outcomes (switched veterinarians, log price, AI success rate) as well as dummies for whether the given observation is missing each baseline mean outcome. In addition, columns (2) through (4) include exact call center script fixed effects and a time trend control. The sample for column (1) is farmers that received a government AI service and were subsequently treated, regardless of whether they then returned. The sample for columns (2) through (4) are farmers that returned after treatment. Note the differences in observations across columns are due to the fact that veterinarian switching can be detected without any successful phone calls, where as log price requires one successful phone call and AI success rate requires two successful phone calls to a farmer. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Returned is a dummy variable equal to one if a farmer that received government AI before treatment subsequently returned for government AI after treatment by the end of the project. Switched veterinarians is a dummy variable equal to one if the veterinarian that a farmer saw for a service provision was different than the last veterinarian seen. Log price is the log price paid for the service provision, as reported by the farmer when called to verify service provision. AI success rate is the rate of success of the AI services provided at a specific service provision upon follow up 60 days later.

Figure 3: Farmer returned added-variable plot—clearinghouse data



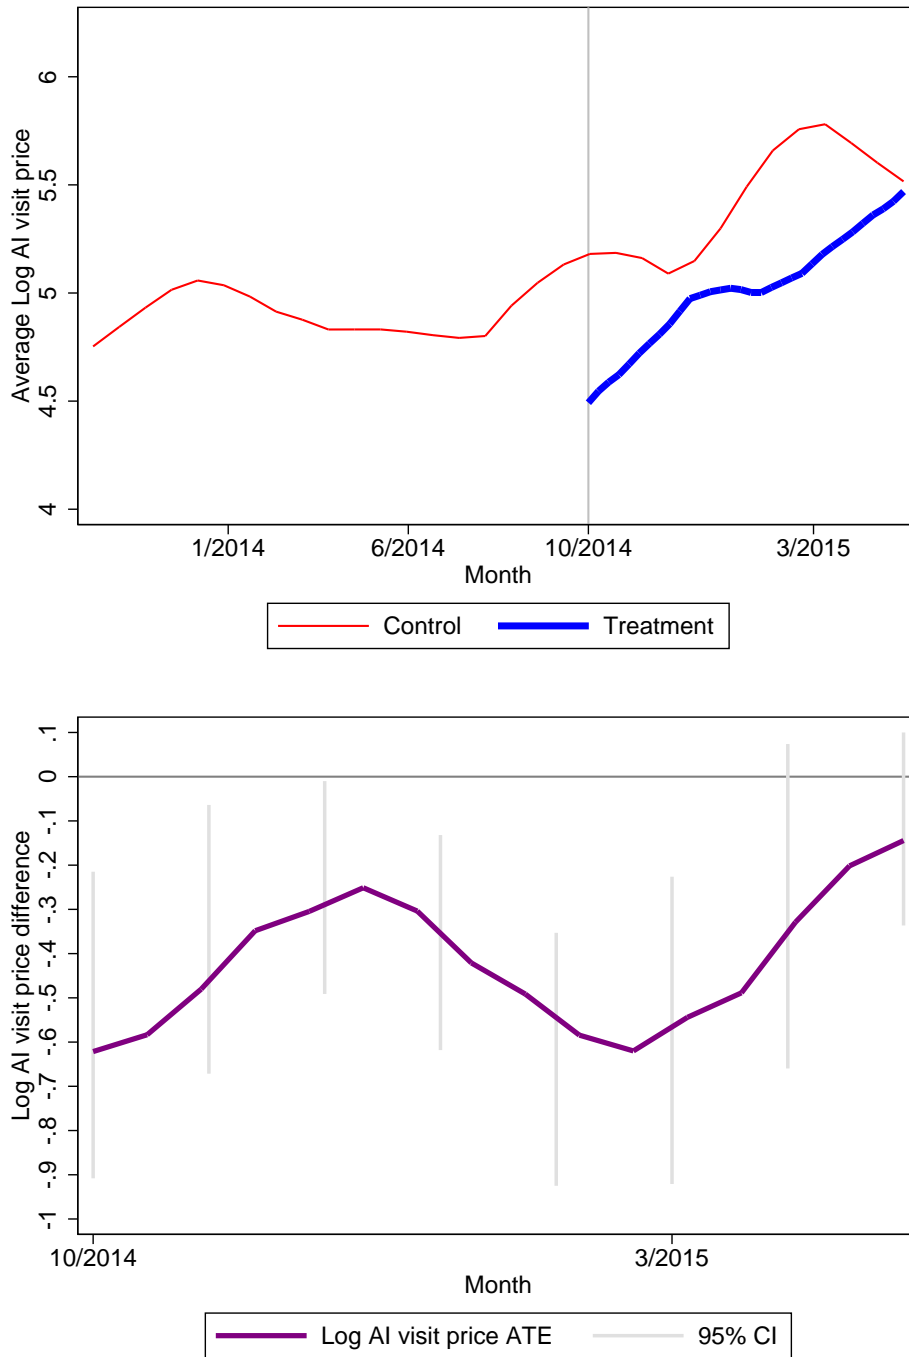
Notes: The sample is farmers that received a government AI service and were subsequently treated, regardless of whether they then returned. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Returned is a dummy variable equal to one if a farmer that received government AI before treatment subsequently returned for government AI after treatment by the end of the project. The covariates used to predict residual values are randomization strata fixed effects and controls for baseline mean outcomes (switched veterinarians, log price, AI success rate) as well as dummies for whether the given observation is missing each baseline mean outcome.

Figure 4: AI success rates in real time—clearinghouse data



Notes: The sample is farmers that received a government AI service and then answered the phone and reported AI success 60 days later. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Lines are smoothed using a kernel-weighted local polynomial regression with the Epanechnikov kernel and bandwidth one. Confidence interval bootstrapped and truncated at 0.4.

Figure 5: Log price per AI visit in real time—clearinghouse data



Notes: The sample is farmers that received a government AI service and then answered the phone and reported price paid one day later. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Lines are smoothed using a kernel-weighted local polynomial regression with the Epanechnikov kernel and bandwidth one. Confidence interval bootstrapped and truncated at 0.1.

Table 5: Does treatment induce a veterinarian reporting bias?

	Treatment	Control	Difference	P-value
Farmer reported AI and veterinarian submitted data to call center (=1)	0.299 [0.459]	0.276 [0.448]	0.023 (0.044)	0.758 .
Farmer reported receiving a call verifying AI service (=1)	0.287 [0.449]	0.240 [0.422]	0.047 (0.041)	0.566 .

Notes: Standard deviations reported in brackets. Standard errors reported in parentheses. Means and differences are unconditional. P-values are from OLS regressions with randomization strata fixed effects and standard errors clustered at the village-cluster level. The sample consists of 730 farmer-visit-level observations from 440 unique farmers across 83 village-clusters from our endline survey, conducted in June 2015. Some regressions have fewer observations due to missing data. Treatment farmers received information about the AI success rates of their local government veterinarians. Treatment calls were conducted in November 2014 and January 2015. “Farmer reported AI and veterinarian submitted data to call center” is a dummy equal to one if a government AI service provision reported in our endline survey was subsequently submitted to the clearinghouse by the veterinarian that performed the service. This is done by verifying survey data with clearinghouse data directly.

Table 6: Treatment balance of returning sample—clearinghouse data

	Treatment	Control	Difference	P-value
Pre-treatment mean satisfaction with AI service provision (1-5)	4.212 [0.684]	4.248 [0.713]	-0.036 (0.080)	0.765
Pre-treatment mean veterinarian switching rate	0.047 [0.218]	0.026 [0.206]	0.020 (0.019)	0.131
Pre-treatment mean log AI visit charges	4.852 [1.356]	4.838 [1.352]	0.014 (0.147)	0.660
Pre-treatment mean AI success rate	0.694 [0.445]	0.669 [0.439]	0.025 (0.051)	0.541
Pre-treatment mean no. of cows	2.770 [2.785]	3.168 [2.349]	-0.398 (0.384)	0.351
Pre-treatment mean no. of buffalo	3.493 [3.243]	3.321 [4.109]	0.173 (0.444)	0.929
Pre-treatment mean distance to closest AI center (km)	2.413 [2.158]	2.007 [2.190]	0.406 (0.245)	0.728

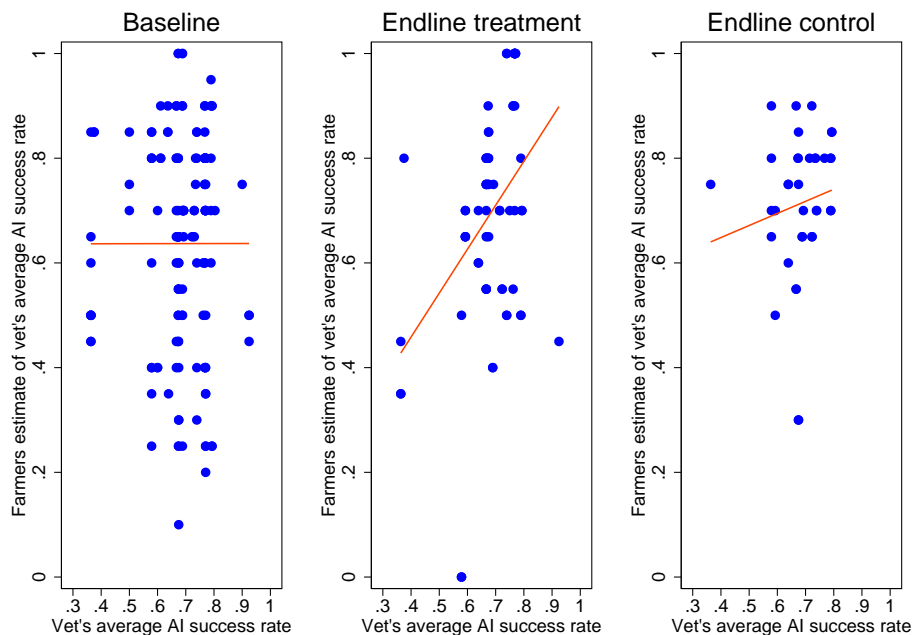
Notes: Standard deviations reported in brackets. Standard errors reported in parentheses. Means and difference are unconditional. P-values are from OLS regressions with randomization strata fixed effects and standard errors clustered at the village-cluster. The sample consists of 300 farmer-level observations across 108 village-clusters of those farmers who received government AI service provisions both before and after receiving a treatment or control phone call. Some regressions have fewer observations due to missing data. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Satisfaction, AI visit charges, and numbers of cows and buffalo are reported by farmers on the phone one day after AI service provision. AI visit success rate is reported by farmers on the phone 60 days after AI service provision. Farmer switched vets and distance to closest AI center are automatically generated administrative data.

Table 7: Treatment effects by veterinarian ranking—clearinghouse data

Outcome:	Returned	Switched veterinarians	Log price	AI success rate
	(1)	(2)	(3)	(4)
Panel A: Farmers told vet. was in top three in area				
Treatment farmer (=1)	0.008 (0.013)	-0.009 (0.035)	-0.169 (0.136)	0.010 (0.115)
Mean of dependent variable	0.091	0.098	4.903	0.654
# Observations	1977	439	169	124
# Village-clusters	174	78	66	56
R-Squared	0.102	0.363	0.717	0.743
Panel B: Farmers told vet. was not in top three in area				
Treatment farmer (=1)	0.039* (0.020)	0.005 (0.079)	-0.994 (1.419)	0.285* (0.161)
Mean of dependent variable	0.067	0.050	5.574	0.429
# Observations	1087	166	82	68
# Village-clusters	161	55	40	34
R-Squared	0.121	0.576	0.819	0.873
Sample	Pre	Post	Post	Post

Notes : * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the village-cluster level reported in parentheses. All regressions include randomization strata fixed effects and controls for baseline mean outcomes (switched veterinarians, log price, AI success rate) as well as dummies for whether the given observation is missing each baseline mean outcome. In addition, columns (2) through (4) include exact call center script fixed effects and a time trend control. The sample for column (1) is farmers that received a government AI service and were subsequently treated, regardless of whether they then returned. The sample for columns (2) through (4) are farmers that returned after treatment. Note the differences in observations across columns are due to the fact that veterinarian switching can be detected without any successful phone calls, where as log price requires one successful phone call and AI success rate requires two successful phone calls to a farmer. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Returned is a dummy variable equal to one if a farmer that received government AI before treatment subsequently returned for government AI after treatment by the end of the project. Switched veterinarians is a dummy variable equal to one if the veterinarian that a farmer saw for a service provision was different than the last veterinarian seen. Log price is the log price paid for the service provision, as reported by the farmer when called to verify service provision. AI success rate is the rate of success of the AI services provided at a specific service provision upon follow up 60 days later. Panels are divided by whether a farmer was told when treated that his/her veterinarian from the last visit was in the top three or not, or would have been if s/he was not selected for control.

Figure 6: Treatment effect on farmer expectations—representative survey sample



Notes: The sample is farmers that received AI from a reported veterinarian that could be matched to our clearinghouse veterinarians. Farmer’s estimates of vet’s average AI success rate reported by farmers in baseline and endline surveys. Vet’s actual average AI success rate is from clearinghouse data before October 2014. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians.

Table 8: Change in farmer expectations—representative survey sample

	Farmer’s estimate of vet’s average AI success rate		
	(1)	(2)	(3)
Vet’s actual average AI success rate	0.001 (0.177)	0.839** (0.385)	0.231 (0.229)
# Observations	145	66	37
# Village-clusters	34	21	20
R-Squared	0.000	0.162	0.020
Sample	Baseline	Endline T	Endline C

Notes : * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the village-cluster level reported in parentheses. The sample is farmers that received AI from a reported veterinarian that could be matched to our clearinghouse veterinarians. Farmer’s estimates of vet’s average AI success rate reported by farmers in baseline and endline surveys. Column (1) limits to baseline responses by eventual treatment and control farmers. Column (2) limits to endline responses by treatment farmers. Column (3) limits to endline responses by control farmers. Vet’s actual average AI success rate is from clearinghouse data before October 2014. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. The null hypothesis that the coefficients in columns (2) and (3) are equal is rejected with a p-value of 0.115 from a regression interacting Vet’s actual average AI success rate with a treatment indicator in the Endline sample.

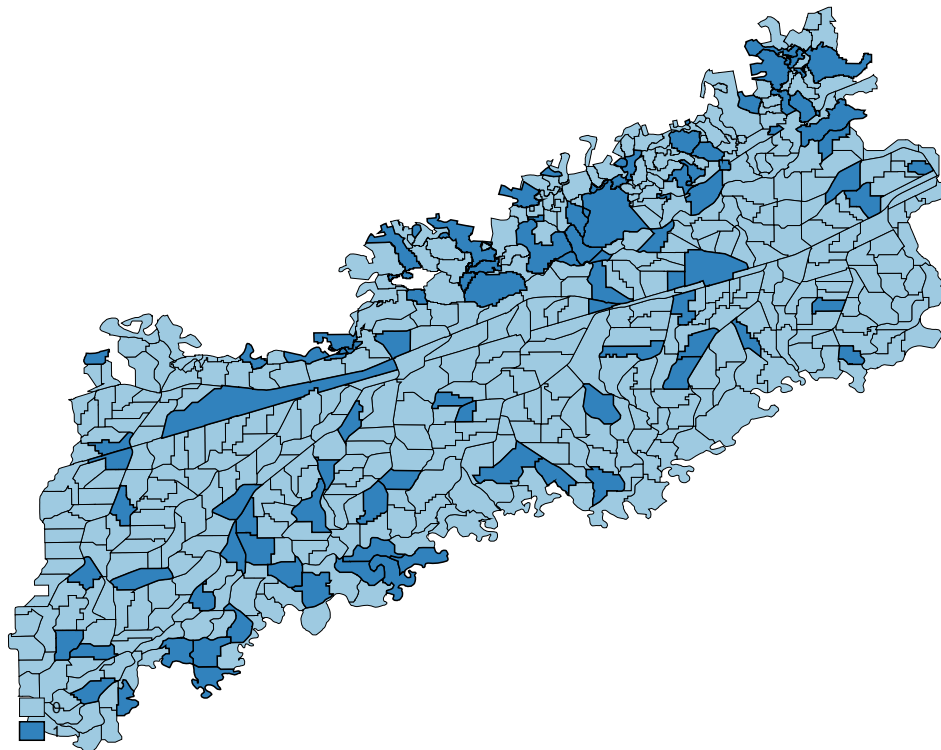
Table 9: Treatment effects by farmer expectations—representative survey sample

Outcome:	Returned	Switched veterinarians	Log price	AI success rate
	(1)	(2)	(3)	(4)
Panel A: Farmers with above median expected-actual AI success				
Treatment farmer (=1)	-0.083 (0.135)	0.049 (0.055)	0.294 (0.493)	0.318 (0.412)
Mean of dependent variable	0.370	0.231	5.688	0.500
# Observations	60	29	29	20
# Village-clusters	28	12	12	9
R-Squared	0.536	0.589	0.738	0.514
Panel B: Farmers with below median expected-actual AI success				
Treatment farmer (=1)	0.113 (0.274)	0.369 (0.329)	-1.399*** (0.385)	0.749* (0.370)
Mean of dependent variable	0.419	0.118	5.939	0.563
# Observations	53	32	28	28
# Village-clusters	29	16	14	16
R-Squared	0.468	0.756	0.898	0.588
Sample	Pre	Post	Post	Post

Notes : * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the village-cluster level reported in parentheses. All regressions include randomization strata fixed effects and controls for baseline mean outcomes (switched veterinarians, log price, AI success rate) as well as dummies for whether the given observation is missing each baseline mean outcome. In addition, columns (2) through (4) include survey wave fixed effects and restricts the sample to those farmers that returned. The sample is limited to post treatment reports of AI service provision from farmers during our endline survey, conducted in June 2015. Returned is a dummy variable equal to one if a farmer that received government AI before treatment subsequently returned for government AI after treatment by the end of the project. Switched veterinarians is a dummy variable coded as one if the veterinarian a farmer saw for a service provision was different than the last veterinarian seen. Log price and AI success rates are recalled by farmers from service provisions two to seven months ago. Panels are divided above and below the median of veterinarian’s farmers’ estimate of their veterinarian’s average AI success rate minus veterinarian’s actual average AI success rate from clearinghouse data before October 2014. Positive values in this difference occur when farmers are told their veterinarian is better than they expected’ negative values occur when farmers are told their veterinarian is worse than they expected. The median is .012.

A Appendix tables and figures

Figure A.1: Representative Survey sample villages



Notes: Sampled villages are dark blue. The sample was stratified by whether or not a government veterinarian center was in the village and on whether the village was a canal colony. It is balanced along the following variables: area, settled area, cultivated area, area of wheat, rice, cotton, sugar cane, pulses, orchards, and vegetables, having a river, distance to the nearest veterinarian center, number of livestock in the village, literacy rates, religion, age, and standard wealth index characteristics. Results available upon request. Within each village, we selected ten households using the well-documented EPI cluster sampling method. In order to be surveyed, households had to report owning at least two livestock (cows and/or buffalo) and having regular access to a cellular phone.

Table A.1: Treatment balance—representative survey sample, additional covariates

	Treatment	Control	Difference	P-value
Head of household education = None (=1)	0.388 [0.488]	0.404 [0.492]	-0.016 (0.038)	0.814
A child in the household attends public school (=1)	0.533 [0.500]	0.525 [0.500]	0.008 (0.038)	0.915
Household has used govt health services in past two years (=1)	0.399 [0.490]	0.466 [0.500]	-0.067 (0.038)	0.045
Amount of land household owns and rents for livestock	1.455 [3.248]	1.417 [2.875]	0.038 (0.273)	0.646
Household owns the house that they live in (=1)	0.926 [0.261]	0.948 [0.223]	-0.021 (0.020)	0.210
Hours of electricity per day	10.458 [3.366]	10.022 [3.573]	0.436 (0.276)	0.214
Household has a cooking stove/range (=1)	0.086 [0.280]	0.121 [0.326]	-0.035 (0.024)	0.119
Household made less than 100k PKR last year (=1)	0.320 [0.468]	0.301 [0.460]	0.019 (0.036)	0.349
Any member of household has bank account (=1)	0.235 [0.424]	0.275 [0.447]	-0.040 (0.034)	0.109
Believed it was likely that last vote was not secret (=1)	0.542 [0.499]	0.582 [0.494]	-0.040 (0.041)	0.396
Is likely to believe information given by gov't employee (=1)	0.776 [0.417]	0.815 [0.389]	-0.039 (0.031)	0.180
Average number of digits recalled	3.308 [0.992]	3.308 [1.129]	0.000 (0.112)	0.818
On a scale fo 0-10, how willing are you to take risks?	4.345 [3.008]	4.715 [6.894]	-0.370 (0.503)	0.332
Agreeableness	4.017 [0.743]	4.033 [0.702]	-0.016 (0.057)	0.756
Conscientiousness	4.071 [0.627]	4.128 [0.656]	-0.057 (0.051)	0.263
Extroversion	4.163 [0.686]	4.096 [0.695]	0.067 (0.056)	0.530
Neuroticism	2.363 [0.845]	2.375 [0.854]	-0.013 (0.066)	0.761
Openness	3.724 [0.711]	3.689 [0.755]	0.034 (0.057)	0.796

Notes: Standard deviations reported in brackets. Standard errors reported in parentheses. Means and differences are unconditional. P-values are from OLS regressions with randomization strata fixed effects and standard errors clustered at the village-cluster. The sample consists of 190 baseline farmer-level observations across 61 village-clusters. Some regressions have fewer observations due to missing data. All data come from baseline surveys fielded in August and September 2013. This sample of farmers was selected to be geographically representative of Sahiwal and is drawn from 90 different villages. The sample is limited to farmers that report receiving services from a government veterinarian at baseline. Treatment farmers received information about the AI success rates of their local government veterinarians. Treatment calls were conducted in November 2014 and January 2015. Agreeableness, conscientiousness, extroversion, neuroticism, and openness are all measures from the Big 5 Personality Index. These traits are each mean responses to statements that represent the trait on a five point likert scale, in which 1 corresponds to disagree strongly, 2 to disagree a little, 3 to neutral, 4 to agree a little, and 5 to agree strongly. Likert responses are given the same direction (5 always being more agreeable, for example, never less).

Table A.2: Treatment effects—representative survey sample

Outcome:	Returned	Switched veterinarians	Log price	AI success rate
	(1)	(2)	(3)	(4)
Treatment farmer (=1)	0.063 (0.062)	-0.058 (0.171)	0.027 (0.407)	0.470** (0.187)
Mean of dependent variable	0.222	0.152	5.852	0.581
# Observations	251	69	70	64
# Village-clusters	72	27	28	30
R-Squared	0.235	0.457	0.633	0.503
Sample	Pre	Post	Post	Post

Notes : * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the village-cluster level reported in parentheses. All regressions include randomization strata fixed effects and controls for baseline mean outcomes (switched veterinarians, log price, AI success rate) as well as dummies for whether the given observation is missing each baseline mean outcome. In addition, columns (2) through (4) include survey wave fixed effects and restricts the sample to those farmers that returned. The sample is limited to post treatment reports of AI service provision from farmers during our endline survey, conducted in June 2015. Returned is a dummy variable equal to one if a farmer that received government AI before treatment subsequently returned for government AI after treatment by the end of the project. Switched veterinarians is a dummy variable coded as one if the veterinarian a farmer saw for a service provision was different than the last veterinarian seen. Log price and AI success rates are recalled by farmers from service provisions two to seven months ago.

Table A.3: Comparing farmers by pre-treatment veterinarian ranking—clearinghouse data

	Vet. in top three	Vet. not top three
Satisfaction with AI service provision (1-5)	4.170 [0.736]	4.142 [0.769]
Farmer switched vets since last AI visit	0.051 [0.220]	0.071 [0.257]
AI visit charges (PKR)	192 [170]	212 [269]
AI visit success rate (pregnancy / AI attempts)	0.628 [0.477]	0.635 [0.476]
No of cows owned by farmer	2.382 [3.154]	2.668 [3.660]
No of buffalo owned by farmer	2.816 [3.165]	3.516 [5.949]
Distance to closest AI center (km)	1.710 [1.572]	3.257 [2.949]

Notes: Standard deviations reported in brackets. The sample consists of 4,788 pre-treatment farmer-visit-level observations from 2,981 unique farmers that received government AI service provision. Some regressions have fewer observations due to missing data. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Satisfaction, AI visit charges, and numbers of cows and buffalo are reported by farmers on the phone one day after AI service provision. AI visit success rate is reported by farmers on the phone 60 days after AI service provision. Farmer switched vets and distance to closest AI center are automatically generated administrative data. Columns are divided by whether a farmer was told when treatment that his/her veterinarian from the last visit was in the top three or not, or would have been if s/he was not selected for control.