Evaluating the benefits of agricultural credit: Evidence from India

Sunil Mitra Kumar*

Abstract

Access to agricultural credit is emphasised in the policies of several developing countries, backed by the assumption that credit can aid investment and thereby farmers' income. Yet it is difficult to evaluate the benefits of agricultural loans at household level because farmers self-select into taking loans. In this study we use survey data from five semi-arid states of India to examine the household-level effects of loans on farm investment and on an index of assets. To account for the self-selection, we use propensity score matching to compare farmers who do and do not avail of loans but are otherwise nearly identical, and we do so using a subset of the data which form a panel, thereby enabling valid before-after inference. We find that loans lead to positive but very small and statistically insignificant effects on both outcomes, and interpret this as ambiguous evidence on the benefits of agricultural credit. While the majority of related literature focuses on the macro-effects of credit, our analysis shows that the borrower-level benefits of agricultural credit may be less apparent, and we explain these findings in relation to a recent literature on an ongoing agrarian decline in India. Our approach also demonstrates how large-scale survey data can be used to infer causal relations at the micro level.

Key Words: agriculture; credit; matching; treatment-effects; Asia; India.

*India Institute, King's College London, Strand, London WC2R 2LS. E-mail: Sunil.Kumar@kcl.ac.uk.
I am grateful to Ragupathy Venkatachalam and Bibhas Saha for comments on this paper, and to Orazio Attanasio and Arjan Verschoor for comments on a previous version. I thank seminar participants at the Indira Gandhi Institute for Development Research, the 2012 Annual Conference of the Royal Statistical Society, and the 2012 Asian Meeting of the Econometric Society for their suggestions and criticisms which have helped improve the paper substantially. The usual disclaimer applies.

1 Introduction

Credit helps economies to run and to grow (Besley, 1995), and its importance for developing countries is underscored by the fact that the poor often lack access to financial institutions (Morduch, 1995; Besley and Burgess, 2003). In particular, the agricultural sector remains a large employer in several developing countries (Awokuse and Xie, 2015), and farmers usually require credit because of the time lag between purchasing inputs and receiving returns after harvest (Conning and Udry, 2007). As a result, government policies have long emphasised better access to agricultural credit as a way to reduce poverty (Adams, 1971; Von Pischke and Adams, 1980).

In this paper we examine the household-level effects of agricultural bank loans in rural India using household survey data. Indian government policy has focused on expanding formal agricultural credit for over half a century (Sriram, 2007). This has been actioned through targeted lending at subsidized rates (Golait, 2007), compulsory sectoral-lending targets for banks (Reserve Bank of India, 2014), and expanding the network of rural bank branches (Burgess and Pande, 2005). As of 2003, nearly half of all farmer households had outstanding debt, of which nearly two thirds was from banks of various kinds (Government of India, 2007).

Even so, it is challenging to evaluate the micro-level effects of such credit because farmers self-select into taking loans. Without suitable data, it is difficult to disentangle the effects of loans from differences in demand, the use to which the credit is put, and the fungibility of credit (Meyer, 1990). Randomised experiments can address these problems of identification, but being unfeasible to implement over long time horizons or at scale, tend to be restricted to studying micro-finance borrowing (e.g. Pitt and Khandker, 1998; Banerjee et al., 2015). Alternatively, observational data can be coupled with strong assumptions to aid identification in a regression-based framework. Carter (1989) and Feder et al. (1990) use this approach, suggesting a theoretical model to link credit with agricultural productivity which is then estimated with the data.

We attempt to address these challenges by using matching techniques based on the propensity score (Rosenbaum and Rubin, 1983) to compare farmers who avail of loans with those who do not. We use a subset of the data which constitutes a panel, wherein the characteristics used to statistically explain loan status are measured a few months before the outcome variables are. This enables valid before-after inference, avoiding the potential confounding that might otherwise arise with purely cross-sectional data (Rosenbaum, 1984). Matching is the preferred method for evaluating large-scale job training programmes which exhibit the same characteristic of self-selection into treatment, and usually lack useful exclusion restrictions (Heckman et al., 1997; Sianesi, 2004; Lechner et al., 2011; Larsson, 2003; Wunsch and Lechner, 2008). Unlike regression, matching avoids extrapolation, and restricts the analysis to households who have similar – and ideally identical – characteristics. This minimises the influence of potential unobservables to the extent that they are correlated with observed characteristics, and we select the optimal matching algorithm by using quantile-quantile plots to evaluate the resulting extent of balance.

We focus on farmer households in five semi-arid states of India, and analyse the effects of credit on two linked outcomes. The first is a binary indicator for investment in farm enterprise, and the second, as a proxy for household wealth, is an index of consumer durable assets. The first outcome is thus a direct measure for the use of credit, while the second is a less direct measure of any resulting changes in household wealth. Our main finding is that the treatment effects of agricultural loans on both outcomes are positive, but very small and statistically insignificant. While raw comparisons suggest that farmers who obtain loans are more likely to invest as well as own more durables at a later date, this association weakens substantially once the comparisons are undertaken on matched samples.

This is a negative finding given the policy emphasis on agricultural credit, but it lends empirical support to a recent literature pointing towards an agrarian decline. This decline is thought to be linked to factors including low output and high input prices, limited access to markets, limited irrigation facilities, and decreasing state investment in agriculture (Mishra, 2008; Reddy and Mishra, 2009; Vakulabharanam and Motiram, 2011), but also to indebtedness, which in extreme cases has contributed to farmer suicides (Government of India, 2007; Vaidyanathan, 2006). While the majority of the literature focuses on macro-effects of credit, our analysis thus shows that the borrower-level benefits of agricultural credit might be more uncertain. Our approach also demonstrates how large-scale survey data can be used to infer causal relations at the micro level.

The remainder of the paper is structured as follows. Section 2 outlines the causal mechanism linking access to credit with agricultural incomes, introduces the outcomes we focus on, and describes the estimation methodology, while section 3 introduces the data. Section 4 presents our main results, and section 5 examines their sensitivity to changes in the matching process. Section 6 concludes.

2 Methodology

2.1 From credit to income

The role that credit plays in agricultural production and thereby the incomes of farmer households can be hypothesised in terms of the following steps.

a) obtaining a loan.

b) investing the loan in agricultural inputs and undertaking agricultural production. Besides the availability of agricultural inputs, this process would be influenced by the farmer's abilities, and the amount and quality of land owned by the household.

c) selling the harvest to obtain earnings, which would depend in part on the availability of markets, transport and storage facilities.

d) using these earnings for one or more of i) paying back the loan, ii) current consumption, iii) adding to savings. In adverse circumstances, the household might borrow yet more or sell existing assets to fund (i) and (ii).

The effects of agricultural loans could manifest in two ways: they could aid step (b), and they could increase household wealth levels through step (d).

By focusing exclusively on the role of agricultural loans, any corresponding causal effects will be averages over unmodeled heterogeneity. For instance, credit and savings are fungible, so that certain farmers might not borrow at all. Likewise in step (d), even with the same earnings from a harvest, some farmers will save more while others might consume more. In common with data from several developing countries, we do not directly observe household income or earnings from agriculture. Instead, we have information on the household's ownership of a number of consumer durable assets, and on any investments undertaken in agricultural business enterprise.

Our estimations therefore focus on two related outcomes. The first is a binary indicator of investment in farm enterprise, corresponding to step (b) above. This includes improvements to land and irrigation, the purchase or maintenance of machinery, and land or buildings. The second outcome is an index of consumer durables, since these are known to proxy households' wealth and income in the absence of direct information on income (Filmer and Pritchett, 2001). Following Kolenikov and Angeles (2009), we estimate this index using polychoric principal components analysis as detailed in appendix A. Again, since farmers would spend only part of their income on consumption, and in particular on the purchase of consumer durables, our estimates based on the ownership of these assets are averages over this unmodelled (and unobserved) heterogeneity.

2.2 Treatment effects

Our aim is to estimate the impact of bank loans on indicators of wealth and investments in farm enterprise. Using the Rubin causal framework (Rubin, 1974), a standard representation of this causal inference problem is as follows (see for instance Todd et al. (2008)). Let T denote treatment status with $T = 1$ if the household has an agricultural bank loan and 0 otherwise, and Y_1 denote the value of some outcome Y if the household had a loan and Y_0 if it did not. For a given household, the difference in outcome due to having a loan is $Y_1 - Y_0$, where only one of Y_1 and Y_0 is observed, depending on treatment status. We focus on the average treatment effect for the treated (ATT):

$$\theta = E(Y_1 - Y_0 | T = 1)$$

Therefore, the aim of our empirical strategy is to construct the counterfactual outcome Y_0 for treatment households. For identification of the ATT, we require one part of the 'strongly ignorable treatment' assumption (Rosenbaum and Rubin, 1983), namely, that the distribution of the counterfactual outcome Y_0 is independent of treatment status conditional on a vector of covariates $\mathbf{X}$ (where the more general assumption of $Y_1, Y_0 \perp T|\mathbf{X}$ would be required if our aim were to estimate the average treatment effect):

$$Y_0 \perp T|\mathbf{X} \tag{1}$$

From this, it follows that $E(Y_0|T = 1, \mathbf{X}) = E(Y_0|T = 0, \mathbf{X}) = E(Y_0|\mathbf{X})$. To rule out the presence of treatment units for whom no matching control units can be found a priori, we also require one part of the 'common support assumption', namely that for any given household, the probability of receiving a loan is strictly less than one:

$$P(T = 1|\mathbf{X}) < 1 \tag{2}$$

Given assumptions (1) and (2), we can use a suitable estimator to impute the missing Y_0 value for each observed Y_1, and use this to estimate the ATT:

$$\theta = E(Y_1 - Y_0|T = 1, \mathbf{X}) = E(Y_1|T = 1, \mathbf{X}) - E(Y_0|T = 1, \mathbf{X}) \quad (3)$$

$$= E(Y_1|T = 1, \mathbf{X}) - E(Y_0|T = 0, \mathbf{X}) \quad (4)$$

The validity of our analysis crucially rests on the untestable assumption (1). We argue in support of its validity through our use of a detailed set of covariates, and by subjecting our results to different sensitivity checks. The covariates used to undertake matching are explained in detail in section 3.2, and these include indicators of economic status, human capital, and region fixed effects. Nonetheless, it is possible that certain unobserved characteristics determine agricultural borrowing and the impact of this credit, such as motivation, and agricultural knowledge and skill. As partial proxies for these we utilise information on education levels and age of the household head, and to the extent that they are reflected in the household's economic status, they are also partially captured by the latter set of variables. However we cannot fully rule out the possibility that these or other unobservables could potentially bias our results. Therefore, we check the robustness of our results through various sensitivity analyses, both by changing different parameters of the matching process and by varying the set of covariates used to calculate the propensity score.

2.3 Matching estimators

Since loans are not distributed at random, the average difference in outcome between households with loans and those without does not yield an unbiased estimate of the ATT. A regression-based approach would attempt to adjust for this non-random treatment allocation by controlling for covariates through a linear model, but this involves extrapolation of the regression function. If the distribution of observed covariates is different by treatment group, the same is likely to apply to unobserved or omitted variables to the extent that they are correlated with observed variables. Treatment effect estimates would then be biased due to the influence of unobservables on selection into treatment and potential outcomes. Therefore, matching methods create a sample where all observed covariates have similar, and ideally identical distributions by treatment group. This minimises any corresponding differences in unobservables to the extent that the latter are correlated with observables.

In other words, matching tries to artificially create an ideal randomised experiment. In the simplest case, this is implemented by matching each treatment unit with a control unit that is ideally identical in terms of a vector of covariates $\mathbf{X}$. We follow a 'doubly robust' approach (Imbens and Wooldridge, 2009; Robins and Rotnitzky, 1995), and adjust the simple difference in mean outcome between matched treatment and control groups for any remaining imbalance in covariates. This is done using weighted regression, where each outcome variable is regressed on treatment status and the covariates used for the matching. This yields the treatment effect estimate as coefficient of the treatment indicator together with the standard error of this estimate.

Given the challenge of identifying matches with a multidimensional set of covariates, Rosenbaum and Rubin (1983) showed that matching on $\mathbf{X}$ was equivalent to matching on the propensity score, the probability of a household receiving the treatment conditional on a vector of covariates $\mathbf{X}$, thus reducing a multidimensional matching problem to a single dimension. In practice, the true propensity score is not known, and is estimated using a logit or probit model. An alternative approach is to use the Mahalanobis distance metric, however we do not use this since it been found to perform worse than propensity score matching when there are many covariates or these are non-Normally distributed (Stuart, 2010).

We employ different matching algorithms including 1-to-n matching, kernel, and radius matching, all based on the propensity score, and choose the best algorithm according to the resulting degree of balance across different covariates. To evaluate this degree of balance, we use summary measures of quantile-quantile plots and standardized bias plots following the recommendations of Imai et al. (2008); Ho et al. (2007); Stuart (2010) and Austin (2008).

The standardized bias before and after matching is defined by Rosenbaum and Rubin (1985) as follows (see, also, the discussion in Lee (2011)):

$$\mathrm{SB}_X^U = \frac{100(\bar{X}_{T,U} - \bar{X}_{C,U})}{\sqrt{\frac{s^2(X_T)+s^2(X_C)}{2}}} \tag{5}$$

$$\mathrm{SB}_X^M = \frac{100(\bar{X}_{T,M} - \bar{X}_{C,M})}{\sqrt{\frac{s^2(X_T)+s^2(X_C)}{2}}} \tag{6}$$

Here SB_X^U (SB_X^M) is the standardized bias for covariate X in the unmatched (matched) sample, and respective $\bar{X}$s refers to the mean of X in groups defined by treatment status (T, C) and the matched (M) or unmatched (U) sample. $s^2(X_T)$ denotes the sample variance for covariate x in the treatment sample, and $s^2(X_C)$ the same quantity for the control sample.

3 Data

We use data from the decadal Debt and Investment Survey carried out by the Indian National Sample Survey Organisation in 2003 (Round 59, Schedule 18.2). We focus on the five states classified as semi-arid tropical by the International Crop Research Institute for the Semi-Arid Tropics, viz. Andhra Pradesh, Tamil Nadu, Karnataka, Maharashtra and Madhya Pradesh (Dinar et al., 1998), and restrict the sample to rural households who report their main occupation as farming. Data were collected through interviews at the household level, and questions focused household demographics (including caste, household size and composition in terms of age and gender, primary occupation, education

Figure 1: Timeline of the survey

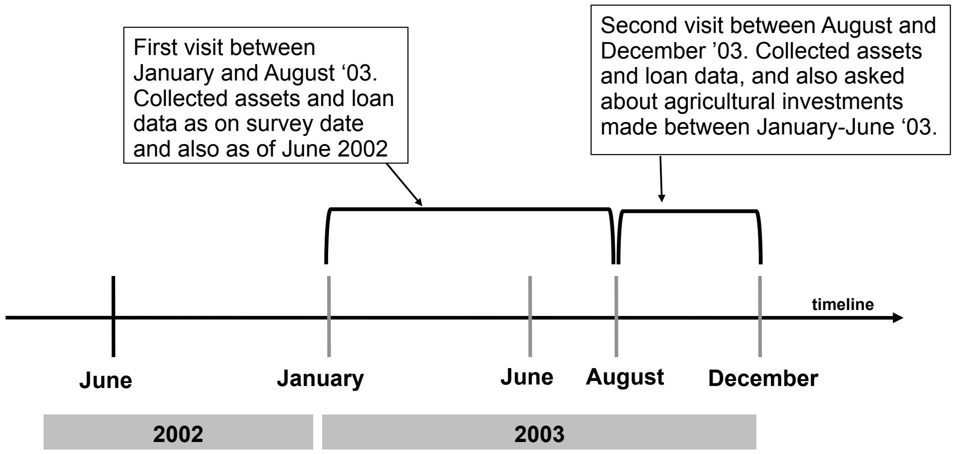

and marital status), land and other assets owned by the household (including consumer durables and farm equipment), and credit transactions in the recent past.

Credible measurement of treatment effects in an observational study requires that the covariates controlled for either be measured before treatment was allocated, or be independent of treatment status (Cochran and Chambers (1965); see also the discussion in Rosenbaum (2002)). To this end, an important feature of the survey is that households were interviewed in two separate visits a few months apart, thus yielding a panel for a subset of the data. The first visit took place between January and July 2003, and gathered data on all variables. This includes the consumer durable assets owned by the household and any outstanding loans as on the date of the survey as well as of June 1, 2002. The second visit took place between August and December 2003, and repeated questions about households' ownership of these assets and their borrowings. At this visit households were also asked whether they had invested in farm enterprise (purchase or upgrade of land and machinery) during the six months from January to June 2003, without specifying the exact date of such investments (see figure 1).

Table 1: Outcomes and definition of treatment

Outcome	**Treatment definition**
Index of consumer durable assets owned by household as on date of second survey visit	T_{assets}=1 iff household obtained an agricultural loan during July 2002 - July 2003
Binary indicator for investments in farm enterprise made by the household during January 2003 -June 2003	T_{invest}=1 iff households obtained an agricultural loan during July 2002 -December 2002

3.1 Definitions and sample restrictions

Our first outcome is an index of consumer durables owned by the household on the date of the second survey visit, while our second outcome is a binary indicator for investments in farm enterprise undertaken between January 2003 and June 2003. The index of consumer durable assets is constructed using polychoric principal components analysis (Kolenikov and Angeles (2009); also see appendix A).

Our definition of treatment status is based on obtaining loans with the stated purpose of capital or current expenditure on farm business from a cooperative or commercial bank. Given the timing at which information on the two outcomes was collected during the survey, we use two corresponding definitions for treatment status. Focusing on consumer durable assets, the outcome variable measured on the second survey visit, T_{assets}=1 for households who obtained a loan between July 2002 and July 2003 and zero otherwise. Similarly, since investments in farm enterprise are observed during January 2003 - June 2003, T_{invest}=1 for households who obtained a loan during July 2002 - December 2002, and zero otherwise. These definitions are summarised in table 1.

In order to minimise the confounding influence of previous borrowings we drop households who had existing loans from informal sources (moneylenders and friends or relatives) in June 2002 or who obtained them over the course of the two survey visits. Excluding these households is a necessary limitation of our analysis, as it is difficult to theorise how households might put loans from respective sources to different uses, or prioritise repayments of one loan over another. And in the absence of more detailed information, it would be difficult to identify the effects of bank loans from those of informal loans.

The same challenge of accounting for previous borrowings also applies to households who had existing agricultural bank loans in June 2002. Yet unlike borrowings from informal sources, the terms imposed by existing bank loans would be comparable to those of new loans, even though we do not have information on how households prioritise existing repayments over fresh investment. Therefore, we use two separate samples in our analysis. Our main results are based on the sample of households who did not have any existing loans as of June 2002, and we present separate estimates for households who already had an agricultural bank loan in June 2002 as part of the sensitivity analysis.

3.2 Selection of covariates

The covariates used for matching are chosen so as to capture three types of attributes. The first attribute is the household's economic status which directly influences the ability to borrow and make use of credit, and the variables used for this are land area owned, (and as a proxy for its quality) the value of this land, households' monthly per capita expenditure, and two indices for the household's ownership of a) consumer durable assets in June 2002, b) farm machinery in June 2002 (details of these indices are provided appendices A and B respectively). In particular, land ownership is known to be an important determinant of access to credit since land is used as collateral (Pal, 2002; Swain, 2007).

The second attribute we seek to capture is skill and experience, both of which are not directly observable but are likely important influences on agricultural production. We use the household head's age and years of education as proxies for these. The third attribute is demographic information to capture the availability of labour within the household, and the age-gender composition which would influence how any change in agricultural income is reflected in changes in consumer durables' ownership. The variables used are household size, sex of the household head, the proportion of children, adult males and adult females, and a dummy for whether a married adult son lives in the household.

Finally, we use dummies for the state in which the household resides to capture inter-state differences due to climate, market conditions, and potential unobservables. We also control for household caste-group, known to be an important determinant of social capital and economic status in India (Deshpande, 2001) and a significant determinant of access to credit (Kumar, 2013). This take on four values: Scheduled Tribes (ST) who are technically not part of the caste-system but have the lowest socio-economic status amongst the four groups, Scheduled Castes (SC) who are the lowest in the caste hierarchy, Other Backward Classes (OBC) who are of middling disadvantage, and Others, the high castes.

3.3 Summary statistics

Table 2 provides summary statistics for the outcome variables and covariates by treatment status. The assets index outcome variable has significantly higher levels in treatment households compared to control households (p=0.000 using a two-sample t-test), and the farm investment outcome has a significantly higher proportion among treatment households (p=0.001 using Pearson's chi-squared test).

Table 2 shows that on average, all households own quite little land. That said, respective groups of treatment households own more land compared to control households, and this trend holds across most correlates of economic status. Similarly, treatment households are more likely to have a male head of the household, to have a (resident) married male child, and to have a higher proportion of adults in the family. As we now discuss, the levels of some of these covariates are also statistically significant predictors of treatment status.

Table 2: Descriptive statistics

	Full sample		T_{assets}=1		T_{assets}=0		T_{invest}=1		T_{invest}=0	
	Means (Standard deviations in parentheses)									
Assets index second survey visit	-0.03	(1.373)	0.385	(1.476)	-0.114	(1.335)	-	-	-	-
Land area owned (Ha)	1.639	(2.01)	2.384	(2.931)	1.487	(1.724)	2.372	(2.937)	1.528	(1.803)
Value of land owned (Rs. '000)	248.91	(419.66)	365.354	(538.809)	225.028	(386.572)	372.743	(567.978)	230.051	(388.868)
Assets index June 2002	-0.035	(1.325)	0.353	(1.422)	-0.114	(1.29)	0.277	(1.432)	-0.082	(1.302)
Farm machinery index June 2002	-0.282	(0.65)	-0.014	(0.785)	-0.337	(0.605)	-0.009	(0.813)	-0.323	(0.611)
Household monthly per capita expenditure (Rs.)	503.606	(221.88)	541.923	(240.019)	495.747	(217.186)	534.502	(236.025)	498.9	(219.307)
Residential area owned	0.017	(0.054)	0.021	(0.058)	0.016	(0.053)	0.021	(0.056)	0.016	(0.054)
Years of education of household head	5.123	(5.787)	5.931	(5.73)	4.958	(5.786)	5.685	(5.621)	5.038	(5.808)
Age of household head	47.649	(13.931)	49.69	(13.736)	47.231	(13.936)	49.957	(13.63)	47.298	(13.945)
Household proportion of children	0.325	(0.233)	0.32	(0.223)	0.326	(0.236)	0.318	(0.223)	0.326	(0.235)
Household proportion of adult males	0.233	(0.165)	0.247	(0.149)	0.23	(0.168)	0.25	(0.152)	0.23	(0.166)
Household proportion of adult females	0.202	(0.196)	0.215	(0.174)	0.2	(0.201)	0.219	(0.171)	0.2	(0.2)
Household size	5.087	(2.602)	5.7	(2.743)	4.961	(2.555)	5.758	(2.709)	4.985	(2.571)
	Proportions									
Invested in farm enterprise during January 2003-June 2003	0.023		-		-		0.048		0.019	
Male household head	0.924		0.954		0.918		0.947		0.92	
Household with married male child	0.229		0.317		0.211		0.34		0.212	
Household caste group										
Other backward classes	0.131		0.106		0.136		0.11		0.134	
Others	0.139		0.09		0.149		0.1		0.145	
Scheduled Caste	0.467		0.468		0.466		0.473		0.466	
Scheduled Tribe	0.264		0.335		0.249		0.317		0.256	
State										
Andhra Pradesh	0.179		0.211		0.172		0.217		0.173	
Karnataka	0.126		0.126		0.126		0.116		0.128	
Madhya Pradesh	0.323		0.28		0.332		0.32		0.324	
Maharashtra	0.258		0.282		0.253		0.26		0.258	
Tamil Nadu	0.114		0.101		0.116		0.087		0.118	
Number of observations	3314		564		2750		438		2876	

4 Results

4.1 Estimating the propensity score

Table 3 shows the estimation results from probit models, where the dependent variables are, respectively, the twin treatment status binary variables. The signs of most covariates reflect the pattern of raw differences in table 2. Of these, the value of land owned is a statistically significant positive predictor of treatment status, as is the level of the assets index in June 2002 and the proportion of adult males in the family. Besides, there are significant inter-state differences in the proportion of households receiving loans. Household caste categories are not statistically significant predictors of treatment status. However, the coefficients themselves reflect known patterns of access to credit (Pal, 2002, e.g.), where, with scheduled tribes (STs) as the base, higher caste 'Others' are more likely to have a loan, and scheduled castes (SCs, the lowest in the caste hierarchy) are least likely to have a loan.

We use these estimated coefficients to predict the respective probabilities of treatment, i.e. the propensity score. Figure 2 shows the distributions of these estimated propensity scores, and establishes that there is substantial overlap in the distribution for treatment and control groups.

4.2 Matching

We use nearest-neighbour matching (n=1 to 10) and kernel matching. Nearest neighbour matching is undertaken with replacement, since this allows good matches to be used multiple times thereby reducing bias (Dehejia and Wahba, 2002), and we also impose a caliper equal to one quarter of a standard deviation of the propensity score to prevent dissimilar matches. The same quantity is used as the bandwidth for performing kernel

Figure 2: Distribution of propensity scores by treatment status

(a) T_{assets}

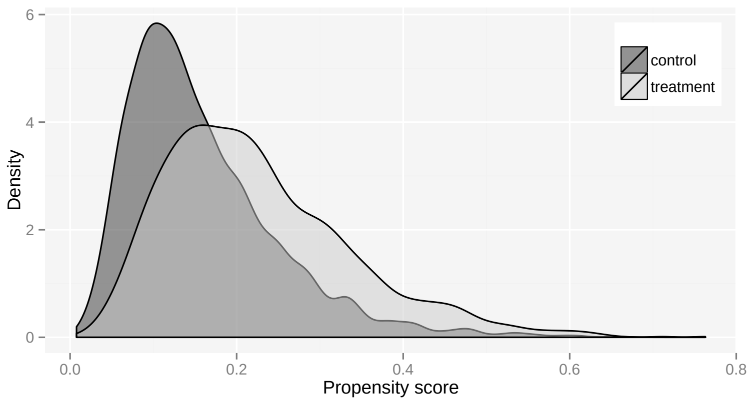

(b) T_{invest}

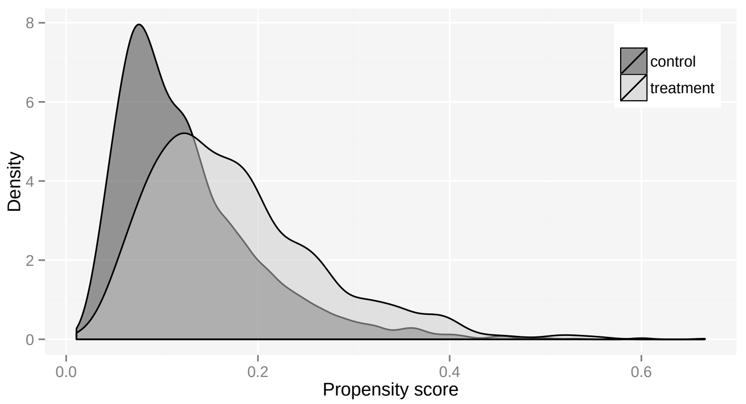

Table 3: Estimation of propensity score

Variable	Dependent variable: T_{assets}		Dependent variable: T_{invest}	
(log) Land area owned	0.100**	(0.0403)	0.0570	(0.0428)
(log) Value of land owned	0.0886**	(0.0388)	0.120***	(0.0414)
Assets index June 2002	-0.00731	(0.0259)	-0.0395	(0.0277)
Farm machinery index June 2002	0.222***	(0.0472)	0.205***	(0.0495)
(log) Household monthly per capita expenditure	0.0635	(0.0847)	0.0713	(0.0899)
Residential area owned	-0.0639	(0.489)	-0.0448	(0.509)
Years of education household head	0.00600	(0.00549)	0.00593	(0.00587)
Age of household head	0.00451	(0.00292)	0.00448	(0.00310)
Male household head	0.199*	(0.119)	0.0901	(0.124)
Household has married male child	0.0802	(0.0840)	0.121	(0.0886)
Household proportion of children	0.251	(0.192)	0.193	(0.204)
Household proportion of adult males	0.471**	(0.208)	0.531**	(0.221)
Household proportion of adult females	0.107	(0.193)	0.142	(0.206)
Household size	0.0180	(0.0148)	0.0201	(0.0155)
Household caste group				
Scheduled Tribe	(base)		(base)	
Scheduled Caste	-0.128	(0.115)	-0.0777	(0.121)
Other Backward Classes	-0.00348	(0.0935)	0.00248	(0.0989)
Others	0.0650	(0.103)	0.0281	(0.109)
State				
Andhra Pradesh	(base)		(base)	
Karnataka	-0.403***	(0.0998)	-0.434***	(0.107)
Madhya Pradesh	-0.466***	(0.0847)	-0.361***	(0.0885)
Maharashtra	-0.285***	(0.0827)	-0.321***	(0.0883)
Tamil Nadu	-0.270**	(0.108)	-0.366***	(0.117)
Constant	-2.834***	(0.642)	-3.354***	(0.682)
Number of observations	3314		3314	
Log likelihood	-1405.2		-1213.8	

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

matching using an Epanechnikov kernel.

A central question surrounding matching methods is on the selection of the optimal matching algorithm. Since the aim of matching is to artificially recreate an ideal randomized experiment, this translates into creating a matched set with identical joint distributions across all covariates, and the extent of balance that actually results from the matching ought to be used as the criterion for choosing the optimal algorithm. Since it is difficult to compare multi-dimensional joint distributions, a compromise is to consider pair-wise distributions by treatment status for each covariate. (Imai et al., 2008) argue why these comparisons should be based not on statistical tests but on quantile-quantile plots (which are independent of sample size) and comparisons of standardised bias (also see Stuart (2010); Ho et al. (2007)).

We follow this approach, and focus on quantile-quantile plots for five covariates that are important determinants of access to and the use of credit: the (log) value of land owned, the total area of land owned, the index of consumer durable assets on the baseline date (June 2002), household monthly per capita expenditure, and years of education of the household head.

In order to construct the quantile-quantile plots, matched sets of treatment and control observations must be constructed after taking into account the weights obtained from the matching process (Joffe et al., 2004). We do this by expanding the matched dataset such that the number of times a given observation is replicated is proportional to the weights placed on it by the matching process. And, we summarise these quantile-quantile plots in terms of the percentage reduction in mean deviation for unmatched relative to matched samples from the 45° line of perfect symmetry. Table 4 shows these results for respective matching algorithms.

For both outcome variables and their corresponding definitions of treatment, kernel matching results in the least decrease in imbalance. Among the 1-n algorithms, matching on three or more nearest neighbours results in a 90% or greater decline in imbalance for most covariates, with the exception of years of education of the household head. In addition, figure 3 shows box plots of the standardized bias across all continuous variables used in the propensity score estimation for each matching algorithm. Unlike table 4, this metric focuses only on the difference in means and not the overall distribution, and so is arguably a less useful metric to evaluate balance. In contrast to table 4, these plots suggest that kernel matching performs best, with no obvious pattern for nearest-neighbour matching. Taken together with 4, this suggests that kernel matching helps reduce the imbalance in means, but not in the overall distribution of respective variables.

Since we prioritise reductions in imbalance across the overall distributions, we select 1-3 matching for the assets index outcome and 1-8 matching for the farm investment outcome. These algorithms result in the highest reduction in imbalance for years of education, and achieve at least a 90% reduction in imbalance for the remaining four covariates in table 4.

4.3 The impact of loans

Table 5 shows the estimated treatment effects. These are calculated using weighted regression such that treatment and control units are almost identical in terms of background covariates (Ho et al., 2007). Matching with the selected 1-n algorithm yields a set of weights, where each matched treatment observation has weight one, and each control observation has weight equal to the number of times it is used as a match, normalised by n. These matching weights are then multiplied with the survey probability weights, and the quantities thus obtained are used as weights for the regression in order that the treatment effect estimate is representative for the survey population.

For both investments in the farm enterprise and the index of durable assets owned by the household, the treatment effects are positive but statistically insignificant. Notwith-

Table 4: Reduction in imbalance for key variables by matching algorithm

(a) T_{assets} (assets index)

	Nearest neighbour matching										Kernel
Variable	1	2	3	4	5	6	7	8	9	10	
(log) Land area owned	95.84	99.45	99.01	98.00	97.87	97.24	99.06	96.94	96.87	96.68	81.42
(log) Value of land owned	92.62	97.68	99.99	97.33	96.07	94.46	95.16	96.18	96.38	96.89	83.86
Assets index June 2002	96.83	94.12	96.90	94.77	90.41	92.14	93.13	94.07	94.72	94.57	80.43
(log) household monthly per capita expenditure	93.20	77.76	90.74	95.77	88.52	86.09	85.44	85.46	88.19	90.08	84.55
Years of education household head	77.15	86.56	91.10	93.37	99.38	97.87	99.82	98.10	99.78	97.50	80.50

(b) T_{invest} (farm investments)

	Nearest neighbour matching										Kernel
Variable	1	2	3	4	5	6	7	8	9	10	
(log) Land area owned	98.03	95.55	91.37	91.41	94.03	91.32	93.15	94.04	93.52	91.96	68.54
(log) Value of land owned	87.62	89.69	95.84	95.89	95.84	96.42	97.43	97.69	99.33	99.49	68.93
Assets index June 2002	66.34	94.87	99.32	98.40	99.57	95.55	97.65	97.62	98.96	99.71	66.69
(log) household monthly per capita expenditure	34.69	74.13	77.36	88.72	85.48	93.14	94.69	95.10	94.34	99.59	71.72
Years of education household head	73.08	75.03	77.78	71.93	75.34	74.44	69.94	77.00	79.05	72.45	66.35

Notes: Numbers shown are the percentage reduction in imbalance for respective matching algorithms. This is measured, using quantile-quantile plots, as the percentage reduction in mean deviation from the 45° line of perfect symmetry, for unmatched relative to matched samples.

Figure 3: Standardised bias plots

(a) Assets index (T_{assets}) (b) Farm investments (T_{farm})

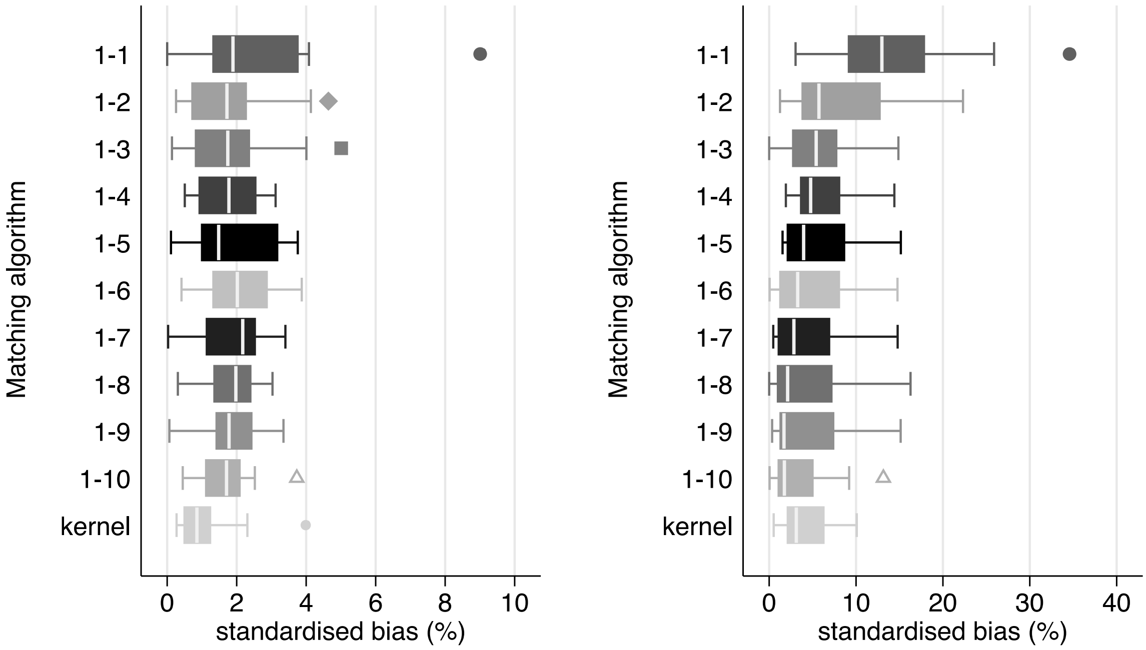

Table 5: Treatment effects for assets and farm investment

	θ	Standard error	p-value
Assets index	0.0491	0.0319	0.1241
Farm investment	0.0198	0.0151	0.1902

standing debates on statistical significance (e.g. McCloskey & Ziliak 1996), the absolute magnitude of the treatment effect for the assets index is also small, corresponding to less than 0.04 of a standard deviation. The treatment effect is larger for the farm investment outcome (0.019 compared to the sample proportion of 0.023), but the accompanying p-value is also large (p=0.19). Compared to the raw differences in table 2, the matching-adjusted differences are thus both smaller in absolute value and statistically insignificant. In other words, we do not find evidence that obtaining an agricultural loan makes it more likely that households will invest in farm enterprise once we account for other attributes, and nor do such loans lead to higher levels of assets ownership.

While these results suggest a negative answer to the overall question this paper seeks to address – whether agricultural credit is useful – the results for the two outcomes have distinct but related implications. As a measure of the overall impact of agricultural borrowing (and the culmination of steps (a)-(d) outlined in section 2), the assets index outcome is arguably the more significant in an economic sense. Yet this metric is also likely insensitive to small changes in income, especially over a relatively short period of approximately one year, as in this data. Alternatively, it could also be the case that the 'true' treatment effect is zero, with only some households benefiting and others incurring losses due to (here unobserved) shocks. Unfortunately the current data do not enable us to differentiate between these scenarios.

There is less ambiguity, however, regarding the second outcome variable, since investments in agricultural enterprise (step (b) in section 2) are a necessary component of the production process. Since the timing of successive observations in the data are suitable for observing agricultural borrowing and subsequent investments, our finding of an insignificant treatment effect is less likely to be a limitation of the data and empirical approach, and more likely evidence of the unobserved variations in the use of agricultural credit (e.g. for consumption instead of investment). While we cannot directly observe such variations in the data at hand, overall our results support the mixed evidence surrounding financial access and indebtedness due to which the net effect of agricultural credit remains ambiguous Vakulabharanam and Motiram (2011); Vaidyanathan (2006). We discuss this issue further in section 6, and before that, examine the sensitivity of our results.

5 Sensitivity analysis

We now discuss the sensitivity of our results to changes in the matching process and the selection of covariates used in the analysis. Table 6 shows respective treatment effects corresponding to changes in parameters of the selected matching algorithms. Part (a) shows the changes in treatment effects if we drop observations in the thinnest 5% and 10% of the distribution of the propensity score. Doing so excludes observations that are outliers in terms of the propensity score, though we would not expect such trimming to change the quality of matches because the use of calipers would in any case exclude poor matches. The results show that the estimated treatment effect(s) remain largely the same even if observations on the relatively thin parts of the propensity score distribution are excluded.

Part (b) of table 6 shows the effects of varying the calipers used for matching to exclude poor matches. The first row corresponds to table 5, where the caliper is set to one quarter of a standard deviation of the estimated propensity score, while rows 2 and 3 use smaller calipers of, respectively, 0.10 and 0.05 of a standard deviation. Since calipers define the maximum permissible difference in propensity scores of any two matched observations, smaller calipers would improve the quality of matches. Correspondingly, the number of unmatched observations also increases. The results show that treatment effects are largely insensitive to the change in calipers, and that even with more exacting matching criteria, the treatment effects for both outcomes remain small and statistically insignificant.

Table 6: Sensitivity to trimming and calipers

(a) Sensitivity to trimming

	Assets index				Farm investment			
% observations dropped from the thinnest support of propensity score	θ	Standard error	p-value	Unmatched treatment observations	θ	Standard error	p-value	Unmatched treatment observations
No trim	0.0491	0.0319	0.1241	0	0.0198	0.0151	0.1902	1
Trim 5%	0.0369	0.0302	0.2225	28	0.0227	0.0144	0.1146	21
Trim 10%	0.0320	0.0315	0.3099	56	0.0160	0.0142	0.2606	43

(b) Sensitivity to calipers

	Assets index				Farm investment			
Caliper as a proportion of standard deviation of the propensity score	θ	Standard error	p-value	Unmatched treatment observations	θ	Standard error	p-value	Unmatched treatment observations
0.25	0.0491	0.0319	0.1241	0	0.0198	0.0151	0.1902	1
0.10	0.0528	0.0327	0.1061	1	0.0207	0.0148	0.1631	2
0.05	0.0510	0.0321	0.1129	4	0.0193	0.0152	0.2059	3

Table 7: Sensitivity to use of survey weights and selection of covariates

	Assets index			Farm investment		
	θ	Standard error	p-value	θ	Standard error	p-value
Excluding survey weights	0.0295	0.0221	0.1824	0.0162	0.0116	0.1642
Excluding age and education	-0.0046	0.0265	0.8610	0.0245	0.0146	0.0924
Excluding economic status covariates[a]	0.3395	0.0948	0.0003	0.0355	0.0134	0.0083
Excluding household demographics[b]	0.0286	0.0244	0.2427	0.0244	0.0149	0.1019

[a] Value and area of land owned, household monthly per capita expenditure, farm machinery index, assets index June 2002, and the residential land area owned.

[b] Sex of the household head, a binary indicator for a married son living with the family, household proportions of children, adult males and adult females, and household size.

Next we examine how the results presented in table 5 are sensitive to the use of survey weights and the exclusion of three categories of covariates in the propensity score estimation: ability and experience, economic status, and household demographics. Table 7 shows that our results are invariant to the exclusion of survey weights in that the results remain statistically insignificant. However, both the absolute magnitude of treatment effect as well as statistical significance change substantially according to the choice of covariates used to undertake the matching. That is, excluding indicators of economic status yields a much larger, positive, and statistically significant treatment effect for both outcome variables ($p<0.01$). This is not surprising, as we would expect households' existing economic status to predict not only access to credit, but also subsequent wealth as captured by the assets index outcome. Excluding these variables from the matching process would then lead to differences in the assets outcome being correlated with treatment status.

Finally, we also calculate treatment effects for households who already had bank loans as of June 2002. For these households, the respective treatment definitions of T_{assets} and T_{farm} are according to whether they obtained a new loan in the corresponding periods under study as described in table 1. There are 998 such households in the sample, and appendix C provides summary statistics, results for estimation of the propensity score(s)

Table 8: Sample of households with existing bank loans in June 2002

	θ	Standard error	p-value
Assets index	-0.0741	0.0664	0.2651
Farm investment	0.0070	0.0175	0.6894

and selection of the optimal matching algorithm for this sample. The treatment effect estimates for this sample are given in table 8, which shows that there is no statistically significant evidence of a treatment effect on either outcome. In fact, the treatment effect for the assets index of negative here (with a large p-value), and the proportional change in farm investment is much smaller in absolute magnitude (0.0070 vs 0.019) compared to the results in table 5 (again, with a large p-value). Thus, the main result remains qualitatively unchanged for this sample as well, demonstrating that *additional* loans too do not have a statistically significant influence on farm investments or assets.

6 Discussion

We have attempted to measure how obtaining agricultural credit influences subsequent investments in farm enterprise and an assets index for farmer households in the semi-arid states of India. We define treatment status according to whether the household obtained a new agricultural loan in the approximately year-long period under study, and use propensity score matching to estimate treatment effects, selecting the best matching algorithm according to the corresponding extent of balance. In order to minimise the effects of unobservable differences between households as a result of pre-existing borrowings, our main results are based on a subset of households who did not have any existing loans. As a sensitivity check, we also present results for households who already had bank loans, some of whom went on to borrow a fresh loan.

Our results are largely negative when viewed in the context of policies that emphasise agricultural credit as key to increasing rural incomes. That is, we find limited evidence that obtaining credit increases the probability of investing in agricultural production or results in higher values of a wealth index, and this evidence is statistically insignificant. Further, these results are robust to various sensitivity checks such as changes in the matching algorithm, the inclusion of survey weights, and the use of regression adjustment. In other words, since we have focused on average treatment effects for the treated, we find no significant evidence that households who availed of agricultural bank credit would have fared worse had they not done so, both in terms of investments in agriculture as well as increases in an assets index.

While these findings appear to undermine some of the main assumptions behind government policies for rural financial inclusion in India, our analysis also has certain important limitations. The primary limitation stems from the nature of data available, in that the period of time over which we can observe the same household and any changes in the situation thereof is approximately one year. The earliest data for loans and assets are for June 2002, which households are asked to list retrospectively at the first survey visit in 2003, while the latest data are from the second survey visit in late 2003 (as figure 1 details). Arguably, a one-year period might limit the extent to which changes in economic status or incomes can manifest, which then limits any analysis into causal links between these and agricultural loans.

Multiple observations from the same households over a longer period of time would also enable more detailed analysis that allows for multiple borrowings from different sources, for instance, to examine how households fare when borrowing from both informal and formal sources. In the absence of such data it is difficult to argue for the valid identification of treatment effects given potentially unobservable heterogeneity in existing loans and commitments on these at any one point in time. To avoid this problem, the main part of our analysis has focused on households without existing loans, who therefore represent only a subset – albeit a significant one – of farmers as a whole. Furthermore, in the absence of information about soil quality and the types of crops grown, our results are effectively averaged over across these characteristics. We have partially addressed this issue by restricting the analysis to the agro-climatically similar semi-arid states of India and allowing for state fixed effects so as to capture some of this heterogeneity. Yet it is possible that credit proves more or less useful for certain combinations of soil and cropping patterns, even if at an average level the net effect is statistically insignificant.

Bearing in mind these limitations, our findings are nonetheless consistent with recent literature on the stagnation of agricultural incomes in India (Vakulabharanam and Motiram, 2011; Reddy and Mishra, 2009). This literature argues that alongside other factors and inputs, credit can support higher productivity and incomes, but that it can also lead to unmanageable levels of debt after which incomes can stagnate or even decline if households sell assets to pay off debts if other agricultural production is not profitable as a whole. The other components needed for this include risk mitigation mechanisms, seed, fertilizer and pesticide technology, and farmers' skills and expertise (Sriram, 2007). Our results then support this mixed view of agricultural credit in a major developing country, suggesting that the focus of agricultural policy ought to be broader and one that emphasises other inputs beyond credit.

References

Adams, D. W. (1971), "Agricultural credit in Latin America: A critical review of external funding policy", *American Journal of Agricultural Economics*, 53, 163–172.

Austin, P. (2008), "A critical appraisal of propensity-score matching in the medical literature between 1996 and 2003", *Statistics in medicine*, 27, 2037–2049.

Awokuse, T. O. and Xie, R. (2015), "Does agriculture really matter for economic growth in developing countries?", *Canadian Journal of Agricultural Economics/Revue canadienne d"agroeconomie*, 63, 77–99.

Banerjee, A., Duflo, E., Glennerster, R. and Kinnan, C. (2009), "The miracle of microfinance? Evidence from a randomized evaluation", *American Economic Journal: Applied Economics*, 7, 22-53.

Besley, T. (1995), Savings, credit and insurance, *in* T.N. Srinivasan and J. Behrman, eds, "Handbook of development economics, vol. 3A", Amsterdam, Elsevier, pp. 2123–2207.

Besley, T. and Burgess, R. (2003), "Halving global poverty", *The Journal of Economic Perspectives*, 17, 3–22.

Burgess, R. and Pande, R. (2005), "Do rural banks matter? Evidence from the Indian social banking experiment", *The American Economic Review*, 95, 780–795.

Carter, M. R. (1989), "The impact of credit on peasant productivity and differentiation in Nicaragua", *Journal of Development Economics*, 31, 13–36.

Cochran, W. G. and Chambers, S. P. (1965), "The planning of observational studies of human populations", *Journal of the Royal Statistical Society Series A (Statistics in Society)*, 128, 234–266.

Conning, J. and Udry, C. (2007), Rural financial markets in developing countries, *in*

R. Evenson, P. Pingali, and T.P. Schultz, eds, "Handbook of agricultural economics, vol. 3", Elsevier, New York, pp. 2857–2908.

Dehejia, R. and Wahba, S. (2002), "Propensity score-matching methods for nonexperimental causal studies", *Review of Economics and Statistics*, 84, 151–161.

Deshpande, A. (2001), "Caste at Birth? Redefining Disparity in India", *Review of Development Economics*, 5, 130–144.

Dinar, A., Mendelsohn, R., Evenson, R., Parikh, J., Sanghi, A., Kumar, K., McKinsey, J. and Lonergen, S. (1998), Measuring the impact of climate change on Indian agriculture, Paper 402, World Bank - Technical Papers.

Feder, G., Lau, L. J., Lin, J. Y. and Luo, X. (1990), "The relationship between credit and productivity in Chinese agriculture: A microeconomic model of disequilibrium", *American Journal of Agricultural Economics*, 72, 1151–1157.

Filmer, D. and Pritchett, L. H. (2001), "Estimating wealth effects without expenditure of data or tears: An application to educational enrollments in states of India", *Demography*, 38, 115–132.

Golait, R. (2007), "Current issues in agriculture credit in India: an assessment", *Reserve Bank of India Occasional Papers* 28, 79–99.

Government of India (2007), "Report of the expert group on agricultural indebtedness". Government of India, Ministry of Finance, New Delhi.

Heckman, J. J., Ichimura, H. and Todd, P. E. (1997), "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme", *The Review of Economic Studies*, 64, 605–654.

Ho, D., Imai, K., King, G. and Stuart, E. (2007), "Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference", *Political Analysis*, 15, 199–236.

Imai, K., King, G. and Stuart, E. (2008), "Misunderstandings between experimentalists and observationalists about causal inference", *Journal of the Royal Statistical Society Series A (Statistics in Society)*, 171, 481–502.

Imbens, G. W. and Wooldridge, J. M. (2009), "Recent developments in the econometrics of program evaluation", *Journal of Economic Literature*, 47, 5–86.

Joffe, M. M., Have, T. R. T., Feldman, H. I. and Kimmel, S. E. (2004), "Model selection, confounder control, and marginal structural models: Review and new applications", *The American Statistician*, 58, 272–279.

Kolenikov, S. and Angeles, G. (2009), "Socioeconomic status measurement with discrete proxy variables: Is principal component analysis a reliable answer?", *Review of Income and Wealth*, 55, 128–165.

Kolenikov, S. (2009), "polychoric: The polychoric correlation package". http://web.missouri.edu/ kolenikovs/stata.

Kumar, S. M. (2013), "Does access to formal agricultural credit depend on caste?", *World Development*, 43, 315–328.

Larsson, L. (2003), "Evaluation of Swedish youth labor market programs", *Journal of Human Resources*, 38, 891–927.

Lechner, M., Miquel, R. and Wunsch, C. (2011), "Long-run effects of public sector sponsored training in West Germany", *Journal of the European Economic Association*, 9, 742–784.

Lee, W. (2011), "Propensity score matching and variations on the balancing test", *Empirical Economics*, 44, 47–80.

McCloskey, D. N. and Ziliak, S. T. (1996), "The standard error of regressions" , *Journal of Economic Literature*, 34, 97–114.

Meyer, R. L. (1990), "Analyzing the farm-level impact of agricultural credit: Discussion", *American Journal of Agricultural Economics*, 72, 1158–1160.

Mishra, S. (2008), "Risks, farmers" suicides and agrarian crisis in India: Is there a way out?", *Indian Journal of Agricultural Economics*, 63, 38–54.

Morduch, J. (1995), "Income smoothing and consumption smoothing", *The Journal of Economic Perspectives*, 9, 103–114.

Pal, S. (2002), "Household sectoral choice and effective demand for rural credit in India", *Applied Economics*, 34, 1743–1755.

Pitt, M. M. and Khandker, S. R. (1998), "The impact of group-based credit programs on poor households in Bangladesh: Does the gender of participants matter?", *Journal of Political Economy*, 106, 958–996.

Reddy, D. N. and Mishra, S., eds (2009), *Agrarian crisis in India*, Oxford University Press, New Delhi.

Reserve Bank of India (2014), "Priority sector lending - Targets and classification". Reserve Bank of India, Mumbai.

Robins, J. M. and Rotnitzky, A. (1995), "Semiparametric efficiency in multivariate regression models with missing data", *Journal of the American Statistical Association*, 90, 122–129.

Rosenbaum, P. R. (2002), *Observational Studies*, Springer Verlag, London.

——— (1984), "The consequences of adjustment for a concomitant variable that has been affected by the treatment", *Journal of the Royal Statistical Society Series A (Statistics in Society)*, 147, 656–666.

Rosenbaum, P. R. and Rubin, D. B. (1983), "The central role of the propensity score in observational studies for causal effects", *Biometrika*, 70, 41–55.

——— (1985), “Constructing a control group using multivariate matched sampling methods that incorporate the propensity score”, *The American Statistician*, 39, 33–38.

Rubin, D. B. (1974), “Estimating causal effects of treatments in randomized and nonrandomized studies.”, *Journal of Educational Psychology*, 66, 688.

Sianesi, B. (2004), “An evaluation of the Swedish system of active labor market programs in the 1990s”, *Review of Economics and Statistics*, 86, 133–155.

Sriram, M. S. (2007), “Productivity of rural credit”, *International Journal of Rural Management*, 3, 245 –268.

Stuart, E. A. (2010), “Matching methods for causal inference: A review and a look forward”, *Statistical Science*, 25, 1–21.

Swain, R. B. (2007), “The demand and supply of credit for households.”, *Applied Economics*, 39, 2681–2692.

Todd, P., Durlauf, S. N. and Blume, L. E. (2008), Matching estimators, *in* S. N. Durlauf and L. E. Blume, eds, “The new Palgrave dictionary of economics”, second edn, Palgrave Macmillan, New York.

Vaidyanathan, A. (2006), “Farmers’ suicides and the agrarian crisis”, *Economic and Political Weekly*, 41, 4009–4013.

Vakulabharanam, V. and Motiram, S. (2011), Political economy of agrarian distress in India since the 1990s, *in* S. Ruparelia, S. Reddy, J. Harriss and S. Corbridge, eds, “Understanding India’s new political economy”, Routledge, London, pp. 101–126.

Von Pischke, J. D. and Adams, D. W. (1980), “Fungibility and the design and evaluation of agricultural credit projects”, *American Journal of Agricultural Economics*, 62, 719–726.

Wunsch, C. and Lechner, M. (2008), "What did all the money do? On the general ineffectiveness of recent West German labour market programmes", *Kyklos*, 61, 134–174.

Appendices

A Assets index

The main question this paper seeks to answer is "Do farmers benefit from access to agricultural loans?". In addition to studying investments in farm enterprise, the 'benefit' we are seeking to measure changes in refers to economic status and the linked notion of permanent income, neither of which are easy to define or directly observable. Further, since it is also difficult to directly measure incomes in rural contexts in developing countries such as India, surveys typically gather information about household expenditure and asset ownership as proxies for economic status. Of the two, asset ownership is generally considered superior, being less prone to measurement errors and less sensitive to price differences across regions.

Information about asset ownership can be converted to a single-dimensional index by assigning weights to each asset category. Filmer & Pritchett (2001) suggest what is now a popular approach to estimating the weights, using principal components analysis (PCA). PCA calculates the weights which will yield linear combinations with the maximum variance, and proceeds by calculating the eigenvectors and eigenvalues for the correlation matrix (of asset vectors). This approach ensures that asset categories with larger variation across the sample receive higher weights, and vice-versa. However, as Kolenikov & Angeles (2009) argue, PCA works best with normally-distributed continuous variables, whereas asset ownership is typically in terms of non-negative discrete numbers. Therefore, they instead suggest using polychoric PCA.

Polychoric PCA assumes that the observed discrete variables are in fact discretised versions of unobserved normally-distributed variables. The correlations between respective asset categories are then calculated as the maximum-likelihood estimates of the correlation between these unobserved variables, and the PCA is performed on the correlation matrix derived using maximum-likelihood. Using simulation and data from Bangladesh, Kolenikov & Angeles (2009) demonstrate that polychoric PCA usually helps explain a greater proportion of variance in the original data compared to ordinary PCA.

Our data are cardinal, in that the survey records the number of units of each category of asset that a household owns in June 2002 and then on the second survey visit. The aim is to convert this information into a ranking of the households, to be able to examine changes in households' ranks between these two dates. We use data on the following assets:

1. beds
2. steel / wooden almirah / dressing table
3. pressure cooker/ household utensils
4. electric fan, clock/ watch, water filter / electric iron/ sewing machine
5. stoves
6. radio, record player/tape recorder/stereo/ musical instruments for household use
7. television, VCR/VCP/VCD, DVD Player, home theatre, multimedia PC
8. refrigerator/ air cooler/ air conditioner/ washing machine

The estimation is carried out with the Stata package *polychoric* written by Kolenikov (2009). Given that the maximum likelihood estimation of polychoric correlation is computationally intensive, the programme treats variables with more than 10 categories as being Normally-distributed. The first four variables in the above list are thus treated as continuous. Pair-wise correlations amongst them are calculated as regular (Pearson) correlations. Correlations between one of them and one of the ($<$10 category) variables are polyserial, with one variable treated as Normal and the other as the discretised version of an unobserved Normally-distributed variable.

Panel (a) in table 9 shows the matrix of polychoric correlations between respective asset categories for the June 2002 data, and panel (a) in table 10 shows the corresponding eigenvalues estimated using PCA. Both bottom panels (b) in tables 9 and 10 show the corresponding statistics for the assets owned by the household as on the second survey visit. Table 10 shows that nearly half the total variation in the assets-related variables is captured by the first (polychoric) principal component score for both sets of data.

Table 9: Polychoric correlations for assets data

a) For June 2002 assets data

	beds	almirahs	cooker/utensils	fan etc	stoves	radio etc	TV etc	fridge etc
beds	1.000							
almirahs	0.068	1.000						
cooker/utensils	0.114	0.294	1.000					
fan etc	0.101	0.453	0.396	1.000				
stoves	0.028	0.163	0.175	0.238	1.000			
radio etc	0.017	0.336	0.361	0.628	0.241	1.000		
TV etc	0.199	0.470	0.432	0.595	0.253	0.492	1.000	
fridge etc	0.230	0.326	0.427	0.693	0.223	0.624	0.473	1.000

b) For assets data on the second survey visit

	beds	almirahs	cooker/utensils	fan etc	stoves	radio etc	TV etc	fridge etc
beds	1.000							
almirahs	0.082	1.000						
cooker/utensils	0.118	0.318	1.000					
fan etc	0.106	0.471	0.387	1.000				
stoves	0.043	0.250	0.196	0.281	1.000			
radio etc	0.016	0.370	0.367	0.603	0.281	1.000		
TV etc	0.198	0.499	0.412	0.581	0.288	0.474	1.000	
fridge etc	0.236	0.361	0.418	0.673	0.235	0.621	0.469	1.000

Table 10: Eigenvalues for correlation matrices in table 9

a) For June 2002 assets data

Component	Eigenvalues	Proportion explained	Cumulative proportion
1	3.523	0.440	0.440
2	1.022	0.128	0.568
3	0.881	0.110	0.678
4	0.782	0.098	0.776
5	0.700	0.087	0.864
6	0.467	0.058	0.922
7	0.370	0.046	0.968
8	0.255	0.032	1.000

b) For assets data on the second survey visit

Component	Eigenvalues	Proportion explained	Cumulative proportion
1	3.565	0.446	0.446
2	1.021	0.128	0.573
3	0.854	0.107	0.680
4	0.739	0.092	0.772
5	0.695	0.089	0.859
6	0.466	0.058	0.918
7	0.384	0.048	0.966
8	0.275	0.034	1.000

B Farm machinery index

Similarly, we also use polychoric PCA to calculate an index of farm machinery owned by the household as of June 2002, taking into account the following categories of machinery and tools:

1. sickle, axe, spade & chopper
2. plough (wooden or iron)
3. harrow, seed-drill, sprayer & duster, chaff-cutter
4. power tiller
5. tractor (excluding trolly)
6. thresher
7. pumps (electric)
8. pumps (other)

Table 11 shows the matrix of polychoric correlations between respective farm machinery categories, and table 12 shows the corresponding eigenvalues estimated using PCA. Table 12 shows that the first (polychoric) principal component score captures 37% of the total variation in these data.

Table 11: Polychoric correlations for farm machinery owned in June 2002

	sickle, axe, spade & chopper	plough (wooden or iron)	harrow, seed-drill sprayer & duster, chaff-cutter	power tiller	tractor (excluding trolly)	thresher	pumps (electric)	pumps (other)
sickle, axe, spade & chopper	1.000							
plough (wooden or iron)	0.281	1.000						
harrow, seed-drill, sprayer & duster, chaff-cutter	0.151	0.289	1.000					
power tiller	0.158	0.023	0.088	1.000				
tractor (excluding trolly)	0.248	0.094	0.130	0.626	1.000			
thresher	0.157	0.245	0.139	0.433	0.641	1.000		
pumps (electric)	0.218	0.235	0.200	0.375	0.474	0.408	1.000	
pumps (other)	0.214	0.171	0.114	0.212	0.444	0.414	0.061	1.000

Table 12: Eigenvalues for correlation matrix in table 11

Component	Eigenvalues	Proportion explained	Cumulative proportion
1	2.962	0.370	0.370
2	1.302	0.163	0.533
3	0.979	0.122	0.655
4	0.836	0.104	0.760
5	0.721	0.090	0.850
6	0.523	0.065	0.915
7	0.417	0.052	0.967
8	0.261	0.033	1.000

C Households with existing bank loans

In this section we present a separate set of results for the sample of households who already had agricultural bank loans as on June 2002. There are 998 such households in the sample. The respective definitions of treatment status (T_{assets}, T_{invest}) correspond to *additional* new loans being taken, and these definitions and the corresponding outcomes we examine are summarised in table 13. Table 14 presents summary statistics for various household characteristics for the overall sample and by treatment status.

Table 13: Outcomes and definition of treatment

Outcome	**Treatment definition**
Index of consumer durable assets owned by household as on date of second survey visit	T_{assets}=1 iff household obtained a new agricultural loan during July 2002 - July 2003
Binary indicator for investments in farm enterprise made by the household during January 2003 -June 2003	T_{invest}=1 iff households obtained a new agricultural loan during July 2002 -December 2002

Table 14: Descriptive statistics

	Full sample		T_{assets}=1		T_{assets}=0		T_{invest}=1		T_{invest}=0	
	Means (Standard deviations in parentheses)									
Assets index second survey visit	-0.002	(1.406)	0.761	(2.12)	-0.06	(1.319)	-	-	-	-
Land area owned	2.275	(2.845)	3.83	(5.294)	2.156	(2.529)	4.262	(5.752)	2.17	(2.567)
Value of land owned (Rs. '000)	336.927	(668.006)	711.592	(1446.855)	308.231	(557.325)	786.162	(1664.306)	313.233	(561.474)
Assets index June 2002	0.011	(1.382)	0.748	(2.077)	-0.045	(1.299)	0.686	(2.253)	-0.024	(1.313)
Farm assets index June 2002	-0.088	(0.739)	0.15	(0.768)	-0.107	(0.734)	0.209	(0.782)	-0.104	(0.734)
Household monthly per capita expenditure (Rs.)	494.631	(217.87)	545.268	(270.751)	490.752	(212.964)	571.917	(302.169)	490.554	(211.937)
Residential area owned	0.017	(0.037)	0.024	(0.038)	0.016	(0.037)	0.03	(0.044)	0.016	(0.037)
Years of education of household head	6.002	(5.781)	6.155	(5.801)	5.99	(5.782)	6.4	(5.852)	5.981	(5.780)
Age of household head	49.546	(13.217)	52.465	(14.227)	49.323	(13.118)	50.26	(13.821)	49.508	(13.191)
Household proportion of children	0.332	(0.221)	0.319	(0.188)	0.333	(0.224)	0.314	(0.207)	0.333	(0.222)
Household proportion of adult males	0.226	(0.151)	0.216	(0.133)	0.227	(0.152)	0.23	(0.13)	0.226	(0.152)
Household proportion of adult females	0.212	(0.184)	0.247	(0.195)	0.209	(0.183)	0.261	(0.210)	0.210	(0.182)
Household size	5.589	(2.771)	6.535	(3.541)	5.517	(2.691)	6.52	(3.694)	5.54	(2.707)
	Proportions									
Invested in farm enterprise during January 2003-June 2003	0.035		-		-		0.04		0.035	
Male household head	0.958		0.972		0.957		0.96		0.958	
Household with married male child	0.309		0.507		0.293		0.48		0.3	
Household caste group										
Other backward class	0.132		0.085		0.136		0.1		0.134	
Others	0.13		0.056		0.136		0.08		0.133	
Scheduled Castes	0.408		0.479		0.402		0.48		0.404	
Scheduled Tribes	0.33		0.38		0.326		0.34		0.329	
State										
Andhra Pradesh	0.112		0.169		0.108		0.22		0.107	
Karnataka	0.104		0.127		0.102		0.1		0.104	
Madhya Pradesh	0.309		0.31		0.309		0.38		0.305	
Maharashtra	0.404		0.282		0.413		0.18		0.416	
Tamil Nadu	0.071		0.113		0.068		0.12		0.069	
Number of observations	998		71		927		50		948	

Table 15 presents the regression results for probit models used to estimate the propensity score. There are two probit regressions corresponding to the two definitions of treatment status. Similar to the results in the main text, the signs of most variables are as expected, even though there are fewer statistically significant variables possibly owing to the smaller sample size. We use these estimated coefficients to predict the respective propensity scores for the two treatments. Figure 4 shows the distributions of these estimated propensity scores, and shows that there is substantial overlap in the distribution for treatment and control groups.

Table 15: Estimation of propensity score for sample of households with existing agricultural loans in June 2002

Variable	Dependent variable: T_{assets}		Dependent variable: T_{invest}	
(log) Land area owned	0.0209	(0.101)	0.101	(0.119)
(log) Value of land owned	0.297***	(0.0977)	0.281**	(0.110)
Assets index June 2002	0.0857	(0.0592)	0.0468	(0.0648)
Farm machinery index June 2002	-0.132	(0.111)	-0.0731	(0.123)
(log) Household monthly per capita expenditure	-0.145	(0.210)	-0.0282	(0.237)
Residential area owned	0.767	(1.530)	1.801	(1.465)
Years of education household head	-0.0173	(0.0133)	-0.0131	(0.0153)
Age of household head	-0.00927	(0.00744)	-0.0201**	(0.00894)
Male household head	0.319	(0.389)	0.110	(0.423)
Household has married male child	0.421**	(0.189)	0.339	(0.214)
Household proportion of children	0.124	(0.519)	0.402	(0.623)
Household proportion of adult males	-0.609	(0.562)	-0.0185	(0.640)
Household proportion of adult females	0.184	(0.542)	1.008	(0.655)
Household size	-0.00650	(0.0304)	0.0165	(0.0330)
Household caste group				
Scheduled Tribe	base		base	
Scheduled Caste	-0.157	(0.314)	0.0622	(0.337)
Other backward classes	0.110	(0.234)	0.0352	(0.263)
Others	0.0824	(0.248)	-0.0603	(0.281)
State				
Andhra Pradesh	base		base	
Karnataka	-0.297	(0.257)	-0.637**	(0.296)
Madhya Pradesh	-0.270	(0.221)	-0.382	(0.236)
Maharashtra	-0.542**	(0.214)	-0.868***	(0.244)
Tamil Nadu	-0.0621	(0.288)	-0.149	(0.313)
Constant	-3.821**	(1.703)	-4.187**	(1.925)
Number of observations	998		998	
Log likelihood	-228.2		-169.5	

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 17 summarises the degree of balance that results following different matching al-

Figure 4: Distribution of propensity scores by treatment status

(a) T_{assets}

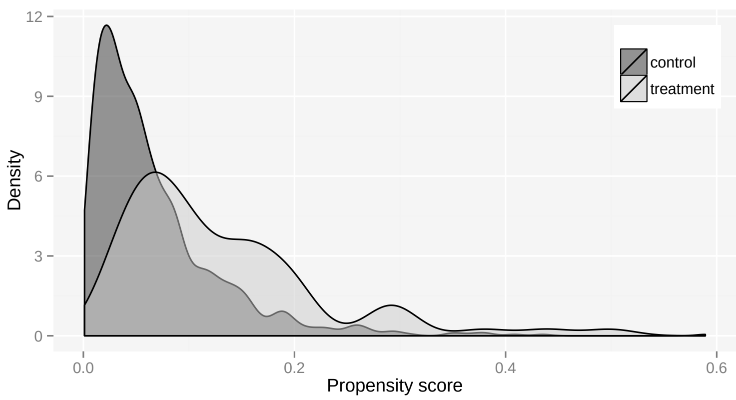

(b) T_{invest}

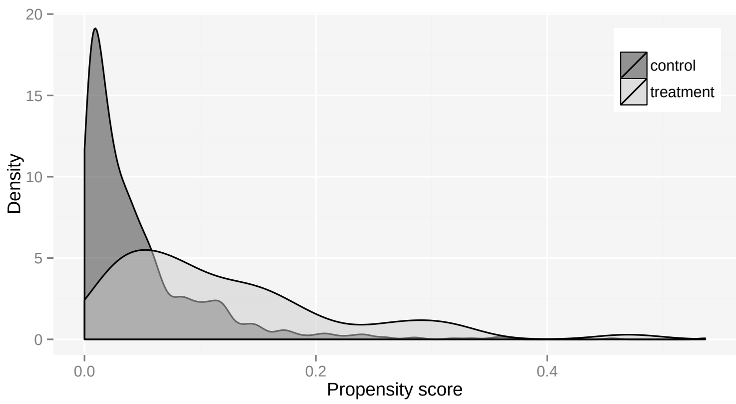

Table 16: Treatment effects for assets and farm investment

	θ	Standard error	p-value
Assets index	-0.0741	0.0664	0.2651
Farm investments	0.0070	0.0175	0.6894

gorithms. As in the main text, kernel matching performs worse than nearest neighbour matching in both cases. Panel (a) of the table shows that balance generally improves as the number of neighbours matched on increases. The exception to this is the household head's years of education, for which balance is actually worse across all matched samples relative to the unmatched samples. This is likely a result of the small sample size on which matching is being undertaken. Notwithstanding, we select nearest neighbour matching with n=10 as the best matching algorithm corresponding to the T_{assets} treatment.

Panel (b) of table 17 does not show as clear a pattern as panel (a). Matching on 3 nearest neighbours results in a reduction in imbalance of atleast 85% across all covariates, and thus we select this as the best matching algorithm corresponding to the T_{invest} treatment.

Table 16 shows the treatment effect estimates that result from the preceding choices of matching algorithm. Similar to the treatment effects for households without existing agricultural loans (in the main text), these effects are statistically (very) insignificant, while the sign of the treatment effect for the assets outcome is negative and that for farm investment is positive.

Table 17: Reduction in imbalance for key variables by matching algorithm

(a) T_{assets} (assets index)

	Nearest neighbour matching										Kernel
Variable	1	2	3	4	5	6	7	8	9	10	
(log) Land area owned	66.43	97.65	91.22	92.72	93.43	97.61	98.12	97.99	97.37	97.00	27.27
(log) Value of land owned	87.04	93.05	90.76	97.11	95.53	99.07	99.81	98.40	98.15	98.60	30.41
Assets index June 2002	62.41	60.41	69.96	67.80	68.54	67.75	68.64	73.22	75.75	79.56	25.69
(log) household monthly per capita expenditure	9.29	69.94	63.97	76.05	87.72	97.09	98.04	94.55	96.55	95.66	24.72
Years of education household head	-207.82	-82.26	-194.32	-171.86	-199.96	-200.11	-264.31	-185.28	-150.40	-62.12	-45.13

(b) T_{invest} (farm investment)

	Nearest neighbour matching										Kernel
Variable	1	2	3	4	5	6	7	8	9	10	
(log) Land area owned	98.84	89.05	87.21	87.16	87.61	89.88	91.28	89.65	89.72	92.74	19.76
(log) Value of land owned	94.13	85.32	99.50	99.60	99.38	99.78	97.92	96.30	96.25	92.19	18.05
Assets index June 2002	85.86	97.32	94.90	93.75	88.74	92.25	93.50	95.01	99.17	94.58	18.95
(log) household monthly per capita expenditure	10.91	72.04	98.49	88.56	82.14	74.58	67.14	63.64	65.43	65.75	3.02
Years of education household head	-198.38	4.09	89.12	70.16	84.12	54.53	55.31	73.50	57.37	55.65	35.01

Notes: Numbers shown are the percentage reduction in imbalance for respective matching algorithms. This is measured using quantile-quantile plots as the percentage reduction in mean deviation from the 45° line of perfect symmetry, for unmatched relative to matched samples.