

No bulls: Experimental evidence on the impact of veterinarian ratings in Pakistan*

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

October 21, 2022

Abstract

We implement a platform to crowdsource information about service provision quality and prices charged and reveal this information to consumers in a market - artificial insemination of livestock in Punjab, Pakistan - where individual signals of quality are noisy. We measure the impact of this information revelation using a randomized controlled trial. Farmers receiving information enjoy 25% higher insemination success and no higher prices than controls. These effects are due to existing veterinarians increasing effort, rather than farmers switching to possibly higher-quality providers. These results illustrate the viability of information clearinghouses successfully aggregating information in low-capacity markets. They also suggest the importance of doing so by implying large welfare benefits from our low-cost information intervention.

1 Introduction

Asymmetric information is ubiquitous across the developing world and often leads to sub-optimal outcomes for the rural poor there (World Bank, 2004; Wild et al., 2012). This asymmetric information can be leveraged by rent seeking government agents (Muralidharan et al., 2021; Ferraz and Finan, 2011; Bandiera et al., 2009; Chaudhury et al., 2006; Reinikka and Svensson, 2004) or private agents (Kelley et al., 2021; Aker, 2010; Svensson and Yanagizawa, 2009; Jensen, 2007). This asymmetric information also means potential gains from trade go untapped, and may lead to sub-optimal outcomes for both poor consumers and

*We thank Eli Berman, Michael Callen, Julie Cullen, Clark Gibson, Craig McIntosh, Edward Miguel, Karthik Muralidharan, Monica Singhal, 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. 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.

[†]Lahore University of Management Sciences. email: hasanain@lums.edu.pk

[‡]University of Pittsburgh. email: myk17@pitt.edu

[§]University of California, Davis. email: abrezaee@ucdavis.edu

agents in steady state. And while the contexts and players vary wildly, this asymmetric information can often be modeled following canonical theories of moral hazard and/or adverse selection (Hölmstrom, 1979; Akerlof, 1970; Klein et al., 2016).

We study asymmetric information in an important developing-country context: livestock service provision. More specifically, we study the market for artificial insemination (AI) of livestock in rural Punjab, Pakistan.¹ Farmers cannot directly observe the quality of the insemination straws that veterinarians apply to their animal, and typically cannot judge the veterinarian’s application technique. Moreover, successful insemination increases the chance of pregnancy, but does not assure it. The farmer’s signal of veterinarian quality is therefore noisy and the market for AI suffers asymmetric information about veterinarian effort. This leads to AI success rates that are lower than technologically possible, costing farmers potential income in calves that are never born and milk that is never lactated. And veterinarians may lose a possible share of this income.

We measure and reveal information about AI service provision to rural households in one district of Punjab, Pakistan, through an information clearinghouse similar to a yelp.com or an angieslist.com. Specifically, the clearinghouse provides households with local (using GPS) government veterinarians’ average success rates at artificially inseminating livestock, an objective measure of veterinarian effort, the average price charged by each of these local veterinarians for AI service, and the number of observations used to estimate these averages. The clearinghouse gathers and disseminates this information from a large base of farmers automatically, in real time, using a call center.

We measure the impact of decreased asymmetric information via this clearinghouse using a randomized controlled trial. We find that farmers treated with information on local government veterinarians’ AI success rates and prices have a 25% higher AI success rate than controls and they do not pay any higher prices.² In addition, treatment farmers who specifically return to a government veterinarian for AI rather than seeking a private provider after treatment selection³ see 61% higher AI success rates. While these results leverage an experiment with a representative sample of farmers that we surveyed before and after treatment, we also find comparable results in a second experiment with those farmers who entered the clearinghouse through its automatic operation.

Interestingly, we find that treatment farmers enjoy these increased AI success rates and no-higher prices *without switching government veterinarians*. This implies that the effect cannot be driven by farmers simply switching to the ‘best vet’ in terms of AI success and/or price. Additional results suggests that the effect cannot be driven purely by changes in farmer behavior on other margins.⁴ To help understand the possible channel(s) for these effects, we

¹AI 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. Livestock agriculture accounts for 12% of GDP in Pakistan, and is a key growth sector for the rural poor (Pakistan Economic Survey 2013-14).

²The estimated treatment effect on log AI price is negative but insignificant.

³In this setting, there are government and private veterinarians but only government veterinarians were rated by our clearinghouse, and only farmers originally contracting government veterinarians were considered for treatment selection.

⁴It is 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

present a simple stylized model based on Klein et al. (2016). While this model and our data allow us to rule out some channels (e.g. decreased adverse selection through *low* quality veterinarians exiting), our results are consistent with some combination of (i) decreased adverse selection of veterinarians (i.e. fewer *low* quality veterinarians masquerading as *high*), (ii) veterinarians’ engaging in less moral hazard (i.e. fewer veterinarians putting in *low* effort when farmers pay for *high*), and (iii) improved bargaining power for farmers (i.e. farmers bargaining a better deal when they learn their current veterinarian is *low* quality and there are other *high* quality veterinarians nearby).⁵

While government veterinarians in Punjab collect a salary and are protected from punishment for poor performance, they are legally allowed to charge a ‘show-up’ fee to farmers 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. This ensures they respond when the market becomes more contested.

Additional results are consistent with our stylized 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 vis-a-vis the channels mentioned above.

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.⁶

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; Annan, 2022). Many other markets have multiple channels for rents and thus expect similar social welfare gains, including education (Andrabi et al., 2017), elections (Ferraz and Finan, 2011), and markets for private restaurants (Jin and Leslie, 2003).

Our study more directly relates to clearinghouses evaluated in Fafchamps and Minten

subsequently choose.

⁵Others have addressed the impact of improving farmers’ bargaining power in their negotiating with service providers. See, e.g., Dreze and Sen (1989); Basu et al. (2009); Muralidharan et al. (2017). This can also be thought of as an increase in contestability, as in Baumol (1986). Note that in the long run increased contestability could even lead to welfare improvements for veterinarians, though such is outside the scope of this paper.

⁶We 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.

(2012) and Mitra et al. (2014), though in both cases, the authors find no treatment effects.⁷ In our case, conservative estimates suggest a 25% higher AI success rate translates into nearly an additional half of one month’s median income per AI provided, a 275% return on the cost of the intervention. These effects hold out hope for improved government accountability as cellular technology improves and becomes cheaper.⁸

The literature on market-based learning provides one lens to understand why our intervention could have been successful. Consumers can learn from each others’ aggregate experiences without intervention (Hubbard, 2002), but this learning is limited by the ability of consumers to experience quality. Learning happens rapidly in markets where quality is easily ascertained, switching costs are low, and volumes are high, such as packaged yogurt (Akerberg, 2003). It happens more slowly when the converse is true, as is the case with car insurance (Israel, 2005). In our case, artificial insemination services are an occasional need for small-scale farmers, so volume is low. At the same time, farmers have difficulty inferring veterinarian quality from outcomes alone, since even when well executed, AI can fail, and animal health and nutrition are co-determinants of insemination success.

This study also relates to the literature on monitoring to improve government service provision. This literature has found mixed results, with research suggesting monitoring may not be effective without complimentary financial incentives (Duflo et al., 2012) and that monitoring’s effects attenuate as agents find alternative strategies to pursue rents (Olken and Pande, 2012).⁹ While we cannot speak to the latter given the time frame of this paper, our results are consistent with the former as veterinarians have a financial incentive to maintain customers. In Pakistan specifically, a literature on health monitoring suggests that the mean zero impacts of smartphone monitoring on the performance of doctors may mask important heterogeneity driven by political competitiveness (Callen et al., 2017b) and individual characteristics (Callen et al., 2017a).

Our work also relates to a connected literature focused on community monitoring specifically. This literature has also found mixed results when citizens take collective action to monitor the performance of their public servants (Olken, 2007; Björkman and Svensson, 2009; Banerjee et al., 2010). While households in our study do not act collectively, it did require a sufficient number of households providing information into a collective information system for it to be useful for anyone.

The paper proceeds as follows: Section 2 provides background on our study district and government AI service provision there, Section 3 presents a simple stylized model of a

⁷Fafchamps and Minten (2012) cite a low take-up rate as the reason for the failure of an sms-based clearinghouse for crop and weather information. While the rate at which farmers answered the phone was nowhere near 100%, we had no problem generating sufficient data for our clearinghouse estimates to be meaningful. Mitra et al. (2014) cite a lack of an outside option as the reason farmers are not able to better leverage information on crop prices. Our clearinghouse directly provided information on outside options. Of course this required contestability could be increased to begin with, which was true in our market (i.e. veterinarians are not monopolists). This will not be the case everywhere.

⁸To further understand the value of the clearinghouse, we investigate whether the data generated by the clearinghouse is biased by either veterinarians (they have to first report providing a service) and/or farmers (they have to answer the phone several times to provide and receive information). We do not find evidence of bias, suggesting the clearinghouse was able to capture and transmit information representative of the truth.

⁹See Finan et al. (2015) for a review of monitoring efforts as apart of a larger review of the growing literature dubbed the personnel economics of the state.

farmer seeking artificial insemination, Section 4 outlines our research design, including providing more information on the clearinghouse and the randomized controlled trial embedded within it, Section 5 provides results, Section 6 discusses the interpretation and social welfare implications of these results, and Section 7 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 of the whole of Punjab, and of similar agricultural districts across the country, though with a slightly higher prevalence of livestock.¹⁰

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.¹¹

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.¹² These veterinarians’ sole job is to provide artificial insemination.¹³

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% of AI services across Sahiwal, with government veterinarians making up the remainder.¹⁴

¹⁰According to the 2010 Punjab’s Multiple Indicator Cluster Survey, households in Sahiwal on average have 1.4 fewer acres of agricultural land and 0.2 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.

¹¹The 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.

¹²Throughout our study period, a total of 77 government veterinarians were active in Sahiwal for any amount of time. Only a handful of veterinarians transferred in or out of Sahiwal.

¹³In 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.

¹⁴We asked farmers for the name and cell phone number of their veterinarians with the goal of precisely mapping the set of veterinarians, government or private, in Sahiwal. Unfortunately, challenges with farmer recall, the facts that some veterinarians go by multiple names, that certain names are very common in Pakistan, and that during this period it was common to switch SIM cards in Pakistan made this exercise incomplete. Thus while we are confident in the percent of services offered by government versus private veterinarians, and while we are confident in the government veterinarian side of the market, we are uncertain

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%, while success rates of 85%-90% are possible with the training and equipment in Sahiwal.

3 A stylized model of a farmer seeking artificial insemination

Revealing to farmers information about government veterinarian quality and prices could affect the market and subsequent outcomes through at least four channels: adverse selection, moral hazard, bargaining, and reputation. In this section, we develop a simple stylized model of the transaction between a farmer seeking artificial insemination and veterinarians who might provide that artificial insemination to specify how we see these channels operating in this particular context. This model largely borrows from Klein et al. (2016).

Imagine one stage in an infinitely repeated game in which a single farmer seeks AI for their cow. This farmer may purchase artificial insemination from one of three veterinarians that live within a short distance of the farmer, two government veterinarians and one private veterinarian.¹⁵ Figure 1 provides an overview of the stage game.

First, veterinarians are endowed with a fixed skill level, s_i , $i \in \{1, 2, 3\}$. Veterinarians can be high-skilled, s_i^H , or low-skilled, s_i^L . This skill relates to the cost veterinarians pay to put in effort toward artificial insemination. Specifically, veterinarians can choose to put in high effort, e_{high} , or low effort, e_{low} where the skill-effort relationship is such that $0 < c_{low}^H < c_{low}^L$, $0 < c_{high}^H < c_{high}^L$, $c_{low}^H < c_{high}^H$, and $c_{low}^L < c_{high}^L$. That is to say, for a fixed effort level, it is less costly for high-skilled veterinarians to put in effort towards artificial insemination than low-skilled veterinarians, and for both skill levels, it is less costly to put in low effort than high effort. Veterinarians are also endowed with a reputation, k_i^j , $j \in \{H, L\}$. This reputation is built from previous transactions with the same farmer.

The farmer with a cow in heat moves first. They choose which veterinarian(s) to call for potential artificial insemination. The farmer pays search cost $c > 0$ for each veterinarian called. The farmer and each called veterinarian immediately engage in Nash bargaining on the phone to decide on an effort-price tuple, $\{e, p\}$. Here $e \in \{high, low\}$ as above. Similarly, $p \in \{high, low\}$. Importantly, each veterinarian has a certain outside option, $\{\bar{e}_i, \bar{p}_i\}$, which determines the bargained effort-price tuple. We assume that the farmer has no option outside of the three veterinarians. Thus the farmer's outside option may only be helpful for bargaining if they choose to call at least two veterinarians.

about the precise number of private veterinarians or about their entry and exit during this period.

¹⁵This distinction will matter once we introduce treatment, which only provides information on government veterinarians.

The farmer then considers each offer along with veterinarians' reputations and selects a veterinarian through expected utility maximization.

The selected veterinarian then comes and provides AI. The agreed-upon price is charged. However, the veterinarian can choose an effort level different from what was agreed upon. This is because the veterinarian's effort level is hidden from the farmer at this time. The veterinarian will select effort to maximize expected utility, where effort today affects future utility through changes in reputation.

Roughly two months after the artificial insemination the farmer finds out if their cow has become pregnant. We can consider this as a noisy signal of the veterinarian's previous effort. The farmer then updates the past veterinarian's reputation given this noisy signal, and if the cow is not pregnant the game begins again immediately. If the cow is pregnant the game begins again in approximately one year.

Adverse selection is present in this model in so far as farmers never know for certain veterinarians' type. Low-skilled veterinarians can thus act opportunistically to charge higher prices and/or put in lower effort than if there was symmetric information. This form of adverse selection could also lead to some low-skilled veterinarians exiting the market, if farmers faced with a so-called lemons market are not willing to pay prices high enough to offset the cost of effort of low-skilled veterinarians.

Moral hazard is present in this model in so far as, regardless of skill, veterinarians take a hidden action when providing artificial insemination. Veterinarians can thus put in lower effort than bargained with farmers and lower effort than they would put in if their actions were not hidden.¹⁶

Bargaining is particularly important to this context as anecdotal evidence suggests farmers and veterinarians do indeed bargain over the phone before any effort costs are paid on the part of veterinarians to travel to a farm. This also allows reputation about quality to affect both prices and effort simultaneously.

Reputation is present in this model in so far as veterinarian reputation does make it more challenging for veterinarians to act opportunistically due to adverse selection and/or moral hazard. But it can never fully remove asymmetric information in this case because even with *high* effort artificial insemination success retains a non-trivial stochastic component (at best a veterinarian can achieve an 85-90 percent success rate). Reputation is also the channel by which our treatment operates. In this model, treatment amounts to providing the farmer with an updated reputation for each of the two government veterinarians.

We can consider what might happen if we treat a farmer with big-N reputation for both of the government veterinarians in the area but not the private veterinarian. While at first glance it may seem farmers would be less likely to select veterinarians reported to be low-quality after treatment, this is actually ambiguous. We can see this by discussing the impact of treatment on each of the four channels outlined in this Section.

The effect of treatment on adverse selection is ambiguous. If low-quality veterinarians have a much harder time masquerading as high-quality after treatment, this should induce farmers to either switch away from such veterinarians or to make more agreeable bargains

¹⁶Note we are not allowing farmers to invest in monitoring veterinarians' effort in this case. We did pilot an intervention to do just this (a detailed brochure showing what proper artificial insemination should look like) but did not end up launching it in this context.

with such veterinarians. This could also lead to low-skilled veterinarians exiting the market if farmers no longer believe low-quality veterinarians would be willing to put in high effort and these veterinarians cannot lower the price enough to compensate without profits going negative. While both of these effects would benefit farmers, farmers could be made worse-off if enough low-quality veterinarians exit the market and high-quality veterinarians gain sufficient market power.

The effect of treatment on moral hazard is equally ambiguous. Veterinarians reported as low-quality may now have their outcomes scrutinized more and may have a harder time keeping customers after lying about effort (assuming a sufficiently high farmer search cost). Veterinarians reported as high-quality could now seize the opportunity to put in low effort and claim it was just bad luck. That is, so long as farmers are not too short-sighted, one new observation of an unsuccessful artificial insemination has less of a negative effect on reputation than before treatment. On the other hand, if treatment causes a farmer to learn that there is another high-quality government veterinarian in the area, that farmer might become more sensitive to decreases in reputation.

The effect of treatment on bargaining is more straightforward. As already discussed, low-quality veterinarians should (weakly) lose bargaining power relative to farmers. High-quality veterinarians should (weakly) gain bargaining power relative to farmers. We say weakly here as there may be no changes in bargaining power if veterinarians retain enough untreated customers to ensure a stable outside option and/or if private veterinarians remain a stable outside option for farmers.

Finally, the effect of treatment on reputation has already been discussed as the mechanism by which all of these other changes would occur.

It is also worth discussing the role of the private veterinarian in this case. Treatment should not affect the reputation of private veterinarians—any asymmetric information related to these veterinarians' quality should remain.¹⁷ Otherwise farmers will treat them like another veterinarian and consider them for artificial insemination provision conditional on their reputation. Whether we see farmers utilizing private veterinarians more or less after treatment will depend on these veterinarians' reputations relative to what is revealed by treatment. Note that at baseline farmers believe the success rate of an average government veterinarian is 6.0 out of 10 while the success rate of an average private veterinarian is 6.3 out of 10. In other words, they do not see one group as systematically better than the other. We thus might think of private veterinarians as offering a stable outside option for farmers.

3.1 Model predictions

Perhaps the primary takeaway from this framework is that, *ex-ante*, the effect of price and quality information revelation on veterinarian, farmer, and cow outcomes in this market is ambiguous. This ambiguity motivates the empirical test we will now turn to. *Ex-post*, we will find several results point toward not necessarily a particular channel as being most

¹⁷It is possible here that farmers learn through treatment that veterinarians are on average better or worse than believed and this could affect private veterinarian reputation. Our treatment would not allow us to detect the effects of such aggregate information shocks, though, as they would be indistinguishable from time trends.

important but to farmers and cows being better off as a result of treatment. We will discuss these results in the context of this framework in Section 6.

4 Research design

4.1 Clearinghouse

To measure veterinarian prices and effort and to subsequently disseminate that information to consumers, we developed a novel cellular-based information clearinghouse. Figure 2 diagrams the four components of this clearinghouse.

Initial service provision data collection: 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.¹⁸ The data was immediately sent to the clearinghouse. We denote this data collection as $t = 0$ in Figure 2.

Follow-up data collection and aggregation: Each service provision generated two subsequent phone calls. First, one day later (denoted $t = +1$ day in Figure 2), 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.

Information provision: The clearinghouse collected and aggregated information from January to September, 2014. Beginning October 2014, once we had sufficient data on veterinarians to have meaningful measures of price and AI success rates, the clearinghouse began providing information back to farmers as well. Information provision took place during the second call (at $t = +60$) for selected farmers. Farmers were selected at random to receive information according to the two experiments outlined in the next two subsections.

Post-Information provision: The clearinghouse allowed us to link farmers over time, so we observe service provision by government veterinarians after the provision of information (if the farmers return; Figure 2 depicts the return of a treatment farmer but not a control farmer). These post-information provision observations also generate two follow-up phone calls, and so on.

4.2 Experiment 1: Representative sample

Parallel to implementing the clearinghouse, we independently surveyed a representative sample of farmers from across Sahiwal. For these surveys, we sampled 90 of Sahiwal’s approximately 500 villages from a district village census.¹⁹ Within each village, we selected ten households using the Expanded Program on Immunization (EPI) cluster sampling method

¹⁸In practice, veterinarians did not always comply. See Section A.1.1 for discussion.

¹⁹We 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.

(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.

Sample villages can be seen in Appendix Figure 7.

We then conducted our primary experiment on this representative sample through our clearinghouse. First, we selected farmers from this sample that had reported using a government veterinarian for their last artificial insemination, to ensure consistency with the clearinghouse and relevance of potential information provision. We then manually entered selected survey farmers' phone numbers into our clearinghouse to generate treatment randomization and a follow-up phone call. For control farmers this phone call included only data collection (verifying our survey data on their last AI visit). For treatment farmers this phone call included data collection *and* information provision.

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 had already been present in our clearinghouse and those who returned for service provision after treatment assignment retained their initial assignment.²⁰

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.²¹ 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.²² The clearinghouse then sent a follow-up SMS with the same information. We also gave farmers veterinarians' phone numbers and, if they requested it, information on average farmer-reported satisfaction with veterinarians on a 1-5 scale, and information on any other veterinarian in our system.

Figure 3 presents a timeline of the clearinghouse and survey data collection and randomization. The baseline survey occurred prior to our clearinghouse implementation, and the endline survey occurred immediately prior to the clearinghouse being shut down.²³

Table 1 reports the balance of our representative survey sample between treatment and control farmers on baseline outcomes and other important covariates. Appendix Table 9 further tests the balance on a large number of characteristics of treatment and control farmers. The study is fairly balanced on the pretreatment co-variates.

4.3 Experiment 2: Clearinghouse sample

The clearinghouse also ran automatically as designed during our study period. Anytime a government veterinarian recorded a service provision on our app, it began the process

²⁰Unfortunately, the coin used for randomization was shaved, due to a glitch in the clearinghouse algorithm. This resulted in 52% of farmers being treated. However, the probability of treatment remained fixed across farmers across time.

²¹When 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.

²²There can be overlap in the most successful veterinarians in terms of cows and buffalo.

²³We 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.

outlined in Figure 2. This resulted in a second ‘clearinghouse’ sample of farmers that were not necessarily in our survey sample (in practice very few were given we only surveyed 900 farmers of hundreds of thousands in the district) but for which we have administrative data on initial service provision, treatment information provision (or not), and subsequent service provision. Calls with these farmers were near identical to those to farmers we manually entered as part of our representative survey sample, and the treatment information provision component was identical.²⁴

We consider our clearinghouse sample experiment to be secondary to our representative survey experiment. This is because, while our clearinghouse sample experiment is very relevant to any policy discussions as it this is the natural sample that the clearinghouse built through automatic operation, the clearinghouse sample is selected on post-treatment outcomes: 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). See Appendix Section A.1 for a discussion of selection into the clearinghouse by veterinarians and farmers.

Table 2 reports the balance of our clearinghouse sample between treatment and control farmers. The study is balanced on pretreatment outcomes and the few other co-variables collected during clearinghouse data collection.

4.4 Empirical specifications

We use the following specification for our primary analysis for both experiments:

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

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. Because we administered treatment at the farmer level, we cluster standard errors at the farmer level.

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 the next day.

²⁴The only difference was that instead of asking about farmers’ last AI service recorded in our survey, up to 6 months prior, we asked about an AI service recorded by the veterinarian app exactly 1 and then 60 days prior. I.e. for our representative survey sample treatment occurred at the same time for all farmers but with different lags since the last service provision, where as treatment occurred at a different time for each clearinghouse sample farmer, but with the same lags since the last service provision. This means that the post-treatment period differs for each farmer.

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.²⁵

5 Results

In this section, we present results. First, we present treatment effects for our two experiments (Sections 5.1 and 5.2). Second, we study the farmers’ choice to return to a government veterinarian after treatment selection (Section 5.3). Third, we explore the possible channels of our treatment effects through heterogeneity analyses (Sections 5.4.1 and 5.4.2).

5.1 Representative sample treatment effects (Experiment 1)

Table 3 presents our primary intent-to-treat effects on prices and AI success rates. Columns (1) and (2) present results for the entire sample of farmers that had been using a government veterinarian at baseline (a requirement for treatment selection) and that returned for AI service provision, government or private, after treatment selection (i.e. have at least one post-treatment outcome). We find no significant impact on prices paid, though the direction of the effect is negative. We do, however, find a large, significant (at 10%) impact of treatment on AI success rates of 17 percentage points, or 25%. In fact, treatment farmers see their AI success rates raise to 84.7% on average, near the theoretical maximum.

Columns (3) through (6) break these treatment effects into those for farmers who either returned to a government veterinarian post-treatment selection (columns (3) and (4)) or instead returned to a private veterinarian (columns (5) and (6)). We see that the impact on AI success rates is being driven entirely by those that returned to government veterinarians after treatment. In this sample, AI success rates increase by 36 percentage points, or 63%.

These results in columns (3) through (6) should be taken with a grain of salt. Because the choice to return to a government veterinarian or not was made post treatment, this analysis could undo the balance created by our randomization in the first place.²⁶ We do show in Section 5.3 that those who returned to a government veterinarian and those who switched to a private veterinarian maintain balance on baseline observables, including baseline AI success rates, prices, and satisfaction, but we cannot rule out this analysis suffers an imbalance on unobservables. We believe, however, that the effect of information provision conditional on farmers returning to a government veterinarian is policy relevant and intellectually quite interesting. These conditional results speak toward the ability for our clearinghouse to increase the effort put in by veterinarians who are being ranked by the system (private veterinarians were not apart of the clearinghouse rankings), and could be

²⁵We 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.

²⁶This is often referred to as the “bad control” problem.

closer to the impacts we might expect if all veterinarians were included in the clearinghouse as would be the case in many markets.

Of course, the results in Table 3 columns (1) and (2) are the true, unconditional intent-to-treat effect of information provision on AI success rates. As such, we will use this for farmer welfare calculations rather than our speculative results conditioning on a post-treatment outcome.

5.2 Clearinghouse sample treatment effects (Experiment 2)

Once conditioning on farmers returning to a government veterinarian after treatment selection, we can also measure treatment effects using our clearinghouse sample.²⁷ Table 4 presents treatment effects using this data, in Panel B. In Panel A, it presents results using our survey sample, where columns (4) and (5) are identical to the previous table to allow for comparisons. In Panel B column (4) we see a much larger negative treatment effect on price, though it is also not significant. In Panel B column (5) we find a large impact of treatment on subsequent AI success of 20.9 percentage points, or 37%. We cannot reject that this effect is the same in magnitude to the effect we find in the survey sample. The fact that we find similar results in two independently drawn samples is particularly reassuring given our small sample sizes. This also supports our findings in Appendix Section A.1 that the clearinghouse sample is representative of farmer reports.

Table 4 also presents two additional results. In columns (1) and (2), it presents the effects of treatment on a farmer returning for any AI before the end of the study (including private) and for government AI specifically, respectively. In the case of returning at all in column (1), we find no significant impact though the coefficient is positive. We find more mixed results in the case of returning specifically to a government veterinarian for AI. In Panel A column (2), we find a treatment effect of 4.7 percentage points. In Panel B, we find a treatment effect of 3.1 percentage points, which is 44% of the control mean in the clearinghouse sample. However, the effect in Panel A is insignificant while the effect in Panel B is significant at 1%.²⁸ In column (3), we present the effect of treatment on a farmer switching which government veterinarian he/she used for AI conditional on returning to a government veterinarian post-treatment selection. In this case, we find a consistent zero impact across both panels. We will discuss this result below as this means the improvement in AI success rates for returning farmers came through specific veterinarians improving rather than farmers switching to better veterinarians.

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) using the clearinghouse data. The top panel illustrates that treatment farmers have higher AI success rates

²⁷By design, all of the clearinghouse sample data is susceptible to post-treatment selection and thus the potential “bad control” problem, though we argue in Appendix Section A.1 that selection into the clearinghouse at various stages of the experiment is not the main driver of the patterns that we see in this section.

²⁸The 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 Appendix Table 10 as well, only roughly 20% of return visits were recorded by veterinarians, so even in five months the true return rate is likely around 40%. Also, the fact that the representative survey does not suffer from recording error could explain the difference between panels A and B.

throughout the entire post period, while the bottom panel traces the size and significance of this treatment effect over the post period with bootstrapped standard errors. These results confirm our regression analyses.

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 over several months in the post period. This serves to strengthen our insignificant price effect result in Table 4.

Of course, all results beyond the initial ITTs in Table 3 columns (1) and (2) are conditional on selection. In the case of the representative survey, they are conditional on farmers returning to a government veterinarian after treatment selection. We will examine this margin of selection in the next subsection. In the case of the clearinghouse data, they are conditional on the same return by farmers as well as conditional on veterinarian and farmer reporting to the clearinghouse itself. We examine these margins of clearinghouse selection at length in Appendix Section A.1.

5.3 Farmers’ choice to return to government veterinarians after treatment selection

Table 5 compares those farmers that chose to return to a government veterinarian after treatment selection with those that chose to instead shift to a private veterinarian using our representative sample (all farmers considered for treatment had been contracting government veterinarians at baseline). Across 23 variables, the means between the two groups of farmers are insignificantly different from each other in 22 cases. In one case, the groups differ at the 5% significance level, though in this case the point estimates are precise and substantively quite similar (4.67 vs 4.47 for risk willingness out of 10) . We would argue this is consistent with sampling error and that the two groups look the same on observables. Importantly, farmers baseline AI success rate, visit charges, farmer satisfaction, and farmers’ estimates of their own government veterinarian’s average AI success rate are all balanced between the two groups. So we find no evidence that farmers that are unhappy with their veterinarians are switching. This is consistent with the results we find in the clearinghouse in Appendix Section A.1.

5.4 Heterogeneous Effects

In order to explore the potential channels driving our treatment effects, we present a series of heterogeneous treatment results.

5.4.1 Clearinghouse sample treatment effects by government veterinarian rank

First, in Table 6, we present treatment effects for two important sub-populations, separated according to the ranking of the last government veterinarian who served them—those for whom this veterinarian was ranked in the top three during their treatment phone call, 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.²⁹

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.^{30,31}

Perhaps the most surprising result in Table 6 is that farmers whose past veterinarian was not ranked in the top three are more likely to return. To investigate this, we show in Appendix Table 11 that farmers in Table 6 Panel B tend to live almost twice as far away from their closest veterinary center. 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.³²

5.4.2 Results using farmer expectations from the representative survey sample

If we are to believe our stylized model, we should expect the level of asymmetric information between veterinarians' baseline reputations and their true skill level 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 very little 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. We see no correlation for this group, suggesting there were no measurable information spillovers.

Point estimates for these two results are reported in Table 7. The null hypothesis that the coefficients in columns (2) and (3) are equal is almost rejected, with a p-value of 0.103. Note these measures were not incentivized. We simply asked farmers, "In general how many times in 10 do you believe that an artificial insemination from [your veterinarian] will be

²⁹Note 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%.

³⁰These results are suggestive because, while the point estimates are qualitatively different, we cannot reject this difference with significance.

³¹We should also expect heterogeneous treatment effects based on whether or not a farmer's past government veterinarian was ranked top 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.

³²In addition, these farmers have more buffalo. They also pay slightly more on average, which is consistent with larger 'show-up' fees due to higher veterinarian travel costs.

successful?” It is possible farmers misreported at baseline despite having accurate beliefs. Though we used the exact same question in the baseline and endline and for treatment and control farmers. So for the estimated treatment effect to be due to a change in misreporting, it would need to be that treatment did not change beliefs but changed something else such as the effort cost of forming an accurate belief on-the-fly (and thus not misreporting) at endline. If farmers are able to form accurate beliefs with less effort that would be sufficient to generate our results.

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 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 8 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 they are able to then bargain away rents from the veterinarian, in this case through prices while both groups benefit from increased AI success rates. Of course these results rely on few observations so we consider them speculative.

6 Discussion

6.1 Interpretation through the lens of our stylized model

Several of our results can be interpreted through the lens of our stylized model presented in Section 3 above. We can start with our primary result that treatment led to higher artificial insemination success rates at no higher prices, driven by farmers that stuck with their pre-treatment government veterinarians. Coupled with this are the suggestive results that it is farmers whose government veterinarian is revealed to be a low rank and/or worse than expected that benefit the most from treatment. While it is tempting to argue that these results point toward moral hazard mattering more than adverse selection as it is not farmers switching to *high* quality veterinarians after treatment seeing the biggest gains, if adverse selection is in terms of effort costs and not maximum achievable success rates as in the case of our model, the channel(s) driving these results is/are ambiguous. These results are consistent with some combination of (i) decreased adverse selection of veterinarians (i.e. fewer *low* quality veterinarians masquerading as *high*), (ii) veterinarians’ engaging in less moral hazard (i.e. fewer veterinarians putting in *low* effort when farmers pay for *high*), and (iii) improved bargaining power for farmers (i.e. farmers bargaining a better deal when they learn their current veterinarian is *low* quality and there are other *high* quality veterinarians nearby).³³

³³Government veterinarians knew we were collecting information on their artificial insemination success and prices (since the first step in this process was them using a smartphone to record service provision), but we did not tell them we were going to give this information to farmers. So any change in bargaining

Anecdotes compliment our finding that veterinarians can improve artificial insemination success rates deliberately, but they do no better to disentangle these channels. Take, for example, 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 veterinarians still perform artificial insemination and charge the farmer. Adverse selection here would be in terms of effort costs to transport semen at the correct temperature (canisters are heavy so strength could affect this cost). Moral hazard would be in so far as farmers do not know to look for canisters versus semen straws. And because farmers call veterinarians before transportation decisions are made, improved farmer bargaining power could manifest as veterinarians exerting more effort in transporting semen.

While it is not ideal that we cannot separate these channels in our research, this is also in keeping with the fact that it would be very hard to design a policy that targets one of these channels and not the others. Encouraging farmer monitoring of veterinarian artificial insemination performance could clamp down on moral hazard, but it could also allow farmers to learn veterinarian types vis-a-vis adverse selection. Prices could be better standardized to remove bargaining over price but it would be difficult to remove bargaining over effort. And so on.

We *can* rule out some channels of impact outlined in our stylized model. Our results are not consistent with adverse selection leading government veterinarians to exit the market (we see no exit on the extensive margin nor on the intensive margin in terms of frequency of artificial insemination provided by veterinarians). They are also not consistent with *high* quality veterinarians seeing better bargains through improved outside options. In other words, while we cannot determine exactly which channel is driving the impact on farmers and their livestock, we see no results that suggest veterinarians are made better off through treatment. We will explore this more through a consideration of the social welfare impacts of treatment in the next subsection.

One result requires some acrobatics to fit our model—that farmers who are treated and then switch to a private veterinarian are no better off. In our model, farmers are expected utility maximizing so it is possible for them to switch veterinarians without seeing gains but what seems less plausible would be farmers switching to a private veterinarian when there are identified *high* quality government veterinarians in the area and even *low* quality veterinarians which farmers have more bargaining power in dealing with. This result could be caused by some farmers interpreting treatment as a negative reputation shock for all government veterinarians (‘they are all so bad’). It could also be caused by some treatment farmers failing to reach a bargain with government veterinarians if they try to overplay their hand. Anecdotes also suggest private veterinarians in particular compete on a third dimension beyond price and effort—variety of semen. It is possible private veterinarians are responding to treatment farmers by offering different varieties of semen (e.g. sex selective,

power must have been farmer-driven, at least at first. Anecdotally veterinarians learned about treatment quite quickly, which is consistent with treatment farmers confronting them during bargaining. This suggests decreased adverse selection and/or less moral hazard would have had to come through the bargaining channel, at least at first.

different breeds, semen from livestock with higher milk production; all government veterinarians have an identical, single variety of semen), though usually such varieties of semen are quite expensive, targeted more toward large farms, and suffer from asymmetric information about quality. Future research could incorporate variety into the information collected by the clearinghouse. There is also no reason a third party could not set up a clearinghouse that rates all veterinarians, public or private.

6.2 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.³⁴

Benefit to farmers: if the treatment effect of 25% on AI success rates (.170 over a control mean of .677) translates into just 2.8% more calves born per year per farmer (i.e. if farmers with a failed 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 3,010 PKR (30 USD) per year, equal to almost half of one month's median income.³⁵ 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.³⁶

Cost to farmers: several results suggest that farmer treatment effects are not due to changes in farmer behavior toward their livestock that could come at a cost, including that treatment farmers see no better outcomes when they switch to a private veterinarian and when their veterinarian is revealed to be *high* quality. There could be changes in farmer effort costs from treatment through additional (or less) search or through additional (or less) bargaining. We did ask farmers at endline "On a scale of 1-10, 1 being as little as possible and 10 being a great amount, how much effort do you put into selecting a [veterinarian for AI]?" We see a small and insignificant treatment effect on this question. We take this as suggestive evidence treatment did not on net change effort put in by farmers toward their veterinarians. Thus we do not consider there to be a meaningful increase in average costs to farmers from this intervention.

Benefit to veterinarians: we have already argued in the subsection above that veterinarians are no better off as a result of treatment since we do not see higher prices or lower success rates (i.e. lower effort) for any sample.

Cost to veterinarians: it is likely that government veterinarians pay increased effort costs

³⁴We do not consider changes in price as such is a transfer with no net social welfare implications.

³⁵This 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 Pakistan, according to the World Bank, is 73.26 USD per month, accessed 10/8/2015.

³⁶Ideally we would have measured actual calf births and sales but unfortunately due to the delay between measured AI success and birth we do not have sufficient data for calf-level analysis. The only reason higher AI success would not translate 1-for-1 to additional calves would be because of livestock abortions. We have not found data on livestock abortion rates in Sahiwal but believe them to be relatively small as this is a setting where livestock is crucial and most are vaccinated against abortion-causing diseases endemic in other parts of the country. A separate concern would be that an influx of additional calves born due to treatment would lower market prices through general equilibrium effects. We do not believe the scale of this experiment would warrant such effects, but we do discuss them in the context of a scale-up at the end of this section.

to improve artificial insemination success rates as outlined in our stylized model. These could include the additional cost of carrying a liquid nitrogen container versus a semen straw, the opportunity cost of coming more quickly to a farm (though we find no change in the frequency of government veterinarian service provision), increased concentration costs, etc. At the same time, we find no evidence of increased travel costs (veterinarians must travel to farms regardless of treatment) and government veterinarians also do not spend any more time visiting treatment farmers to perform services (our smartphone application data collects this information). On net, we suspect the increased effort costs to government veterinarians are low in this context, and it is exactly this fact that allowed our treatment to be so successful.

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 30 USD per farmer from an intervention that cost 8 USD per farmer. This suggests a 275% return. While this does not include increased effort costs paid by government veterinarians, we believe them to be much smaller than this estimated return, and it also 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 which would increase the estimated return.

Another approach to understanding the potential social welfare effects of the clearinghouse is to look at price dispersion. The link between reduced asymmetric information, decreased price dispersion, and increased social welfare is outlined in Jensen (2007). In this case, however, farmers and veterinarians bargain on two dimensions simultaneously—price and quality. In Appendix Table 12, we thus examine the dispersion of price, quality (artificial insemination success), and price-adjusted quality (the number of artificial insemination successes per 100 PKR charged). While we do not find impacts on dispersion of price or quality separately, we find a strong decrease in dispersion of price-adjusted quality at the village-market level post-treatment. Unfortunately, this analysis is only suggestive as it is a simple pre-post comparison. We are not able to estimate average treatment effects for a market-level variable since our randomization was at the farmer level and at any point in time there was an equal proportion of treatment and control farmers in a given market. Nonetheless, these results serve as affirmation of our positive social welfare estimates through benefit-cost analysis and they again point towards the importance of considering both price and quality when understanding welfare.

Potential general equilibrium effects from scale-up: the accounting in this section has to this point been limited to the short-term effects from the evaluated intervention itself. But what might we expect if the intervention was scaled up to cover every government veterinarian in Punjab, Pakistan? Assuming such a scale-up does not induce entry and/or exit of government veterinarians and/or farmers in the market,³⁷ we would expect the costs and benefits to scale linearly with two exceptions. First, the average cost of the intervention would decrease to near-zero as it was scaled up. Second, we might expect a decrease in the market price for calves as many more are born and sold.³⁸ Specifically, if every farmer

³⁷We see no evidence of this during our evaluation.

³⁸We might also expect an increase in livestock fodder prices, which can be quite important to the returns of livestock (Anagol et al., 2017). This would benefit net-sellers of fodder and harm net-buyers. During our experiment, farmers report no additional money spent on fodder nor a change in behavior towards feeding

saw 2.8% more calves born per year, we would expect the total calves born across Punjab to increase by 1.2% (only 43% of farmers use government veterinarians). According to the 2018 Livestock Census of Punjab there was a near zero change in the number of livestock in Sahiwal between 2006 and 2018. Meanwhile, population (demand for meat and milk) grew rapidly (Punjab Livestock & Dairy Development Department, 2018). It is thus possible that there is excess demand and prices would not drop. But we can consider what would happen if they did. Appendix Table 13 presents estimates of farmer benefits to the intervention accounting for different general equilibrium price decreases. We see for price decreases up to 6 percent, this intervention would still be net positive for affected farmers, and likely net positive as a whole given average costs would be near zero. This would correspond to a price elasticity of greater than three which we believe is quite unrealistic. Of course, this is assuming all households are net sellers of calves and lose from a price decrease in calves. Net buyers of calves would gain and the overall social welfare implications are thus unclear.

We might also consider scale-up not in space but in time. The results in this study are over a relatively short period of time. In the long-run, we might imagine effects increase as farmers slowly change veterinarians and push the lowest quality to exit. Or, we might imagine effects decrease as veterinarians re-optimize and cut corners in behavior not measured by the clearinghouse. Long-run effects would also depend on the information quality of and trust in the clearinghouse.³⁹ While we have limited post data to test such hypotheses, we can see in Figure 4 that our result on AI success persists during our study. Using clearinghouse data, we can also test whether our results are localized to farmers that only show up in our clearinghouse once (one-time) versus those that show up multiple times (repeat). We might expect farmers that show up multiple times to be closer to the long-term steady state in terms of behavior. While an under-powered test, in Appendix Table 14 we see that in magnitudes the effects on these two groups in AI success rates are very similar.⁴⁰

7 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 and associated adverse selection and moral hazard. 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. While this confirmation is neither novel nor heartening in and of itself, it allows us to fit the livestock sector in Punjab into a context that is much more general. Adverse selection and moral hazard have been documented in numerous sectors, public and private, across the developing world. We might

their livestock.

³⁹We discuss selection of data in the clearinghouse during our study in Appendix Section A.1.

⁴⁰Given that being a one-time versus repeat farmer is conditional on coming back at all, we can not study this heterogeneity on returning to government AI as an outcome.

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 275% return to farmers on the cost of the intervention. This is driven by a 25% 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.

Finally, 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

- Akerberg, Daniel A (2003). “Advertising, learning, and consumer choice in experience good markets: an empirical examination.” *International Economic Review*, 44(3), 1007–1040.
- Aker, Jenny C (2010). “Information from markets near and far: Mobile phones and agricultural markets in Niger.” *American Economic Journal: Applied Economics*, 2(3), 46–59.
- Akerlof, George A (1970). “The market for” lemons”: Quality uncertainty and the market mechanism.” *The quarterly journal of economics*, pp. 488–500.
- Anagol, Santosh, Alvin Etang, and Dean Karlan (2017). “Continued existence of cows disproves central tenets of capitalism?” *Economic Development and Cultural Change*, 65(4), 583–618.
- Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja (2017). “Report cards: The impact of providing school and child test scores on educational markets.” *American Economic Review*, 107(6), 1535–63.
- Annan, Francis (2022). “Misconduct and reputation under imperfect information.” *Available at SSRN 3691376*.
- Bandiera, Oriana, Andrea Prat, and Tommaso Valletti (2009). “Active and Passive Waste in Government Spending: Evidence from a Policy Experiment.” *American Economic Review*, 99(4), 1278–1308.
- Banerjee, Abhijit V, Rukmini Banerji, Esther Duflo, Rachel Glennerster, and Stuti Khemani (2010). “Pitfalls of participatory programs: Evidence from a randomized evaluation in education in India.” *American Economic Journal: Economic Policy*, 2(1), 1–30.
- Basu, Arnab K, Nancy H Chau, and Ravi Kanbur (2009). “A theory of employment guarantees: Contestability, credibility and distributional concerns.” *Journal of Public economics*, 93(3), 482–497.
- Baumol, William J (1986). “Contestable markets: an uprising in the theory of industry structure.” *Microtheory: applications and origins*, pp. 40–54.
- Björkman, Martina and Jakob Svensson (2009). “Power to the people: evidence from a randomized field experiment on community-based monitoring in Uganda.” *The Quarterly Journal of Economics*, 124(2), 735–769.
- Callen, Michael, Saad Gulzar, Ali Hasanain, Muhammad Yasir Khan, and Arman Rezaee (2017a). “Personalities and Public Sector Performance: Evidence from a Health Experiment in Pakistan.” Nber working paper #21181.
- Callen, Michael, Saad Gulzar, Ali Hasanain, Muhammad Yasir Khan, and Arman Rezaee (2017b). “The Political Economy of Public Sector Absence: Experimental Evidence from Pakistan.” Nber working paper #22340.

- Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, and F Halsey Rogers (2006). “Missing in action: teacher and health worker absence in developing countries.” *The Journal of Economic Perspectives*, 20(1), 91–116.
- Dreze, Jean and Amartya Sen (1989). *Hunger and public action*. Oxford University Press on Demand.
- Duflo, Esther, Rema Hanna, and Stephen P Ryan (2012). “Incentives work: Getting teachers to come to school.” *The American Economic Review*, pp. 1241–1278.
- Fafchamps, Marcel and Bart Minten (2012). “Impact of sms-based agricultural information on indian farmers.” *The World Bank Economic Review*, 26(3), 383–414.
- Ferraz, Claudio and Frederico Finan (2011). “Electoral accountability and corruption: Evidence from the audits of local governments.” *American Economic Review*, 101(4), 1274–1311.
- Finan, Frederico, Benjamin A Olken, and Rohini Pande (2015). “The personnel economics of the state.” Tech. rep., National Bureau of Economic Research.
- Henderson, Ralph H and T Sundaresan (1982). “Cluster sampling to assess immunization coverage: a review of experience with a simplified sampling method.” *Bulletin of the World Health Organization*, 60(2), 253.
- Hölmstrom, Bengt (1979). “Moral hazard and observability.” *The Bell journal of economics*, pp. 74–91.
- Hubbard, Thomas N (2002). “How do consumers motivate experts? Reputational incentives in an auto repair market.” *The Journal of Law and Economics*, 45(2), 437–468.
- Israel, Mark (2005). “Services as experience goods: An empirical examination of consumer learning in automobile insurance.” *American Economic Review*, 95(5), 1444–1463.
- Jensen, Robert (2007). “The digital divide: Information (technology), market performance, and welfare in the South Indian fisheries sector.” *The Quarterly Journal of Economics*, pp. 879–924.
- Jin, Ginger Zhe and Phillip Leslie (2003). “The effect of information on product quality: Evidence from restaurant hygiene grade cards.” *The Quarterly Journal of Economics*.
- Kelley, Erin M., Gregory Lane, and David Schönholzer (2021). “Monitoring in Small Firms: Experimental Evidence from Kenyan Public Transit.” Working paper.
- Klein, Tobias J, Christian Lambertz, and Konrad O Stahl (2016). “Market transparency, adverse selection, and moral hazard.” *Journal of Political Economy*, 124(6), 1677–1713.
- Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria (2014). “Asymmetric information and middleman margins: An experiment with west bengal potato farmers.” Working paper.

- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar (2017). “General equilibrium effects of (improving) public employment programs: Experimental evidence from india.” Tech. rep., National Bureau of Economic Research.
- Muralidharan, Karthik, Paul Niehaus, Sandip Sukhtankar, and Jeffrey Weaver (2021). “Improving last-mile service delivery using phone-based monitoring.” *American Economic Journal: Applied Economics*, 13(2), 52–82.
- Olken, Benjamin A (2007). “Monitoring corruption: evidence from a field experiment in Indonesia.” *Journal of political Economy*, 115(2), 200–249.
- Olken, Benjamin A. and Rohini Pande (2012). “Corruption in Developing Countries.” *Annual Review of Economics*, 4, 479–509.
- Punjab Livestock & Dairy Development Department (2018). “Livestock Census Punjab 2018.”
- Reinikka, Ritva and Jakob Svensson (2004). “Local Capture: Evidence from a Central Government Transfer Program in Uganda.” *The Quarterly Journal of Economics*, 119(2), 679–705.
- Svensson, Jakob and David Yanagizawa (2009). “Getting prices right: the impact of the market information service in Uganda.” *Journal of the European Economic Association*, 7(2-3), 435–445.
- Wild, Lena, Vikki Chambers, Maia King, and Daniel Harris (2012). “Common Constraints and Incentive Problems in Service Delivery.” Tech. rep., Overseas Development Institute.
- World Bank (2004). *World Development Report 2004: Making Services Work for the Poor*. World Bank.

8 Tables and figures

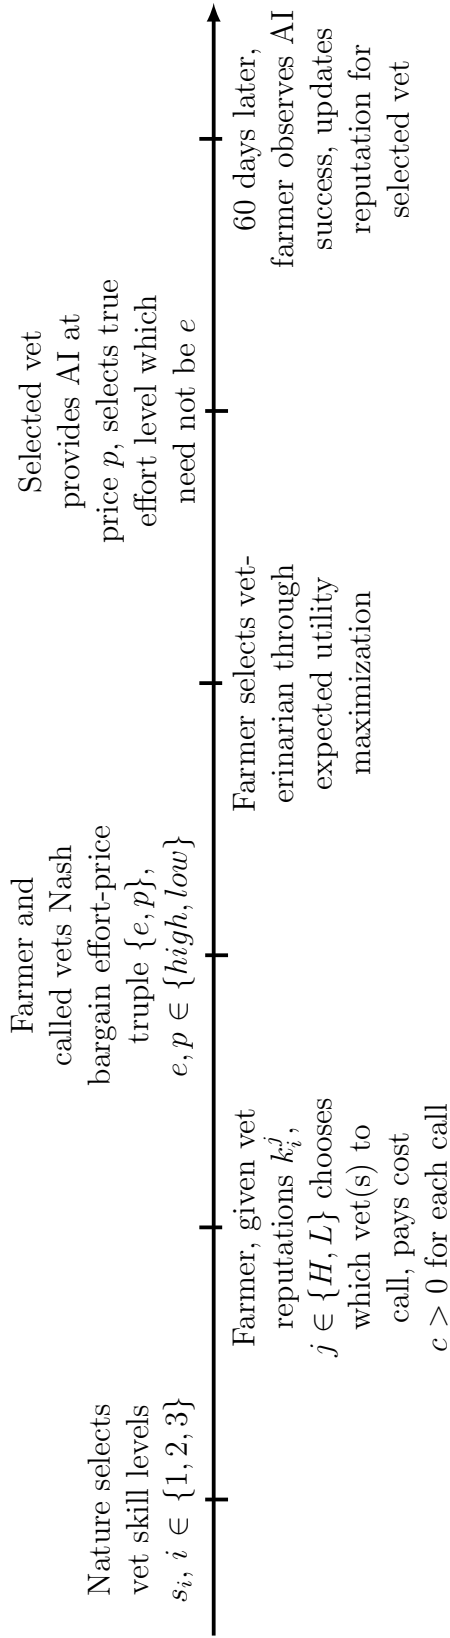


Figure 1: Sequence of decisions of a typical artificial insemination transaction

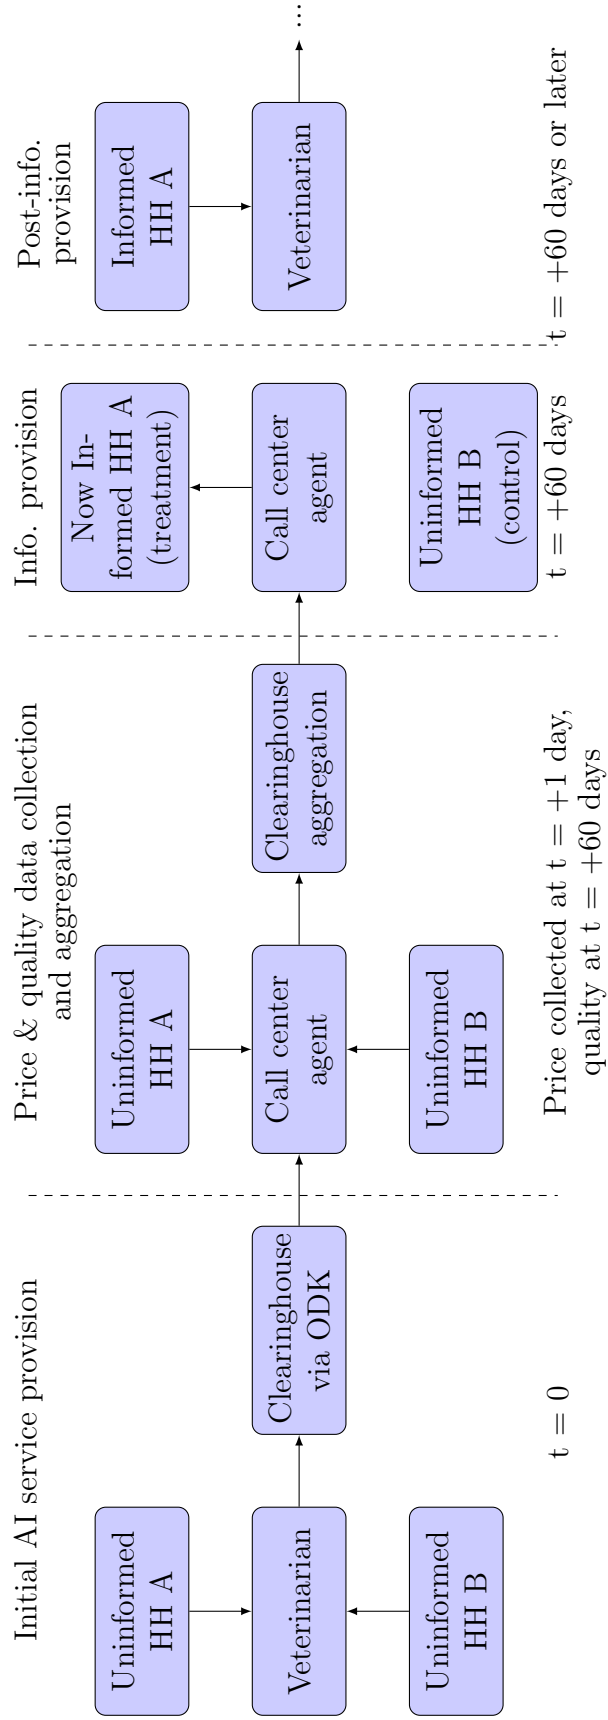


Figure 2: Clearinghouse flowchart . *Notes:* Arrows indicate the flow of information. The collection of quality data and information provision (for treatment farmers only) occur during the same follow-up phonecall 60 days after service provision. Beginning in October 2014, treatment farmers received information about the AI success rates of their local government veterinarians.

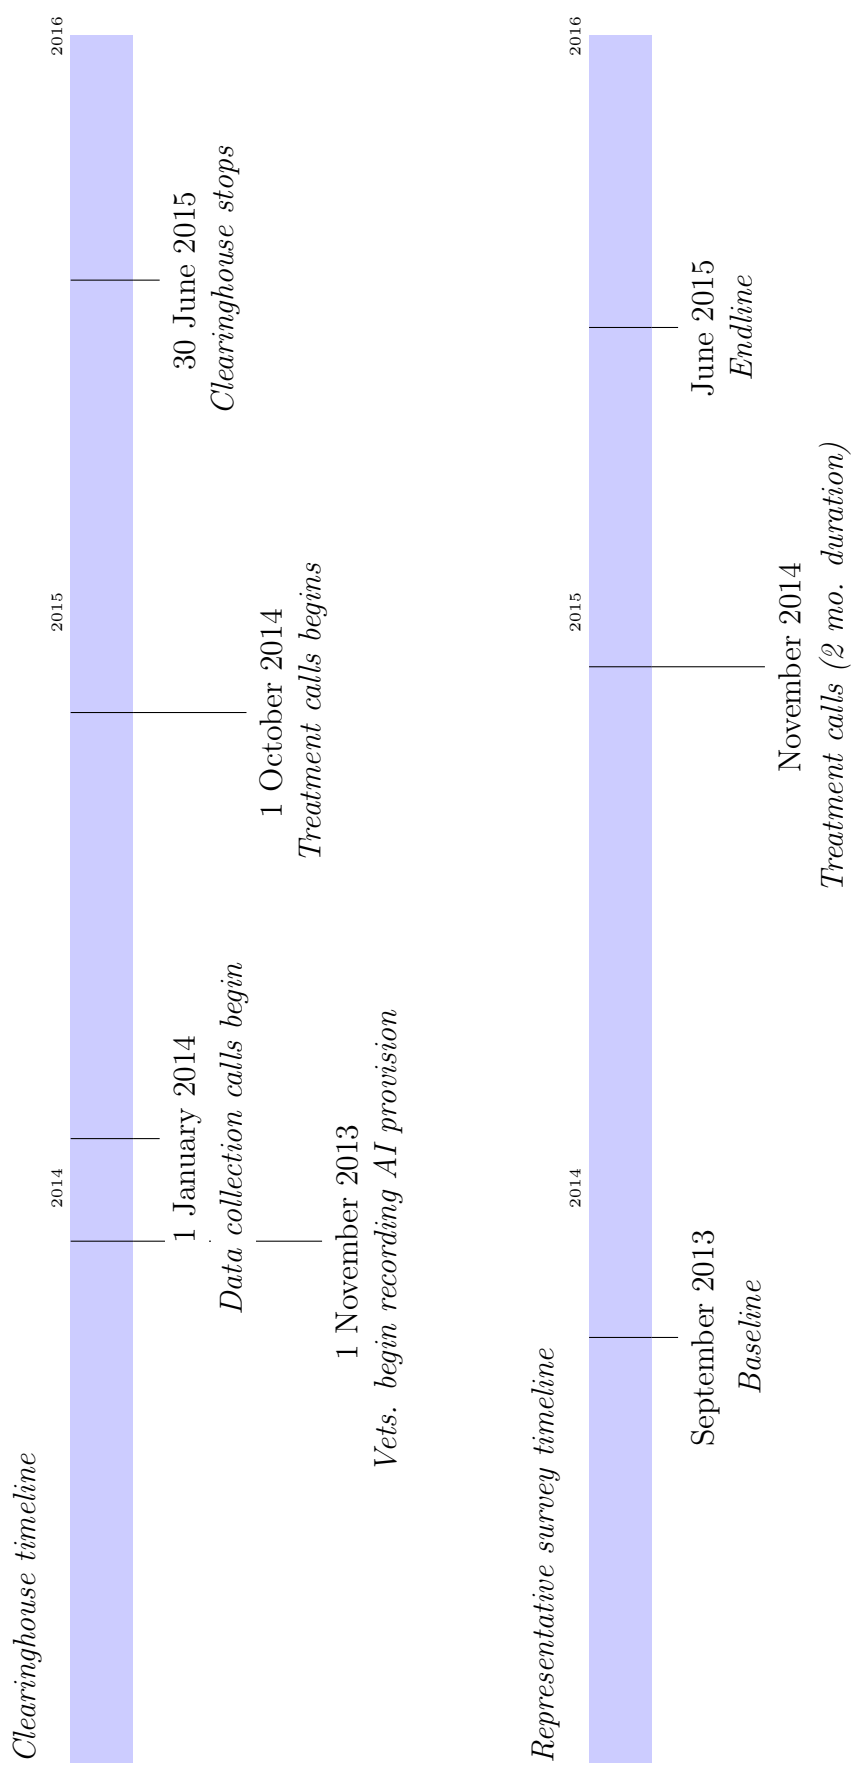


Figure 3: Clearinghouse and representative survey timelines

Table 1: Treatment balance—representative survey sample

	Treatment	Control	Difference	P-value
<u>Farmer-level baseline variables—190 observations</u>				
Livestock is primary source of HH's income (=1)	0.099 [0.300]	0.079 [0.271]	0.020 (0.038)	0.589
1-10 effort HH puts into selecting a vet. off.	6.195 [2.294]	5.490 [2.024]	0.705 (0.457)	0.236
Farmer attrited from in-person endline	0.018 [0.134]	0.026 [0.161]	-0.008 (0.020)	0.507
<u>Farmer-visit-level variables—356 pre-treatment observations from 190 farmers</u>				
Farmer switched vets since last recorded AI visit (=1)	0.182 [0.387]	0.198 [0.400]	-0.016 (0.051)	0.651
AI visit charges	374 [396]	364 [367]	10 (45)	0.981
AI visit success rate	0.689 [0.453]	0.760 [0.423]	-0.071 (0.042)	0.068
1-10 AI visit farmer satisfaction	7.689 [2.086]	9.091 [20.373]	-1.403 (1.456)	0.250
Farmer estimated AI visit vet success rate	6.556 [1.848]	6.586 [1.987]	-0.029 (0.242)	0.679

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 farmer. An F-test of joint significance of all covariates excluding primary outcomes has a p-value of 0.52. 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 2: Treatment balance—clearinghouse data

	Treatment	Control	Difference	P-value
Satisfaction with AI service provision (1-5)	4.185 [0.736]	4.138 [0.758]	0.048 (0.029)	0.187
Farmer switched vets since last AI visit	0.055 [0.228]	0.049 [0.216]	0.006 (0.010)	0.087
AI visit charges (PKR)	195 [171]	202 [245]	-7 (9)	0.484
AI visit success rate (pregnancy / AI attempts)	0.683 [0.459]	0.683 [0.459]	0.000 (0.016)	0.456
No of cows owned by farmer	2.510 [3.440]	2.414 [3.053]	0.096 (0.155)	0.275
No of buffalo owned by farmer	3.073 [3.761]	3.268 [6.364]	-0.195 (0.367)	0.744
Distance to closest AI center (km)	2.172 [2.255]	2.268 [2.257]	-0.096 (0.114)	0.747

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 farmer. An F-test of joint significance of all covariates excluding primary outcomes has a p-value of 0.27. The sample consists of 6,462 pre-treatment farmer-visit-level observations from 3,088 unique farmers. 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 3: Treatment effects—representative survey sample

Outcome:	Log price	AI success rate	Log price	AI success rate	Log price	AI success rate
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment farmer (=1)	-0.086 (0.126)	0.170* (0.097)	-0.106 (0.227)	0.357** (0.148)	-0.179 (0.171)	-0.005 (0.161)
Mean of dependent variable	5.872	0.677	5.852	0.581	5.888	0.765
# Observations	158	143	71	64	87	79
# Farmers	119	112	54	51	65	61
R-Squared	0.517	0.267	0.614	0.456	0.582	0.242
Sample	All	All	Returned gov't	Returned gov't	Left gov't	Left gov't

Notes: Standard errors clustered at the farmer level reported in parentheses. All regressions include randomization strata fixed effects and a control for the baseline mean outcomes of the dependent variable. 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. All indicates farmers that received government AI before treatment and subsequently returned for either government or private AI by the end of the project. Returned gov't indicates farmers that received government AI before treatment and subsequently returned for government AI after treatment by the end of the project. Left gov't indicates farmers who received government AI before treatment and instead subsequently received private AI by the end of the project. Price and AI success rates are recalled by farmers from service provisions two to seven months ago.

Table 4: Treatment effects for farmers who return for government AI

Outcome:	Returned at all	Returned to gov't	Switched veterinarians	Log price	AI success rate
	(1)	(2)	(3)	(4)	(5)
Panel A: Representative Survey					
Treatment farmer (=1)	0.092 (0.073)	0.047 (0.061)	-0.046 (0.112)	-0.106 (0.227)	0.357** (0.148)
Mean of dependent variable	0.518	0.246	0.152	5.852	0.581
# Observations	225	225	70	71	64
# Farmers	225	225	51	54	51
R-Squared	0.200	0.268	0.431	0.614	0.456
Panel B: Clearinghouse Data					
Treatment farmer (=1)		0.031*** (0.010)	0.001 (0.027)	-0.275 (0.245)	0.239** (0.107)
Mean of dependent variable		0.070	0.091	5.083	0.563
# Observations		3108	565	213	174
# Farmers		3108	262	143	122
R-Squared		0.095	0.305	0.703	0.514
Sample	Pre	Pre	Post	Post	Post

Notes: Standard errors clustered at the farmer level reported in parentheses. All regressions include randomization strata fixed effects and a control for the baseline mean outcome. 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. In Panel A, columns (2) through (4) 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. In Panel B, 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.

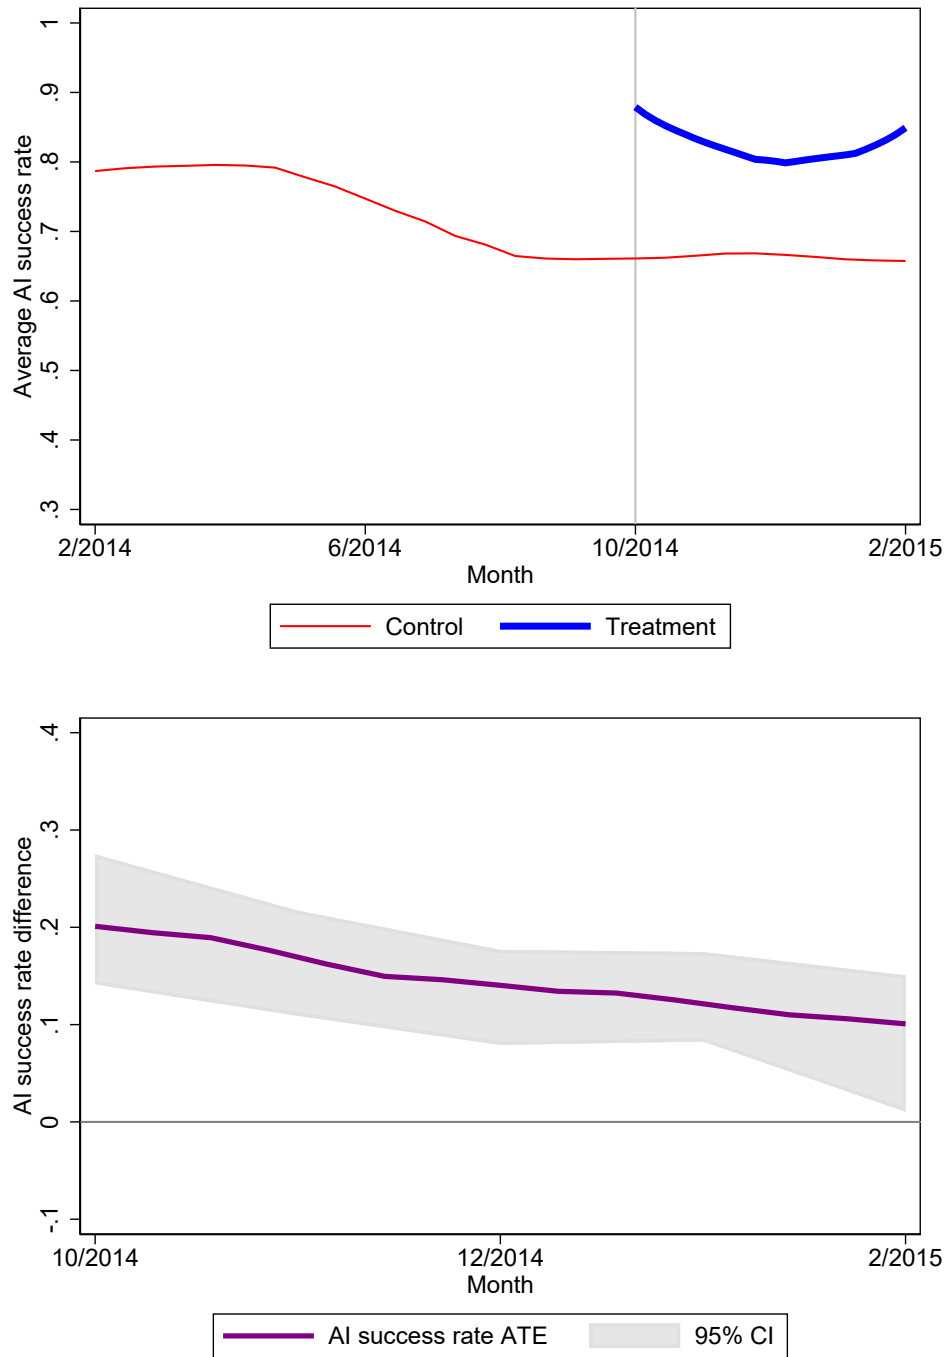


Figure 4: AI success rates in real time—clearinghouse data . *Notes:* The sample is farmers that received a government AI service and then answered the phone and reported AI success 60 days later. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Lines are smoothed using a kernel-weighted local polynomial regression with the Epanechnikov kernel and a rule-of-thumb bandwidth estimator. Confidence intervals are bootstrapped.

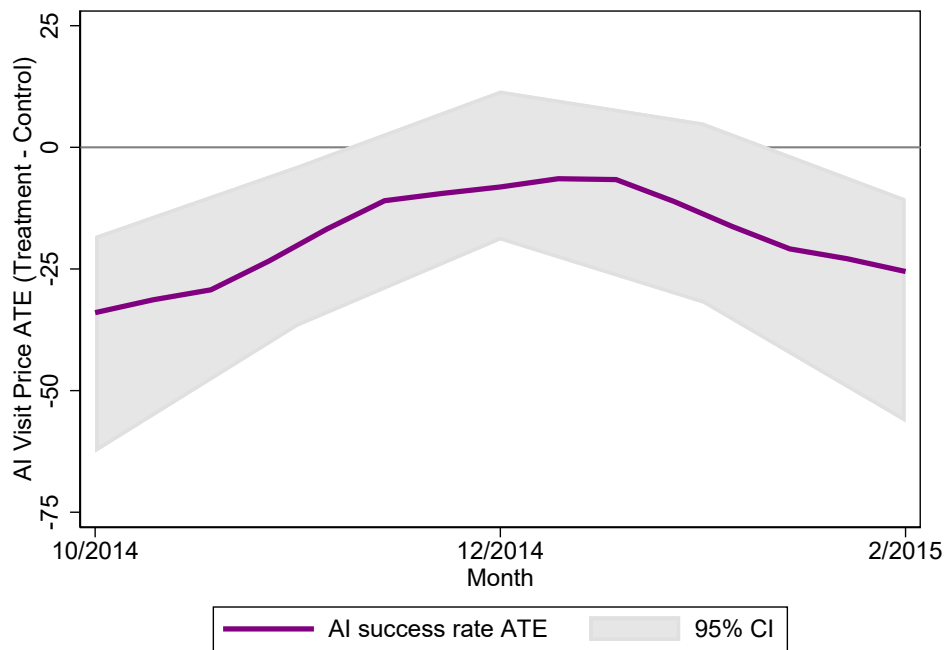
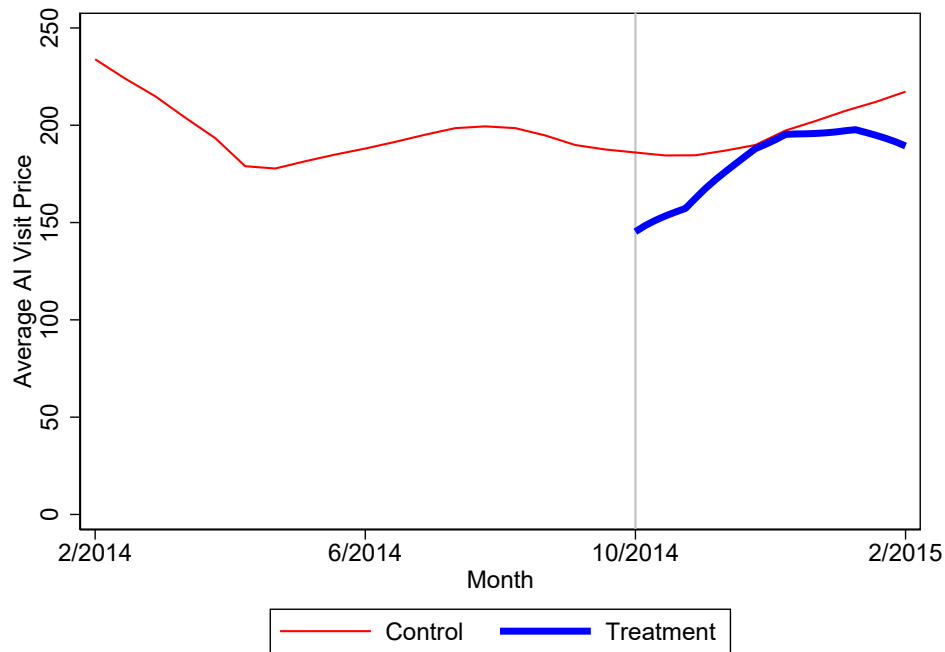


Figure 5: Price per AI visit in real time—clearinghouse data . *Notes:* The sample is farmers that received a government AI service and then answered the phone and reported price paid one day later. Beginning in October 2014 treatment farmers received information about the AI success rates of their local government veterinarians. Lines are smoothed using a kernel-weighted local polynomial regression with the Epanechnikov kernel and bandwidth one. Confidence intervals are bootstrapped.

Table 5: Comparing farmers that returned to a government veterinarian with those that returned to a private veterinarian—representative survey sample

	Returned to gov't	Returned to private	Difference	P-value
<u>Farmer-level baseline variables—124 observations</u>				
HoH education = None (=1)	0.400 [0.494]	0.323 [0.471]	0.077 (0.087)	0.793
A child in the HH attends public school (=1)	0.383 [0.490]	0.662 [0.477]	-0.278 (0.087)	0.217
HH has used govt health services in past 2 yrs (=1)	0.533 [0.503]	0.477 [0.503]	0.056 (0.090)	0.790
Amount of land HH owns and rents for livestock	1.269 [2.695]	1.310 [2.667]	-0.041 (0.532)	0.112
HH owns the house that they live in (=1)	0.983 [0.129]	0.923 [0.269]	0.060 (0.037)	0.168
Hours of electricity per day	10.322 [3.466]	10.354 [3.511]	-0.032 (0.627)	0.332
HH has a cooking stove/range (=1)	0.136 [0.345]	0.123 [0.331]	0.013 (0.061)	0.279
HH made less than 100k PKR last year (=1)	0.217 [0.415]	0.283 [0.454]	-0.067 (0.079)	0.121
Any member of HH has bank account (=1)	0.350 [0.481]	0.369 [0.486]	-0.019 (0.087)	0.453
Believed it was likely that last vote was not secret (=1)	0.685 [0.469]	0.509 [0.504]	0.176 (0.092)	0.557
Is likely to believe information given by govt employee (=1)	0.817 [0.390]	0.797 [0.406]	0.020 (0.071)	0.986
Average number of digits recalled	3.587 [0.985]	3.162 [1.092]	0.425 (0.278)	0.825
On a scale fo 0-10, how willing are you to take risks?	4.667 [2.541]	4.469 [3.267]	0.197 (0.601)	0.056
Agreeableness	4.279 [0.547]	4.031 [0.716]	0.248 (0.113)	0.293
Conscientiousness	4.221 [0.504]	4.185 [0.603]	0.036 (0.099)	0.776
Extroversion	4.217 [0.694]	4.112 [0.692]	0.105 (0.124)	0.106
Neuroticism	2.225 [0.682]	2.354 [0.792]	-0.129 (0.132)	0.137
Openness	3.750 [0.621]	3.615 [0.709]	0.135 (0.119)	0.209
1-10 effort HH puts into selecting a vet. off.	5.960 [1.837]	5.760 [2.204]	0.200 (0.574)	0.873
<u>Farmer-visit-level variables—239 pre-treatment observations from 124 farmers</u>				
AI visit success rate	0.743 [0.423]	0.703 [0.440]	0.040 (0.051)	0.695
AI visit charges	459 [507]	456 [546]	3 (82)	0.669
AI visit farmer satisfaction	7.926 [2.065]	7.771 [2.031]	0.156 (0.285)	0.184
Farmer estimated AI visit vet success rate	6.651 [2.006]	6.634 [1.994]	0.017 (0.290)	0.287

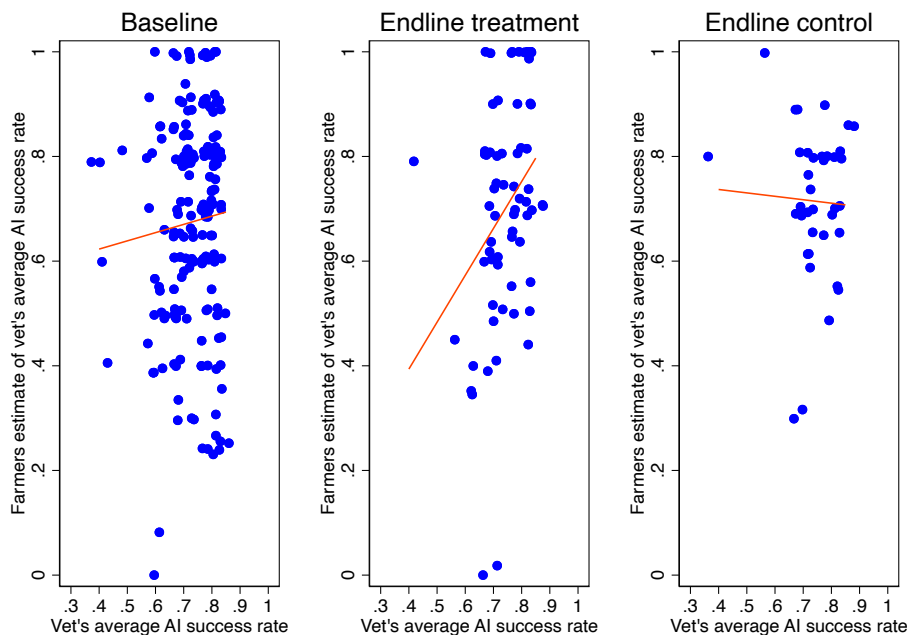
Notes: Standard deviations reported in brackets. Standard errors reported in parentheses. Means and differences are unconditional. P-values are from OLS regressions with randomization strata fixed effects and standard errors clustered at the farmer level. Some regressions have fewer observations due to missing data. An F-test of joint significance of all covariates excluding primary outcomes and those with above 20 percent missing values (livestock land, likely to believe gov't info, digit recall, risk willingness, effort into selecting vet.) has a p-value of 0.13. All data come from baseline surveys fielded in August and September 2013. 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.

Table 6: 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)	0.018 (0.012)	-0.006 (0.032)	0.013 (0.082)	0.093 (0.130)
Mean of dependent variable	0.072	0.108	4.895	0.661
# Observations	1970	404	132	105
# Farmers	1970	167	88	76
R-Squared	0.086	0.346	0.892	0.648
Panel B: Farmers told vet. was not in top three in area				
Treatment farmer (=1)	0.039** (0.017)	-0.004 (0.060)	-0.905 (1.337)	0.440** (0.209)
Mean of dependent variable	0.058	0.038	5.690	0.273
# Observations	1083	152	73	62
# Farmers	1083	88	48	40
R-Squared	0.114	0.499	0.751	0.792
Sample	Pre	Post	Post	Post

Notes: Standard errors clustered at the farmer level reported in parentheses. All regressions include randomization strata fixed effects and a controls for the baseline mean outcome. 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 7: 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	0.157 (0.204)	0.896* (0.496)	-0.065 (0.307)
# Observations	191	67	37
# Farmers	109	46	28
R-Squared	0.005	0.095	0.001
Sample	Baseline	Endline T	Endline C

Notes: Standard errors clustered at the farmer 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 0.103 from a regression interacting Vet’s actual average AI success rate with a treatment indicator in the Endline sample.

Table 8: 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 below median expected-actual AI success				
Treatment farmer (=1)	-0.375 (0.273)	-0.065 (0.086)	0.256 (0.303)	0.516*** (0.171)
Mean of dependent variable	0.643	0.214	5.720	0.545
# Observations	29	34	32	25
# Farmers	29	22	21	17
R-Squared	0.689	0.492	0.696	0.489
Panel B: Farmers with above median expected-actual AI success				
Treatment farmer (=1)	-0.140 (0.145)	-0.022 (0.257)	-1.109*** (0.316)	0.615** (0.271)
Mean of dependent variable	0.800	0.125	5.893	0.533
# Observations	29	28	26	23
# Farmers	29	23	22	20
R-Squared	0.769	0.344	0.729	0.516
Sample	Pre	Post	Post	Post

Notes: Standard errors clustered at the farmer level reported in parentheses. All regressions include randomization strata fixed effects and a control for the 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 farmers' estimate of their veterinarian's average AI success rate minus veterinarian's actual average AI success rate from clearinghouse data before October 2014. Negative values in this difference occur when farmers are told their veterinarian is better than they expected. Positive values occur when farmers are told their veterinarian is worse than they expected. The median is .012.

A Online appendix

A.1 Selection into the clearinghouse

In this section, we test for a number of ways in which clearinghouse sample data can be biased by treatment. We first explore whether treatment induced veterinarian reporting bias (Appendix Section A.1.1) and we then study if farmers selection and reporting is biased (Appendix Section A.1.2).

The goal of this section is twofold. First, we seek to establish that our clearinghouse data is not subject to selection (on observables) and thus treatment effects estimated using this data can be believed. Second, and more importantly given that we have the representative survey to ensure unbiased ITTs, we seek to understand farmer and veterinarian interaction with the clearinghouse system, which we believe is quite interesting in and of itself.

A.1.1 Treatment does not induce a veterinarian reporting bias

We might worry that veterinarians selectively report artificial insemination service provision to our clearinghouse through their smartphone app. For example, knowing that farmers of reported artificial insemination service provisions will be called and asked about their service, veterinarians may only report when they put in high effort and their artificial insemination is likely to be successful. Or veterinarians might not report service provision when they bargain a higher price from a farmer so as to not scare away potential customers (or a lower price if they are hoping to appear high quality and thus expensive). If treatment induces changes in bargained price and/or artificial insemination success, then this selective reporting would be correlated with treatment and it would bias treatment effect estimates.⁴¹

In Appendix Table 10, we find that 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 20% of the time, they are as likely to report for treatment and control farmers. And there is not a statistically significant difference in the likelihood that farmers that had their service provision data submitted to the call center report receiving a phone call to verify that service (those not reporting receiving a call may have not picked up, may have refused the call, or may not remember). These results suggest that veterinarians are not biasing the clearinghouse data by choosing only to report when they are sure they were successful. Though of course they are often not reporting at all. Anecdotally many veterinarians struggled to maintain the phone charge and/or often left it at home with their family as this was the first smartphone they had ever had.

⁴¹This is distinct from the external validity concern that data collected by the clearinghouse is not accurate, which we discuss in the following subsection.

A.1.2 Ruling out farmer selection and reporting biases in the clearinghouse sample

We might imagine that farmers that have experienced higher AI success rates and/or lower prices with a government veterinarian could be more likely to select back into government AI after treatment selection, biasing the clearinghouse sample. Similarly, even after farmers return for government AI, they have to answer the phone multiple times to report to the clearinghouse on the prices they paid and whether the AI was successful. We could imagine certain farmers may be more or less likely to report—e.g. those that were less happy might be more likely to answer and to complain, or vice versa.

In Appendix Table 15, we show that the characteristics of control farmers change little with each additional selection margin, the exception being that there is a drop off in response rates by farmers who live farther away from the closest AI center both in answering the phone for the first time and then in subsequently returning for AI (and having their return service recorded by their government veterinarian). This is not surprising, as farmers living far away from AI centers may not use the public AI as frequently as the ones living nearby. It is also possible these farmers far away differ on economic fundamentals such as their unemployment rate.⁴²

We, of course, might also worry that treatment itself biases the returning clearinghouse sample. In Appendix Table 16 we show treatment balance on all measured pre-treatment outcomes for those farmers who return for government AI and answer subsequent phone calls. 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.

B Online appendix tables and figures

⁴²While we are confident that the information we provided farmers in this setting was valuable (and our results affirm this), selection into who answers the phone, especially if it is correlated with economic fundamentals, could affect information “quality” in so far as we consider quality as distance away from the true mean. In the short-run, and so for the analysis done in this paper, we would argue what matters is not the true quality but perceptions of quality by farmers and veterinarians as this will affect bargaining and decisions around adverse selection and moral hazard. In the long-run we might expect true quality to matter more as farmers’ own experienced N could be larger than the N for our clearinghouse. Of course, the average N at the time of treatment is 65 and this number is only increasing in time at a much faster rate than farmers’ own N.

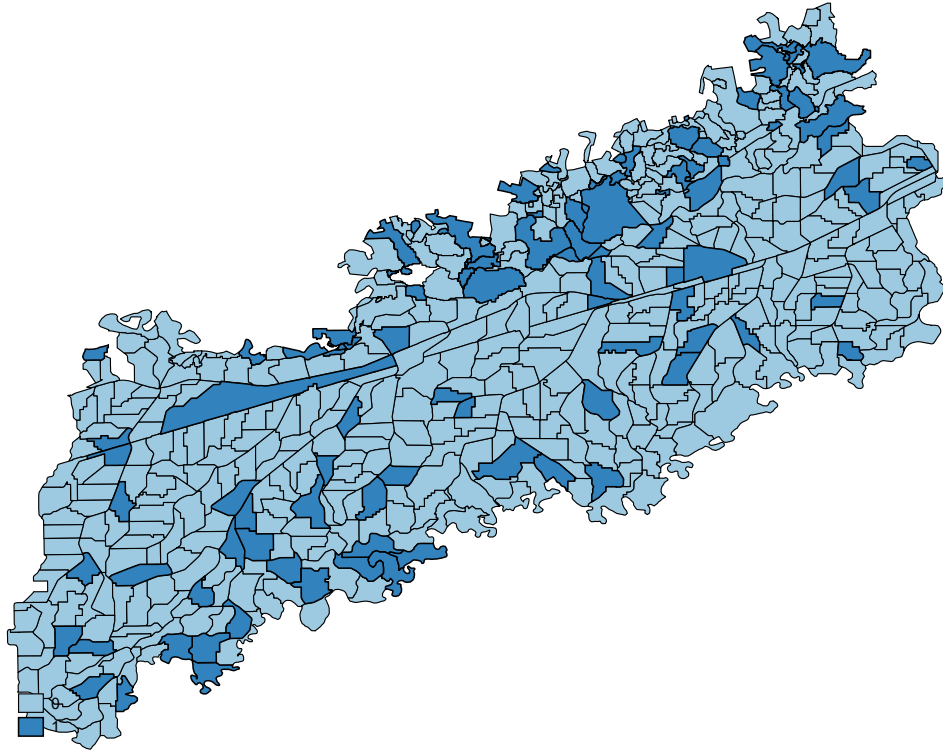


Figure 7: Representative Survey sample villages . *Notes:* Sampled villages are dark blue. The sample was stratified by whether or not a government veterinarian center was in the village and on whether the village was a canal colony. It is balanced along the following variables: area, settled area, cultivated 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 9: Treatment balance—representative survey sample, additional covariates

	Treatment	Control	Difference	P-value
Head of household education = None (=1)	0.387 [0.489]	0.430 [0.497]	-0.042 (0.066)	0.486
A child in the HH attends public school (=1)	0.532 [0.501]	0.500 [0.502]	0.032 (0.067)	0.622
HH has used govt health services in past 2 yrs (=1)	0.495 [0.502]	0.482 [0.502]	0.013 (0.067)	0.889
Amount of land HH owns and rents for livestock	1.795 [3.795]	1.166 [2.164]	0.629 (0.490)	0.161
HH owns the house that they live in (=1)	0.964 [0.187]	0.965 [0.185]	-0.001 (0.025)	0.568
Hours of electricity per day	11.046 [3.201]	10.106 [3.579]	0.940 (0.455)	0.030
HH has a cooking stove/range (=1)	0.081 [0.274]	0.150 [0.359]	-0.069 (0.043)	0.148
HH made less than 100k PKR last year (=1)	0.275 [0.449]	0.336 [0.475]	-0.061 (0.062)	0.748
Any member of HH has bank account (=1)	0.279 [0.451]	0.316 [0.467]	-0.037 (0.061)	0.441
Believed it was likely that last vote was not secret (=1)	0.531 [0.502]	0.583 [0.496]	-0.052 (0.070)	0.666
Is likely to believe information given by govt employee (=1)	0.809 [0.395]	0.772 [0.421]	0.037 (0.055)	0.743
Average number of digits recalled	3.306 [0.929]	3.317 [1.045]	-0.011 (0.189)	0.874
On a scale fo 0-10, how willing are you to take risks?	4.469 [3.107]	4.060 [2.729]	0.409 (0.457)	0.524
Agreeableness	4.072 [0.725]	4.011 [0.750]	0.061 (0.099)	0.548
Conscientiousness	4.086 [0.570]	4.164 [0.587]	-0.079 (0.078)	0.279
Extroversion	4.212 [0.657]	4.115 [0.649]	0.097 (0.088)	0.659
Neuroticism	2.372 [0.788]	2.383 [0.773]	-0.011 (0.105)	0.938
Openness	3.689 [0.683]	3.669 [0.652]	0.020 (0.090)	0.772

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 farmer level. The sample consists of 190 baseline farmer-level observations. Some regressions have fewer observations due to missing data. An F-test of joint significance of all covariates excluding primary outcomes and those with above 20 percent missing values (livestock land, likely to believe gov't info, digit recall, risk willingness, effort into selecting vet.) has a p-value of 0.22. 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 10: Does treatment induce a veterinarian reporting bias?

	Treatment	Control	Difference	P-value
Farmer reported AI and vet submitted data to call center (=1)	0.234 [0.425]	0.159 [0.367]	0.075 (0.067)	0.164
Farmer reported receiving a call verifying AI service (=1)	0.379 [0.494]	0.385 [0.506]	-0.005 (0.207)	0.281

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 farmer level. The sample consists of 316 farmer-visit-level observations from 175 unique farmers from 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. "Farmer reported AI and vet 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. "Farmer reported receiving a call verifying AI service" is conditional on a vet submitting data to the call center.

Table 11: Comparing farmers by pre-treatment veterinarian ranking—clearinghouse data

	Vet. in top three	Vet. not top three	Difference	P-value
Satisfaction with AI service provision (1-5)	4.171 [0.735]	4.144 [0.768]	0.026 (0.030)	0.952
Farmer switched vets since last AI visit	0.054 [0.226]	0.076 [0.265]	-0.022 (0.013)	0.894
AI visit charges (PKR)	191 [170]	211 [269]	-19.597 (11)	0.021
AI visit success rate (pregnancy / AI attempts)	0.624 [0.478]	0.635 [0.477]	-0.010 (0.017)	0.382
No of cows owned by farmer	2.344 [3.155]	2.612 [3.659]	-0.268 (0.156)	0.532
No of buffalo owned by farmer	2.757 [3.148]	3.442 [5.957]	-0.684 (0.259)	0.038
Distance to closest AI center (km)	1.710 [1.573]	3.253 [2.951]	-1.543 (0.120)	0.000

Notes: Standard deviations reported in brackets. The sample consists of 4,784 pre-treatment farmer-visit-level observations from 2,978 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.

Table 12: Price and success rate dispersion after treatment

	SD(AI Success Rate)	SD(Charges per AI)	SD(AI Successes / 100 PKR)
Post treatment (=1)	-0.050 (0.030)	2.674 (49.051)	-0.066*** (0.024)
Mean of dependent variable	0.424	247.550	0.271
# Observations	127	129	129
# Village clusters	69	69	69
R-Squared	0.019	0.000	0.049

Notes: Standard errors clustered at the village cluster level reported in parentheses. Mean of dependent variable calculated pre-treatment. Observation is at the village cluster level. While we administered treatment at the farmer level, treatment information provision was localized at the village-cluster level. 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.

Table 13: Cost-effectiveness and general equilibrium price effects

Decrease in prices	Per-farmer benefit of intervention
None	30 USD
1 percent	25.29 USD
2 percent	20.58 USD
3 percent	15.88 USD
4 percent	11.17 USD
5 percent	6.46 USD
6 percent	1.75 USD
7 percent	-2.96 USD

Notes: Estimates explained in Section 6.2.

Table 14: Treatment effects for one-time vs repeat customers—clearinghouse data

Outcome:	Switched veterinarians	Log price	AI success rate
	(1)	(2)	(3)
Panel A: One-time farmer customers			
Treatment farmer (=1)	0.007 (0.041)	0.298 (0.225)	0.262 (0.286)
Mean of dependent variable	0.060	5.025	0.455
# Observations	176	89	72
# Farmers	176	89	72
R-Squared	0.501	0.802	0.660
Panel B: Repeat farmer customers			
Treatment farmer (=1)	0.005 (0.037)	0.035 (0.092)	0.273** (0.133)
Mean of dependent variable	0.105	5.143	0.694
# Observations	389	124	102
# Farmers	86	54	50
R-Squared	0.392	0.965	0.654
Sample	Post	Post	Post

Notes: Standard errors clustered at the farmer level reported in parentheses. All regressions include randomization strata fixed effects and a controls for the baseline mean outcome. The sample for columns (1) through (3) 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. 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 only received services recorded by the clearinghouse once during our post period or whether they received services two or more times.

Table 15: Farmer selection in to clearinghouse at various stages

	All control	Answered once	Answered twice	Returned for AI	Returned & answered once	Returned & answered twice
No of cows owned by farmer	2.03 [2.43]	2.06 [2.47]	2.08 [2.53]	2.32 [2.51]	2.35 [2.84]	2.37 [2.87]
No of buffalo owned by farmer	2.61 [3.73]	2.68 [3.92]	2.71 [3.80]	3.07 [4.13]	2.86 [3.50]	2.90 [3.60]
Distance to closest AI center [km]	6.22 [143.63]	3.36 [54.90]	3.55 [60.02]	2.23 [2.30]	2.16 [2.22]	2.14 [2.25]
Satisfaction with AI service provision [1-5]		4.16 [0.77]	4.17 [0.76]	4.20 [0.73]	4.16 [0.76]	4.16 [0.77]
Log AI visit charges [PKR]		4.91 [1.40]	4.92 [1.39]	4.88 [1.35]	4.90 [1.24]	4.91 [1.20]
AI visit success rate [pregnancy / AI attempts]			0.68 [0.46]	0.71 [0.45]	0.69 [0.45]	0.70 [0.45]
# Observations	16,345	4,263	3,591	1,039	557	462

Notes: Sample: Non-treatment farmers from the clearinghouse sample. Standard deviations in parenthesis.

Table 16: Treatment balance of returning sample—clearinghouse data

	Treatment	Control	Difference	P-value
Pre-treatment mean satisfaction with AI service provision (1-5)	4.190 [0.685]	4.258 [0.709]	-0.068 (0.087)	0.473
Pre-treatment mean veterinarian switching rate	0.063 [0.235]	0.034 [0.216]	0.029 (0.025)	0.110
Pre-treatment mean log AI visit charges	4.791 [1.347]	4.805 [1.405]	-0.014 (0.172)	0.732
Pre-treatment mean AI success rate	0.689 [0.446]	0.671 [0.441]	0.018 (0.056)	0.707
Pre-treatment mean no. of cows	2.745 [2.790]	3.488 [2.361]	-0.743 (0.445)	0.117
Pre-treatment mean no. of buffalo	3.489 [3.371]	3.737 [4.102]	-0.248 (0.517)	0.392
Pre-treatment mean distance to closest AI center (km)	2.394 [2.171]	1.949 [2.177]	0.444 (0.259)	0.946

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 farmer level. The sample consists of 262 observations of 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. An F-test of joint significance of all covariates excluding primary outcomes has a p-value of 0.53. 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.