The (In)efficiency of Share-Tenancy Revisited: The Role of Supervision

Hanan G. Jacoby* Ghazala Mansuri*

September 2004
(*preliminary; please do not cite*)

Abstract

Sharecropping, particularly its effect on agricultural efficiency, has long fascinated economists. With a couple of notable exceptions, past studies find small or insignificant productivity differentials between sharecropped and owner-cultivated land. This paper provides more conclusive evidence using a large, nationwide, micro-data set from rural Pakistan. Our estimates show that the average yield differential between share-tenanted and owner-cultivated plots is highly unlikely to exceed 8 percent. An analysis of tenant labor allocation corroborates this conclusion.

To understand *why* sharecropping does not lead to substantial productivity losses on average, we use unique data on monitoring frequency collected directly from tenants. We find that "unsupervised" tenants are significantly less productive than their "supervised" counterparts. We show that the coexistence of these two types of tenants is consistent with an agency model in which landlords have different costs of supervision. To assess the model, we investigate whether a landlord's decision about the form of incentive contract and degree of supervision is driven, in part, by variation in supervision costs.

*Development Research Group, The World Bank. Contact Information: Jacoby: e-mail: hjacoby@worldbank.org; Mansuri: e-mail: gmansuri@worldbank.org.

1 Introduction

Sharecropping, particularly its effect on agricultural efficiency, has fascinated economists since Adam Smith. Writers in the "Chicago" tradition (e.g., Johnson, 1950; Cheung, 1968; Reid, 1977) argued that the incentive problem inherent in share-tenancy is largely obviated by landlord supervision of tenant effort. But, with the advent of agency theory, in which sharecropping was a leading example (Stiglitz, 1974), the notion that effort could be effectively monitored began to be regarded with skepticism.

In a landmark study, Shaban (1987) vindicated the moral hazard paradigm by finding lower labor intensity and yields on sharecropped land as compared with owner-cultivated land. Sharecropping, at least in the six South Indian villages investigated by Shaban, did lead to substantial productivity losses. Recent evidence from a tenancy reform in West Bengal (Banerjee, et al., 2002) appears to reinforce this view, showing that a modest reallocation of property rights in favor of share-tenants had a dramatically positive productivity effect. Notwithstanding these two prominent studies from India, the idea that share-tenancy entails serious inefficiency remains controversial. Much, if not most, other evidence points to small or insignificant productivity differentials, though clearly some of these findings can be attributed to small sample sizes or to other methodological shortcomings (Otsuka, et al., 1992; Binswanger, et al., 1995).[1]

This paper aims to provide more conclusive evidence on the (in)efficiency of sharecropping using a large, nationwide, micro-data set from rural Pakistan, a country with an agriculture similar to that of India, but where tenancy is even more common. Our empirical work is divided into three parts. In the first part, presented in Section 3, we estimate the average yield differential between share-tenanted and owner-cultivated plots. Based on these estimates, we are able to state with a high degree of confidence that this differential does not exceed 8 percent and is most likely smaller. Thus, we can easily rule out yield differentials of, respectively, 16 percent (Shaban) and more than 50 percent (Banerjee, et al.) in the Pakistani context.[2] Our analysis of tenant labor allocation, simi-

[1] The well-known study of Laffont and Mattousi (1995), for example, is based on fewer than 200 plots from a single Tunisian village. Possibly as a consequence, none of the differences in productivity between share-tenants and owner/renters that they estimate appear to be statistically significant. Moreover, their point estimates indicate that, *at the sample mean of tenancy contract duration,* productivity on sharecropped land is actually *higher* than on owner-cultivated land!

[2] The impressive productivity gains on sharecropped land found by Banerjee et al. are all the more remarkable considering that the tenancy reform in West Bengal fell far short of providing share-tenants the same incentives as owner-cultivators. Were moral hazard in current production effort the only source of inefficiency, their estimates could thus be viewed as a *lower* bound on the effect of converting sharecropped

lar to that of Shaban (except, in our case, broken down by agricultural task), corroborates this conclusion.

The second part of the empirical analysis asks *why* sharecropping does not lead to substantial productivity losses in Pakistan. In Section 4, we reconsider the question of landlord supervision. First, we develop a model in which tenant effort is observable and enforceable, but at a cost to the landlord. In this model, otherwise identical tenants can receive different levels of supervision because supervision costs vary across landlords. Using unique information on monitoring frequency collected directly from tenants,[3] we then show that "unsupervised" tenants are significantly less productive than their "supervised" counterparts. There are simply not enough of these unsupervised tenants, however, to make sharecropping unproductive in the aggregate.

In the third part of the empirical work, we assess the assumptions underlying the supervision costs model by investigating the landlord's choice between providing incentives and supervision. Section 5 shows that, consistent with the model, this choice is driven by variation in the cost of supervising a tenant on a particular plot of land. Before turning to these results, we set out the context for our study in the next section along with the framework for the empirical analysis. Section 6 concludes the paper.

2 Analytical Framework

2.1 Data

Our empirical analysis draws mainly upon a new nationally representative rural household survey. The Pakistan Rural Household Survey (PRHS), which was completed in late 2001, collected data from about 2,800 households sampled across 17 districts and 150 villages. Roughly 60 percent of the households surveyed were farm households and a considerable fraction of these operated multiple plots. The survey was designed to provide detailed information at the plot level on land characteristics (soil type, irrigation, and so forth), land tenancy contracts, and production activities. These data were collected for the two major agricultural seasons, *kharif* (May-November) and *rabi* (November-May).

The agricultural module of the 2001 PRHS was modeled on an earlier survey (round

land into owner-cultivated land. However, as they point out, the tenancy reform may also have raised land productivity by improving investment incentives for share-tenants.

[3] Ai et al. (1997) is the only other study that we are aware of utilizing similar information. Based on plot-level data from a single Tunisian village, they find that crop yield is a *decreasing* function of the frequency of landlord visits to the tenant.

14 of the IFPRI panel) which also collected plot level data. This smaller survey, fielded in 1993, was carried out in four districts and 52 villages. All of the households surveyed in this earlier survey were purposively included in the PRHS. Because much of the plot-level agricultural production and tenancy data are comparable across the two surveys, part of our empirical work pools the 2001 and 1993 surveys to improve the precision of our estimates.

2.2 Context

Pakistan has a very unequal distribution of landownership. Consequently, the fraction of tenanted land is high (more than a third), and about two-thirds of tenanted land is under sharecropping. While various land and tenancy reforms have been attempted over the years, none has been particularly successful. There are, however, numerous tenancy laws on the books which protect tenants from expropriation and set lower bounds on inputs shares provided by landlords. These laws are rarely used to settle actual disputes, though, and their enforcement is questionable at best. Large landowners continue to wield tremendous political power in the country, while the bulk of share tenants belong to landless households with little education or opportunities for employment outside agriculture.

There are important regional variations in the form and prevalence of land tenancy in Pakistan. Sharecropping is the predominant form of tenancy in Sindh province.[4] Based on our data, almost two-thirds of all cultivated plots in Sindh are under share tenancy, while only 3% are under fixed rent tenancy. The distribution of landownership is also particularly skewed in Sindh. According to reports by surveyed tenants, the median landlord owns 28 acres, while nearly 80% of the share tenants are landless farmers. Landlords in this region often employ labor supervisors (*kamdars*) and also have significant coercive power over their tenants. In Punjab province, by contrast, 77 percent of sampled plots are cultivated by their owners, and of the remainder, over half are cultivated by fixed rent tenants. Landlords in Punjab are also much smaller, with a median holdings of only 7 acres, and are more likely to be resident in the same village as their tenants.

Tenancy contract terms are quite diverse in Pakistan, although nearly three-quarters of our sample share-tenants report a nominal output share of 50%.[5] Many tenants,

[4]Pakistan has four provinces, Sindh, Punjab, the North West Frontier and Baluchistan. The bulk of agricultural households fall in the first two provinces which together account for more than 80 percent of the total agricultural production of the country.

[5]Output shares vary considerably by province, though. In Sindh, 17 percent of sample tenants report a

however, also borrow from their landlords. These loans typically cover an agricultural season and the loan, plus any implicit interest, is deducted from output at the time the crop is harvested and divided. Thus the relative uniformity of nominal output shares may conceal significant variation in effective shares. Traditionally, the tenant has been solely responsible for all labor, but shared hiring of labor at harvest time has emerged more recently. Other major inputs (fertilizer, tractor hire, seeds, and pesticides/fungicides) are typically shared between the landlord and the tenant in accordance with the output share, although this also varies, as discussed below.

2.3 Moral hazard and yields

To clarify the empirical issues, we consider a simple tenancy model with a constant returns to scale technology. This avoids having to model the choice of plot size, which we set to unity. Gross output, or yield, is given by the production function $y = f(e, x) + \varepsilon$, which depends on unmonitorable tenant effort e and a purchased input (e.g., fertilizer) x, as well as a random shock ε. Net productivity is given by $\pi = y - px$, where p is the price of the purchased input normalized by the price of output. The tenant's disutility of total effort is given by the convex function $v(e)$. The share-contract specifies an output share β and possibly also a fixed component α, which may be negative. Suppose also that the cost of the purchased input is shared between landlord and tenant at the same rate as output.

A risk neutral tenant chooses e and x such that $\beta f_e = v\prime$ and $f_x = p$,[6] from which we may write $y = g(\beta, p) + \varepsilon$. Given that any tenancy contract must be incentive-compatible, moral hazard delivers the unambiguous prediction $\frac{de_s}{d\beta} > 0$, where (e_s, x_s) denotes the tenant's optimal choice. The marginal effect of share-tenancy on yields, however, is given by

$$\frac{dy}{d\beta} = f_e \frac{de_s}{d\beta} + f_x \frac{dx_s}{d\beta}$$

which, in general, is ambiguous.[7] If e and x are perfect substitutes, for example, pro-

25 percent share, whereas in Punjab the comparable figure is less than 5 percent. In Baluchistan, output shares of 25 and 33 percent are the most common, and in the NWFP a share of 2/3 is reported by 36 percent of tenants.

[6] Precisely the same first-order conditions would hold for a risk-averse tenant. However, since we later model share-tenancy as arising from financial constraints, tenant risk aversion is an unnecessary complication.

[7] The sufficient condition for a positive yield effect is that $f_{xe} > f_e f_{xx} / f_x$

viding the tenant better incentives does not actually increase yields at all (of course, in this case, the landlord could dispense with the tenant altogether and use only fertilizer). More generally, a test of the null of zero moral hazard using yield data only has power against alternatives in which moral hazard is present *and* tenant effort is an input without very close purchased substitutes.[8] Also, the extent to which a purchased input, such as fertilizer, is substitutable with tenant effort (i.e., $\frac{dx_s}{d\beta} = \frac{dx}{de}\frac{de_s}{d\beta}$) is indicated by just how much more of this input is used on tenanted plots than on owner cultivated plots.[9]

2.4 Empirical model and the selection problem

Our regression model for yields realized by cultivator c on plot i is

$$y_{ci} = \gamma s_{ci} + \omega\prime x_{ci} + \nu_c + \eta_{ci} \tag{1}$$

where s_{ci} is an indicator of whether the plot is sharecropped, $\beta \in (0,1)$, and x_{ci} is a vector of exogenous plot characteristics. Thus, γ estimates the yield differential between sharecropped plots on the one hand and owner-cultivated and rented plots on the other. One component of the error, ν_c, captures all unobserved factors common to a given cultivator that determine productivity; e.g., access to credit, farming knowledge, average land quality, and ownership of non-marketed assets more generally. Included in this term would also be the effect of input prices, p, to the extent that these are not observed. The error component η_{ci} contains plot-specific unobservables, such as soil fertility, that are not captured by x_{ci}.

In general, the decision to enter into a sharecropping contract will depend upon the cultivator's unobserved characteristics (i.e., $E[\nu_c \mid s_{ci} = 1] \neq 0$), which leads to a selection (or endogeneity) bias in 1. All of the major theories of share-tenancy that have thus far been proposed in the literature, however, imply that sharecroppers have *lower* unobserved productivity than owner-cultivators or fixed renters, which imparts a particular direction to the selection bias. For example, models with either ex-post or ex-ante finan-

[8] For this reason, it might seem more attractive to compare net rather than gross productivity. In particular, $\frac{d\pi_s}{d\beta} = f_e\frac{de_s}{d\beta}$ is unambiguously positive. The problem, in practice, is that there are usually several purchased inputs, each measured with considerable noise, so that net productivity tends to be less precisely measured than gross productivity. Moreover, estimates based on yield data are more comparable to those from previous studies.

[9] Note that we do not need to assume that x is *contractible*. Although fertilizer, for example, can be readily purchased, it may be difficult to enforce a contract in which a tenant uses a quantity that is incompatible with his first-order conditions. The key feature that distinguishes e and x is that the cost of effort is not observable to the landlord and hence cannot be shared with his tenant.

cial constraints (Shetty, 1998; Basu, 1992; Mookherjee, 1997; Laffont and Mattoussi, 1995) imply that wealthier cultivators would be less likely to take land on share, but wealthier cultivators, if anything, will tend to have higher productivity on a given piece of land. An analogous argument would apply if share-tenancy is motivated by risk aversion (e.g., Stiglitz, 1974); i.e., share tenants are more risk averse, but greater risk aversion reduces productivity by, for instance, encouraging the use of safer but less effective production techniques. In the double-sided moral hazard model of Eswaran and Kotwal (1985), both tenant and landlord supply a noncontractible input, which in the landlord's case may be farming know-how. Here, again, sharecroppers tend to come from the lower tail of the productivity distribution. The same holds when farming ability is private information to the tenant, as in Halligan's (1978) screening model – less able cultivators become sharecroppers rather than fixed renters.

If sharecroppers are indeed, on average, less productive farmers, then we expect that estimates of γ that fail to account for this selectivity will overstate the disincentive effects of share-tenancy. In other words, the OLS estimate of γ will tend to be negative even if the true γ is zero. A corollary to this observation is that if the OLS estimate of γ is in fact zero, then the selection problem is not likely to be empirically important. This is because, under either the null of no moral hazard or under the alternative of moral hazard, the true γ cannot plausibly be positive.

Our strategy for correcting the selectivity bias is essentially the same as that of Shaban (1987) and Bell (1977). We use household fixed effects to purge ν_c, a procedure that requires a sufficient number of *owner-cum-sharecropper,* households that cultivate at least one sharecropped plot and one plot of their own (or one plot on fixed rent). Even after taking out household fixed effects, there is no guarantee that $E[\eta_{ci} \mid s_{ci} = 1] = 0$. Highly fertile plots may be more (or less) likely to be sharecropped out than those cultivated by their owners or given on fixed rent. Unfortunately, finding good instruments for within household variation in s_{ci} is difficult, so we will have to rely on indirect methods to assess this potential endogeneity problem.

3 Results: Is Sharecropped Land Less Productive?

3.1 Yields

To obtain precise estimates of productivity differentials, we focus on yield from the five major crops: wheat, rice, cotton, sugarcane, and maize. We thus exclude the value of outputs from fodder and a number of minor crops, which are difficult to measure accurately. In the 1993 IFPRI sample, major crops account for 66% of cultivated area (71% for sharecropped land), whereas in the nationally representative 2001 PRHS, 80% of cultivated area is devoted to major crops (83% for sharecropped land). Yield is defined as the value of output from these five crops (evaluated at median prices for that year) divided by the area of the plot planted to major crops over *kharif* and *rabi* seasons. Plots growing none of the major crops are dropped from the sample, as are cases of yield above the 99th percentile (the presence of which inflates the standard errors). Table 1 breaks down the number of plots available, by year. For the household fixed effects analysis of yields, the sample consists of the 1,718 plots belonging to households with multiple plots.[10] Of these, 403 belong to owner-cum-sharecropper households; these plots directly identify the share-tenancy effect.[11] The remaining 1,315 plots, belonging to pure sharecroppers or pure owner-cultivator households, contribute to greater precision when we control for plot characteristics.

Table 2 presents household fixed effects results for yields. The dependent variable is scaled so that the coefficients can be interpreted as percentage deviations relative to owner-cultivators/fixed renters. Unconditional on plot attributes, we see that yields are about 3% lower on share-tenanted plots, a difference which is not remotely significant. Controlling for plot characteristics (value, area, irrigation, soil, etc.; see Appendix) hardly changes this result.[12] Evidently, these important *observed* characteristics are not highly correlated with the tenancy status of the plot. In light of this, it would be surprising if the presence of *unobserved* plot characteristics (e.g., soil fertility) would seriously bias our results. Adding crop composition variables (i.e., the fraction of area cultivated with major crops devoted to a given crop) also has little effect, except to slightly improve the precision of the sharecropping coefficient estimate.[13] While, in theory, crop choice may be specified in the tenancy contract, and hence endogenous, in Pakistan, share-tenants generally have autonomy over crop choice and grow basically the same mix of crops as

[10] Although most households from the 1993 round are resurveyed in 2001, we do not make use of the panel element in the estimation. Grouping together the plots of the same household in different survey rounds results in less precise estimates as compared to treating them as separate households in each round.

[11] Owner-cum-sharecropper households are over-represented in the IFPRI survey by geographical accident; these households tend to be concentrated in central Punjab.

[12] Shaban (1987), by contrast, finds that the yield differential falls by more than one third, from 25% to 16%, once he controls for irrigation and other plot characteristics.

[13] We allow the effect of crop composition to vary by survey year to capture any change in relative prices among the major crops.

owner-cultivators.[14]

Failure to reject the null hypothesis of equal gross productivity across tenure types is informative only to the extent that moderate productivity differences would be detectable in our data. Andrews (1989) has devised a statistic, the inverse power function, that allows one to quantify the set of alternatives against which a given test has power. Table 2 reports, next to each estimate, two points along the inverse power function. The first point is the percentage yield differential against which our test has low power. Based on the household fixed effects estimates, we would be equally likely as not to reject the null if the true yield differential were around 7%; this figure demarcates the region of low power. On the other hand, if the true yield differential were 13%, we would be 95% certain of rejecting the null. Thus, our test has high power against yield differentials exceeding 13%.

Although 13% is a respectable number, we can do even better by imposing restrictions in the estimation. To this end, we report analogous estimates using village, rather than household, fixed effects. These estimates are not robust to the selection problem outlined above, but they do control for input price variation across villages and may be more efficient than household fixed effects estimates.[15] As it happens, all of the village fixed effects estimates are within a single standard error of their household fixed effects counterparts. This finding suggests that selection into share-contracts on cultivator-specific unobservables is not particularly strong.[16] Moreover, the standard error on the sharecropping dummy coefficient falls by almost 40% as we move to the village fixed effects estimator. As a consequence, the inverse power interval (which is proportional to this standard error) shifts to the left. We can now be 95% certain that the yield differential is no greater than 8%, whereas our test has low power against only fairly trivial yield differentials, below 4%.

[14] For all five of the major crops in our yield measure, there is no significant difference in the proportion of area cultivated between share tenanted and other plots, once we control for *tehsil* (there are 23 *tehsils* in our sample).

[15] Note that this estimator uses all 2,807 plots in the sample, including those belonging to single-plot households. In the estimation, we also allow for household random effects to deal with the correlation across plots within multi-plot households.

[16] Again using the inverse power function calculation, we can be 95% certain that the household fixed effect estimate of the yield differential is within about 11 percentage points of the village fixed effect (household random effect) estimate. In other words, we would be very likely to detect moderate selection bias if it existed in our data.

3.2 Nitrogen

As discussed above, to test for moral hazard using yield data one needs to maintain the assumption that tenant effort has no close purchased substitutes. However, if such substitutes exist, then they must respond to marginal incentives in the opposite direction as tenant effort. In other words, these inputs must be used *more* intensively on sharecropped plots than on owner-cultivated plots. A relatively clean test-case is provided by chemical fertilizer, a major purchased input in Pakistan. We focus on nitrogen, which is more widely used than phosphate. Nitrogen has the advantage of being a very homogeneous input compared to, say, seed and even tractor services. Moreover, the costs of fertilizer are generally, but not universally (see below), shared between landlord and tenant at the same rate as output. Since share-tenancy, consequently, does not distort the fertilizer margin, differences in fertilizer intensity across tenure types are driven solely by complementarity or substitutability between fertilizer and tenant effort.

The bottom panel of Table 2 presents estimates of equation 1 with yields replaced by the quantity of nitrogen in kg per acre. We expand the sample to include all cultivated plots, not just those growing major crops (see Table 1). In about one quarter of the share-tenancies, fertilizer costs are either the sole responsibility of the tenant or the tenant bears a larger share of the cost than he receives in output. In these cases, the share-contract may indeed distort the fertilizer margin. But when we include an indicator variable to capture this fact in the fertilizer regressions, it never attains statistical significance. To maximize precision, we omit this variable from the results reported below.

The point estimate of the sharecropping dummy coefficient for nitrogen use in Table 2 is indistinguishable from zero in every specification. Inverse power function thresholds are, in this case, positive numbers because we are only interested in alternatives wherein sharecroppers use nitrogen *more* intensively than owner-cultivators. Based on the household fixed effects standard errors, we can be confident that the nitrogen differential is no greater than 16%. Using the village fixed effects estimates instead, this threshold drops to just 10%. In short, there is little chance that the lack of a yield effect is due to the substitutability of nitrogen for tenant effort. Of course, this evidence is not decisive in itself, as there may be other purchased inputs that substitute for effort.

3.3 Labor

A more direct test for moral hazard involves comparing the cultivator's family labor input on sharecropped versus owned plots. Indeed, perhaps the most compelling piece of evidence from Shaban's (1987) study is that owner-cum-sharecroppers allocate substantially less family labor to their tenanted plots than they do to their owned plots. Naturally, the question arises as to how well reported labor hours on a plot correspond to cultivator 'effort' when the quality of labor is variable, but the fact remains that Shaban's finding has been widely viewed as indicating moral hazard in effort.

The 1993 IFPRI survey (but not the 2001 PRHS) provides information on plot-level labor inputs, both hired and family. These data are disaggregated by type of worker (adult male and female, and male and female child), but since the vast majority of farm labor is supplied by adult men in this sample, we combine the hours of all worker types. Unlike the ICRISAT data used by Shaban, the IFPRI survey collects farm labor data by major task (plowing/irrigation, sowing, weeding, harvesting, and threshing). To analyze labor use, we modify our previous procedure in one respect by taking logs of labor hours instead of using levels, as the former gives us considerably lower standard errors.[17]

Table 3 shows results with alternative definitions of labor hours per acre, and with each regression including the full set of controls (plot characteristics and crop composition). Based on the household fixed effects estimates, we find that family labor, aggregated across all tasks, is about 6% lower on sharecropped plots than on owner-cultivated plots. Once again, this difference is not significant, but the standard errors are noticeably larger than in Table 2. The village fixed effects estimate is a bit more precise, but in this case a Hausman test rejects equality between the (household) fixed and random effects coefficients (p-value = 0.006). Conservatively, then, our test has high power against family labor use differentials of around 21%. Shaban (1987) estimates a family labor use differential of exactly 21% for males (47% for females). Based on our findings, we can all but rule out such high differentials in the case of Pakistan.

The second specification in Table 3 combines family and hired labor. Shaban's analysis found that hired labor use was modestly lower on sharecropped plots than on owned plots cultivated by the same household. It is also conceivable, as discussed earlier, that hired labor substitutes for tenant effort and is thus used more intensively on share-tenanted land. If anything, our evidence suggests the latter scenario, as the percentage differential in total (family + hired) labor use is closer to zero than that of family labor alone, although the relatively high standard errors preclude firm conclusions.

Certain tasks are inherently easier to monitor than others. For example, in contrast to other activities, much of hired harvest labor in Pakistan is paid on a piece-rate (see Jacoby

[17] Standard errors of percentage changes are calculated using the approximation formula given by van Garderen and Shah (2002).

and Mansuri, 2003), which suggests that monitoring and enforcement are feasible for this task. Probably for this very reason, the cash costs incurred for harvesting/threshing (principally hired labor) are more frequently shared between landlord and tenant than in the case of land preparation (plowing, sowing, and weeding). Combining time spent on land preparation and on harvesting may obscure any moral hazard effects present in the former but not in the latter.

To investigate this issue, we re-run the labor-use regressions excluding hours spent harvesting and threshing. These two activities account for almost half of all family labor hours devoted to the average plot. Surprisingly, the family labor-use differential between sharecropped and owner-cultivated plots does not increase when we focus solely on land preparation tasks. To the contrary, the differential based on the household fixed effects estimates is actually closer to zero, although the corresponding results for total labor are practically identical.

Overall, then, the findings do not favor the implications of moral hazard in tenant effort. What paltry evidence there is of lower family labor intensity on sharecropped land is belied by the failure of this effect to strengthen when we concentrate on tasks that are likely to be particularly susceptible to moral hazard.

4 Does Landlord Supervision Matter?

The results of the previous section rule out a sizeable yield shortfall on share-tenanted land *vis a vis* owner-cultivated/rented land. Faced with similar, albeit less conclusive, evidence, Otsuka, et al. (1992) surmise that supervision (and enforcement) of share-tenant effort is generally effective. Sharecropping, they argue, is adopted only when monitoring costs are low enough to make it worthwhile relative to fixed rental. They go on to suggest that "significant inefficiency of share-tenancy is expected to arise only when the scope of contract choice is institutionally restricted" (p. 2007). Where fixed rent tenancy is legally discouraged, landlords without a comparative advantage in supervision are forced to enter into sharecropping contracts. Thus, rather than evidence of the *general* inefficiency of sharecropping, these authors view Shaban's (1987) findings as a peculiarity arising from India's legal environment. Specifically, the argument is that land-to-the-tiller legislation effectively penalized landlords who entered into longer-term tenancy contracts. Sharecroppers, however, were easier to disguise as permanent farm laborers than were fixed rent tenants. As a consequence, many landlords switched over to share-contracts, ill-equipped with the requisite supervision technology.

As attractive as this argument may sound, the proposition that sharecropper inefficiency declines with the degree of supervision has yet to be subjected to formal empirical testing. In this section, we provide such a test, but, before doing so, we need to understand why two otherwise identical tenants might receive different levels of supervision. We begin, then, by setting out a model of landlord supervision that forms the basis for our empirical work in the remainder of this section. The model has implications for the relationship between yields and the level of supervision as well as for how supervision costs determine the form of the tenancy contract.

4.1 A tenancy model with supervision

Returning to the set-up of subsection 2.3, consider first the extreme case in which the landlord can costlessly monitor and enforce, and thus effectively choose, tenant effort. Restricting ourselves to linear contracts, the landord must select a fixed payment α, output share β, and tenant effort e to

$$\begin{aligned} &Max \quad (1-\beta)\pi(e,x)+\alpha \qquad s.t. \\ &\beta\pi(e,x)-v(e)-\alpha\geq 0 \end{aligned} \tag{2}$$

where the second line is the tenant's participation constraint with outside utility normalized to zero.[18] The solution to the landlord's problem sets $f_e = v_e$ and $f_x = p$, the latter condition provided by the tenant's incentive compatibility constraint. Given this first-best optimum, say (e^*, x^*), the landlord can choose any combination of α and β such that the tenant's participation constraint binds. Obviously, with full enforcement of effort, the choice of these contract terms has no effect on either the tenant's effort level or on the landlord's payoff.

Next, we introduce a financial constraint along the lines of Laffont and Matoussi (1995). In particular, the tenant cannot be made to pay more up front than he has in net liquid wealth w (which may be negative, especially at the beginning of a season). Up front payments include α and the tenant's share of the input costs βpx, so that $\alpha + \beta px \leq w$. Wealthier tenants, for whom this constraint does not bind, can be offered a fixed rent contract with $\beta = 1$, whereas low wealth tenants who cannot afford upfront payments end

[18] The fixed payment α may be negative, in which case the contract has a credit element to the extent that such payments are made up front. This 'loan' is repaid at harvest time through the output share β. Formally, the model abstracts from the timing of payments.

up with share contracts. In either case, tenant effort remains first-best.[19]

At the other extreme, tenant effort is prohibitively costly to monitor, delivering the traditional moral hazard view of sharecropping discussed in subsection 2.3. In this case too, the binding financial constraint implies a direct mapping from tenant wealth into different contractual forms, as illustrated in Figure 1. At high levels of tenant wealth ($w \geq w_1$), the landlord offers a fixed rent contract, which yields first-best effort e^*. Low wealth tenants, meanwhile, are offered share (S) contracts, which induce less than first-best effort, $e_s < e^*$. As tenant wealth declines, and with it the amount that the landlord can extract upfront, the output share β must also fall to keep the tenant on his participation constraint.[20] Beyond a point, however, the landlord may have to let the tenant keep some rents to induce sufficient effort. In Figure 1, w_0 is the wealth threshold at which the tenant's participation constraint ceases to bind.[21] Below this point, the landlord's payoff declines one-for-one with tenant wealth, as the landlord transfers rents to the tenant.

Now consider an intermediate case between the extremes of costless and prohibitively costly monitoring; the landlord can choose tenant effort, but at a cost. Let $c(e)$ represent the cost of implementing a given level of effort, where $c_e > 0$ and $c_{ee} \geq 0$.[22] The landlord's problem is now

$$\begin{aligned} &Max \quad (1-\beta)\pi(e,x) + \alpha - c(e) \qquad s.t. \\ &\beta\pi(e,x) - v(e) - \alpha \geq 0 \\ &\alpha + \beta px - w \leq 0. \end{aligned} \tag{3}$$

The interesting situations involve share-contracts, because if the tenant is wealthy enough to afford a fixed rent contract monitoring is redundant.[23] Under a monitored share (MS)

[19] Allowing the landlord to set the tenant's share of the input cost independently from the output share has no effect on this or on subsequent analyses. When the financial constraint is binding, the input share drops out of both the landlord's objective function and the tenant's participation constraint; only the total amount the tenant pays up front is relevant to the contracting parties.

[20] Although many empirical studies find that output shares are quite inflexible, as Otsuka et al. (1992) note, the tenant's 'effective' share may depend on the implicit interest charged by landlords in credit advances that are deducted at the time output is shared. In our data, too, landlord loans to tenants are pervasive and the bulk of these are settled at harvest in the form of output retained by the landlord.

[21] We ignore the possibility that landlord's may switch to pure wage contracts ($\beta = 0$) for sufficiently low wealth tenants.

[22] More generally, we could assume that there are two types of effort, e_1 and e_2, the first of which is prohibitively costly to monitor and the second of which can be monitored at a cost. As long as e_1 and e_2 are perfect substitutes in production as well as in the disutility of effort function, all the results of our model go through.

[23] Demougin and Fluet (2001) model supervision in a limited liability setting, which does not allow for

contract, which solves problem 3, it is straightforward to show that the landlord chooses tenant effort according to $f_e = v_e + c_e$. The optimal effort level, e_m, must be less than first-best, since the cost to the landlord of inducing tenant effort is always higher than the tenant's marginal disutility of effort. Less obvious is how e_m differs from e_s.

Figure 1 shows how the introduction of landlord supervision affects contractual choice. In the MS contract, the tenant's participation constraint is binding at all wealth levels and the landlord's return is constant in tenant wealth at $r_m = \pi(e_m, x_m) - v(e_m) - c(e_m)$. We now introduce a cost shift parameter θ such that $dc(e;\theta)/d\theta > 0$, and consider two landlords, one with high supervision costs θ_H and one with low costs θ_L, where $\theta_L < \theta_H$ and thus, correspondingly, $r_m^L > r_m^H$. For landlord L, it is optimal to offer the MS contract to any tenant with wealth below w_L and to offer the S contract to any tenant for whom $w_L \leq w < w_1$. The same applies to landlord H at wealth level w_H, except that in his case the tenant earns information rents in the S contract.

We are now ready to address our main question: Does landlord supervision increase the productivity of a share-tenancy, *holding tenant wealth constant*? The way to think about this in the case of, say, landlord L is to fix tenant wealth at some w that is just epsilon above w_L. The landlord would thus choose contract S with a return of $r_s^L = \pi(e_s^L, x_s^L) - v(e_s^L)$, recalling that the tenant's PC binds here. Now lower θ (raise r_m) by some tiny amount delta so that the landlord is just willing to switch the tenant from the S to the MS contract. The question is whether tenant effort is higher after the switch. Since MS is now chosen, we have that $r_m^L > r_s^L \Longrightarrow [\pi(e_m^L, x_m^L) - v(e_m^L)] - [\pi(e_s^L, x_s^L) - v(e_s^L)] > c(e_m) > 0$. Further, since the function $\pi - v$ attains a maximum at (e^*, x^*) and both e_m^L and e_s^L are less than e^*, $\pi - v$ must be increasing over this range of effort. It follows, therefore, that $e_m^L > e_s^L$. The landlord must get more effort out of the tenant in the MS contract in order to justify his added supervision cost.

Interestingly, this reasoning does not apply to the case of landlord H. Performing the same thought experiment around w_H leads to a very different conclusion. This is because here the alternative to an MS contract is an S contract that leaves the tenant with rents. For this contract, the landlord's return is $r_s^H = (1-\beta)f(e_s^H, x_s^H) - px_s^H + w_H$. Could tenant effort remain unchanged (say) after the switch from the S to the MS contract? That is, would the assumption that $e_m^H = e_s^H$ lead to a contradiction? Note that the condition $r_m^H > r_s^H$ implies, after rearranging terms, that $\beta f - v - w > c$, which is not *necessarily* a contradiction. If information rents are high enough relative to supervision

fixed rent contracts. As a consequence, principals always monitor their agents to some extent.

costs, tenant effort under MS may be equal to (and, possibly, even lower than) effort under S. Intuitively, the tenant can compensate the landlord for bearing the monitoring cost in this case by relinquishing his information rents rather than working more diligently.

We can summarize these arguments as follows

Proposition 1 *If tenant effort is an input without very close purchased substitutes (see footnote 7), then yield will be higher in a share-contract with monitoring than in one without monitoring, provided the tenant does not earn rents in the latter contract. Otherwise the yield effect is ambiguous.*[24]

Our next step, then, is to make the distinction between a monitored and unmonitored share-contract operational.

4.2 Quantifying landlord supervision

In the 2001 PRHS, each share-tenant was asked "during [*kharif/rabi* season] how many times did the landlord meet with you to discuss or supervise your activities on this plot." In case the landlord employed labor overseers, or *kamdars*, the same question was asked about meetings between the tenant and these individuals as well. Exactly analogous questions were also asked of landlords about each of their share-tenants. None of these supervision questions were posed in the 1993 IFPRI survey, hence this round of data is excluded from the analysis of this section.

Very few share-tenants in our data (less than 4%) report never having had supervisory meetings with their landlord or with a *kamdar* during the year. On half of all sharecropped plots, the tenant reports having had more than 30 meetings per year with his landlord, and, on half of these plots, tenants claim to have had at least 90 meetings. To be sure, many of these conversations may have occurred during non-crucial periods or were not otherwise intended to elicit or enforce effort on the part of the tenant. Nevertheless, these numbers belie the notion that landlords are aloof, let alone absent, from their tenants' cultivation activities in Pakistan (Nabi, 1986, provides similar evidence in a smaller scale survey).

We certainly do not want to treat supervision intensity as linear in the number of meetings, since there must be diminishing returns beyond a point, and possibly increasing returns at very low numbers of meetings as well. The simplest empirical approach, and the one we adopt here, is to assume a threshold number of annual meetings above which

[24] Note that yield is always higher in a fixed rent contract than in a share-contract.

a tenant can be considered "supervised". But, what should this threshold be? This is a question on which we prefer to let the data speak.

To this end, we estimate a version of Hansen's (1999) threshold regression model for panel data. Let m_{ci} be the number of meetings that cultivator c on plot i had with his landlord (defined only for share-tenanted plots). Our modified yield regression is then

$$y_{ci} = \gamma s_{ci} + \delta s_{ci} I(m_{ci} > k) + \omega\prime x_{ci} + \nu_c + \eta_{ci} \tag{4}$$

where $I(\cdot)$ is the indicator function and k is the threshold, which is treated as a parameter to be estimated. For ease of interpretation, we demean the supervision indicator using $E\left[I(m_{ci} > k)|s_{ci} = 1\right]$, so that γ continues to estimate the mean difference in yields between sharecropped and owner-cultivated/rented plots.

Like the choice of share-tenancy itself, supervision may be endogenous to the extent that cultivators differ in farming ability or in other unobserved productive attributes that landlords use as the basis for their supervision decisions. If, for example, low productivity cultivators are monitored more intensively, then we would find a spurious negative relationship between yields and supervision. Estimation of equation 4 with household fixed effects can correct for this problem provided that the unobservable component of productivity is constant across tenanted and owned plots. However, it is possible that particular attributes a cultivator exhibits on his sharecropped plot(s), but not on his owned plot(s), induce a supervision response by his landlord. In other words, there may be an additional error component of the form $s_{ci}\mu_c$. One can think of this as a "bad tenant" effect versus the above "bad cultivator" effect. Fortunately, our data allow us to obtain an estimate of the supervision parameter, δ, that is robust even to this problem. Using a subsample of tenant households with multiple sharecropped plots ($s_{ci} = 1\ \forall i$), we can estimate by household fixed effects a regression of the form

$$y_{ci} = \delta I(m_{ci} > k) + \theta\prime x_{ci} + \mu_c + \nu_c + \eta_{ci}, \tag{5}$$

which purges both μ_c and ν_c (γ is absorbed in the constant term).

4.3 Results: Supervision and yields

Our analysis of supervision and yields is based on a sample of 1256 plots cultivated by multi-plot households in the 2001 PRHS (see Table 1). Replicating the third specification in Table 2 on this smaller sample, we obtain a yield differential of -4.2% (4.8), which is very similar to, but less precise than, our earlier result. To estimate the monitoring threshold k, we search over values of m_{ci} within a reasonable range and find the $\widehat{k}$ that minimizes the sum of squared residuals ($\widehat{\eta}_{ci}$) from equation 4 (see Hansen, 1999, for details).[25] Although conventional standard errors on the coefficients in equation 4, which treat $\widehat{k}$ as the true value of k, are asymptotically valid, the test of the null hypothesis $\delta = 0$ is non-standard.[26] We thus implement the bootstrap F-test proposed by Hansen (1999) for this purpose.

Household fixed effects regressions, including plot characteristics and crop composition, are reported in Table 4. For baseline specification (1), the estimation algorithm produces an optimal threshold value of 10 meetings. In other words, the definition of supervision that best fits the data is one in which the tenant meets his landlord at least 11 times per year, or about once each month. Notice that the average yield differential between sharecropped plots and owner-cultivated/rented plots remains about -4% after including this supervision variable. Supervised tenants, however, achieve 28% higher yields than unsupervised ones, and this difference is significant (albeit only just so, when we use the more conservative bootstrap F-test; p-value=0.051). According to our data-derived definition, about two-thirds of the 351 share-tenanted plots in this sample receive supervision from their landlords and/or the landlord's *kamdar*.[27] Viewed in comparison to owner-cultivated or fixed rent plots, plots cultivated by supervised tenants realize 3.0% (5.6) higher yields, a trivial difference. By contrast, land cultivated by unsupervised tenants is 17.8% (7.2) *less* productive than owner/renter cultivated land.[28] This latter figure is very close to the -16% yield differential relative to owner-cultivated land found by Shaban (1987) for all share-tenanted plots in his Indian sample.

To test the robustness of this remarkable finding, we fix the value of $\widehat{k}$ at 10 and add alternative sets of extra controls in Table 4. In specification (2), we are concerned with the possibility that our supervision variable is picking up characteristics of the tenancy that may have independent effects on yields. For example, one could argue that newer tenants, whose abilities are less familiar to the landlord, are more heavily supervised and are also less productive. However, a dummy variable indicating that the share-tenancy

[25] We restrict the search for k between the 10th and 50th percentiles of m_{ci} among the 351 sharecropped plots in this sample. The 50th percentile is 21 annual meetings, which is already fairly intensive supervision.

[26] The problem with the standard t-test is that the threshold parameter k is not identified under the null hypothesis.

[27] Note that this subsample of tenants is somewhat unrepresentative as it excludes sharecroppers of single plots, who are more likely to be from Sindh province and have large landlords. Supervision is considerably higher (75%) in the full sample of share-tenants.

[28] To obtain these estimates, we rescale the dependent variable by the average yield of owner-cultivated/rented plots and redefine the contract dummy variables.

has lasted no more than 3 years does not attract a significant coefficient. The number of landlord-tenant meetings could also reflect the social relationship between the two, which may have independent productivity effects as well. Again, this does not appear to the case, as a dummy for whether the landlord and tenant are related (including membership in the same caste/clan) is insignificant. The inclusion of these two variables also has no appreciable effect on our estimate of δ. Supervision may reflect landlord characteristics as well. It is conceivable, for example, that wealthier landlords supervise more and are also able to provide more or better quality inputs to their tenants. In specification (3), we control for the land, tractor, and tubewell ownership of the landlord, which again has a negligible effect on the supervision coefficient. Specification (4) includes both sets of extra control variables together, also with no discernible impact on $\widehat{\delta}$.

In Table 5, we replicate this sequence of regressions using the sub-sample of 113 households that cultivate at least two plots on share contracts (264 plots in all). Thus, we are no longer comparing yields on sharecropped plots, supervised or not, with yields on owned or rented plots cultivated by the same households. Rather, we are comparing yields on supervised sharecropped plots to those on unsupervised ones cultivated by the same households. This allows us to control for unobservables that are common to a household on its sharecropped land, but do not necessarily affect productivity on its owned land (e.g., a "bad tenant" effect). A potential downside of this approach is that it requires variation in supervision levels across plots within the same household and, often, tenants who sharecrop multiple plots do so from the same landlord. Whether we have enough variation in our sample to identify a supervision effect is, of course, an empirical question.

Using the threshold estimation algorithm on this new sample, we obtain specification (1) of Table 5 with $\widehat{k} = 9$, practically the same value as in previous case. But the supervision effect is now extremely large – landlord supervision raises yields by about 73% (versus our earlier 28%)– and very significant, with the bootstrap F-test p-value coming in at 0.006. As before, none of the extra control variables in specifications (2)-(4) put much of a dent in the estimate of δ.

Figure 2 plots the bivariate regression between yields and the supervision variable, in deviations from the respective household means. The supervision parameter δ is identified off of the 28 plots (11% of the sample) for which $I(m_{ci} > k)$ deviates from its household mean. Despite this small number of observations, the regression line (slope = 0.74) does not appear to be driven by extreme outliers. While one may question the size of the estimated yield difference between supervised and unsupervised plots on this sample, it is hard to argue that no statistical relationship exists between sharecropper yields and landlord supervision.

5 Supervision Costs and the Form of Tenancy Contract

In our model of subsection 4.1, differences in landlord supervision, and, ultimately, in yields, are driven by variation in the costs of monitoring tenant effort. In this section, we develop and implement a direct test of this assumption. The test is based on an analysis of the *landlord's* choice among alternative tenancy contracts. For this we use data collected from landlords rather than tenants.

5.1 Contract choice and matching

To understand the effect of supervision costs on contractual choice, return to Figure 1. Here we see that an increase in θ (fall in r_m) lowers the wealth threshold below which the landlord prefers to monitor, but has no effect on the wealth threshold above which he prefers to offer a fixed rent contract (w_1). This is only part of the story, however, because it does not account for how landlords and tenants are matched in equilibrium. In particular, landlords with lower supervision costs would tend to pair up with lower wealth tenants. To see why, suppose that we have two tenants with wealth levels w_L and w_H and two landlords with supervision costs θ_L and θ_H. From the landlords' perspective, the pairing $\{(w_L\ ,\theta_L),(w_H\ ,\theta_H)\}$ is preferred to $\{(w_L\ ,\theta_H),(w_H\ ,\theta_L)\}$ because for the landlord with low supervision costs the relatively high payoff from the MS contract is invariant to any decline in his tenant's wealth. Therefore, this landlord is indifferent among the two tenants, whereas the landlord with high supervision costs clearly prefers the high wealth tenant, to whom he can offer the S contract.

Based on this intuition, we can sketch out a simple landlord search model. Prior to contracting, potential tenants randomly drawn from the wealth distribution arrive sequentially before landlords who can either accept or reject them. If rejected, the tenant reenters the candidate pool. Since searching for tenants is costly, landlords will adopt a "reservation wealth" strategy; that is, hire the tenant if his (the tenant's) wealth exceeds a given threshold. It is obvious from the previous discussion that landlords with higher supervision costs will have higher reservation wealth. In equilibrium, therefore, tenant wealth and landlord supervision costs are positively correlated.

5.2 Econometric Implementation

An econometric model of landlord choice across the S, MS and FR (fixed rent) contracts can be built up from the following equations:

$$
\begin{aligned}
r_m &= a\theta + \psi_m{\prime}x && \text{(6a)}\\
r_s &= bw + \psi_s{\prime}x && \text{(6b)}\\
w_1 &= \sigma_1{\prime}x && \text{(6c)}\\
w &= c\theta + \xi && \text{(6d)}
\end{aligned}
$$

As we have already seen, the landlord's return on the MS contract, r_m, depends on supervision costs but is invariant to tenant wealth, whereas the converse is true for the return on the S contract, r_s. We assume that both return functions are linear in θ, keeping in mind the theoretical restrictions $a < 0$ and $b > 0$. We also let the returns to each contract, as well as the fixed rental wealth threshold, w_1, depend on plot characteristics x (which includes a constant). Equation 6d is the matching function (see Ackerberg and Botticini, 2002) with match error ξ being the sole source of random noise in our model. If $c = 0$, the model reduces to one in which tenant wealth is the only unobservable. The theory, however, implies that $c > 0$.

Putting together these equations, we find that the landlord chooses

$$
\begin{aligned}
MS &\iff r_m > r_s \iff \xi < d_1\theta + \sigma_2{\prime}x\\
S &\iff r_m \leq r_s \text{ and } w < w_1 \iff d_1\theta + \sigma_2{\prime}x \leq \xi < d_2\theta + \sigma_1{\prime}x\\
FR &\iff w \geq w_1 \iff d_2\theta + \sigma_1{\prime}x \leq \xi
\end{aligned} \tag{7}
$$

where $d_1 = (\frac{a}{b} - c)$, $d_2 = -c$, and $\sigma_2 = (\psi_m - \psi_s)/b$. Contract choice probabilites can thus be estimated using an ordered logit model, generalized to allow the slope coefficients to vary with the categorical value of the dependent variable. Our model provides three testable restrictions: $d_1 < 0$, $d_2 < 0$, and $d_1 - d_2 < 0$. Intuitively, higher supervision costs affect contract choice both by directly lowering the landlord's return to monitoring a tenant as well as by increasing the attractiveness to him of wealthier tenants, who can be given stronger incentives and concomitantly less supervision. The combination of these two effects, however, operates only at the margin between the MS and S contract, whereas the matching effect alone operates at the margin between the S and FR contract.

5.3 Results: Supervision costs and contract choice

The 2001 PRHS not only asks landlords about each of their tenants and the terms of their contract, but also about the number of supervisory meetings that they or their *kamdars* had with their tenant. Thus, we can use the threshold number of annual landlord-tenant meetings (10) estimated from the yield data and the tenants' responses on supervision to construct a like supervision indicator on the landlord side. Based on this definition, out of 609 leased plots in our landlord sample, 29% are given on fixed rent, 25% on standard share-contracts, and 46% on monitored share-contracts.

To implement the econometric model, we need a variable that shifts the cost of supervision. A natural candidate for θ is a measure of the accessibility of the plot to the landlord. We use the location of the plot relative to the landlord's village; 14% of plots are "outside" the landlord's village of residence. Although we do not have the actual distance between the plot and the landlord's house, in most cases a plot in a different village will not be within walking distance, whereas a plot inside the village will be no more than a kilometer or two away.

Table 6 reports the generalized ordered logit estimates based on 7. Included in x are the plot characteristics used previously, plus dummies for whether the contract applies to the *kharif* or *rabi* seasons only and dummies for the four provinces. The first two columns under specification (1) report the estimates of the threshold parameters for, respectively, the margin between the S and MS contracts and the margin between the FR and the S contracts. Thus, for example, the coefficient -1.59 in the first column is an estimate of d_1 (scaled by the error variance) and the coefficient -0.766 in the second column is the corresponding estimate of d_2. Both of these estimates are significantly different from zero and from eachother, all of which is consistent with the implications of the model. In short, accessibility of the plot to the landlord does appear to influence the tradeoff between incentives and supervision.

A possible caveat, however, is that only a select group of landlord's may own plots outside their village. These landlords, in turn, may prefer particular types of contracts for reasons unrelated to supervision costs on the plot. For example, large and wealthy landlords may be more likely to own distant plots and to supervise their share-tenants. To address this problem, specification (2) of the contract choice model controls for landlord characteristics; namely, ownership of land, tractors and tubewells. While these asset variables do significantly affect the choice between S and MS contracts (although not in a consistent direction), their inclusion has little effect on the results of interest. Indeed, our key finding that greater supervision costs lower the returns to tenant monitoring (i.e., $d_1 - d_2 = \frac{a}{b} < 0$) is even somewhat strengthened.

6 Conclusions

Recent empirical evidence implies that tenancy (or land) reform may be a rare example of a "win-win" policy. Redistributing property rights over land from wealthy landlords to poor tenants clearly has attractive equity implications, at least in principle. But Banerjee, et al.'s (2002) study of the tenancy reform in West Bengal makes a much stronger case, suggesting that such redistributions can have positive efficiency effects as well. This paper delivers a less sanguine conclusion. Our evidence establishes that gross productivity of land cultivated by sharecroppers differs little from that of land cultivated by owners and fixed renters. At most, this yield differential can be 8%, and our point estimates are less than half this magnitude. In Pakistan, at least, giving higher powered incentives to share-tenants would not have a dramatic impact on agricultural productivity.[29]

The overall efficiency implications, though, are more difficult to assess. Maintaining share-tenant productivity requires fairly heavy landlord supervision. Our evidence shows that yield on land cultivated by sharecroppers who are monitored by their landlords – the majority of share-tenants in Pakistan – is on the order of 30% higher than yield on land cultivated by unmonitored sharecroppers. Since this supervision is costly to the landlord, either in his own time or in money (for hiring labor overseers), redistributing land rights would generate an efficiency gain by eliminating the need to supervise. Indeed, Ai, et al. (1997) provide evidence from Tunisia that output on the *landlord's* own cultivated land suffers the more frequently he meets with his tenants. Even if this were true in Pakistan, it is unlikely that there would be much impact on aggregate agricultural productivity, since the proportion of land cultivated by landlords is small in comparison to the proportion cultivated by share-tenants. More generally, it is improbable that the total resources expended on supervising tenants would come close to matching the gross return on supervision uncovered in this paper, although a definitive answer to this question must await further research.

[29] In Jacoby and Mansuri (2003), we reach similar conclusions regarding land-specific investment. Although such investment is significantly lower on tenanted land, the productivity effects are very small, on the order of 1-2% of yields.

References

[1] Ackerberg, D. and M. Botticini (2002): "Endogenous Matching and the Empirical Determinants of Contract Form," *Journal of Political Economy*, 110(3):564-90.

[2] Ai, C., J.-L. Arcand, and F. Ethier (1997): "De l'Efficacite Allocative des Contrats Agricoles: Cheung avait-il Raison?" *Revue d'economie du developpement*, 2: 103-27.

[3] Andrews, D. (1989): "Power in Econometric Applications," *Econometrica*, 57(5):1059-90.

[4] Banerjee, A., P. Gertler, and M. Ghatak (2002): "Empowerment and Efficiency: Tenancy Reform in West Bengal," *Journal of Political Economy*, 110(2):239-280.

[5] Basu, K. (1992): "Limited Liability and the Existence of Share Tenancy," *Journal of Development Economics*, 38(1):203-220.

[6] Bell, C. (1977): "Alternative Theories of Sharecropping: Some Tests Using Evidence from Northeast India," *Journal of Development Studies*, 13(July):317-46.

[7] Binswanger, Hans P.; Deininger, Klaus; Feder, Gershon (1995): "Power, Distortions, Revolt and Reform in Agricultural Land Relations," in *Handbook of development economics*. Volume 3B, 2659-2772.

[8] Cheung, S. (1968): "Private Property Rights and Sharecropping," *Journal of Political Economy*, 76.

[9] Demougin, D. and C. Fluet (2001): "Monitoring versus Incentives," *European Economic Review*, 45:1741-64.

[10] Johnson, D. G. (1950): "Resource Allocation Under Share Tenancy," *Journal of Political Economy*, Vol 58, Issue 2: 111-123.

[11] Eswaran M., and Kotwal, A. (1985): "A Theory of Contractual Structure in Agriculture," *American Economic Review*, v75, n3: 352-67

[12] Halligan, William S. (1978): "Self-selection by Contractual Choice and the Theory of Sharecropping'" *Bell Journal of Economics*,9: 344-54.

[13] Hansen, Bruce (1999) "Threshold Effects in Non-Dynamic Panels: Estimation, Testing and Inference," *Journal of Econometrics*, 93, 345-368.

[14] Jacoby, H. and G. Mansuri (2003): "Incomplete Contracts and Investment: A Study of Land Tenancy in Pakistan," manuscript, World Bank.

[15] Laffont, J.-J. and M. Matoussi (1995): "Moral Hazard, Financial Constraints and Sharecropping in El Oulja," *Review of Economic Studies*, 62(July): 381-99.

[16] Mookherjee, D. (1997): "Informational Rents and Property Rights in Land," in *Property Rights, Incentives, and Welfare*, edited by J. Roemer. London: Macmillan Press.

[17] Nabi, I. (1986): "Contracts, Resource Use and Productivity in Sharecropping," *Journal of Development Studies*, 22(2):429-42.

[18] Otsuka, Chuma and Yujiro Hayami (1992) "Land and Labor Contracts in Agrarian Economies: Theories and Facts," *Journal of Economic Literature*, Vol. 30, No. 4, 1965-2018.

[19] Reid, Josph D., (1977): "The Theory of Share Tenancy Revisited-Again," *Journal of Political Economy* v85, n2: 403-07.

[20] Shaban, R. (1987): "Testing between Competing Models of Sharecropping," *Journal of Political Economy*, 95(5):893-920.

[21] Shetty, S. (1988): "Limited Liability, Wealth Differences and Tenancy Contracts in Agrarian Economies," *Journal of Development Economics*, 29:1-22.

[22] Stiglitz (1974) "Incentives and Risk Sharing in Sharecropping'" *Review of Economic Studies*, vol (41), No. 2, 219-255.

[23] Kees Van Garderen and Shah (2002) "Exact Interpretation of Dummy Variables in Semilogarithmic Equations," *Econometrics Journal*, volume 5, 149-159.

Figure 1: Tenant Wealth and Contractual Choice

Figure 2: Check for Outliers

Table 1: Number of Plots in Estimation Samples

	IFPRI-1993	PRHS-2001	Total
Plot growing any crop[a]			
O-C-S households[b]	234	273	507
	(119)	(132)	(251)
All multi-plot households	660	1461	2121
(hh fixed effect sample)	(228)	(402)	(630)
All households	885	2305	3190
(village fixed effect sample)	(323)	(728)	(1051)
Plots growing five main crops[c]			
O-C-S households	163	240	403
	(81)	(113)	(194)
All multi-plot households	462	1256	1718
(hh fixed effect sample)	(174)	(351)	(525)
All households	719	2088	2807
(village fixed effect sample)	(279)	(670)	(949)

Note: Number of sharecropped plots in parentheses.

[a]Sample for fertilizer analysis.

[b]Owner-cum-sharecroppers (39 plots cultivated by renter-cum-sharecroppers).

[c]Sample for yield analysis.

Table 2: Effect of Share-Tenancy on Yields and on Nitrogen Use

	Household fixed effects		Village fixed effects[a]	
	Coeff. (S.E.)	IP interval[b]	Coeff. (S.E.)	IP interval[b]
Yield (Rps/acre)				
no controls	-2.7 (4.2)	[-6.9,-13.8]	0.6 (2.6)	[-4.3,-8.5]
plot characteristics	-2.8 (4.2)	[-6.9,-13.9]	0.6 (2.6)	[-4.2,-8.5]
plot characteristics + crops composition	-2.3 (4.0)	[-6.5,-13.0]	-0.0 (2.5)	[-4.0,-8.1]
Nitrogen (kg/acre)				
no controls	0.6 (5.0)	[8.1,16.2]	-0.8 (3.1)	[5.1,10.3]
plot characteristics	2.3 (5.0)	[8.2,16.5]	-0.4 (3.1)	[5.2,10.3]
plot characteristics + crops composition	1.1 (4.9)	[8.0,16.0]	-1.5 (3.1)	[5.1,10.2]

Notes: Coeff. and s.e. normalized by average yield on owner-cultivated plots ×100
See Appendix Table A.1 for details on the control variables.
[a]Household random effects.
[b]Inverse power function for one-sided test evaluated at 0.5 and 0.95, respectively.

Table 3: Effect of Share-Tenancy on Labor Use

	Household fixed effects		Village fixed effects[a]	
Type of Labor (hrs/acre)	Coeff. (S.E.)	IP interval[b]	Coeff. (S.E.)	IP interval[b]
All agricultural tasks				
Family	-6.4 (7.0)	[-10.4,-20.7]	5.3 (5.4)	[-9.0,-18.1]
Family + hired	0.8 (7.7)	[-12.2,-24.4]	8.0 (5.5)	[-9.4,-18.8]
Excluding harvesting/threshing				
Family	-1.9 (8.3)	[-12.8,-25.6]	6.1 (6.2)	[-10.4,-20.9]
Family + hired	0.2 (8.5)	[-13.3,-26.6]	6.4 (6.3)	[-10.6,-21.3]

Notes: Coefficients and standard errors from logarithmic specifications are converted to percentage changes. All specifications include plot characteristics and crop composition.

[a]Household random effects.

[b]Inverse power function for one-sided test evaluated at 0.5 and 0.95, respectively.

Table 4: Landlord Supervision and Yields

	Mean (s.d.)	**(1)**	**(2)**	**(3)**	**(4)**
Sharecropped plot[a]	0.28	-4.4	-5.0	-2.5	-3.7
		(4.8)	(4.9)	(4.9)	(5.0)
Supervised ($m_{ci} > \widehat{k}$)	0.65	28.1	27.4	25.8	24.9
		(10.9)	(11.0)	(11.0)	(11.0)
Tenant characteristics					
Recent ($\leq$ 3 yrs.)	0.29		-5.7		-1.8
			(10.2)		(10.3)
Relative of landlord	0.34		10.0		16.7
			(10.3)		(11.0)
Landlord characteristics					
Log of landholdings	5.78			0.050	0.071
	(1.64)			(0.040)	(0.043)
Owns a tractor	0.32			12.7	10.4
				(12.4)	(12.5)
Owns a tubewell	0.21			13.6	13.1
				(14.8)	(14.8)

Notes: All regressions include plot characteristics, crop composition variables, and household fixed effects. Coeff. (s.e.) normalized by average yield on unsupervised sharecropped plots $\times 100$ unless otherwise noted. Sample size is 1256 plots.
[a]Coefficient renormalized by average yield on owner-cultivated plots. Mean is taken over all plots in the sample; means of other variables are taken over 351 sharecropped

Table 5: Landlord Supervision and Yields in Pure Sharecropper Households

	Mean (s.d.)	**(1)**	**(2)**	**(3)**	**(4)**
Supervised ($m_{ci} > \widehat{k}$)	0.66	73.1	78.1	69.4	74.1
		(22.4)	(22.9)	(23.2)	(23.8)
Tenant characteristics					
Recent ($\leq$ 3 yrs.)	0.26		19.5		21.0
			(18.3)		(18.6)
Relative of landlord	0.30		12.3		16.8
			(19.7)		(20.7)
Landlord characteristics					
Log of landholdings	5.87			0.038	0.067
	(1.71)			(0.081)	(0.085)
Owns a tractor	0.35			6.4	2.9
				(20.1)	(20.3)
Owns a tubewell	0.22			6.4	7.5
				(31.6)	(31.7)

Notes: All regressions include plot characteristics, crop composition variables, and household fixed effects. Coeff. (s.e.) normalized by average yield on unsupervised sharecropped plots $\times 100$. Sample size is 264 plots.

Table 6: Generalized Ordered Logit Model of Contract Choice

		(1)		(2)	
	Mean (s.d.)	**S - MS**	**FR - S**	**S - MS**	**FR - S**
Plot outside landlord's village	0.14	-1.59**	-0.766*	-1.73**	-0.727
		(0.367)	(0.382)	(0.427)	(0.389)
Landlord characteristics					
Log of landholdings	3.96			0.648**	0.192
	(1.43)			(0.242)	(0.165)
Owns a tractor	0.11			-1.33*	0.511
				(0.593)	(0.567)
Owns a tubewell	0.10			1.10*	-0.656
				(0.519)	(0.505)
$H_0 : d_1 = d_2$ (p-value)		0.024		0.004	

Note: Robust standard errors adjusted for village-level clustering in parentheses (* denotes p-value<0.05; ** denotes p-value<0.01). Each equation also includes plot characteristics (see appendix), dummies for seasonal leases, and province dummies. $N = 609$.

Appendix

Table A.1: Descriptive Statistics for Plot Characteristics

	IFPRI -1993		PRHS - 2001	
	Owned/Rented	Sharecropped	Owned/Rented	Sharecropped
Area (acres)[a]	10.8	9.9	8.3	7.4
	(15.4)	(8.9)	(14.1)	(8.2)
log(value/acre)	10.6	10.3	9.5	9.1
	(1.1)	(1.1)	(1.0)	(1.0)
Outside village	0.07	0.06	0.07	0.11
Year-round canal irrigation	0.43	0.58	0.31	0.53
Seasonal canal irrigation	0.09	0.10	0.23	0.21
Tubewell access	0.18	0.10	0.57	0.46
Sandy soil	0.11	0.19	0.18	0.24
Maira soil	0.24	0.24	0.29	0.20
Chikni soil	0.03	0.03	0.24	0.36
Cotton[b]	0.02	0.00	0.13	0.09
Rice[b]	0.19	0.38	0.15	0.30
Sugarcane[b]	0.11	0.12	0.07	0.12
Maize[b]	0.08	0.05	0.08	0.06
N	440	279	1418	670

Notes: Means (Std. Dev.). Omitted categories: canal irrigation = none; soil type = clay.

[a]Entered in logs in the regressions.

[b]Fraction of area planted to five major crops (omitted crop is wheat).