

# RISK MANAGEMENT IN THE COOPERATIVE CONTRACT

ETHAN LIGON

## 1. INTRODUCTION

Agricultural cooperatives have long played an important role in helping their members manage risk. Yet the typical cooperative (co-op) does a much better job of helping its members manage some sorts of risk than it does others. In particular, co-ops are good at helping members manage *marketing* risk, or idiosyncratic variation in prices observed within the course of a single season. However, agricultural cooperatives seem not to be particularly good at helping their members to manage *production* risk, which involves variation in yield over the course of several years. This paper argues that by taking advantage of the multi-year nature of most members' relationship with the cooperative, the cooperative can also provide a useful (though limited) form of insurance against crop shortfalls.

Agricultural producers face various sorts of risk; some of these are difficult to manage using standard institutions. Agricultural cooperatives can play an important role in helping their members to manage these sources of risk. For example, marketing co-ops traditionally help to reduce price risk by pooling sales across time and space, and could reduce production risk by making some payments to members on the basis of predetermined shares, rather than on actual delivery or "patronage."

However, just because a cooperative *can* help its members manage risk does not mean that membership in a cooperative *will* serve as a useful means for dealing with risk. The details of a cooperative's by-laws, management of equity, mechanisms governing the transfer of delivery rights, methods of capitalization, pooling rules, and expansion options can all have a big impact on the usefulness of the cooperative structure as a risk-management tool for users.

Here, we provide an enumeration of some different sorts of shocks which may be important to agricultural producers and then describe the manner in which a typical cooperative market pooling mechanism

can reduce the risk associated with these shocks. However, market pools do not help producers manage yield risk at all—the cooperative must deploy a different mechanism if it is to help members manage shocks to output.

Accordingly, we describe a simple way of structuring an agricultural marketing cooperative’s operations so as to best help members manage yield risk, in addition to the other sources of risk they face. Marketing cooperatives offer the greatest scope for mutual insurance among members and so will serve as the principal focus of the proposed paper, but some approaches to risk management will apply to other forms of cooperatives.

### RISKS FACED BY THE PRODUCER

We consider four different sources of risk faced by agricultural producers. These are yield risk, quality risk, basis risk, and price risk—together, these will determine the total revenue generated by the farmer for a particular crop.

More specifically, suppose that at the beginning of period  $t$ , farmer  $i$  decides to devote  $m_{it}$  acres to the production of some particular commodity. The farmer invests  $a_{it}$  in inputs. The farmer harvests at the end of the period and realizes a yield (averaged across his  $m_{it}$  acres of production) of  $q_{it}$  having quality (also averaged across acres) of  $\theta_{it}$ . The quality here should be thought of as an index which can directly reduce the value of output.

Shocks to yield and quality having been realized for all farmers, aggregate supply and demand yield a market price for the commodity in question of  $p_t$ —variation in these aggregates gives rise to price risk.<sup>1</sup> However, farmer  $i$  will typically not be able to sell his output at exactly this price—rather, the price he receives will depend not only on aggregate demand and supply but also on his “quality shock”  $\theta_{it}$  as well as on a “basis shock”  $b_{it}$ , which (while random) may vary depending

---

<sup>1</sup>The term “price risk” here may be something of a misnomer (after all, the real shocks are to supply and demand), but we do have a reason for using it. For commodities for which a futures market exists, there will be some “spot” price which reflects aggregate supply and demand conditions for the market. Importantly for the producer, variation in this price can be hedged, using those same futures markets. The difference between the market price and the price actually received by the producer who does his own marketing is what is sometimes called the “basis” in the literature on commodity markets; hence, this variation in the price actually received by the producer (netted of the variation in the spot price) is what we term the producer’s basis risk.

on the farmer's location, transportation costs, and other idiosyncratic factors.

We assume that farmer  $i$ 's total revenue can be written:

$$y_{it} = p_{it}q_{it}m_{it} = (p_t + b_{it} + \theta_{it})q_{it}m_{it}.$$

Although the farmer has some control over this risk via his choices of  $m_{it}$  and inputs  $a_{it}$ , idiosyncratic variation in basis ( $b_{it}$ ), quality ( $q_{it}$ ), quality ( $\theta_{it}$ ), and yield ( $q_{it}$ ) implies that variation in the farmer's revenue will not be perfectly correlated with that of other farmers.

Now, under very modest assumptions regarding the distribution of the idiosyncratic variables ( $q_{it}, b_{it}, \theta_{it}$ ) (e.g., so long as these are not perfectly correlated across farmers), the variation in *average* revenue across  $n$  farmers will be smaller than the average variation for a single farmer. Total revenues for the cooperative will be:

$$\bar{y}_t^n = \sum_{i=1}^n p_{it}q_{it}m_{it},$$

and as the size of the cooperative increases we may expect (under only somewhat less modest assumptions) that a law of large numbers will apply to this sum, and that  $\text{plim}_{n \rightarrow \infty} \bar{y}_t^n/n = \bar{y}_t$ . As a consequence, by pooling revenues, the cooperative can reduce the risks faced by every one of its members. However, typically agricultural cooperatives distribute their revenues in proportion to "current patronage," or current year deliveries to the cooperative, so that member  $i$  will receive  $(q_{it}m_{it}/\sum_{j=1}^n q_{jt}m_{jt})\bar{y}_t^n$ . While pooling within the cooperative effectively reduces variation in  $\bar{y}_t^n$ , it has no such effect on the variation of the share, which depends on  $q_{it}$ . Thus, an effect of relying on current patronage to divide revenues is to make the co-op ineffective at sharing yield risk.

#### BENCHMARK: FULL RISK SHARING IN A COOPERATIVE.

In principle a marketing cooperative could completely insure its members against risks associated with idiosyncratic shocks to yield or production as well as risks associated with variation in prices, providing a sure 'home' for members' production at a price determined in advance. For example consider a closed marketing cooperative. A simple mechanism which would fully insure the members of this example cooperative would have four elements:

- (1) Each member would be assigned (the assignment could be accomplished via negotiation at the time the member joined the cooperative) a *delivery target* in the cooperative. Member  $i$ 's

delivery target divided by the sum of all members' delivery targets would determine his or her *share* in the cooperative.

- (2) Regardless of their negotiated delivery target, members would commit to deliver all of their production to the coop—they would have, in effect, unlimited delivery rights, but not an *obligation* to deliver in the event of a production shortfall.
- (3) The cooperative would commit to distribute net revenues from the sale of all members' deliveries in direct proportion to members' initial shares.

Inducing members to commit to deliver their entire production to the cooperative would require that shares in the cooperative be allocated in rough proportion to members' expected production, but with this assignment in place, these two elements would suffice to ensure that all members of the cooperative collectively shared any production risk as well as collectively sharing price risk associated with variation in prices over the period (e.g., year) and spatial area over which production was sold. But the cooperative could do even better by:

- (4) The cooperative could further insure its members against variation in aggregate prices by using futures markets (if those markets existed for the commodities in question) or forward sales to hedge aggregate price risk.

In principle by pooling shocks to yield as well as to basis and quality, the mechanism described above works to provide full mutual insurance. However, in practice providing full insurance may be impossible because *ex post* more successful farmers cannot be compelled to always share with their less successful brethren—when a farmer's yield is very high, he will be tempted to use some alternative marketing mechanism which does not require him to subsidize other members of the co-op. Following a literature on risk-sharing which contemplates a similar problem, we term this a problem of *limited commitment*.

#### LIMITS TO POSSIBLE INSURANCE: FAILURES OF COMMITMENT

In our description above of the elements of a scheme which would permit our example marketing cooperative to provide full insurance to its members, the word “commit” appeared in two key places—members must commit to deliver all their production to the coop, and the cooperative must in turn commit to distributing net revenues in proportion to initial shares. But what if this commitment is not feasible? It may not be possible to induce a member with unusually high production to share his windfall with other cooperative members; he may instead simply opt to market some of his production outside the cooperative.

In the academic literature, Kimball (1988) was the first paper to consider the consequences of this sort of “limited commitment” in a producers’ cooperative. Kimball shows that although limited commitment (on the part of the members) reduces the ability of the cooperative to insure its members, it nonetheless can provide *some* useful insurance. A mechanism which would implement the risk-sharing that Kimball envisioned would differ from the full insurance mechanism described above in its second element. In particular, rather than asking members to deliver *all* of their production to the cooperative, members would only be asked to deliver up to their delivery targets. A member who had a production shortfall (i.e., who produced less than their delivery rights) would deliver all of their production to the cooperative; a member who produced more than their delivery rights would deliver just enough to exhaust their delivery rights and could then market the surplus outside the cooperative (or alternatively sell the surplus to the cooperative on a cash basis). In Kimball’s scheme a failure to deliver the minimum of total production or delivery rights would be punished by permanently expelling the offender from the cooperative. Net revenues would be distributed to members on the basis of shares, as before.

One way to understand Kimball’s chief contribution is that he was able to show that under the mechanism described there is some set of levels of member delivery targets which makes it so that no member ever has an incentive to ‘short’ the cooperative, and yet which provides enough reduction in yield risk so that members are willing to participate in the cooperative.

More recent developments in the general theory of risk-sharing under limited commitment suggest some ways in which Kimball’s mechanism could be improved upon. First, in Kimball’s mechanism shares are not allowed to change (except in the event of member default). But it is possible to show that if members who exceed their delivery targets are rewarded by promises of a larger future share in the cooperative’s revenues, then the value of the cooperative as a risk-sharing institution can be dramatically improved. Second, Kimball was only able to show that the agricultural cooperative could serve as a useful risk-management tool provided that no other risk-management tools (such as savings or credit markets) were available. But research by Gauthier et al. (1997) and Ligon et al. (2000) suggests that, far from “crowding out” cooperatives, with careful design the existence of outside credit institutions could in fact greatly increase the value of agricultural cooperatives to their members as an insurance mechanism.

## THE OPTIMAL CONTRACT WITH LIMITED COMMITMENT

In this section we pursue three main extensions to the standard dynamic model of limited commitment (e.g. Kehoe and Levine, 1993; Kocherlakota, 1996; Ligon et al., 2002, notation here is drawn from this last). What we've termed the "standard" model features two risk-averse producers, at least one of which faces some yield risk. In our first extension we show how to treat the case in which there are more than two producers. Second, we replace the simple stochastic yields of the standard model with a general, producer-specific intertemporal technology, permitting investment. Third, we permit each producer to choose the technology it will operate.

Consider, then, a cooperative with  $n$  infinitely lived producers. We index these producers by  $i = 1, 2, \dots, n$ ; we suppose that  $i$  also denotes the location of the producer in some metric space  $(\Lambda, d)$ , with  $\Lambda$  the set of possible locations and  $d$  a distance metric on  $\Lambda$ . Time is discrete and indexed by  $t$ . At any date  $t$  some state  $s \in \mathcal{S}$  is realized (with  $\mathcal{S}$  finite); given that the current period's state is  $s$ , the probability of the state next period being  $r \in \mathcal{S}$  is given by  $\pi_{sr} > 0$ . Producer  $i$  derives momentary utility from consumption according to some function  $u^i : \mathbb{R} \rightarrow \mathbb{R}$  and discounts future utility at a common rate  $\beta \in [0, 1)$ .

Each producer chooses some stochastic production technology from a set  $F$ . In particular, at each date let each producer  $i$  choose a stochastic intertemporal technology such that if the current state is  $s$ , and the producer invests  $a$ , then next period the technology returns some quantity  $f_{sr}^i(a)$  in the event that the subsequent state is  $r$ . We assume that each of the functions  $f_{sr}^i$  is non-decreasing, concave, and continuously differentiable.

Producers can agree to participate in a scheme involving mutual insurance, but the scope of this insurance is limited by the fact that after any history, each producer has the option of renegeing on any proposed insurance transfers. In the event that a producer  $i$  which has saved  $a^i$  units of the consumption good reneges in state  $s$ , he is assumed to obtain a discounted, expected utility given by the continuously differentiable function  $Z_s^i(a^i)$ .<sup>2</sup> Accordingly, any 'sustainable' insurance scheme must guarantee that in state  $s$  every producer  $i$  having saved  $a^i$  obtains at least  $Z_s^i(a^i)$  utils under the proposed insurance scheme.

---

<sup>2</sup>In the standard model, a producer who reneges is assumed to be forced into autarky in all future periods. There is no need to make this assumption here, but were we to do so, we could think of this function as corresponding to the autarky utility of producer  $i$  in state  $s$  if its resources available at the beginning of the period are  $a^i$ .

When there is an intertemporal technology which allows investments in future production, consumption allocations depend on the claims producers have to these assets in the event of default<sup>3</sup> as well as the sequence of technology choices made by the producer, whether under the continuation contract or in the event of a deviation. For our application, we'll assume that the distribution of possible outcomes under a given technology depends on the producer's location (e.g., some areas may be particularly suitable for cultivation, and others less so).

Denote the discounted expected utility for producers  $i$  in state  $s$  by  $U_s^i$ . We set up a dynamic programming problem so that the current state is  $s$ , treat the discounted expected utilities for producers as additional state variables in the dynamic program. Thus,  $\mathbf{U}_s^{-n} = \{U_s^i\}_{i=1}^{n-1}$  are utilities that have been promised to the first  $n - 1$  producers. As in the standard model, these promised utilities will be treated as state variables in our dynamic program. In addition a state variable  $z$  denotes the resources available to all the members of the collective at the beginning of the period, which can be divided into consumption and investment. Choice variables in the programming problem will be consumption assignments  $c^i$  for  $i = 1, \dots, n$ , the continuation utilities  $U_r^i$  for each possible state  $r$  in the next period, and an assignment of both technologies  $\{f_{sr}^i\}$  and investments  $a^i$  for each producer. The value function for producer  $n$  can now be written to depend on the current target utilities and collective resources:  $U_s^n(U_s^1, \dots, U_s^{n-1}; z)$ . Then, the dynamic programming problem is:

(1)

$$U_s^n(\mathbf{U}_s^{-n}; z) = \max_{(\mathbf{U}_r^{-n})_{r \in \mathcal{S}}, (c^i, (f_{sr}^i)_{r \in \mathcal{S}}, a^i)_{i=1}^n} u_n(c_s^n) + \beta \sum_{r \in \mathcal{S}} \pi_{sr} U_r^n \left( \mathbf{U}_r^{-n}; \sum_{i=1}^n f_{sr}^i(a^i) \right)$$

subject to an aggregate resource constraint  $\sum_{i=1}^n (a^i + c^i) \leq z$ , with associated Lagrange multiplier  $\mu$ , and subject also to a set of promise-keeping constraints  $u_i(c_s^i) + \beta \sum_{r \in \mathcal{S}} \pi_{sr} U_r^i \geq U_s^i$  with associated multipliers  $\{\lambda^i\}$ , which must hold for all  $i \neq n$ . The solution must also be

---

<sup>3</sup>For example contrast the treatment of storage in Gobert and Poitevin (1998) in which defaulting producers forfeit stored assets with the treatment in Ligon et al. (2000), where holding a large store of assets can provide an incentive for a producer to renege on existing arrangements. It is important to note that the latter treatment (which is similar to our approach here) may introduce nonconvexities, since the value associated with autarky now depends on choice variables. We ignore this interesting difficulty here but refer the reader to Ligon et al. (2000) for a more satisfactory treatment.

sustainable, and so satisfy the sustainability constraints:

$$(2) \quad \beta \lambda^i \pi_{sr} \phi_r^i: \quad U_r^i \geq Z_r^i(a^i)$$

for all  $r \in \mathcal{S}$ , for all producers  $i \neq n$ , and

$$(3) \quad \beta \pi_{sr} \phi_r^n: \quad U_r^n \left( \mathbf{U}_r^{-n}; \sum_{i=1}^n f_{sr}^i(a^i) \right) \geq Z_r^n(k^n) \quad \text{for all } r \in \mathcal{S}.$$

The expressions on the left of the colons in the equations above are the Kuhn-Tucker multipliers associated with the corresponding constraints. The first-order conditions from the dynamic program yield

$$(4) \quad \frac{u'_n(c_s^n)}{u'_i(c_s^i)} = \lambda^i, \quad \forall i \neq n,$$

$$(5) \quad \lambda_r^i = \lambda^i \frac{1 + \phi_r^i}{1 + \phi_r^n}, \quad \forall r \in \mathcal{S}, \forall i \neq n,$$

where  $\lambda_r^i \equiv \partial U_r^n / \partial U_r^i$  (by the envelope condition this is equal to next period's ratio of marginal utilities between producers  $n$  and  $i$ ), and:

$$(6) \quad u'_i(c_s^i) = \beta \sum_{r \in \mathcal{S}} \pi_{sr} [f_{sr}^i(a_s^i) u'_i(c_r^i)] + \beta \sum_{r \in \mathcal{S}} \pi_{sr} \phi_r^i [f_{sr}^i(a_s^i) u'_i(c_r^i) - Z_r^i(k_s^i)].$$

Of central interest is the behavior of the Lagrange multipliers  $\{\lambda^i\}$ ; by characterizing the evolution of these we can compute the sharing rule across all producers and states. Together, (4) and (5) imply a simple updating rule for utility ratios, where producer  $n$ 's marginal utility is treated as a numeraire. Equation (6) is analogous to the usual Euler equation; the left-hand term is the marginal cost of increased investment associated with foregone contemporaneous consumption, while the first term on the right-hand side is the usual marginal benefit. However, the equation differs from the usual case in that there is a second term. This term reflects both additional marginal benefits measured by the terms  $\phi_r^i f_{sr}^i(a_s^i) u'_i(c_r^i)$  and additional marginal costs measured by the terms  $\phi_r^i Z_r^i(k_s^i)$ . The former terms capture the feature that additional resources can help to relax sustainability constraints; the latter terms have to do with the problem that if too many resources are assigned to a producer with low surplus, then autarky may become relatively more attractive, and thus make the sustainability constraints more binding, actually reducing welfare (see Ligon et al. (2000) for an illustration).

Although the sign of the contribution that the additional terms in (6) make is generally ambiguous, in many situations optimal assignment

of the  $a^i$  to producers who are unlikely to have binding sustainability constraints means that the sign will be positive. In this case we can interpret the additional terms as a sort of endogenous “liquidity constraint,” since current consumption will be lower relative to future consumption than predicted by the usual Euler equation. Intuitively, we can think of current consumption being reduced due to some producers (those who would otherwise be likely to have binding sustainability constraints in the subsequent period) posting “bonds” with other producers (save for timing, this mechanism is similar to one considered by Gauthier et al. (1997) which they term “ex ante payments.”).

#### IMPLEMENTING THE OPTIMAL CONTRACT WITH LIMITED COMMITMENT

In the previous section we considered the problem of devising an optimal intertemporal sharing rule which would provide maximal risk-sharing within a cooperative. However, the rule we devised is specified in terms of consumption and investment allocations and in terms of promised utilities. It may not be practical or natural to write the membership agreement, by-laws, and so on for the cooperative in these terms; so, in this section we make an effort to recast the optimal contract in terms more closely related to existing standard cooperative agreements.

The key to mapping between the customary jargon of agricultural cooperatives and the model we have outlined involves making each farmer’s share of cooperative revenue depend, not on current patronage, but on what we will call “accumulated patronage points.” This is simply an accounting mechanism which would allow the cooperative to keep track of the history of a member’s deliveries and in particular to keep track of the extent to which a given member has subsidized others in the past, so as to reward that same member in the future. Some key points:

- (1) Accumulated patronage (or “patronage points”) for farmer  $i$  in state  $s$  corresponds to the quantity  $\lambda_s^i$  in the model. When (for example) farmers have logarithmic utility functions, farmer  $i$  will receive a share of total cooperative revenue in state  $s$  equal to

$$\sigma_s^i = \frac{\lambda_s^i}{\sum_{j=1}^n \lambda_s^j}.$$

- (2) Anyone can join the cooperative simply by delivering output, but a “new” producer has an “accumulated patronage” which

will be somewhat less than the total share of his deliveries to the cooperative in the year he joins. Since he thus provides an initial subsidy to existing members, he will be welcomed. In turn a new member has an incentive to join (even though he'll be compensated for less than his full deliveries) because of the future benefits of risk reduction that he receives by virtue of joining the cooperative.

- (3) Note from the sharing rule that producers with more accumulated patronage receive higher compensation for delivery of some fixed amount than do producers with less accumulated patronage, regardless of current deliveries.
- (4) Every farmer has some 'delivery target;' the value of this target depends on his accumulation of 'patronage points,' which in turn depends on historical deliveries. New members start with a delivery target of zero; so their initial delivery is immediately rewarded with some patronage points.
- (5) If the cooperative has enough members and no single member is "too big," then every farmer is fully insured (in the current period) against failure to reach his delivery target. Farmers who are "too big" will have some limited insurance against failure to reach their delivery targets.
- (6) A farmer receives additional 'patronage points' whenever his deliveries exceed his delivery target (and receives no additional points otherwise).
- (7) The cooperative markets total deliveries  $\bar{q}_t = \sum_i m_{it}q_{it}$ , realizing an average price  $p_t$ .
- (8) The cooperative distributes  $\tau_{it}$  to the  $i$ th farmer; this distribution is equal to total revenue  $p_t\bar{q}_t$  times the farmer's share of total patronage points.
- (9) Because each farmer's share of current revenue depends on his *accumulated* patronage, he is protected against current production shortfalls. His past deliveries will have sometimes subsidized other members when they had a shortfall and resulted in an accumulation of patronage points. Since the division of cooperative revenues depends on these accumulated points, he will not be seriously hurt by a bad shock. However, since the subsidy he receives from others may result in new patronage points for *them*, his share of total accumulated patronage points will fall, resulting in a smaller share of total revenues for the farmer in the future.

## CONCLUSION

Agricultural marketing cooperatives are typically organized so as to do a very good job of helping their members manage some sources of risk but not others. In particular cooperatives do a good job of reducing risk associated with variation in price, basis, and quality, but do almost nothing which helps members to manage yield risk.

Agricultural cooperatives could perfectly share yield risk among their members by pooling total revenues from all their members and then distributing these revenues in proportion to members' shares—the key would be to specify these shares in such a way that the shares did not depend on current production. But there may be a good reason that agricultural cooperatives typically fail to insure their members' yield risk—in the scheme we have just described, a farmer with particularly high yields in a given period would wind up subsidizing other less fortunate farmers in that same period. He might be better off by leaving the cooperative altogether rather than sharing his bounty.

Although this problem of limited commitment may doom *perfect* sharing of production risk within the cooperative, it is possible to modify the sharing mechanism so as to provide at least some sharing, a point first made by Kimball (1988), who described a simple mechanism which would provide farmers with high yields an incentive to share at least *part* of their high yields with their fellows. The basic trick to create some scope for sharing yield risks is to adopt a sharing rule for cooperative revenues which does not only depend on random current yields but which instead is at least partly predetermined. Kimball's idea was to assign *fixed* shares, but to not require members with unusually high yields to market all of their product through the cooperative.

In this article we describe an alternative, dynamic scheme which improves on Kimball's mechanism in terms of the scope for risk-sharing that it creates. The key improvement is that our shares are made a function of the history of deliveries to the cooperative. In this scheme farmers never have an incentive to market outside the cooperative—instead, a farmer who makes an unusually large delivery benefits a little bit immediately, but in exchange for the current subsidy that he provides other farmers, he is rewarded by being assigned a larger share of future revenues. Conversely, a farmer who consistently brings in less patronage than expected will tend to see his share fall over time.

Allowing shares to vary dynamically with the history of deliveries not only improves risk-sharing but should also make it possible for the cooperative to deal gracefully with the arrival of new members, changes in the scale of operations of existing members, and the retirement of

old members. There is considerable evidence (and some litigation) surrounding the last point in particular—for many existing agricultural marketing cooperatives, there is an “equity redemption problem” having to do with easing cooperative members out of the cooperative when those members retire. The dynamic mechanism described here offers a promising approach to solving this and related issues but is left for future research.

#### REFERENCES

- Gauthier, C., M. Poitevin, and P. González (1997). Ex ante payments in self-enforcing risk-sharing contracts. *Journal of Economic Theory* 76, 106–144.
- Gobert, K. and M. Poitevin (1998). Non-commitment and savings in dynamic risk-sharing contracts. manuscript.
- Kehoe, T. J. and D. K. Levine (1993). Debt-constrained asset markets. *Review of Economic Studies* 60(205), 865–888.
- Kimball, M. S. (1988). Farmers’ cooperatives as behavior toward risk. *American Economic Review* 78, 224–236.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies* 63(4), 595–610.
- Ligon, E., J. P. Thomas, and T. Worrall (2000). Mutual insurance, individual savings, and limited commitment. *Review of Economic Dynamics* 3(2), 216–246.
- Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies* 69(1).