Selective and Mutlidimensional Learning
A case study of bt cotton farmers in India

Srijita Ghosh

July 31, 2018

Abstract

After the introduction of genetically modified bt cotton seeds cotton farmers in India overwhelmingly switched to the more profitable bt seed variety but failed to adjust the level of pesticide use leading to lower profit. Since over use of pesticides cannot be explained by financial or learning constraint we consider a decision problem involving learning constraint for a rational producer. The producer needs to choose more than one input to maximize payoff. He faces the choice of learning about each input unconditionally, i.e, learning about best seed, or learn conditionally about an input combination, i.e, mix of seed and pesticide, subject to two different costs of learning. We find that it is optimal for the producer to learn unconditionally when uncertainty of the belief is higher and conditionally when uncertainty is lower. This implies the optimal learning strategy is to choose either only conditional or unconditional learning or to start with unconditional learning, e.g., best seed variety, then switch to conditional, i.e, given the best seed variety optimal use of pesticide. Finally we show reducing the cost of only one type of learning can reduce the amount of learning.

1 Introduction

Most production technologies require multiple inputs. For example in agriculture, a farmer needs to decide seed variety, fertilizer, timing of sowing and harvesting, irrigation, pesticide and so on. In context of developing economy this feature of agricultural technology has significant impact on policy making.

One of the major policy goals for these countries is to increase agricultural productivity by technology adoption and/or optimal use of inputs. But previous literature has found that in many contexts farmers do not use the optimal level of inputs even when there is possibility of significant increase of profitability (see Mullianathan et al(2014)).

In this paper we will study one particular case of sub-optimal choices by bt cotton farmers in India. Cotton is one of the most important cash crop in India and is cultivated in 9 states in West and South India with significant regional heterogeneity in practices and productivity. It is also the major raw material of the textile industry which the second largest industry(in terms of GDP) in India.

Bt cotton was introduced as an alternative to pesticide for cotton production. This is a type of transgenic cotton seed that is resistant to cotton bollowrm, the major pest for cotton. Use of bt cotton significantly reduced the cost of producing cotton and increased average profitability.

Before the introduction of bt cotton in 2002 for commercial purposes, one of the main challenges the cotton farmers across India used to face is the risk pest infestation. In 1990's when only 5% of total cropped area was used for cotton production, 45% of pesticides were being used for cotton to protect the plants against pests. The average productivity was also significantly low compared to other cotton producing countries in the world [1].

The adoption rate for bt cotton was significantly high in India despite debates over impact of bt cotton in media and policy discussion. After being commercially introduced in 2002 when only about 1% of total cotton production was from bt cotton by 2014-15 95% of cotton production was by bt seeds. This success of adoption however did not translate well into increased productivity across the country.

One important observation is that farmers who switched to bt cotton kept using high level pesticide. Kiersur and Ichangi (2011) found that in Karnataka farmers under the guidance of extension workers used *less* pesticide which lead to reduction of total cost by 11.6% on average.

We want to explain the rationale behind the sub-optimal over use of pesticides. Even though several explanations are plausible for sub-optimal behavior over use of an input contradicts the assumption of financial constraint or supply side constraint. Hence, we explore the role of learning constraints on the farmers in this paper.

In this paper we focus on two major features of bt cotton farming. One, cotton farmers overwhelmingly adopted the new technology but did not optimize their choice of other inputs and two, the inefficiency in choice of inputs conditional on adoption of technology cannot be explained by financial or supply-side constraints.

The second observation leads us to a learning constraint type model for a rational DM. The learning constraint can be due to costly information or the farmers cognitive ability to process information. In the Indian context, before bt cotton was introduced commercially in 2002 the department of agriculture (DoA) of Governement of India conducted several small and large scale field experiment between 1995 (when Monsanto the bt cotton seed company introduced bt cotton in world market) and 2001 and have already concluded about the optimal use of different inputs, especially pesticide use and length of cycle (two major factors affecting productivity). Taking into consideration the caveat of external validity the farmers should have had access to the relevant information, thus the learning failure is not due to lack of information.

This implies we need to consider a model of learning where the DM's cognitive constraint is the main obstacle to learning similar to the recent literature of attention [2]. In all these models DM faces a cognitive cost for updating beliefs.

In this paper we assume the DM faces a similar cognitive cost for updating but with a major difference. In all these attention models the DM does not have an option to learn about only one component unconditionally neither would it be optimal. But to capture the feature of the farmers' problem in this paper we assume the DM can learn about one input unconditionally and cost of doing that is different from that of learning about input combination.

We write the farmer's problem as follows: he wants to maximize his payoff from farming but does not know the true production function. There are two possible ways to learn about the production function: one, to learn about the payoff generated by a combination of inputs, which gives him precise information about the payoff from the combination but does not inform him about any other combination, we would call this *depth strategy*. The other way is to learn about the payoff from an input unconditionally, though this strategy does not give precise information about any input combination (in general), it gives some aggregate information (e.g. average) about a set of input combinations, i.e., all combinations involving the input in question. We would call this *breadth strategy*

In our bt cotton case, the farmer has two possible learning strategies. One he can learn about

[1] While India was the third largest cotton producer in the world in 1990 only after China and US, yield per acre was 2.5 times higher in US compared to India. (source: National Cotton Council of America)

[2] See Sims(2001), Caplin and Dean 2014, Caplin et al 2017, Herbert and Woddford (2017), Mullianthan et al(2014, 2017), Schwarzstein(2014), Gabaix(2014) etc.

one particular combination of seed, pesticide etc and other to learn about seeds or pesticides separately, i.e, whether *on average* bt seeds are better than conventional seeds, what level of pesticide is best *on average*. The cost of these two strategies can be very different.

The main question we ask in this paper is what is the optimal learning strategy for the DM? More specifically what is the trade-off between the two types of learning strategy. More generally, we want to study how does the multi-dimensionality of the problem affects the learning choice of the DM when he has to choice to learn selectively about some features of the model.

The main trade off between the two types of learning are as follows: depth strategy uncovers one combination at a time. In the case the particular observation leads to low payoff not much information is revealed about other possible combination but this strategy guarantees reaching to the best combination if DM keeps learning for sufficiently many periods. Breadth strategy may not lead to the input combination with highest payoff since the farmer is learning about each inputs separately unconditionally but observing the unconditional payoff can reduce uncertainty significantly even when the observed payoff is low.

This trade-off between depth and breadth of learning strategies gives us the following result: at any possible belief and value of expected payoff for higher level of uncertainty the DM prefers learning unconditionally, once uncertainty reduces sufficiently the DM would want to learn conditionally about combinations of inputs.

With this result we characterize the optimal learning strategy which involves one of the three possible strategies for a given prior belief and cost of learning (where learning always decreases uncertainty, see assumption 2). First, the DM only learns unconditionally, DM learns only conditionally and DM starts with learning unconditionally first and then switch to learning conditionally.

In our bt cotton case this result would imply the relative cost of learning the payoff from input combinations are significantly high and also the uncertainty about the production function is sufficiently high. This is partly supported by the observation that in districts where farmers had guidance from extension program workers they used less fertilizer and saved on cost of production without affecting the productivity.

In the theory model we further show that making one of the two types of learning cheaper may lead less learning by the DM even though it would increase the welfare of the farmer. The optimal policy would thus be a combination of reduction of both types of cost that does not distort the incentives for the farmers for learning.

The rest of paper is organized as follows: in the next section we describe the case study of bt cotton. Section 3 describes the model, section 4 sets up and discusses the features of the decision problem. Section 5 states the main result about optimal learning strategy and discusses the policy implications for reducing the cost of the two types of learning. Section 6 discusses the relevant literature, section 7 discusses the future direction on policy analysis and section 8 concludes. All proofs are given the the appendix (A).

2 Case Study: Bt cotton

Previous studies on the condition of farmers in India have discussed the importance of information(see Birthal et al(2015), Ferroni and Zhou(2012)). In additional to learning about output price and other market information about crops the farmers also wants to learn about input choices and technology.

Agriculture sector contributes to 17% of Indian GDP but employs 50% of the labor force in India, thus policies targeting the condition of farmers are of extreme importance for reducing poverty. One of the major agricultural policy is that of agricultural extension programs that helps farmers maximize their productivity by providing required information, training and increasing the ease of access of relevant input. However, in recent decades the efficacy of the existing extension programs that has been successful during green revolution has been questioned.

Policy debates have centered around the question of the provision of information. Main criticism of the traditional extension policy is that it does not consider the information demand of the farmer and entirely driven by supply side effort of the agricultural universities and KVKs(Krishi Vikas Kendra). New extension policies have tried to address this issue however the impact these new policies is not yet conclusive.

The new extension policy desires to decentralize the information provision mechanism and want to give more control to the farmers about what type of information should be provided. Also, private input dealers and other private entities are encouraged to provide information to the farmers.

However, no systematic analysis is available that compares the impact of different type of information provision. If the government extension workers provide different type of information compared to their private counterparts the final choice of the farmers might be different depending on the information source they use.

As Gladdening et al (2012) showed that the type of information acquisition depends on farm size and other farm level factors there might be some undesirable impact of the new extension policy. For example, marginal farmers (land holding below 1 ha.) are more likely to get information from input dealers than large farmers. If this type of information is less effective than the information provided by extension workers which is mostly used by large farmers this type of policy can actually increase inequality in productivity for different groups of farmers.

2.1 Data

In this paper we want to focus on the impact of the type of information on input choice and productivity. We focus only on the cotton farmers in India using the NSS 2013 dataset (Round 70, schedule 33). Cotton is one of the major cash crop in India and also it the main input for the textile industry which is the second largest industry in India in terms of contribution to GDP.

The major producer of cotton are ten Indian states divided in three regions, Northern, Central and Southern. The Northern region consists of the states of Punjab, Haryana and Rajasthan, Central region of Gujarat, Madhya Pradesh and Maharastra and Southern region of Karnataka, Telengana, Andhra Pradesh and Tamil Nadu. The three regions differs significantly in terms of irrigation, type of cotton produced and the length of the crop cycle. The following summarizes the differences across the regions.

The schedule 33 is the "Situation Assessment Survey of Agricultural Households" that has been conducted across India in two rounds for the agricultural year 2012 -13. Round 1 is conducted in July-December 2012 and round 2 is conducted in January - June 2013. Since cotton is a *Kharif* crop, i.e, cultivated during the rainy season, in India we only consider the visit 1 of schedule 33 survey.

Considering only the farmers who produces cotton as only crop or one of the major crops we have 2334 households. For these households we have information about expenditures on different inputs, value of output, access to various sources of information, knowledge about minimum support price, outstanding loans, insurance and crop loss etc.

2.2 Learning failure

Even though India has been one of the major cotton producing and exporting country for several decades, India lagged behind in terms of yield. One of the major reason for that was susceptibility

Particulars	Northern	Central	Southern
States	Punjab, Haryana, Rajasthan	Gujarat, Madhya Pradesh, Maharashtra	Karnataka, Telengana, Andhra Pradesh, Tamil Nadu
Irrigation	Mostly Irrigated	Irrigated and Rainfed	Irrigated and Rainfed
Planting Season	April-May	June-July	July-August
Harvesting Season	October-November	November- April	December- March
Variety	Medium and Short Staple	Medium and Long Staple	Long and Extra-Long Staple

Table 1: Regional Variation in Cotton Production in India

of the cotton crop to pest attacks. In 1990's only 5% of cultivated area was under cotton but 45% of total pesticides was used for cotton production. This severely affected the profitability of the major cash crop.

In 1995 a new variety of cotton seed was introduced that was genetically modified to be immune to the major pests for cotton, namely varieties of bollworms. After several trials on both small and large farms agricultural department of Indian government approved of this new variety, namely Bt cotton for commercial use in 2002.

Various studies have shown already that bt seeds increased the profitability of the cotton production significantly. In 2002, the year of introduction only around 1% of total cotton producing area was under bt cotton. In ten years time, i.e, during the agriculture year 2012-13 93% of total cotton producing area was under bt cotton.

However we do not observe the seed variety for these households but since bt cotton is genetically modified, farmers cannot reuse seeds for earlier season. Hence we include only those households that has bought seeds in the 2012-12 agricultural year for cotton. This reduces our sample size 2311 households which also shows that most households are likely to not use traditional seed varieties.

The major reason behind using this genetically modified seed variety was to reduce the use of pesticides. Since the bt cotton is resilient to the major pests for cotton the farmers should optimally use less pesticide and focus on minor pests. Since we do not observe the type of pesticides use by the farmers we would only consider the expenditure on plant protection chemicals (PPC).

Several studies have shown that even after switching to bt seeds farmers tend to overspend on pesticides leading to lower profitability (see Kiresur and Ichangi 2011). We observe a similar pattern in the NSS data as well. We define net profit as the difference between total value from the crop minus the total cost of all inputs[3] per unit of land. Instead of considering the total net profit we consider per unit profit because earlier studies in context of Indian agriculture have already shown than that farmer behavior and farm quality changes with the size of the farm.

Table 2 shows the impact of pesticide on net profit, total quantity and crop loss. The variable pesticide measures the cost of pesticide per unit of land (ha.). Average level of pesticide use for

[3] Other than seeds all input expenditures are recorded in the aggregate level for all crops but since we are considering only those farmers who either produces only cotton or cotton is the major crop, this reported cost is a good proxy for cost of inputs.

the sample is Rs. 5856.858 with a standard deviation of Rs 9942.464. The net profit variable is constructed by subtracting total cost of all inputs from total value of output both measured in Rs per ha. All control variables are listed in the appendix (ref table 14 for details).

	Dependent variable:		
	Net profit	Loss (=1)	Quantity
	felm	*logistic*	*felm*
	(1)	(2)	(3)
Pesticide	-2.086^{***}	-0.00002^{***}	0.011^{**}
	(0.177)	(0.00001)	(0.005)
Farm	Yes	Yes	Yes
Information	Yes	Yes	Yes
Household	Yes	Yes	Yes
District FE	Yes	No	Yes
Observations	2,186	2,186	2,185
R^2	0.403		0.355
Adjusted R^2	0.365		0.314
Log Likelihood		−1,448.209	
Akaike Inf. Crit.		2,936.418	
Residual Std. Error	54,721.700 (df = 2052)		1,399.294 (df = 2051)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2: Effect of Pesticide on Net payoff, Crop loss and Total Quantity produced

Note that one rupee increase in cost of pesticide per ha of land reduces net profit by Rs 2.086. In the second column we regress the probability of crop loss and obtain that pesticide use reduces the probability of crop loss. The third column uses the total quantity of output produced measured in Kg/ha. The average value of total output is 1383.096 kg/ha with a standard deviation of 1687.591 kg/ha.

Table 2 suggests that use of pesticide is probably helpful for reducing crop loss, however it does not increase profitability which implies the farmers are most likely using more pesticide than the most profitable level. We want to claim that this is the feature of bt cotton technology.

2.2.1 Panel 1: pre and post cotton farmers

To that extent next we construct a panel data (not same households) by combining the data from both NSS 59 and NSS 70 round for cotton farmers. Note that, in 2002-03 only about 1% of all cotton farmers were using bt seeds whereas in 2012-13 season around 97% of them were using bt seed. The post variable is an indicator that the household data is from 2012-13 agricultural year.

As shown in table 3 on average pesticide positively impacts profitability but for 2012-13

	Dependent variable:		
	Net Payoff = Total Value - Total Cost		
	(1)	(2)	(3)
Pesticide	0.106	0.013	−0.023
	(0.729)	(0.769)	(0.772)
Pesticide* post	−2.468***	−2.353***	−2.309***
	(0.732)	(0.771)	(0.774)
District FE	Yes	Yes	Yes
Post FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	2,984	2,962	2,960
R^2	0.390	0.391	0.391
Adjusted R^2	0.361	0.360	0.360
Residual Std. Error	50,204.080 (df = 2845)	50,379.810 (df = 2821)	50,420.530 (df = 2815)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 3: Panel 1 : Impact of pesticide on net payoff

farmers, i.e., bt cotton farmers use of pesticide in fact reduces profitability. We can see the same pattern when instead of using net profit we use total value of production (in Rs/ha) or total quantity produced (in kg/ha) (see table 4).

	Dependent variable:	
	Total Value	Total Quantity
	(1)	(2)
Pesticide	2.737**	0.049**
	(1.108)	(0.022)
Pesticide* post	−2.306**	−0.038*
	(1.112)	(0.022)
Observations	2,792	2,780
R^2	0.409	0.302
Adjusted R^2	0.375	0.262
Residual Std. Error	72,819.850 (df = 2639)	1,449.823 (df = 2628)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 4: Panel 1: Impact of pesticide on total value and total quantity of output

2.2.2 Panel 2: Cotton producing households in *Kharif* and *Rabi* seasons

In the next panel we consider the same households that produced cotton in 2012-13 agriculture year and use their pesticide use for other crops during the *Rabi* (winter) season. The "other" variable is a dummy for any crop other cotton produced by same household during winter season. This variable varies considerably across different regions.

Table 5 shows controlling for household fixed factor indeed pesticide has a worse impact on both net profit and total value of output compared to other crops produced by the same household.

2.2.3 Panel 3: Household and time fixed effect

Finally we combine the two panel, namely cotton farmers in 2002-03 and 2012-13 and also the same cotton farmers in *Rabi* season. In the full panel Pesticide*post denote the impact of pesticide on all crops in 2012-13 agricultural year whereas pesticide*post*other separates the impact for crops other than cotton.

As shown in table 6 use of pesticide has a significantly negative impact on net profit but this effect is due to cotton mostly since the coefficient for pesticide*post*other is statistically significantly positive. The same result holds true if we consider the total value of output per ha instead of net profit.

This suggests that the bt cotton farmers are most likely using more pesticide than optimal level. The diagram below shows the relationship between the pesticide use and net profit earned by the household for 2012-13 cotton farmers from NSS round 70.

	Dependent variable:	
	Net profit	Total Value
	(1)	(2)
Pesticide	−1.886***	−0.667
	(0.507)	(0.898)
Pesticide*other	2.383*	9.857***
	(1.348)	(2.390)
District FE	Yes	Yes
Farm FE	Yes	Yes
Input	Yes	Yes
Household	Yes	Yes
Information	Yes	Yes
Observations	2,380	2,527
R^2	0.933	0.934
Adjusted R^2	0.260	0.270
Residual Std. Error	57,696.340 (df = 214)	102,247.400 (df = 214)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 5: Panel 2: Net profit and total value of output)

	Dependent variable:		
	Net profit		
	(1)	(2)	(3)
Pesticide	0.015	−0.090	−0.135
	(0.704)	(0.741)	(0.743)
Pesticide* post	−2.445***	−2.325***	−2.272***
	(0.708)	(0.745)	(0.748)
Pesticide*post*other	4.792***	4.781***	4.763***
	(0.813)	(0.817)	(0.819)
District FE	Yes	Yes	Yes
Post FE	Yes	Yes	Yes
Other FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Farm	No	Yes	Yes
Information	No	No	Yes
Observations	3,115	3,092	3,089
R^2	0.342	0.338	0.338
Adjusted R^2	0.309	0.304	0.304
Residual Std. Error	49,375.130 (df = 2966)	49,535.990 (df = 2942)	49,566.940 (df = 2935)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 6: Panel 3: Net profit

	Dependent variable:		
	Total Value of Output		
	(1)	(2)	(3)
Pesticide	3.284***	2.821**	2.904**
	(1.195)	(1.256)	(1.260)
Pesticide* post	−3.016**	−2.520**	−2.594**
	(1.202)	(1.263)	(1.267)
Pesticide*post*other	13.823***	13.652***	13.576***
	(1.381)	(1.384)	(1.388)
District FE	Yes	Yes	Yes
Post FE	Yes	Yes	Yes
Other FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Farm	No	Yes	Yes
Information	No	No	Yes
Observations	3,548	3,526	3,522
R^2	0.469	0.470	0.472
Adjusted R^2	0.447	0.447	0.449
Residual Std. Error	97,954.400 (df = 3405)	98,118.810 (df = 3382)	98,026.160 (df = 3374)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 7: Panel 3: total value of output

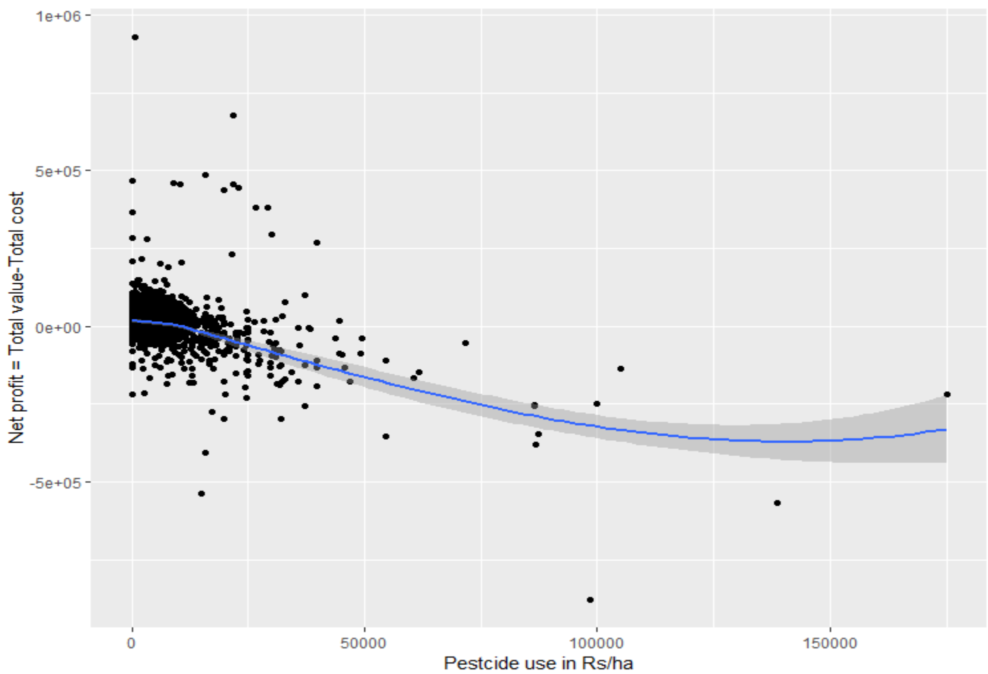

Figure 1: Pesticide use and net profit

2.3 Role of Prior

Next observation we make about the variables that affect the choice of pesticide. In most learning models if agents do not choose to update their beliefs they would choose according to their prior beliefs. We want to show that is also true for this dataset. Note that, we do need to use the entire prior distribution of belief to show the dependence. We have assumed that the farmers are rational, hence given their beliefs they should choose optimal level of pesticide.

However since we do not have a panel dataset, we cannot use the pesticide use by same farmer in a previous period, so we construct a variable at the district level. A similar survey on situation of agricultural households have been done in 2003 (NSS round 59, schedule 33). From this dataset we calculate the mean level of pesticide use in each district.

Most districts are very similar in terms of geography and weather hence the optimal input use would be similar across the district. Also, information through public extension service is decentralized at the district level, hence households in same district is likely to have similar prior belief about the efficacy of different inputs.

In table **??** we show the relationship between prior belief about pesticide use and actual pesticide use in the 2012-13 agricultural year. In the first column we do not include the information choice by the household and in the second column the information choice of the household and the interaction terms are also included. In both cases the prior variable is statistically significantly positively related to current pesticide use.

	Dependent variable:		
	Pesticide		
	(1)	(2)	(3)
Prior	0.767***	0.803***	0.817***
	(0.125)	(0.124)	(0.126)
Region FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	2,169	2,165	2,165
R^2	0.413	0.421	0.423
Adjusted R^2	0.410	0.417	0.417
Residual Std. Error	7,666.834 (df = 2159)	7,516.171 (df = 2148)	7,511.311 (df = 2144)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 8: Pesticide use and Prior

	Dependent variable:		
	Net Profit		
	(1)	(2)	(3)
Prior	−1.809*	−2.278**	−2.153**
	(1.009)	(1.019)	(1.033)
Region FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	2,169	2,165	2,165
R^2	0.197	0.210	0.211
Adjusted R^2	0.194	0.204	0.204
Residual Std. Error	62,098.510 (df = 2159)	61,517.410 (df = 2148)	61,531.710 (df = 2144)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 9: Net profit and prior

2.3.1 Alternate definition of prior

Next we consider an alternate definition of prior belief. For every household in the 2012-13 NSS dataset we find the average the pesticide use for all other crops in the winter season and use that as the prior belief for the household. In absence of extra information about bt seeds the household should use it's knowledge from other crops to decide the optimal level of pesticide.

Table 10 shows that is indeed the case. The use of pesticide for other crops is a strong indicator for the pesticide for bt cotton as well controlling for all other inputs and household level variables. Table 11 shows the negative impact of the prior belief about pesticide on net profit even though the value of the coefficient is not statistically significant.

3 Model

Based on the observations from the case study we want to construct a decision problem for a rational DM who needs to choose as least two inputs to maximize his payoff. The main feature of the data we want to explore is the following: optimal choice of one input depends on other inputs, e.g., switching to bt cotton requires change in pesticide use *but* DM updates his belief about only one of the input and this suboptimal choice is payoff relevant for him.

3.1 Primitives

Consider a DM producing output Y with two inputs A_1 and A_2. Each of the two inputs has n possible levels, namely, $\{a_{i1}, a_{i2}, \dots a_{in}\}$ for $i = 1, 2$. Output Y can take two possible values

	Dependent variable:		
	Pesticide		
	(1)	(2)	(3)
Prior	0.208*	0.199*	0.198*
	(0.110)	(0.110)	(0.110)
District FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	815	813	813
R^2	0.558	0.571	0.572
Adjusted R^2	0.497	0.507	0.505
Residual Std. Error	8,339.592 (df = 715)	8,263.816 (df = 706)	8,279.761 (df = 702)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 10: Alternate prior: Pesticide use

	Dependent variable:		
	Net Profit		
	(1)	(2)	(3)
Prior	−0.278	−0.228	−0.212
	(0.801)	(0.795)	(0.800)
District FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	815	813	813
R^2	0.565	0.580	0.580
Adjusted R^2	0.505	0.517	0.514
Residual Std. Error	60,637.800 (df = 715)	59,925.010 (df = 706)	60,077.400 (df = 702)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 11: Alternate prior: Net payoff

$Y \in \{0,1\}$. [4]

Suppose the DM has one unit of indivisible land and he needs to choose a combination of two inputs to maximize his payoff from producing Y. Let us define $\mathbf{A} = A_1 \times A_2$ as the set of all possible input combinations, i.e the action/decision space. The DM can choose a mixed strategy over input choices, i.e., he can choose several combinations of inputs with positive probability. In case of a mixed strategy only one of the input combination is actually realized according to the strategy chosen by the DM.

Also we assume that the DM can choose only one of the inputs in which case the other input is chosen according to some known distribution $G(.)$. For simplicity we assume that G is uniform. i.e., if the DM chooses only input A_1 then every level of input A_2 is chosen with equal probability.

Let $\pi : \mathbf{A} \to Y$ denote the unknown payoff function that maps each combination in the action space to the corresponding value of output. If $\mathbf{a}_{ij} = (a_{i1}, a_{j2})$ is a typical element of the action space $\mathbf{A}$ then $\pi(\mathbf{a}_{ij})$ would denote the payoff from choosing $\mathbf{a}_{ij}$. We can also define the payoff from each level of each input a_{ij} separately, namely $\pi(a_{ij})$ defined as the average of choosing level j of input i along with all possible level of input $-i$. For example $\pi(a_{11})$ is the average of choosing (a_{11}, a_{21}) and (a_{11}, a_{22}).

We can also represent the payoff function as a payoff matrix where the $(i,j)^{th}$ cell denote the payoff from $\mathbf{a}_{ij}$. For example, if $n = 2$, the payoff matrix can be written as follows:

	a_{21}	a_{22}	A_1
a_{11}	$\pi(\mathbf{a}_{11})$	$\pi(\mathbf{a}_{12})$	$\pi(a_{11})$
a_{12}	$\pi(\mathbf{a}_{21})$	$\pi(\mathbf{a}_{22})$	$\pi(a_{12})$
A_2	$\pi(a_{21})$	$\pi(a_{22})$	

Table 12: 2×2 model

We will denote $\pi(\mathbf{a}_{ij})$ as *conditional payoff* or payoff from *cell* (i,j) and $\pi(a_{i,j})$ as *unconditional* or *average* payoff for a_{ij}.

The state space Ω is defined as the set of all possible realizations of the payoff matrix π. Since there $n \times n$ many combinations and each combination can take two possible values, Ω contains $2^{n \times n}$ possible states. As the state space is defined is terms of payoff matrices, any typical state $\omega \in \Omega$ would denote a particular payoff matrix and the notation $\pi(\mathbf{a}_{ij}|\omega)$ would denote the payoff from $\mathbf{a}_{ij}$ in state ω.

3.2 Information

3.2.1 Prior

We assume that there exists a "true" payoff function which is chosen according to the data generating process μ^*, that is unknown to the DM. Instead let $\mu_0 \in \Delta(\Omega)$ be the prior belief of the DM over Ω. We do not assume that $\mu_0 = \mu^*$, but rather $supp(\mu^*) \subset supp(\mu_0)$, i.e., if DM learns perfectly given his prior belief he would be able to uncover the "true" payoff function.

Since we do not impose any restriction on possible payoff matrix $\pi(.)$ in a given decision problem μ_0 can denote any possible relationship between the two inputs . Let us consider two special cases, namely the two inputs are complements/substitutes and the two inputs are orthogonal, i.e, the payoff of one input does not affect the payoff from other input.

[4]For simplicity we normalize the cost of both the inputs to zero hence net payoff would be same as value of output Y.

Using the matrix we can clearly show the difference between the two possible type of payoff function. In case of complement/substitute all rows and all column have same average payoff whereas in case of orthogonal inputs all cells in a row or column has the same payoff. The two tables below illustrates:

	a_{21}	a_{22}	a_{23}	
a_{11}	1	1	1	1
a_{12}	0	0	0	0
a_{13}	0	0	0	0
	1/3	1/3	1/3	

(a) Example: Orthogonal

	a_{21}	a_{22}	a_{23}	
a_{11}	0	0	1	1/3
a_{12}	0	1	0	1/3
a_{13}	1	0	0	1/3
	1/3	1/3	1/3	

(b) Example: Substitute/Complement

3.2.2 Learning Technology

The DM can learn about the payoff matrix in two possible ways, namely by observing conditional payoff of an input combination or a *cell* in the payoff matrix or by observing the unconditional payoff of an input by observing an average value from the matrix. We would call the former as learning by *cell* and the later as learning by *average*.

In both cases when the DM observes the cell he perfectly observes the value of output in the cell. For example if there is only one cell in a row that gives payoff $Y = 1$ and the DM observes the average cell then he would observe the payoff $1/n$ which is indeed the average of payoff of all cells in that row.

The learning strategy is sequential, i.e., once the DM observes the payoff from one cell he can decide whether and how to learn the next cell. . Let $\mathbf{P}$ be the set of all choices available to the DM which includes $\varnothing$, i.e., no learning as well.

Learning is costly in this model. We assume that the cost of learning can possibly be different for the two types of learning. Let c_a denote the constant cost of learning by averages and c_l be the constant cost of learning by cell. [5]. We will assume two features of the cost of learning.

Assumption 1. *1. Learning about averages is weakly cheaper than learning about cells, i.e.,* $c_a \leq c_l$.

2. No learning is costless, i.e, $c(\varnothing) = 0$.

A learning strategy is defined as a conditional sequence of cells that the DM can choose to observe consecutively where the choice of k^{th} cell can depend on the outcome of the history of cells observed before k^{th} round. Let $\mathcal{P}$ denote a sequence of cells chosen from $\mathbf{P}$ then $\gamma(\mathcal{P})$ would denote learning strategy, which is a distribution over possible sequences of cells $\mathcal{P}$.

3.2.3 Updating procedure

We assume the DM is Bayesian, i.e, given μ_0 and observed cell from $\mathbf{P}$ he would update according to the Bayes law. Since the posterior belief is defined over Ω but the DM observes only one cell/average at a time the updated posterior need to be consistent with Bayes law. This implies

[5] Note that we assume that learning about any cell (or averages) would be equally costly which may not be realistic. For example different types of inputs may require different learning techniques as a result can have different costs. However, in this model we assume away from that.

if μ_t is the t^{th} round belief and $\gamma(\mu_{t+1})$ denote the distribution of all possible posteriors in $t+1^{th}$ round then by Bayes consistency,

$$\mu_t = E(\mu_{t+1}|\gamma_t)$$

Bayes updating rule here implicitly assumes perfect memory for the DM, i.e., if he observes one cell and updates his belief based on the observation, then he would never need to observe the same cell again for any further updating.

The belief generated by observing any cell or average can be summarised by a partition of the state space. For example if the DM chooses to observe cell (i, j) he is partitioning the state space into two blocks, namely one where (i, j) generates $Y = 0$ and where (i, j) generates payoff $Y = 1$. Suppose ω_k and ω_l denote two states, in both of which the observed cell produces either $Y = 0$ or $Y = 1$, then these two states would belong the same block of the partition. Since the DM is only learning about (i, j) the following consistency condition must hold true:

$$\frac{\mu_t(\omega_k)}{\mu_t(\omega_l)} = \frac{\mu_{t+1}(\omega_k)}{\mu_{t+1}(\omega_l)}$$

where μ_t and μ_{t+1} denote the t^{th} and $t+1^{th}$ round beliefs resp. We would denote this condition as *partition consistency* condition.

4 Decision Problem

We consider a static choice problem for the DM who enters the period with prior belief μ_0. Once entering the DM chooses a learning strategy which specifies a conditional sequence of cells $\gamma(\mathcal{P})$ that the DM would observe. After learning is done the DM chooses an action, namely, mixed strategy defined over the action space $\mathbf{A}$. Then final payoff is realized.

Given the DM is rational, for any posterior belief the DM would choose the mixed strategy that maximizes his expected payoff thus the only choice we need to consider to derive the DM's value function is the learning strategy chosen by the DM.

The value function of the DM is given as follows:

$$W(\mu_0) = \max_{\gamma(\mathcal{P})} E(\pi(\mathbf{a}_{ij}) - c(\mathcal{P})|\mu_0|\gamma, \mu_0) \tag{DP}$$

Lemma 1. *Given any prior μ_0 DM's problem (DP) can be written in recursive form as*

$$V(\mu_t) = \max_{p \in \mathbf{P}} -c(p) + E_{\mu_t} V(\mu_{t+1})$$

4.1 Recursive Problem: two extreme cases

Before we solve the full model where the DM has the choice over both types of learning we will consider two extreme cases, one where he can only observe cells in the matrix and two, where the DM can observe only averages.

4.1.1 Learning by Cells

When the DM has access to only cells his decision problem can be described as follows: at each round of belief μ_t he decides whether to observe another cell from the matrix or stop learning. Also, if he decides not to learn any further then he needs to choose the optimal cell to learn about.

Note that, if the DM finds a cell that generates a payoff of $Y = 1$ since the learning is fully revealing he would no longer observe any more cells and choose the cell with payoff $Y = 1$ for sure. Also, by assumption of rationality, in case the DM does not uncover a cell with $Y = 1$ he would choose the cell with highest probability of generating $Y = 1$ and if there are more than one cell with the same probability he would be indifferent between choosing any of them.

Lemma 2. *Learning by cells is optimal when the uncertainty of the belief μ_t, as measured by the entropy of the distribution, lies between an interval $(H_l(\pi_t), H_u(\pi_t))$. The size of the interval decreases with increase in expected payoff.*

Proof of the lemma is given in Appendix. The main argument is as follows; observing a cell always partitions the state space into two blocks, namely, the cell generates $Y = 0$ or $Y = 1$, irrespective of the uncertainty of the belief. This implies for high level of uncertainty the reduction of uncertainty by this partitioning is not sufficient to cover the cost. Also when uncertainty is too low then the value of information becomes too low for learning to be optimal.

4.1.2 Learning by Average

Next, we consider the scenario where the DM can only learn about averages. The relevant problem in round t for the DM is to choose whether to learn or not and if learning is optimal then which average to uncover. If he decides to stop learning he can choose any mixed strategy over possible input combination or only one of the two inputs in which case the other input is chosen according to the uniform distribution.

Note that, learning averages only will not always lead to the true state. For example if the prior belief is that the two inputs are substitutes then observing averages does not reduce uncertainty at all.

Lemma 3. *The optimal learning strategy is a cutoff strategy defined on uncertainty (measured by the entropy of the state space Ω), i.e, only if the entropy is higher than the cutoff $\bar{H}_a(\pi_t)$ the DM would choose to learn. The cutoff $\bar{H}_a(\pi_t)$ is increasing in the expected payoff.*

The proof of the lemma is given in the Appendix. The main difference in case of learning by average is that the reduction in uncertainty by observing one more average increases with uncertainty in state space Ω itself. Thus even for higher values of uncertainty learning is actually optimal as that is when the benefit from learning by reducing uncertainty is largest.

4.1.3 Full Model

Finally we consider the full model where the DM has access to both types of learning. In this case at any round the relevant problems for the DM are, first whether to learn or not, second, which type of learning to choose given learning is optimal and third, which cell or average to uncover given a choice of learning.

Lemma 4. *There exists three cutoff values of uncertainty of the belief μ_t, namely $\bar{H}(\pi_t) \geq \hat{H}(\pi_t) \geq \underline{H}(\pi_t)$ such that*

1. *If uncertainty is higher than $\bar{H}(\pi_t)$ then it is optimal for the DM to uncover an average,*
2. *If uncertainty is in between $(\underline{H}(\pi_t), \hat{H}(\pi_t))$ then DM optimally chooses to uncover a cell in the matrix*
3. *Everywhere else no learning is the optimal strategy.*

The proof of the lemma is given in the Appendix. The learning strategy is obtained combining the two earlier lemmas. For higher uncertainty if averages are informative then optimal learning is to uncover averages. As uncertainty reduces the averages become less informative but the gain in expected payoff from observing cells starts to increase and hence learning cells becomes optimal.

However since the cost of two types of learning are different as $\delta > 0$ even when learning about averages become not profitable the cost of learning cells may not become optimal which can generate a gap difference between $\hat{H}(\pi_t)$ and $\bar{H}(\pi_t)$.

5 Results

5.1 Solution of DM's problem

In the last section we discussed the recursive version of the DM's problem. Now we will consider the entire learning strategy for any given prior μ_0. The main difference is in each stage the level of expected payoff and the level of entropy changes as the DM learns more about the payoff matrix.

Before we prove our main theorem about the nature of learning strategy let us introduce a condition that we will assume to be true for the prior belief.

Assumption 2. *Given the prior there does not exist any cell or average such that uncovering it can increase the level of uncertainty of the belief.*

Note that, assumption 2 is only false in the following case; there exists a small subset of cells that contains $Y = 1$ with high probability but in the event of the subset containing only $Y = 0$, the belief that any other cell will contain $Y = 1$ is distributed uniformly.

Theorem 1. *Given any prior μ_0 if assumption 2 holds true then there are three possible strategies:*

1. *Observe only cells.*
2. *Observe only averages.*
3. *Observe averages first then observe cells.*

Proof. For any given value of prior, in round $t = 0$ there are three possibilities, namely, no learning is optimal, learning about averages are optimal, learning about cells are optimal. The first case would occur if the cost of both types of learning are sufficiently high. In this case the DM chooses an input combination based on his prior belief μ_0 and payoff is realized. So we would focus on the two other cases.

If at μ_0 the uncertainty is below $\hat{H}$, then learning about cells are optimal. In the next round when DM has observed one cell since by assumption 2 the entropy of the belief μ_1 reduces observing cells remains optimal. This is true for every round of learning until the entropy becomes sufficiently low that no learning becomes optimal given the level of expected payoff. Since the expected payoff can both increase or decrease with uncovering more cells the conditional sequence may have different cutoff values of entropy for different realizations of the payoffs of the cell.

If at μ_0 the level of uncertainty is in $(\hat{H}(\pi_0), \bar{H}(\pi_0))$ the no learning is optimal in first round. Since the entropy or the expected payoff does not change from first round, no learning would be the optimal strategy in that case.

If at μ_0 the level of uncertainty is above the cutoff $\bar{H}(\pi_0)$ then in first round learning averages would be optimal. Since observing averages would reduce the uncertainty (by assumption 2) in the subsequent rounds one of the two possibilities can happen. First if the cost of learning cells are sufficiently high such that uncertainty is reduced to be in the region $(\hat{H}(\pi_t), \bar{H}(\pi_t))$ for some round t then no further learning would take place.

On the other hand if the reduction is uncertainty is such that for round t the uncertainty in the belief is in the region $(\underline{H}(\pi_t), \hat{H}(\pi_t))$ then the DM would switch to learning cells. No other cases are possible under assumption 2. □

If assumption 2 does not hold true then it is possible that the DM would start by learning about cells then switch to learning about averages. But this can be the case only if the cost of learning cells is sufficiently small compared to learning averages otherwise learning cell would not be the initial strategy.

The first type of learning where the DM only observes the average values would be called as *selective learning* since the DM only selects to learn about one or more input unconditionally. Note that, under selective learning the DM does not necessarily learn about the true state as discussed in previous section.

5.1.1 Special Cases

Next we will consider two special cases of the prior introduced in the earlier section, namely, substitute/complement and orthogonal inputs.

We define two inputs to be substitute or complement if given a level of input A_1 the choice of A_2 that generates $Y = 1$ and this choice is different for each level of A_1. Thus the payoff matrix would have $Y = 1$ only along the diagonals for either the original or reordered matrix.

If the two inputs are complement/substitutes then learning about averages can never be optimal, hence there are only two choices, learning about cells along a row (all rows generate same expected payoff) or no learning at all. Following lemma 2 learning is only optimal for a intermediate range of uncertainty.

If the two inputs are orthogonal then also observing cells can be the optimal strategy. Orthogonal inputs are defined as the prior belief where either each row including averages or each column including averages are identical but the DM does not know one it is.

In this case the DM has two possible strategies, one to start with observing average row or column and then switch to the other average if the first average is not for the payoff relevant input. The other strategy requires learning cells along the diagonal. The second strategy requires smaller number of cells to be uncovered. If δ is sufficiently small then learning cells would be the optimal strategy in this case as well.

Note that for every other type of prior belief that describes the relationship between the two inputs the level of complementarity or orthogonality lies in between these two extremen cases, unless only one input is payoff relevant. Hence for any belief if the cost of learning cells is not sufficiently high then learning cells can be the optimal strategy.

Finally if the prior belief of the DM is such that only one of the input is payoff relevant then only learning averages is optimal since $\delta \geq 0$. Thus selective learning would happen *for sure* only when only one input is payoff relevant. In case of every other prior belief there exists values of $\delta > 0$ such that learning about cells would be optimal.

5.2 Policy Analysis

In this model, the parameters that are subject to policy are the two types of cost of learning. Any extension program can change the two types cost of learning. We can also consider changing the prior belief of the DM, but since we haven't modeled how the DM forms his prior it would not be possible to analyze such a policy in context of this model.

5.2.1 Reduction in Cost of Learning Cells

If the cost of learning cells c_l is reduced then even though the net expected payoff, i.e., the welfare of the DM increases but the impact on the learning strategy remains ambiguous.

Corollary 1. *For sufficiently high value of c_a, reducing the cost of learning cells, c_l can increase the probability of mistake, i.e, probability of choosing an input combination generating $Y = 0$ and hence reduce the expected gross payoff.*

The proof of the corollary is given in the Appendix. The following example would explain the main argument of the proof. Consider the decision problem where $n = 3$ and the prior belief of the DM is that exactly five out nine cells contain $Y = 1$ where each cell is equally likely to generate $Y = 1$. If cost of observing cell, $c_l \in (.278, .3)$ then it is optimal for the DM to start with observing average and switch to observing cell from a different row or column if observed average is .33.

Now if the cost of learning cells decrease to a value between $c_l \in (.267, .277)$ then the DM would optimally choose to observe at most one cell for sufficiently high c_a. The resulting uncertainty in the second case would be higher and would lead to lower expected gross payoff.

5.2.2 Reduction in Cost of Learning Averages

Similar to the case of c_l, even though the reduction in cost of learnig averages c_a leads to higher net payoff, it has ambiguous impact of learning and gross payoff.

Corollary 2. *For sufficiently high value of c_l, reducing the cost of learning averages, c_a can increase the probability of mistake, i.e, probability of choosing an input combination generating $Y = 0$ and hence reduce the expected gross payoff.*

The proof of the corollary is given in the Appendix. The following example would explain the main argument of the proof. Consider the following decision problem where $n = 3$ and the prior belief of the DM is that exactly six out of nine cells contain $Y = 1$ and all cells are equally likely to generate $Y = 1$.

Suppose initially $c_a < .133$ and $c_l \leq .33$ but $\delta < .2$. In this case optimal strategy is to observe the average and then observe a cell if the value of average is .67. The final payoff in this case is 1 and there is no probability of mistake.

If now c_a reduces significantly whereas the c_l remains the same such that $\delta > .2$ then observing only averages become optimal. But in that case if there are exactly two cells in each row and column that gives $Y = 1$ the expected gross payoff is less than 1 and the probability of mistake is .33.

5.2.3 Reduction in both costs

Reduction in both cost would definitely increase the net payoff and hence the welfare of the DM but does not necessarily increase the gross expected payoff.

Similar to the corollary 1 is c_l reduces more than proportionally such that the DM reduces the amount of reduction in entropy by switching to observing cells then it would increase the probability of mistake.

Also, similar to corollary 2 if c_a reduces more than proportionally such that DM switches to learning only averages then the probability of mistake may increase.

Thus the optimal policy would be the case where none of the two costs are reduces more than proportionately such that the DM uses a weakly more informative learning strategy but with a lower cost of learning leading to both increase in net payoff and reduction in probability of mistake.

Note that, even though reducing cost always makes the DM better off the decrease in amount of learning(or increase in probability of mistakes) can be suboptimal for the reasons of externality. For example, as we have shown in our bt cotton study around 31% of the farmers learn from other progressive farmers. If these progressive reduce their learning it may create negative externality on other farmers, the discussion of which is beyond the scope of this model.

6 Literature

Since low productivity in agriculture is a major concern in most developing countries understanding the reason for sub-optimal use of inputs and/or technology has been an important research question in development economics.

The main reasons for the sub-optimal use of inputs as discussed in the literature are: financial constraints[6], supply side constraints and learning constraints(see Duflo et al 2008, 2011, RIRDC report 2003). See de Janvry et al (2017) for a survey of field experiment targeted to affect the sub-optimal choice sof farmers in developing countries [7].

Constraints to learning can take many forms starting from lack of information or costly experimentation/processing or cognitive limitation faced by the farmers. If information is costly to acquire or requires some costly experimentation the farmers may optimally choose not to learn about the optimal practices. One of the major focus of agricultural extension services is to successfully transmit information to the farmers (see e.g. Kondylis(2017), Meijer et al (2015) and Glendenning(2000) for Indian context).

To mitigate learning constraints farmers often learn from the experience of their neighbours. Social learning can change adoption/optimal choice behavior for new farmer based on the outcome of the more experienced farmers [8]. However, social learning mechanisms are often not effective in case of agriculture due to idiosynccratic nature of the agricultural production function (see Munshi(2014)).

Constraint to learning can be a result of cognitive limitation of the farmers. As dicussed in Mullianathan et al (2014, 2016) the poorer farmers face a stricter cognitive constraint as a result of poverty. Also especially in India where the average rate of literacy is around 80% a significant proportion of farmers do not have the necessary education to systematically learn via experiments or from other formal resources (govt websites, newspaper etc.).

The main mechanism of the decision problem in this paper is in the strand of recent attention literature. In the *Rational Inattention* literature the DM faces a decision problem where he has the choice of learning which is costly. The cost of learning/attention is treated as a cognitive cost that affects the net utility of the DM. Several types of deviations of behavior from rationality has been explained as a result of the cognitive cost restricting the possible choices of the DM, which is precisely the same approach we follow in this paper as well. (See Matejka and Mckay (2015), Caplin, Dean and Leahy (2017), Woodford (2014), Herbert and Woodford(2017), Matejka, Steiner and Stewart (2017), Matejka and Tabellini(2015), Gabaix (2014) etc.)

Apart from Matejka and Tabellini(2015) no other paper consider the multidimensional nature of the decision problem which is the key focus of this paper. In Matejka and Tabellini(2015) even though there are multiple dimensions, they are additively combined to derive the final payoff. We do not require any such structure of payoff function.

[6]Lambrecht et al(2014), Karlan(2014)

[7]See Foster and Rosenszeig 2010 for a detailed explanation of factor affecting technology adoption.

[8] Conley and Udry (2010), Beaman et al (2014), Cai et al (2015)

Also, even though the DM in this model faces a one-time decision problem the recursive approach of the problem is similar to Woodford (2014), Herbert and Woodford(2017). However, our emphasis is on the decision what to learn also as opposed to when to stop learning only.

Besides the rational inattention literature that mainly focuses on deviations of behavior from a rational standard other literature on financial decision making has discussed the role multiple dimensions of a decision making problem. See Nieuwerburgh and Veldkamp(2009 a, b), Mondria (2010), Peng and Xiong (2009) for the portfolio choice problem where the DM also needs to learn about the payoff from several assets.

Every model in this literature assume some specific form of relationship between the assets (in terms of correlation) and special cost functions for learning. We have assumed away from that. Also, since the impact of each asset can accessed separately the same trade-off of breadth and depth is not present in these models.

The feature of the trade-off between breadth and depth learning strategy in this model is however similar to *coarse categorization* models of prediction (See Fryer and Jackson(2008), Mullianathan, Schwarzstein and Schleifer(2008), Mohlin(2014) etc). In this models the DM wants to predict an outcome variable based on a set of input variables. The DM partitions the set of possible input variables and predicts the outcome variable for the entire block in the partition.

There are two major differences between the coarse categorization/partition models and our model. First, Instead of choosing a finer or coarser partition as in the coarse categorization models, the DM in this paper is choosing between two types of partition, one based on unconditional learning and the other on conditional learning which generates the depth-breadth trade-off.

Second, in all these categorization models the objective of the DM is to predict the outcome but in our model the DM wants to maximize the payoff and learning about the relationship between input and output variable is instrumental in maximizing payoff. As noted by Nieurwerburgh and Veldkamp (2010) in an otherwise similar environment the learning strategy would be very different if the DM is maximizing payoffs instead of predicting only.

Finally this paper is closest to Schwarzstein(2014) where the author introduces the concept of selective attention, i.e., paying attention to only input in a production process. First major difference is that in Schwarzstein(2014) the objective of the DM is to correctly predict the state rather than maximizing the expected payoff. The second major difference is that the DM in Schwarzstein(2014) does not choose whether to learn about both inputs or not since the cognitive constraint he faces is exogeneously given. This leads to different policy implications. For example reducing the cost of learning in Schwarzstein(2014) can never increase the probability of mistake.

7 Future Directions: Experiment

As discussed in the introduction the main objective of this paper is to understand the optimal learning strategy of the producers who faces a multi-dimensional choice problem and can learn about each input separately beside learning about the production function. Analyzing the learning strategy helps us to understand how the farmers make mistake in learning and results in lower payoff.

In the results section we have shown the optimal strategy would be reduce both types of cost for the farmers in a way that does not distort the incentives of the farmers to learn. Several extension practices both public and private already provides input specific information and also information about best practices. Even though NSS 70 records where the farmers get their information from due to huge difference in regional difference it is not possible to clearly identify the impact of these extension policies.

For example, in 2005 there was a major shift in the public extension services by introduction of Agricultural Technology Management Agency (ATMA) which shifted the focus of the extension program to more demand driven information provision, i.e, providing information as demanded by the farmer. In 2007 this program rolled out to all district in India but no recent study is available discussing the impact of the policy. Since ATMA emphasized on district level decentralization, the real implementation of this program has huge degree of heterogeneity across districts, which makes it hard to analyze the impact of such a policy.

According to our model this policy of demand-driven extension service is not necessarily beneficial since it distorts the incentives for learning. However, ATMA was introduced with many other features facilitating the technology dissemination for the farmers through financial support and increases visit and training programs. Thus the impact of just the policy of affecting the cost of learning cannot be analyzed in this framework.

To do so, we want design an experiment where we ask the subject to choose inputs in a multi-dimensional choice problem. The prior and the true state being known to the experimenter we can collect the state dependent stochastic choice data(SDSC) [9], i.e., choice of learning in given a true state and prior. Using the SDSC we can measure the amount of learning defining is as a difference between final stochastic choice based on posterior and the known prior.

In this experiment if we now only change the cost of one type of learning at a time by offering the subject option to buy some information of one type or other we can analyze the impact of policy intervention.

8 Conclusion

We document suboptimal behavior of cotton farmers in India, who when switching the seed variety does not fully adjust the level of pesticide and this substantially affects their profitability. Since this behavior is due to over use of an input we introduced a model of choice with learning constraint rather than financial or supply constraints.

To explain the behavior of the farmers we set up a decision problem where the DM has choose more than one input but he does not perfectly know the production function that relates the two inputs to the output. The DM can choose to learn about each input unconditionally or learn about the conditional payoff from the combination of the two inputs.

We find that for high level of uncertainty learning unconditionally is optimal and as uncertainty reduces learning conditionally becomes optimal. This generates three possible learning strategies, namely only conditionally, only unconditionally and start with unconditionally and then switch to conditional learning.

The main result generates the policy implication that reducing only one type of cost is not unambiguously better for learning even when is the welfare improving. Finally we want run a controlled experiment to analyze the impact of this policies on learning behavior.

Bibliography

Beaman, L., Magruder, J., and Robinson, J. (2014). Minding small change among small firms in kenya. *Journal of Development Economics*, 108:69–86.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *The Quarterly Journal of Economics*, 128(1):1–51.

[9]See Caplin et al (2014, 2017) etc.

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *The American Economic Review*, 105(7):2183–2203.

Caplin, A., Dean, M., and Leahy, J. (2017). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. Technical report, National Bureau of Economic Research.

Francetich, A. and Kreps, D. M. (2016a). Choosing a good toolkit, i: Formulation, heuristics, and asymptotic properties.

Francetich, A. and Kreps, D. M. (2016b). Choosing a good toolkit, ii: Simulations and conclusions.

Hanna, R., Mullainathan, S., and Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *The Quarterly Journal of Economics*, 129(3):1311–1353.

Hébert, B. and Woodford, M. (2017). Rational inattention and sequential information sampling. Working Paper 23787, National Bureau of Economic Research.

Karlan, D., McConnell, M., Mullainathan, S., and Zinman, J. (2016). Getting to the top of mind: How reminders increase saving. *Management Science*, 62(12):3393–3411.

Matejka, F. and McKay, A. (2014). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *The American Economic Review*, 105(1):272–298.

Matějka, F. and Tabellini, G. (2016). Electoral competition with rationally inattentive voters.

Mullainathan, S. and Shafir, E. (2013). *Scarcity: Why having too little means so much*. Macmillan.

Schwartzstein, J. (2014). Selective attention and learning. *Journal of the European Economic Association*, 12(6):1423–1452.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690.

Steiner, J., Stewart, C., and Matějka, F. (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553.

Woodford, M. (2014). Stochastic choice: An optimizing neuroeconomic model. *The American Economic Review*, 104(5):495–500.

A Appendix

A.1 Control Variables

Variable	Description	Unit	Summary statistics
	A. Input		
Seed	Cost of seed per ha	Rs	mean = 4584.465, sd = 6281.258
Irrigation	Cost of irrigation per ha	Rs	mean = 478.7396, sd = 2001.767
Fertiliser	Cost of chemical fertiliser per ha	Rs	mean = 11269.87, sd = 15881.04
Manure	Cost of manure per ha	Rs	mean = 1140.944, sd = 5098.996
Area_irr	Percentage of irrigated area	%	mean = 0.4034814, sd = 0.488216
Land	Total cultivated land (in logs)	Ha.	mean = 1.322927, sd = 1.367415
	B. Household		
Land Possession	Total land possession (own or rented)	Ha.	mean = 2.657479 , sd = 2.614432
Ration Card	1 = Antodaya (poorest) 2= BPL(below poverty line), 9 = other	%	1: 3.1%, 2: 40.5%, 9: 48.6%
MSP	Aware of Minimum Support Price	%	aware : 21.7%
Insurance	1: insured only with loan, 2: insured, 3: not insured	%	1: 10%, 2: 1.6%, 3: 88.4%
	C. Information		
Extension Agents	Accessed information at least once	%	8.2%
Private	same	%	13.3%
Other progressive farmers	same	%	30.8%
Media	same	%	28%

Table 14: Control Variable

A.2 Crop loss and Prior

	Dependent variable:		
	Experienced Crop Loss (=1)		
	(1)	(2)	(3)
Prior	-0.0001^{***}	-0.0001^{***}	-0.0001^{***}
	(0.00003)	(0.00004)	(0.00004)
Region FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	2,166	2,165	2,165
Log Likelihood	−1,425.115	−1,416.329	−1,412.779
Akaike Inf. Crit.	2,870.229	2,866.657	2,867.558
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 15: Crop loss and prior

	Dependent variable:		
	Crop Loss (=1)		
	(1)	(2)	(3)
Prior	−0.00000	−0.00000	−0.00000
	(0.00002)	(0.00002)	(0.00002)
Region FE	Yes	Yes	Yes
Input	Yes	Yes	Yes
Household	No	Yes	Yes
Information	No	No	Yes
Observations	813	813	813
Log Likelihood	−507.516	−500.254	−497.654
Akaike Inf. Crit.	1,035.031	1,034.509	1,037.308
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 16: Alternate prior: Crop loss

A.3 Proofs

A.3.1 Proof of lemma 1

Proof. We want to show that given any prior belief μ_0 the DM's optimal choice of observing a cell at any stage t where the DM's belief is given by μ_t depends only on μ_t and not the history of μ_s for $s \leq t$. This is sufficient because if μ_s alone determines the choice of cell p then the DM's problem can be written as optimal choice of p given μ_t which is described in the value function $V(\mu_t)$.

The strategy of observing any cell generates an unique partition of the state space Ω. For example, if the DM chooses to observe a cell in the matrix, it can have two possible values, namely 0 or 1. Thus observing the cell is equivalent to partitioning the state space into two block, one where the cell takes value 1 and the other where it takes 0. Whereas if the DM observes an average it can take values in $\{0, 1/n, 2/n, \dots, 1\}$.

Once the DM observes the realization in a given cell he chooses the block of cells that contains the true state. For example, if he observes an average where there are only 1 cell that generates $Y = 1$ then he would learn that the true state belongs to the block of states where the average takes the value $1/n$. If now he observes another cell from the same row (or column) then that cell can take value 0 or 1, which would further divide the new subset into two blocks and the DM would learn which block contains the true state.

Thus any learning strategy can be described as a sequence of partitions where the DM learns that true state belong to the intersection of relevant blocks. Given the partition consistency condition the DM cannot further update his belief about the states in the block without further observation.

Now suppose there are two histories that generate the same belief μ' but the optimal strategy following them would be different. Note that, for both two histories, say μ^1 and μ^2 the resulting subset that contains the true state under μ^{prime} must be the same otherwise the belief μ' cannot be same.

This implies there are three possible ways in which the two histories are different. First the same cells are observed but in different order. Since intersection of sets is commutative and each cells generates an unique partition of Ω the order of partition is irrelevant for decision making.

Second, under one of the histories, say μ^1 one cell has been observed that was not payoff relevant, i.e, the updating did not reduce any uncertainty, but this would be imply the strategy chosen under μ^1 was not optimal.

Third, there exisis multiple sequence of cells observing which would generates same partition. In that case the DM must be indifferent about choosing any of two sequence which again implies it cannot be relevant for future decision making. Otherwise ex ante the DM would not choose one of the two histories optimally.

Thus any t^{th} round belief μ_t denotes an unique subset of states that contains the true state and hence the optimal policy only depends on the belief μ^t and not on the entire history.

□

A.3.2 Proof of lemma 2

Proof. Since the DM would only learn if he has not uncovered a cell with payoff $Y = 1$ let us define(abusing the notation) the expected payoff given his belief π_t at round t. By, assumption of rationality π_t is obtained by choosing the cell(s) with highest probability of obtaining payoff $Y = 1$, i.e., π_t is the probability that this cell would generate payoff $Y = 1$. Once the DM decides to stop learning he would choose the optimal input combination that generates this highest payoff π_t.

Observing a cell can result in one of the two possibilities, either $Y = 1$ or $Y = 0$, hence irrespective of the level of uncertainty and expected payoff the strategy would always partition the state space in two blocks according to the prior belief of the cell being $Y = 1$ or $Y = 0$. By partition consistency condition, the ratio of probability in the chosen block, revealed by the true of the observed cell does not change. Hence observing a cell with lower probability of generating $Y = 1$ can never be beneficial. Hence the optimal strategy would always be to observe the cell(one of the cells) with highest probability of $Y = 1$.

This implies if the expected payoff in round t is μ_t, then the net value from not learning would be

$$-c(\varnothing) + \pi_t = \pi_t$$

The net value from learning would be,

$$-c_l + E(V(\mu_{t+1})|\mu_t) = -c_l + (\pi_t * 1 + (1 - \pi_t)V(\mu_{t+1})).$$

If the observed cell indeed generates $Y = 1$ then DM would stop learning for sure and if not then he can continue learning. Thus learning is optimal only if

$$\begin{aligned} -c_l + (\pi_t * 1 + (1 - \pi_t)V(\mu_{t+1})) &\geq \mu_t \\ (1 - \pi_t)V(\mu_{t+1}) &\geq c_l \end{aligned}$$

The value function $V(\mu_{t+1})$ denotes the expected value when the observed cell is $Y = 0$. Given any value of expected payoff μ_t let us consider two possible cases of uncertainty of the belief μ_t. In the first case the subset of possible sets are much bigger compared to the second case such that the uncertainty of the first case is higher than the second case. We can measure the level of uncertainty by the entropy of the distribution defined as

$$H(\mathbf{p}) = -\sum_{p \in \mathbf{p}} p \ln p.$$

In the first case suppose the entropy is $H_1 > H_2$ where H_2 is the entropy of the smaller subset. Since the information content of a given cell is fixed in the first case the reduction of entropy would be smaller than the second case. Thus, for a given c_l, the DM is more likely to learn under the lower entropy H_2. This generates the upper cutoff for entropy H_u, such that for any belief with uncertainty higher than H_u the reduction of uncertainty by observing only one cell would be higher than c_l and hence learning would not be optimal.

On the other hand given an expected payoff value if the entropy of belief μ_t is very small, learning about one more cell would not reduce uncertainty sufficiently. This would imply the gain in expected payoff would not be higher than the cost of learning c_l. This generates the lower cutoff H_l such that learning is only optimal when uncertainty is at least higher than H_l.

This generates an interval of entropy values $H(\mu)$ such that learning in round t is only optimal if the uncertainty in round t belongs to this interval given an expected payoff. If the expected payoff increaes, then the gain from learning decreases, since the highest possible payoff is given by $Y = 1$ for sure. Thus the size of interval increases with lower expected payoff.

□

A.3.3 Proof of lemma 3

Proof. Before we prove the lemma describing the optimal learning strategy let us define the uncertainty of the average as measured by $H^a(.)$. For any given prior belief we can define a distribution of possible values of averages. The uncertainty of this distribution is measures by the H^a, i.e., the entropy of the average distribution.

Observing an average can lead to partitioning the state space into more than two blocks. For example if in a given any number of cells can take the value $Y = 1$ then observing the average implies partitioning the state space into $n + 1$ possible blocks corresponding to the possible number of cells that can generate $Y = 1$.

There are two featutes of the learning by averages problem that are different than learning about cells. First, for any given prior belief μ we can define separately define the level of unceratinty for the averages and when measured by entropy this value would be lower than the entropy of the prior over original state space.

Also, since learning about averages is not always fully revealing there is no one-to one relationship between the two entropies. Reducing the entropy of the state space would not necessarily reduce the entropy of the average but reducing the entropy of average definitely reduces the entropy of the state space. This happens when the prior belief is such that the relationship between the two inputs are substitute or complements. In that case uncertainty of the average is zero and it is uninformative about true state but the uncertainty of the state space need not be zero.

Second, the possible partitions depends on the value of uncertainty of the state space. For example if the DM learns that there are only $k < n$ cells that can produce $Y = 1$ in a given row then observing the average would partition the state space into only $k + 1 < n + 1$ many blocks. Thus when the uncertainty of state space Ω is higher the possible reduction in uncertainty is also higher.

Similar to the learning by cell, learning is optimal only if

$$E(V(\mu_{t+1})) - \pi_t \geq c_a. \tag{LA}$$

Since under learning by averages the payoff need not be $Y = 1$ we cannot further simplify the expected value function from learning.

Learning about average is different as learning does not always perfectly reveal the true state even after observing all possible avearges, hence the main objective of the DM is to reduce the probability of choosing a wrong input combination. Suppose there are two possible average value such that the expected payoff from choosing both the rows (or columns) but the uncertainty in one average is higher than the other average then learning the former average reduces the probability of making mistake more. Thus learning about the average value that has the higher uncertainty is better.

Since when the uncertainty of the average is high the reduction of uncertainty by observing any single average would be high as well. This implies given the expected payoff there exists a cutoff value of average uncertainty, namely $\bar{H}^a$ such that if the uncertainty of the average is higher than this cutoff the expected gain through reduction in unceratinty is higher than cost of learning and hence learning becomes optimal. Since higher entropy of average imply higher uncertainty of the state space as well, we can rewrite the cutoff in terms of entropy of state space. However, when average is uninformative the cutoff would be defined at the highest possible value of uncertainty of Ω since learning is nowhere optimal.

Finally the cutoff increases with increase in expected payoff. This is evident from equation LA as LHS reduces with increase in μ_t making no learning more profitable. □

A.3.4 Proof of lemma 4

Proof. First we note that for both types of learning there exists a value of uncertainty such that no learning becomes optimal if uncertainty goes below this value. However for learning by cell this cutoff is defined over entropy of Ω and for for learning by average this is defined over entropy of the averages.

In case learning about averages does not reveal the true state the reduction of uncertainty of avearge does not correspond one-to-one with the reduction of uncertainty. This is true because a reduction in entropy may not lead to any reduction of average entropy if averages are uninformative.

First we will consider the cutoff in this case where averages are not fully informative. As entropy of the state space reduces the informatin content of the averages reduces with it, but the information content of cells always remain the same irrespective of the level of uncertainty as a cell can only take two values.

The major difference about learning by cells and learning by averages is that reduction is uncertainty decreases with level of uncertainty but learning cells always reduce unceratinty in the same proportion. Thus for lower value of uncertainty the averages become less profitable compared to the cells. If learning about cells become not profitable then learning about averges cannnot remain profitatble if it was not profitable before because as uncertainty has reduced averages are less informative. Thus the lower cutoff of uncertainty of Ω would be decided by the cutoff of the learning by cell problem.

Next, we consider the case when learning averages perfectly reveals the true state. In this case the reduction in uncertainty on average has a one-to-one relationship with uncertainty of the state space. In that case learning about average can be more profitable since $c_a \leq c_l$. In this the lower cutoff would be determined by the averages but it has a one-to-one relationship with the uncertainty in state space. Thus the lower cutoff can be described in terms of entropy of state space Ω.

Next, we will find the cutoff $\hat{H}(\mu)$. As the value of entropy starts to decrese the information content of averages decreases but learning about cells become more profitable. This implies there exist a cutoff value of entropy where the benefit from learning about cells net of the cost c_l is higher than the benefit from learning about averages net of c_a. Note that, if the cost of learning cells is too high the two cutoffs, $\hat{H}$ and $\underline{H}$ may become the same i.e, learning about cells is never optimal.

Finally below the cutoff $\hat{H}$ it would never be the case that learning about averages become more profitable. This is because of the two following reasons: one, reduction in uncertainty of state space necessarily reduces uncertainty of the average unless uncertainty of the average beomes zero in which case learning averages is no more. Second, as uncertainty of state space the information content of he averages reduce, thus if for a lower value of uncertainty it is optimal to observe averages then it would also be optimal to observe for higher value of entropy where reduction in entropy was higher for averages.

Finally we consider the highest cutoff $\bar{H}(\mu)$. This is derived from lemma 3. Learning is only optimal when uncertainty of the prior is high enough.

Note that $\bar{H}$ and $\hat{H}$ would be same if the averages are informative and cost of learning are not too high. As uncertainty reduces averages become less informative but the increase in expected payoff by learning cells increases and hence the DM would switch from learning about averages to learninga about cell.

Finally the reason that there can be a gap in between $\bar{H}$ and $\hat{H}$ given any μ_t is because cost of observing cells can be sufficiently high. If the averages no longer remain profitable to learn about but the cost of learning cells is too high then learning cells would not become optimal even when learning averages is more profitable. □

A.3.5 Proof of corollary 1

Proof. Consider the following learning strategy: the DM chooses to observe the averages first and for some realizations where the uncertainty is sufficeintly reduced the DM switches to learning about cells. This policy is optimal since c_l is high enough such that at μ_0 the DM would not want to observe any cell.

Now if the cost c_l is reduced substantially such that learning about cells at μ_0 becomes profitable and the value of c_a is sufficiently high such that the DM prefers the new strategy due to higher net expected payoff.

If it is the case that under the first strategy the reduction of uncertainty was substantially higher than the reduction in entropy by observing any cell (which only partitions the state space into two blocks) then the resulting unceratinty can become higher under the second strategy even though it is welfare improving. □

A.3.6 Proof of corollary 2

Proof. Since learning about averages does not necessarily reveal the tru state, using only learning by average may lead to expetced payoff less than $Y = 1$. In case the learning strategy switches from learning cells to learning average for the same prior due to reduction in c_a can thus lead to higher probability of mistake.

Consider the following case, where the uncertainty of the prior was such that the optimal strategy was to observe the average and the as the uncertainty reduecs learning cells, i.e., for some round t the uncertainty is in $(\underline{H}(\pi_t), \hat{H}(\pi_t))$ making learning cells optimal.

When c_a reduces the cutoff $\bar{H}$ reduces as a result since learning averages become more profitable but the other two cutoffs do not change. If the reduction in cost is significantly high such that $\bar{H}(\mu_t) = \underline{H}(\mu_t)$ then only averages would be learned about and this would incresae the probaility of choosing a cell with payoff $Y = 0$ if avearage is not perefectly revealing. □