Risky Income or Lumpy Investments? Evidence on Two Theories of Under-Specialization

Ajay Shenoy*

University of California, Santa Cruz

November 30, 2016

First Version: October 08, 2012

*Email at azshenoy@ucsc.edu. Special thanks to Raj Arunachalam, David Lam, Jeff Smith, David Weil, Anja Sautmann, Chris Udry, Silvia Prina, Stefanie Stantcheva, seminar participants at Michigan and Brown, and conference attendees at PacDev and MWIEDC in 2013 for helpful comments and advice. I am grateful to Marcel Fafchamps and two annonymous referees for their guidance and suggestions. I am also grateful to Robert Townsend and the Townsend Thai Project for collecting and distributing the data I use for this project. File formatting based on a stylesheet by Chad Jones.

Abstract

Why do the poor have so many economic activities? According to one theory the poor do not specialize because relying on one income source is risky. I test the theory by measuring the response of Thai rice farmers to conditional volatility in the international rice price. Households expecting a harvest take on 1 extra activity when the volatility rises by 21 percent. I confirm the decrease in specialization costs households foregone revenue. I find no evidence to back a second theory in which households under-specialize because they cannot afford lumpy business investments. (JEL Codes: D13, O12)

Keywords: Risk, specialization, occupation, lumpy investment, Thailand

1 Introduction

> *To take...the trade of the pin-maker...One man draws out the wire, another straights it, a third cuts it...ten persons, therefore, could make among them upwards of forty-eight thousand pins in a day...But if they had all wrought separately and independently...they certainly could not each of them have made twenty, perhaps not one pin in a day...*

-Adam Smith, *Wealth of Nations*

The idea that specialization is efficient is as old as economics itself. The puzzle, then, is to explain why households in poor countries rarely specialize in a single business or a single job (Banerjee and Duflo, 2007). If entering multiple economic activities is costly, why would the world's poorest people fail to specialize?

This paper tests two well-known but unproven theories for why the poor have so many economic activities: the theory of risky income and the theory of lumpy investments. The theory of risky income compares a poor household choosing economic activities to an investor choosing stocks. Like stocks the activities of the poor have risky returns, driving households to diversify their portfolio even though expanding it is costly. Whereas this theory blames under-specialization on a lack of insurance, the theory of lumpy investments blames a lack of credit. The theory posits that households must make a large investment—tailors must buy a sewing machine and bakers must buy an oven—before expanding any business to its optimal scale. Households that cannot borrow enough to create one large business must cobble together income from many small businesses.

From a simple model I derive several tests of the theory of risky income. Each household has a primary activity and pays a fixed opportunity cost to enter any side activity. The returns to these activities are random and not perfectly correlated. Therefore the theory's first test is that a rise in the riskiness of the primary activity causes a risk-averse household to self-insure by entering more side activities. But since labor spent on side activities is labor taken from the primary activity, a rise in the average return to the primary activity raises the cost of self-insurance. The theory's second test is that a rise in the return to the primary activity causes the household to exit side activities.

The third test, which uses revenue from side activities rather than total revenue to check whether specialization is efficient, is critical for two reasons. First, my empirical approach rules out any test using total revenue because it studies households that do not yet have the revenue from their primary activity. Second, the household reallocates labor between activities when it enters new activities. For both reasons I cannot identify the fixed cost of entering a new activity, the clearest sign of inefficient under-specialization. I can, however, derive the optimal allocation of labor as a function of the number of activities and use this allocation to find the change in side revenue caused when the household enters a new activity. I show that the reallocation of labor works against finding a decrease in revenue, giving a sufficient but not necessary condition. If there is a decrease in side revenue then specialization is inefficient, though the converse need not hold.

To run these tests I study how rice farmers in Thailand respond to volatility in the international price of rice. Using a monthly panel I identify the households who expect a rice harvest in the next three months. Higher volatility in the price of rice raises the riskiness of these farmers' income. I effectively compare these farmers when they face high volatility to themselves at times when they face low volatility, measuring their response to risk. I then compare their response to the response of rice farmers who do not expect a harvest and thus for whom price volatility does not equal income volatility. My identifying assumption is that these two groups would have the same response to high or low volatility. Under this assumption, the difference in their responses is the causal effect of riskier income on specialization. By likewise comparing how the two groups respond to changes in the expected price I identify the causal effect of greater returns on specialization.

My first two tests confirm the theory of risky income. Greater risk drives households into more activities while higher returns tempt households out of activities. After adjusting for how well international prices predict local prices, my baseline estimates suggest a 21 percent rise in volatility causes a household to enter 1 extra activity.

A household that expects a harvest next month sells no rice this month. Though existing research (e.g. Jacoby and Skoufias, 1997) suggests risk lowers investment in physical or human capital, changes in investment take months or years to affect revenue. Since I examine volatility within a short window (the three months before the harvest), any changes to investment will have little effect on my estimates. In short, the mean and variance of the rice price change the number of activities without directly affecting current revenue. I use this change to instrument for the number of activities. Households expecting a harvest do not yet have the revenue from their primary activity, ruling out any test of whether additional activities decrease total revenue. But my third test shows that if additional activities decrease revenue from side activities then a failure to specialize is inefficient. Two-stage least squares confirms exactly that. I confirm that these results are not driven by changes in household labor or composition, by negative shocks rather than volatility, by changes in the aggregate village economy (proxied by wages), or by a correlation between the price of rice and the prices of other crops.

Finally, I test the theory of lumpy investments. The theory predicts that households with easier access to credit can afford the lumpy investments that let them specialize in one business. I test whether households exit activities after a government program creates random variation in the availability of credit. I find no evidence that credit increases specialization. The confidence intervals let me rule out that an increase in credit causes as much specialization as an equal reduction in risk. These results suggest that an intervention meant to increase specialization would be most effective if it insured the poor against risk.

Existing work links risk to under-specialization but lacks the exogenous variation needed to show that risk causes under-specialization. [1] Morduch (1990) shows that households more vulnerable to income shocks tend to diversify their crops, and Bandyopadhyay and Skoufias (2012) find that households in areas with riskier weather tend to have spouses with different occupations. Reardon et al. (2007) study both risk and other factors that might drive farm households into nonfarm activities. But since vulnerability and weather risk are not exogenous, households who endure these problems may endure other problems unrelated to the riskiness of their income. If these other problems also cause under-specialization then estimates of the effect of risk will be biased.

More recent work, on the other hand, uses exogenous variation but does not study the effect of risk on specialization. Cole, Giné, and Vickery (2013) run a field experiment to show that weather insurance causes household to shift production towards riskier crops, and Emerick et al. (2014) find similar results in an experiment that distributes drought resistant seeds. These studies focus mainly on the intensive margin, whereas my aim is to test the two theories on the extensive margin. Adhvaryu, Kala, and Nyshadham (2013) study whether households expand their number of activities in response to shocks, but entering activities in response to shocks is not the same as entering activities in anticipation of risk. The first is a way to cope with a bad outcome whereas the second is a way to insure against that outcome.

The paper most similar to mind is Karlan et al. (2014), which tests whether providing insurance and loans causes households to shift out of non-farm work. Whereas they study the response to a one-time intervention that re-

[1] Many more papers study how imperfect insurance drives households to make other inefficient choices. Those most relevant to this paper study whether farmers with riskier profits marry their daughters to men in different occupations (Rosenzweig and Stark, 1989); choose safer but less profitable bundles of investments (Rosenzweig and Binswanger, 1993; Bliss and Stern, 1982); or delay the planting of their crops (Walker and Ryan, 1990, p. 256).

duces risk, I study the response of farmers to monthly changes in a form of risk they face every day. I also study a risk that evolves rapidly over time. Price volatility changes from month to month. By showing that households respond even to monthly changes in volatility, my results suggest households are deeply aware of the risks they face and respond to them exactly as theory predicts. By contrast, I rule out that access to credit causes a response of similar magnitude.

By testing the theory of lumpy investment alongside the theory of risky income I assess which theory has more merit. This is not to say that imperfect insurance and imperfect credit are the only reasons or even the most important reasons that poor households under-specialize—indeed, the results suggest much of the variation remains unexplained. My tests aim only to reveal whether either imperfection matters, and which matters more.

2 Theory: A Model of Risky Income

2.1 Deriving Tests for the Theory

Each household has one primary economic activity and may enter any number of side activities. These activities could be new businesses or new crops or new jobs. The household pays a fixed cost for each side activity (where, in the case of a job, this cost might be the time spent finding the job). For simplicity I model the cost as a literal cash payment, but it might be more realistic to think of it as the opportunity cost of whatever labor is wasted while switching between activities. The household allocates one unit of labor between all activities. Labor produces a constant return, and the household does not know the return to any activity until after it has made its choices.

The household must first choose the number of side activities. Then it chooses the allocations of labor. Then the returns to the side activities are realized. Finally, the household learns the return to its primary activity and consumes.

Suppose for simplicity that the household has constant absolute risk-averse preferences. The household solves

$$\max_{M,L_p,\{L_{s,m}\}} \mathbb{E}[-e^{-\alpha C}]$$

subject to

$$C = Y = w_p L_p + \sum_{m=1}^{M} w_{s,m} L_{s,m} - MF$$
$$L_p + \sum_m L_{s,m} = 1$$

where α is the coefficient of absolute risk aversion, $M \geq 0$ is the number of side activities, and L_p and $\{L_{s,m}\}_{m=1}^{M}$ are the labor allocated to the primary and each side activity. The household consumes its revenue, which is the sum of revenue from primary (p) and side (s) activities minus fixed costs. The primary and side activities yield returns w_p and $\{w_{s,m}\}_{m \in M}$, which are independent normal random variables with $w_p \sim N(\bar{w}_p, \sigma_p^2)$ and $w_{s,m} \sim N(\bar{w}_s, \sigma_s^2)$ for each m. [2] Assume the side activities yield weakly lower expected returns: $\bar{w}^p \geq \bar{w}^s$. Also assume the average premium to the primary activity, $\bar{w}_+ = \bar{w}_p - \bar{w}_s$, is not too large: $\bar{w}_+ < \alpha\sigma_p^2$. If this assumption fails the household will specialize despite the risk. If $\bar{w}^p = \bar{w}^s$ the household is no better at the primary activity than any other, but even then it is efficient for the household to specialize.

I make many simplifying assumptions about functional forms, but the important results rest on four crucial assumptions. First, the household is risk-averse. Second, the household cannot perfectly smooth its consumption through insurance or savings. (To sharpen the model's predictions I assume

[2] If the returns to side and primary activities were not independent, the properties of normal random variables let me write the returns to each side activity as $w_{s,m} = \rho^m w_p + \xi^m$ for some correlation coefficient ρ^m and an independent error ξ^m. If I then re-label variables accordingly, the qualitative results should hold as long as $\rho^M < 1$.

the household has no insurance or savings.) Third, the returns to side activities are not perfectly correlated with returns to the primary activity. Fourth, each activity has (locally) increasing returns. The first two assumptions force the household to insure itself against risk without using financial markets. The third assumption makes under-specialization a form of insurance. The fourth assumption makes under-specialization costly.

To get the intuition of the model, consider the simple case where the household either specializes ($M = 0$) or has one side activity ($M = 1$). The household chooses between two "bundles" of average consumption $\bar{C}$ and variance of consumption V:

	$\boldsymbol{M = 0}$	$\boldsymbol{M = 1}$
$\boldsymbol{\bar{C}}$	$\bar{w}_p$	$\bar{w}_p - \bar{w}_+(1 - L_p) - F$
$\boldsymbol{V}$	σ_p^2	$(L_p)^2\sigma_p^2 + (1 - L_p)^2\sigma_s^2$

Since $L^p < 1, \bar{w}_+ > 0$ and $F > 0$ the household can lower the variance of its consumption by entering a side activity if it accepts a lower expected consumption.

Suppose the household enters a side activity and must now choose how much labor to shift from the primary activity. Since consumption is a normal random variable, expected utility is (the negative of) a log normal random variable. The household now solves

$$\max_{L^p} -e^{-\alpha\bar{C} + \frac{\alpha^2}{2}V}.$$

The first-order condition is

$$
\begin{aligned}
& & 0 &= -e^{-\alpha\bar{C} + \frac{\alpha^2}{2}V} \cdot \left(-\alpha\frac{\partial\bar{C}}{\partial L_p} + \frac{\alpha^2}{2}\frac{\partial V}{\partial L_p}\right) \\
&\Rightarrow & 0 &= -\bar{w}_+ + \alpha L_p\sigma_p^2 - \alpha(1 - L_p)\sigma_s^2 \\
&\Rightarrow & L_p &= \frac{\alpha\sigma_s^2 + \bar{w}_+}{\alpha\left(\sigma_p^2 + \sigma_s^2\right)}
\end{aligned}
$$

To derive predictions about aggregate statistics, suppose the fixed cost of entering the side activity varies across households because some find it easier to enter activities. For example, two rice farmers might differ only in how closely they live to a construction site where they can find part-time work. For simplicity suppose $F \sim U[0, \mathcal{F}]$ for some upper-bound $\mathcal{F}$.

For any amount of risk there is a household whose fixed cost makes it indifferent between zero and one side activity. Call that household's fixed-cost $\bar{F}_0$. Let $C(M)$ and $V(M)$ be the mean and variance of consumption as functions of the number of side activities. Then $\bar{F}_0$ is defined as the fixed cost that makes this equation hold:

$$
\begin{aligned}
& & -e^{-\alpha\bar{C}(0)+\frac{\alpha^2}{2}V(0)} = -e^{-\alpha\bar{C}(1)+\frac{\alpha^2}{2}V(1)} \\
\Rightarrow & & -\alpha\bar{C}(0) + \frac{\alpha^2}{2}V(0) = -\alpha\bar{C}(1) + \frac{\alpha^2}{2}V(1) \\
\Rightarrow & & \frac{\alpha}{2}[V(0) - V(1)] = \bar{C}(0) - \bar{C}(1)
\end{aligned}
$$

Substitute the expressions from the table above and from the optimal labor allocation:

$$
\bar{F}_0 = \frac{\left(\alpha\sigma_p^2 - \bar{w}_+\right)^2}{2\alpha\left(\sigma_p^2 + \sigma_s^2\right)}
$$

Households who pay fixed costs above the threshold $\bar{F}_0$ will specialize while those below enter a side activity. The threshold rises with the variance of the primary activity σ_p^2, and Figure 1 shows the effect on the number of households with a side activity. When their primary activity becomes riskier, households are willing to pay a bigger fixed cost to make their revenue less risky. The threshold $\bar{F}_0$ rises, and the mass of households with fixed costs between the old and new thresholds enter a side activity. The change in the average number of activities in the sample is

Figure 1
Intuition of the Simplified Case

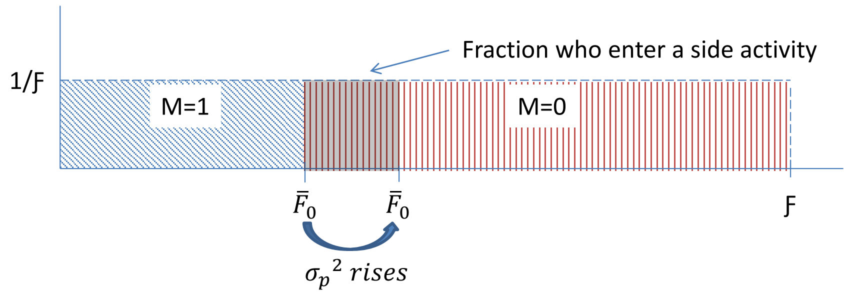

Note: M is the number of side activities; $\bar{F}_0$ the threshold fixed cost for moving from zero to one side activity; σ_p^2 is the variance of the primary economic activity. A rise in the variance raises the threshold fixed cost, which represents the amount households are willing to pay for insurance. In response the highlighted mass of households switches from specialization to having a side activity.

$$\frac{\partial \mathbb{E}_F[M]}{\partial \sigma_p^2} = \frac{\partial \mathbb{E}_F[M]}{\partial \bar{F}_0} \cdot \frac{\partial \bar{F}_0}{\partial \sigma_p^2} = \frac{1}{\mathcal{F}} \cdot \frac{\sigma_p \left(\alpha \sigma_p^2 - \bar{w}_+\right) \left(\alpha \sigma_s^2 \left(\sigma_p^2 + 2\sigma_s^2\right) + \bar{w}_+ \sigma_s^2\right)}{\alpha \sigma_s^2 \left(\sigma_p^2 + \sigma_s^2\right)^2} > 0.$$

(Recall that, by assumption, $\alpha\sigma_p^2 - w_+ > 0$). Similarly we can derive the change in the average number of activities when the average return to the primary activity rises. Since a rise in the expected return makes under-specialization more costly, the threshold will fall and the average number of activities will fall:

$$\frac{\partial \mathbb{E}_F[M]}{\partial \bar{w}_p} = -\frac{\alpha \sigma_p^2 - \bar{w}_+}{\alpha \left(\sigma_p^2 + \sigma_s^2\right)} \cdot \frac{1}{\mathcal{F}} < 0$$

The intuition of the case where $M \in \{0, 1\}$ holds for any number of activities $M \in \{0, 1, 2, \ldots\}$, and the simple but tedious proof is left for Appendix A.1.

On average each additional activity will lower total revenue. However, the empirical approach of Section 3 studies rice farmers who expect but have not yet collected a harvest. These farmers do not have their total revenue, ruling out any test based on total revenue. I must instead derive the model's predictions about what under-specialization does to revenue from side activities; that is, what happens to the rice farmer's revenue from growing cassava if he starts baking bread.

Consider the revenue of the household just before it gets the output from its primary activity. Its revenue at this stage is simply the revenue from its side activities:

$$y_s = \sum_{m=1}^{M} w_{s,m} L_{s,m} - MF$$

For simplicity treat the number of activities M as continuous. [3] Holding a household's cost of additional activities fixed, a small increase in the number of activities changes side revenue on average by

$$\begin{aligned} \mathbb{E}_F\left[\frac{\partial y_s}{\partial M}\right] &= \mathbb{E}_F\left[\mathbb{E}_{w_p,\{w_{s,m}\}}\left[\frac{\partial y_s}{\partial M} \mid F\right]\right] \\ &= \mathbb{E}_F\left[\frac{\partial}{\partial M}\left[-MF + (1-L_p)\bar{w}_s\right]\right] \\ &= -\mathbb{E}[F] + \left(-\bar{w}_s \frac{\partial L_p}{\partial M}\right) \end{aligned}$$

The average change in side revenue, which corresponds to the instrumental variables coefficient estimated in Section 5.1, has two parts: the average fixed cost of a side activity and the effect on side revenue of shifting labor to the side activities. Since an all-else-equal increase in the number of activities makes the portfolio of side activities less risky, the household wants to shift labor away from its primary activity. Then $\frac{\partial L_p}{\partial M} < 0$ and the second term is positive. If large enough it will swamp the cost of under-specialization and make the derivative (and thus the instrumental variables estimate) positive.

[3] I could keep the number of activities discrete and compute the average conditional difference, but given that a regression coefficient is meant to capture an average marginal change the simplification seems justified.

To see why, suppose the household starts with no side activities and thus no revenue from side activities. If the variance of the primary activity rises sharply and the cost of entering a side activity is small, then the household will want to enter the side activity. Then revenue from side activities will have increased, and though the increase might be small compared to what the household loses from its primary activity, the coefficient I estimate will be positive. Thus a negative estimate is sufficient evidence that an additional activity (and thus under-specialization) is costly, but it is not necessary evidence. This argument ignores the direct effect that my instruments, the variance and the average returns, have on the labor allocation. But as I show in the proof in Appendix A, the direct effect only strengthens the result.

The model also makes a prediction about the ordinary least squares coefficient, which estimates the average effect of increasing the number of activities without holding their cost fixed. That is, it estimates the average total derivative

$$\begin{aligned}
\mathbb{E}\left[\frac{dy^s(M,F)}{dM}\right] &= \mathbb{E}\left[\frac{\partial y^s}{\partial M} + \frac{\partial y^s}{\partial F}\cdot\frac{\partial F}{\partial M}\right] \\
&= \mathbb{E}\left[\frac{\partial y^s}{\partial M}\right] + \mathbb{E}\left[\frac{\partial y^s}{\partial F}\cdot\mathbb{E}\left[\frac{\partial F}{\partial M}\mid M\right]\right] \\
&= \mathbb{E}\left[\frac{\partial y^s}{\partial M}\right] + \mathbb{E}\left[\frac{\partial y^s}{\partial F}\cdot\frac{\partial}{\partial M}\mathbb{E}\left[F\mid M\right]\right]
\end{aligned}$$

The term $\frac{\partial y^s}{\partial F}$ is clearly negative; a higher fixed cost will lower revenue. The term $\frac{\partial}{\partial M}\mathbb{E}\left[F\mid M\right]$ gives the selection bias. It captures the difference in fixed cost paid by households who select into many versus few activities. As Figure 2 illustrates, it is also negative. Since a household takes up a large number of activities only if it pays a small fixed cost, the number of activities is informative about their cost. This gives the final test of the model:

$$\begin{aligned}\beta_{OLS} &= \mathbb{E}\left[\frac{\partial y^s}{\partial M}\right] + \mathbb{E}\left[\frac{\partial y^s}{\partial F} \cdot \frac{\partial}{\partial M}\mathbb{E}\left[F \mid M\right]\right] \\ &> \mathbb{E}\left[\frac{\partial y^s}{\partial M}\right] \\ &= \beta_{IV}\end{aligned}$$

To summarize, the model gives four tests for the theory of risky income:

Test 1 (Risk) *Households enter activities when the returns to their primary activity get riskier.*

Test 2 (Return) *Households exit activities when the (expected) returns to their primary activity rise.*

Test 3 (Cost) *The average effect of more activities on revenue is negative only if under-specialization is costly.*

Test 4 (OLS Bias) *Compared to the IV estimate, the OLS estimate of the effect of more activities on side revenue is biased positively.*

2.2 Modeling and Measuring Expectations about Risk

To run these tests I must model farmers' expectations about the returns and volatility of the price of rice. Suppose the household makes its choices at the beginning of period t. It has not yet observed the price w_{pt} and must form its expectation $\bar{w}_{pt}$ using only information from the past. Suppose the monthly price follows the Autoregressive Conditional Heteroskedasticity (ARCH) model of Engle (1982) with one modification: I assume the level of the price follows a random walk. The assumption reduces the number of parameters I must

Figure 2
Why is OLS Upward-Biased?

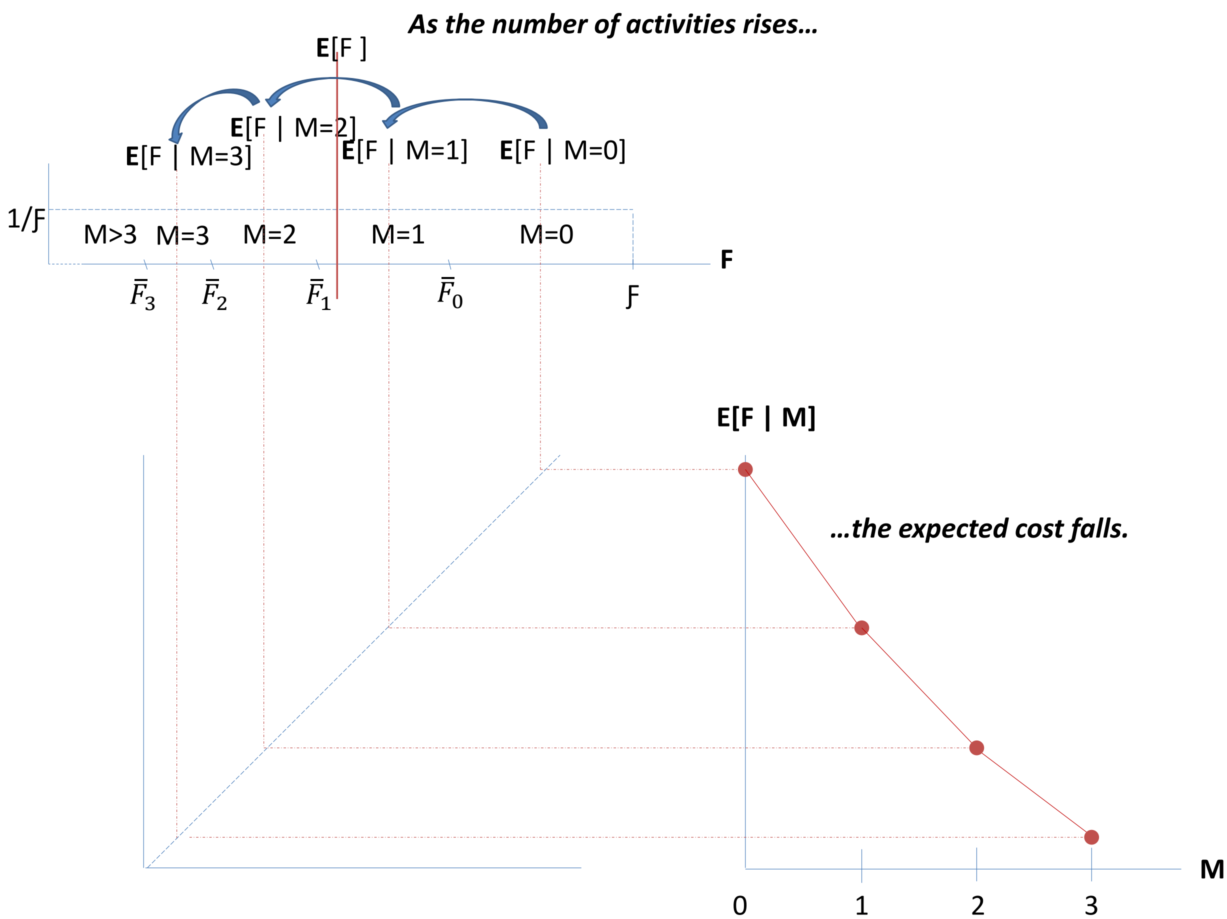

Note: M is the number of side activities; $\bar{F}_m$ is the threshold fixed cost below which a household moves from m to $m+1$ side activities. A household only enters many activities if the fixed cost it must pay for each is low. Then the number of activities predicts a household's costs. The cost of these extra activities appears in the error term of a regression of side revenue on number of activities. Thus the coefficient on the number of activities is biased upwards.

estimate and, as I show below, matches the true series well. [4] Then

$$w_{pt} = w_{p,t-1} + \varepsilon_t$$
$$\varepsilon_t = z_t\sqrt{h_t}, \qquad z_t \sim N(0,1)$$
$$h_t = \tau_0 + \tau_1 \varepsilon_{t-1}^2.$$

where z_t is white noise. At the beginning of period t, the household expects a return of $\bar{w}_{pt} = \mathbb{E}[w_{pt}] = w_{p,t-1}$. The variance of the return is $\sigma_{pt}^2 = V(w_{pt}) = V(\varepsilon_t) = h_t = \tau_0 + \tau_1 \varepsilon_{t-1}^2$. I estimate the model using conditional maximum likelihood. [5] The predicted value $\hat{h}$ is a consistent estimate of the true conditional variance.

In practice I must make several simplifications when I use this measure. I cannot use the actual expected volatility of the price at harvest because the empirical design in Section 3 compares farmers expecting a harvest to non-farmers and farmers who do not expect a harvest. Since I cannot define the volatility at the time of harvest for non-farmers I must use the current volatility. This creates measurement error and may bias my estimates towards zero. I also measure volatility using the conditional standard deviation $\sqrt{h}$ rather than the conditional variance to make the coefficients on the volatility of the price and the expected price comparable.

Figure 3 plots the actual price of rice, the predicted mean, and the predicted standard deviation. Simple though it is, the random walk assumption makes very accurate predictions about the mean. A regression of price on its lag gives a coefficient of .995. The estimated equation for volatility is $\hat{h}_t = 53.3 + .39\varepsilon_{t-1}^2$. The red lines mark the start and end of the time period covered in the monthly panel data. The sample spans a time when prices are relatively

[4] I show in Online Appendix F that a regression of the price on its lag yields a coefficient of 0.995.

[5] The true distribution of z_t need not be normal; the (quasi) maximum likelihood estimator based on a normal distribution is still consistent.

Figure 3
Rice Price and Predicted Mean and Conditional Standard Deviation

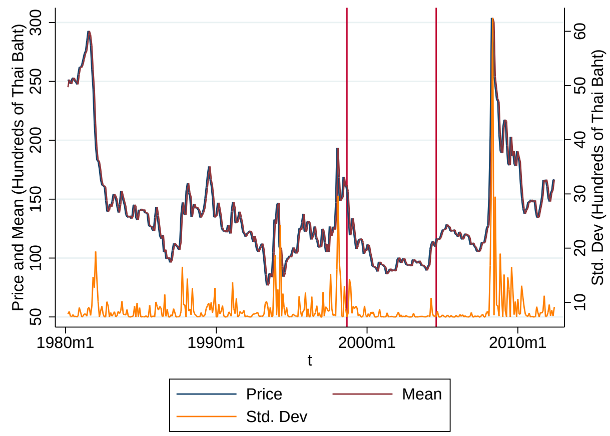

Note: I plot the actual rice price next to both the predicted rice price and the predicted volatility (square root of the predicted conditional variance) from the Autoregressive Conditional Heteroskedasticity (ARCH) model. The red lines mark the start and end of the panel data.

stable, ending well before the massive food price spike of 2008. [6]

It is important to my empirical design that these movements in the mean and volatility are irregular. If there is seasonality, some farmers might time their planting to ensure they harvest when the price is high or the volatility is low. There are strong theoretical reasons to expect prices cannot be regular—otherwise commoditiy traders could reap spectacular profits. I confirm this by running a regression of both the predicted mean and volatility on a set of month-of-year dummies. An F-test confirms that for both mean and volatility these dummies are insignificant. [7] There is no evidence of seasonality.

[6]The reader may worry if regressions on a regressor generated from a time series model are consistent. Pagan (1984) confirms that the ARCH predicted value (though not the residual) will give consistent estimates, and I have confirmed in monte carlo simulations that panel estimators are consistent as well.

[7]The p-values are 0.99 for the mean and 0.60 for the volatility.

3 Empirical Design: Implementing the Tests

3.1 Estimating Risk Response

Changes in the international price of rice—and the responses they evoke in Thai rice farmers—create exogenous variation in risk. I use this variation to test the model. Between planting and harvest the price can change drastically, and anecdotal evidence suggests farmers follow it closely in newspapers, radio broadcasts, and television reports. Since most of my sample grows at least some of the white rice and jasmine rice that make Thailand the world's biggest rice exporter, the international price matters. [8] In Column 1 of Table 1 I report the correlation between the sample-wide average price farmers receive and the international price. The correlation is imperfect, likely because local prices depend on distribution costs and may be subsidized by the government. Nevertheless, the correlation is significant and large enough to make following the international price worth a farmer's time. If prices become more volatile, the farmer knows it and knows the value of her harvest has become riskier. [9]

A response to volatility need not be a response to risky income unless it

[8]As expected, I find in unreported regressions that farmers harvesting only sticky rice, which is not exported, have a lower response. The size of the difference is too large to be the all-else-equal effect of growing rice that will not be exported. As households who grow only sticky rice are unusual, their response may differ from that of other farmers for reasons beyond the type of rice.

[9]Though it would be ideal to test whether volatility as measured in Section 2.2 is correlated with volatility in the local price, directly measuring that correlation is difficult. The household data have prices for only 62 months, with several gaps between those months (as compared to an uninterrupted 390 months for the data on the international rice price). Since an observation can only be used to estimate volatility if the current price and two of its lags are observed, the estimation sample would be further restricted to 47 (compared to 388 for the international price). It is difficult to make reliable estimates of volatility with such a small sample. But it is straightforward to show that a correlation in the price implies a correlation in the conditional volatility. Suppose the local price P_t^L equals a scaled version of the international price ρP_t^I. (According to Table 1, $\hat{\rho} = .33$.) If $P_t^I \sim N(\bar{P}_t, \sigma_t^2)$ then $P_t^L \sim N(\rho \bar{P}_t^L, \rho^2 \sigma_t^2)$. The local volatility, as defined in Section 2.2, is simply $\rho\sigma_t$, which is a rescaled version of the international volatility.

Table 1
Rice Prices and Sales

	(1) Avg. Transaction Price	(2) Rice Sold
Int. Rice Price	0.333**	
	(0.14)	
Rice Harvested		0.856***
		(0.01)
Constant	1.500	-2043.744***
	(1.53)	(70.44)
N	62	2126

Note: **Column 1 —** The dependent variable is the sample-wide average price of a kilogram of rice based on actual transactions, and the independent variable is the international price of rice in baht per kilogram. Not all survey rounds include any sales of rice—hence the number of observations is smaller than the number of survey rounds. Standard errors are robust to heteroskedasticity. **Column 2 —** The unit of observation is the household-month conditional on positive rice harvest.

comes from a specific group of farmers: those who harvest soon. Simply comparing the response of a rice farmer to someone who does not farm rice might just measure how rice farmers differ in their attitude to risk. By contrast, observing a household with rice planted but not yet harvested—a farmer expecting a harvest in the next three months—isolates the effect of risky income.

Farmers harvest rice roughly four months after planting and cannot hasten or delay the date. Harvesting too soon yields immature grains while harvesting too late risks losses to pests. The International Rice Research Institute states that "the ideal harvest time lies between 130 and 136 days after sowing for late" varieties and gives similarly narrow windows for other varieties (Gummert and Rickman, 2011). To be precise, the rice must be harvested when "Grain moisture is between 20-22% which is normally about 30 days after flowering" (International Rice Research Institute, 2015). Harvesting too early "will have many unfilled and immature grains which will break easily" whereas harvesting too late will cause "heavy losses…through shattering and

Figure 4
Time of Planting is a Strong Predictor of Time of Harvest

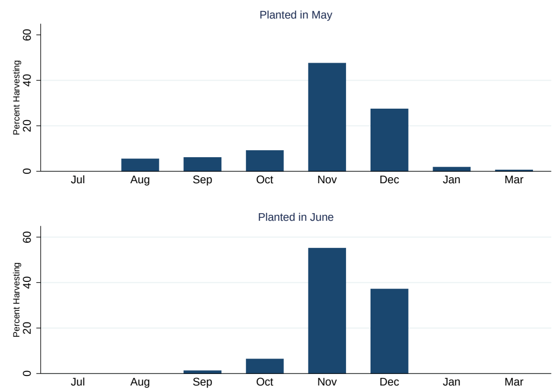

Note: The graph shows the fraction of households that report harvesting rice in each month of 1999, conditional on starting to plant in May or June. The sample is households who harvest only rice.

bird attacks." Leaving rice on the stalk to wait out low prices is not an option.

Figure 4 shows the month in which farmers start harvesting conditional on starting their planting in May or June, which is when most farmers plant. [10] Assuming it takes roughly a month to prepare and sow the land, most of these farmers should harvest between October and December. This is exactly what the graph shows.

With a growing period of over 3 months, it would also be difficult for farmers to strategically plant by predicting the volatility at harvest. My estimates from Section 2.2 suggest a 1 dollar rise in volatility at planting predicts only a 2 to 6 cent rise at harvest. And as described in Section 2.2 I find no evidence of seasonality in either the level or volatility of the rice price. A household cannot avoid volatility by aiming to harvest in a month that always has stable prices—no such month exists.

Although in principle a farmer might store rice after the harvest, in prac-

[10]The graph shows harvest dates for the year 1999, the first complete calendar year in the data. Since the public release of the survey does not identify which crops are being planted, I restrict the sample to farmers who harvest only rice.

Figure 5

Response to Conditional Volatility

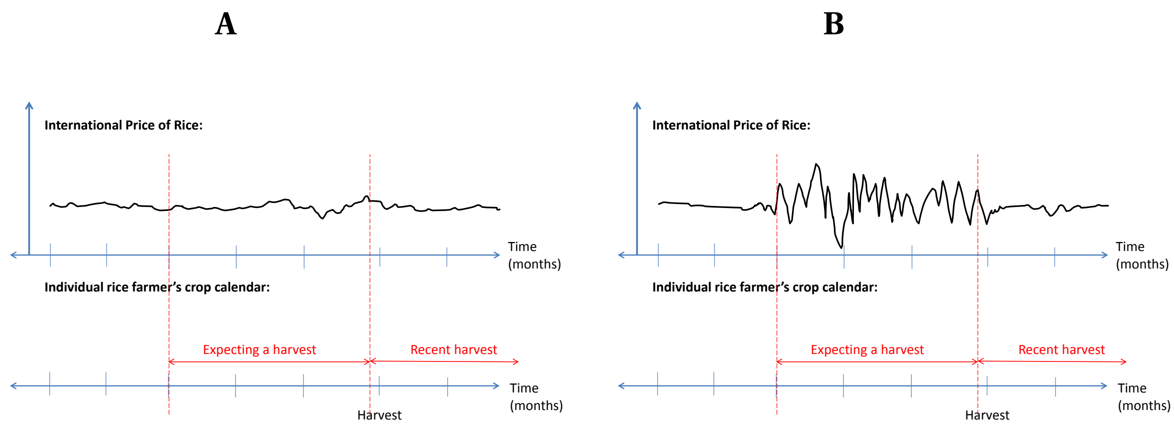

Note: Among rice farmers expecting a harvest I compare the response when rice prices are (**A**) stable to when they are (**B**) volatile. Since I use household fixed-effects I effectively compare each farmer to himself.

tice the farmers in my sample sell most of their rice right away. Colum 2 of Table 1 reports the correlation between how much rice a household sells and how much it harvests conditional on harvesting any during the month. For every kilogram harvested, 0.85 are sold. Why do farmers not simply store their rice when prices are low and wait for higher prices? Recall from Section 2.2 that the price follows a random walk. The farmer who stores her rice at harvest rationally expects to receive the same price next month and forever after. If she stores the rice, the price she gets is just as likely to fall as to rise. Given this belief it is rational to sell the rice immediately.

Could changes in the international price be driven by, say, bad weather in the villages I study? That is unlikely for two reasons. First, bad weather would have to hit the entire country and not just my sample. Second, though Thailand is a big exporter of rice, it is far from the biggest producer, which is China. Bad weather in China and India is more likely to drive prices than bad weather in Thailand. I confirm in Section 5.1 that controlling for rainfall shocks does not change the results.

It is critical to my identification strategy that the farmers who harvest dur-

Figure 6
An Impending Rice Harvest Requires Labor

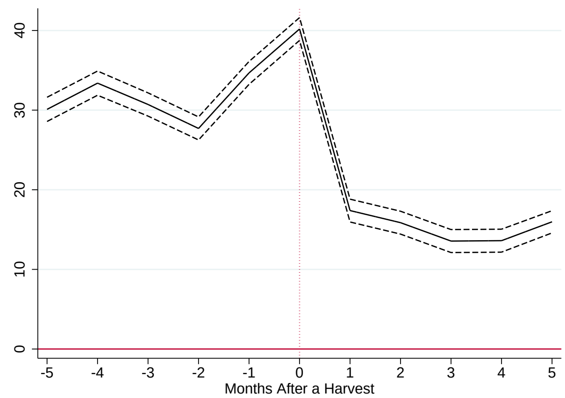

Note: The figure shows how many days the average household works in its fields in the months before and after a rice harvest. More precisely, I plot the coefficients of a regression of the number of days worked in the fields on dummies for periods before and after the harvest. The dashed lines cover 95 percent confidence intervals.

ing times of high volatility would respond similarly to those that farm at other times. Variation in the time of harvest comes from two sources: the time that farmers plant and the time the plant takes to mature. Most rice farmers in the same village plant within the same season, and most farmers in my sample harvest between October and December. We may worry that the farmers who plant at other times are larger farmers who plant many harvests per year. In fact the overwhelming majority of farmers grow only one harvest. During the calendar years 1999 to 2003 (the years for which I observe every month of the year), of the households that ever plant rice roughly 80 percent planted no more than one harvest in any of the five years. Excluding those that ever planted more than one harvest and controlling for whether the farmer harvested outside the typical range from October to December has little effect on my estimates of the effect of risk on under-specialization. [11]

Among farmers who do plant in a season, the variation in planting may arise because it makes sense for farmers to stagger their planting. Transplant-

[11] See Section 5.2.

ing seedlings takes a lot of hired or shared labor. If farmers were to plant and transplant at the same time they would find labor scarce, whereas they avoid this problem by staggering their plantings. This also ensures they will not all need harvest labor at the same time. Finally, even two farmers who plant at the same time may not harvest at the same time because the time it takes rice to mature is not constant. The ideal time to harvest is when “Grain moisture is between 20-22%,” which typically happens 140 days after planting but in practice will vary.

Since some of these decisions depend on the farmer, it is important I control for household fixed-effects. Even if farmers who plant in August are different from those that plant in June, household fixed effects ensure that I compare the households that plant in August to themselves when prices are more volatile. Given the fixed effects, the farmer faces shocks to volatility outside her control. By the time these shocks arrive, the time of harvest is also outside her control. The key identifying assumption is that, after controlling for fixed effects, a farmer whose harvest time exposes her to volatile prices responds to volatility exactly as a farmer who harvests at a different time would have responded had he been exposed to the same volatility.

To summarize, the farmers in my sample are too small to affect the international price and they cannot delay their harvest. Though some plant at different times than others, it is unlikely that the decision of when to plan is systematically correlated with volatility. Then after controlling for fixed effects, the responses of non-rice farmers, and the responses of rice farmers not expecting a harvest, any additional response must be caused by riskier income. [12] The regression I run will actually compare the farmer to herself at times when prices are volatile but she expects no harvest, and times when she expects a harvest but prices are not volatile. Figure 5 illustrates the specification.

When prices become volatile the farmer must decide whether to shift her

[12] I can drop the non-rice farmers from my regressions and still get consistent (albeit noisier) estimates. I confirm in Online Appendix D that estimating Equation 1 with only rice farmers does not change the results.

efforts away from maximizing the upcoming harvest. Figure 6 graphs the average household labor that rice farmers devote to their fields in the months before and after harvest. Bringing a rice crop to harvest requires ceaseless effort. Working as a laborer or planting cassava on the side detracts from rice farming. Like in the model, side activities detract from the primary activity.

Define $[Expecting\ Harvest]$ as a dummy for whether the household expects a rice harvest. (By definition, it is always zero for non-rice farmers.) I run the regression

$$\begin{aligned}[Activities]_{it} = {} & [FE]_i + \beta_M[Mean]_t + \beta_V[Volatility]_t \\ & + \beta_E[Expecting\ Harvest]_{it} + \beta_H[Had\ Harvest]_{it} \\ & + \beta_{RM}[Rice\ Farmer]_i \times [Mean]_t + \beta_{RV}[Rice\ Farmer]_i \times [Volatility]_t \\ & + \beta_{EM}[Expecting\ Harvest]_{it} \times [Mean]_t + \beta_{EV}[Expecting\ Harvest]_{it} \times [Volatility]_t \\ & + \beta_{HM}[Had\ Harvest]_{it} \times [Mean]_t + \beta_{HV}[Had\ Harvest]_{it} \times [Volatility]_t + \varepsilon_{it}.\end{aligned} \tag{1}$$

Aside from the responses of non-farmers and farmers who do not expect a harvest, I must also control for the responses of farmers who just had a harvest. Having had a harvest is negatively correlated with expecting a harvest and cannot be left in the error term. In some regressions I replace the main effects [Mean] and [Volatility] with time dummies—dummies for each of the 73 months I observe. Time dummies eliminate much of the variation in volatility but produce more conservative estimates. Finally, I also estimate specifications that include the interaction $t \times [Expecting\ Harvest]_{it}$, which allows for differential trends between farmers expecting and not expecting a harvest. Since the volatility is generated, I use a two-stage bootstrap for all inference in the results I report in Section 5.1. The details of the bootstrap are in Online Appendix C.

The coefficient β_{EV} on $[ExpectingHarvest] \times [Volatility]$ measures the differential response to volatility of a farmer who expects a harvest—that is, the additional response relative to the responses of non-farmers and farmers who do not expect a harvest. Since the number of activities is my measure of specialization, β_{EV} measures the causal effect of risk on under-specialization. Test 1 predicts it should be positive. The coefficient β_{EM} on $[ExpectingHarvest] \times [Mean]$ measures the response to higher average prices, and Test 2 predicts it should be negative.

3.2 The Costs of Under-Specialization

Risk may drive households into extra side activities, but do these extra activities have an opportunity cost? It is hard to imagine why else the household would diversify only when risk increases. If the extra activities were costless the household ought to have as many as possible. Test 3, however, suggests a direct approach: to check whether revenue from side activities falls as the farmer adds more activities.

Rises in volatility will cause farmers expecting a harvest to increase their number of activities, but by construction these farmers have not yet sold their harvest and collected their primary revenue. I cannot run any test on total revenue. Test 3 solves the problem by showing that if revenue from side activities falls when the household enters activities, then under-specialization is costly. Test 3 says that if a rice farmer's revenue from cassava falls when he starts baking bread, and the loss to cassava outweighs the gain from bread, then extra activities are costly.

Since total household revenue before the harvest does not include revenue from rice, movements in the rice price cannot affect revenue directly. They may cause the household to reallocate labor away from rice farming, but Test 3 already accounts for the change in labor (see Appendix A.2). Greater risk might cause a household to invest less in physical and human capital, or cause couples to change their decisions about child-bearing. But the effect of any change in investment or fertility will not appear for years, whereas my regressions measure changes within a three-month window. Moreover, the question is not whether risk causes the household to invest less in rice farming, as revenue from rice will not be included in its current income, but whether it causes less investment in other activities. There is no reason to expect a rice farmer to invest less in tailoring when the rice price gets riskier—if anything, he will invest more, which again is why Test 3 is necessary.

The key identifying assumption is that, aside from the effect on labor supply (see Section 2), the mean and volatility of the rice price has no differential effect on rice farmers expecting a harvest except through its effect on the number of activities. It is possible that changes in the mean and volatility of rice prices could be correlated with movements in the prices of other goods or have a general equilibrium effect on the economy. Such changes might bias the IV estimates if they have a differential effect on rice farmers expecting a harvest (that is, an effect beyond that on rice farmers that do not expect a harvest). I show in Section 5.2 and Online Appendix D that dropping revenues from crops and controlling for village wages does not change the main results. The identifying assumption is that there is no direct effect on any other source of revenue important only for farmers expecting a harvest.

I run the following first-stage regression:

$$\begin{aligned}[Activities]_{it} = {} & [FE]_i + \textstyle\sum_{m=2}^{73} \beta_{D,m}[Time\ Dummy]_t \\ & + \beta_E[Expecting\ Harvest]_{it} + \beta_H[Had\ Harvest]_{it} \\ & + \beta_{RM}[Rice\ Farmer]_i \times [Mean]_t + \beta_{RV}[Rice\ Farmer]_i \times [Volatility]_t \\ & + \beta_{EM}[Expecting\ Harvest]_{it} \times [Mean]_t + \beta_{EV}[Expecting\ Harvest]_{it} \times [Volatility]_t \\ & + \beta_{HM}[Had\ Harvest]_{it} \times [Mean]_t + \beta_{HV}[Had\ Harvest]_{it} \times [Volatility]_t + \varepsilon_{it}.\end{aligned} \tag{2}$$

The second-stage regression excludes $[ExpectingHarvest] \times [Mean]$ and

$[ExpectingHarvest] \times [Volatility]$ like so:

$$
\begin{aligned}
[Revenue]_{it} = [FE]_i + \gamma_A[\widehat{Activities}]_{it} + \textstyle\sum_{m=2}^{73} \gamma_{D,m}[Time\ Dummy]_t \qquad (3) \\
+ \gamma_E[Expecting\ Harvest]_{it} + \gamma_H[Had\ Harvest]_{it} \\
+ \gamma_{RM}[Rice\ Farmer]_i \times [Mean]_t + \gamma_{RV}[Rice\ Farmer]_i \times [Volatility]_t \\
+ \gamma_{HM}[Had\ Harvest]_{it} \times [Mean]_t + \gamma_{HV}[Had\ Harvest]_{it} \times [Volatility]_t + u_{it}.
\end{aligned}
$$

In some specifications I control for the interaction $t \times [Expecting\ Harvest]_{it}$, which allows for differential trends between farmers expecting and not expecting a harvest.

Test 3 states that if γ_A is negative then under-specialization is costly. The final test, Test 4, predicts the coefficient on $[Activities]$ in the simple OLS regression

$$
[Revenue]_{it} = \kappa_A[Activities]_{it} + \sum_{m=2}^{73} \kappa_{D,m}[Time\ Dummy]_t + \varepsilon_{it} \qquad (4)
$$

should be biased upward relative to the IV regression. In practice I estimate this equation both with and without household fixed effects. [13]

4 Data

I build my sample using annual and monthly surveys collected by the Townsend Thai Project. In May of 1997 the Project surveyed over two thousand rural households in four provinces. The annual survey followed the households from one-third of the baseline districts up through 2010 (Townsend et al., 1997).

[13] Perhaps simple OLS is not the true empirical version of Test 4, but rather OLS that controls for everything in Equation 3 without instrumenting for the number of activities. Running this alternative regression does not change the outcome of the test.

The monthly survey followed the baseline households plus several new additions from four of the remaining districts (Townsend, 2012). The monthly survey records changes in household income, crop conditions, and many other features of the household. I combine the survey with the monthly international price of rice from January 1980 to June 2012, taken from the IMF's commodity price dataset. [14]

I use the monthly data to test the theory of risky income. My final sample contains all 743 households that responded to at least two of the seventy-two monthly rounds the project has released. Table 2 summarizes the sample characteristics. I observe the average household for 65 months, but have the full five years of data for over three-quarters of households. I mark a household to be a rice farmer if it harvests rice at any point in the sample. I mark a household as expecting a harvest if it harvests rice in the next three months; I mark it as having had a recent harvest if it harvested rice in the current month or the previous three months. [15] Table 2 shows that households expected a harvest one-fifth of the time.

The survey asks each household about its economic activities one-by-one. The interviewer walks through each of several possible activities, asking the household if it earned revenue from the activity, and if so how much. I define the number of economic activities as the sum of the number of "large" businesses, crop-plots cultivated, types of livestock raised, number of jobs held by all members, number of miscellaneous or small businesses, and an indicator for whether the household engages in aquaculture (raising fish or shrimp). I define household revenue as the sum of revenue from each economic activity. [16] I define total consumption as total weekly and monthly household expen-

[14] I treat a household-month surveyed in the first half of the month as though observed in the previous month when I merge with time series data and define time dummies. Since the rice price and consumer price index are monthly averages, my convention best matches the survey response period to the horizons of the aggregate prices.

[15] Some fraction of households claim to sell rice during months when they still expect a harvest. In Appendix F I show that dropping these observations does not change the results.

[16] Each planting of any crop, as separated by space or time, is a crop-plot. A farmer who

diture. Net transfers, which I use to classify households as insured in Section 5.2, are the total incoming transfers minus total outgoing transfers. I deflate revenue, consumption, and transfers to be in May 2007 Thai baht. [17] Despite its benefits the dataset has some limitations. The Townsend Thai Project has released only part of the monthly survey. I do not observe how much land or wealth a household owns; I do not observe all of its farm expenses; and I cannot link the monthly survey to the baseline survey collected in 1997. [18]

Table 2 shows that the average monthly revenue is 620 U.S. dollars per month at May 2007 exchange rates. This figure is skewed upward because revenue is bounded below by zero but spikes during rice harvests; hence the high standard deviation. Consumption is less seasonal and the mean of 194 dollars is less skewed.

The monthly data do not contain the information needed to test the theory of lumpy investments; for that I turn to the annual panel. In addition to the four provinces and roughly 1000 households followed from baseline, the project added two more provinces and roughly 500 more households several years into the survey (both from the new provinces and from the original villages to counter attrition). My final annual sample for the lumpy investment tests is 1502 households. I construct the number of activities as closely as possible to my monthly measure: the sum of the number of large businesses, crop-plots, jobs, herds, an indicator for aquaculture, and a subset of the miscella-

---

plants rice and cassava side-by-side and a farmer who plants two fields of rice that are physically separate both have two crop-plots. The assumption is that even planting the same crop on a new field comes with a fixed cost—that of traveling to the plot or calculating a different mix of fertilizer. This definition lets me detect when a household that already has a rice crop in the ground starts growing additional rice or cassava on a new plot of land. Though some households may have more fragmented land (and thus more crop-plots) this fixed difference will be captured by the fixed effects.

[17] For more details on how I construct the variables, see Appendix B.

[18] Since I do not observe expenses I cannot compute income, which would also be of interest. But since the short-run costs of under-specialization will likely be the opportunity cost of wasted time, revenue should capture most of the useful variation.

Table 2
Descriptives of the Monthly Sample

Household-Month Mean and Standard Deviation				Fraction of Households or Household-Months	
Number of activities:	4.6	Revenue:	21352.8	Rice Farmers:	0.48
	(3.3)		(79854.7)		
				Of Whom	
Household size:	5.3	Consumption:	6692.2	Fraction of time expecting harvest:	0.23
	(2.4)		(24449.2)		
Total Labor:	80.0	Net Transfers In:	667.0	Fraction of time just had harvest:	0.31
	(75.6)		(35274.8)		
Households:	743	Avg. Obs/HH:	65.0	Observations:	48329

neous income sources. [19] The annual average of 4.6 activities is almost identical to the monthly average in Table 2, but it varies less because the annual measure wipes out within-year variation in activities. Though this sample is technically different from the monthly sample, by the design of the survey it is nearly identical in location and characteristics. The main difference is that it contains more households, which if anything means my tests of the theory of lumpy investment should be more likely to yield significant results.

The histogram in Figure 7, which shows the distribution of the number of activities in an arbitrary month, confirms that households in Thailand have many economic activities. Rice farmers are particularly under-specialized. Figure 8 graphs the top seven spontaneous responses to "What did your household do in the worst year [for income] of the last five to get by?" The most popular response was to take on an extra occupation, followed by working harder

[19] Miscellaneous income sources in the annual survey often include remittances and other sources that do not meet my definition of economic activities (namely, revenue generating activities that require labor). I filter these unwanted sources using regular expressions on the textual descriptions of sources. The 1999 survey unfortunately does not contain textual descriptions, but the year dummies in the annual regressions should account for any 1999-specific measurement error.

Figure 7
Number of Economic Activities, Rice Farmers and Non-Rice Farmers

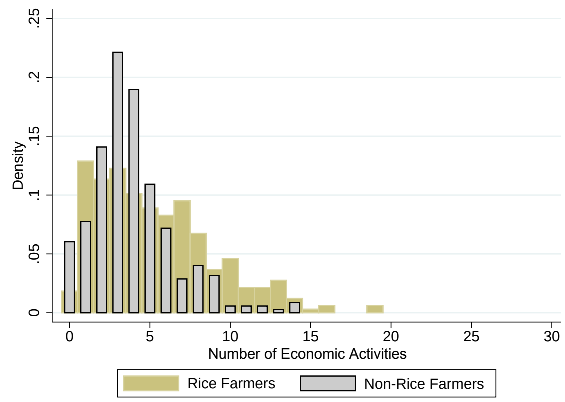

Note: The histogram shows the fraction of households with any number of economic activities in an arbitrary month. Rice farmers are more likely to have many activities.

Figure 8
Household Response to Negative Income Shocks

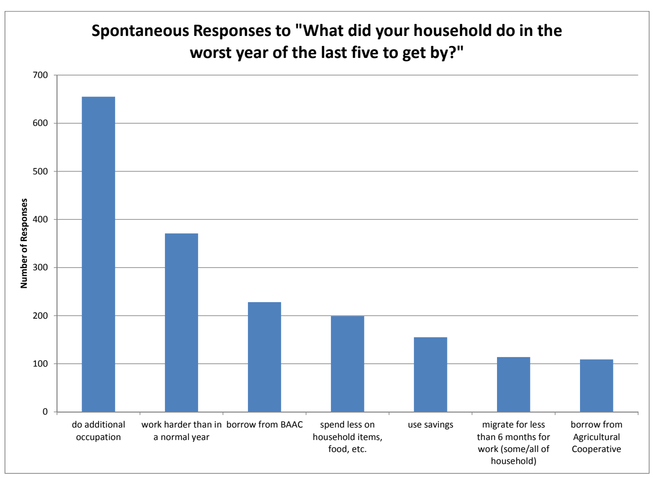

Note: The 1997 round of the Townsend Thai annual survey asks households how they coped during the the worst income year of the last five. They first gave spontaneous responses, which the project classified into categories. The graph reports the frequencies of the seven most popular responses. Many households work more or spend less to absorb income shocks rather than borrowing or using savings.

Figure 9
What Activities do Households Enter and Exit Most Often?

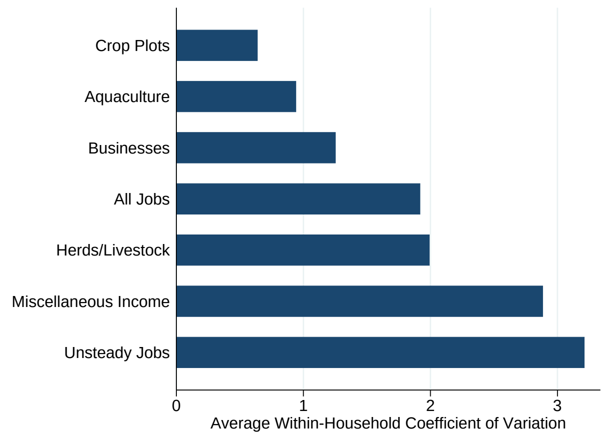

Note: I graph the average within-household coefficient of variation in the number of each type of activity. This shows which activities the household takes up and drops most readily. "Aquaculture" is an indicator for being involved in fish or shrimp farming. "All jobs" refers to all wage or salaried work, and "Unsteady Jobs" to work that lasts for no longer than five months.

than usual. These responses do not necessarily mean that households avoid risk through under-specialization, only that they cope with shocks through under-specialization. But if households must smooth their consumption by working harder, then they must have no better option. Borrowing money is only the third most popular response and using savings only the fifth. The fourth most popular response is to consume less, meaning many households lack even second-rate insurance.

What sorts of activities can households use to weather these shocks? For each of several categories, Figure 9 graphs the average within-household coefficient of variation in the number of activities. The greatest variation is in "Unsteady jobs," which I define as wage work that lasts for five months or less. Households transition in and out of these jobs often. In part this is mechanical—by definition an unsteady job must end. Moreover, that these jobs are entered and exited many times does not prove they are used to hedge against shocks.

Figure 10
Correlation Between Monthly Revenue and Consumption

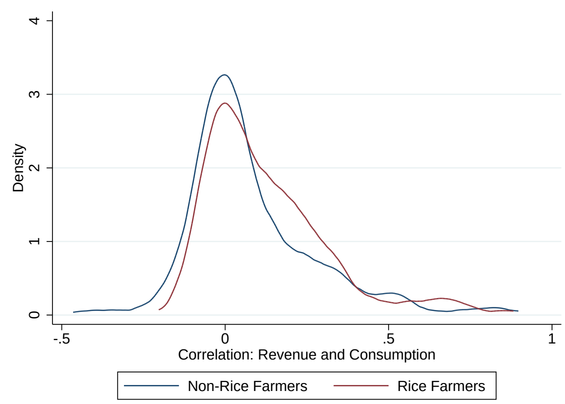

Note: For each household I compute the monthly correlation between total consumption and total revenue. I plot the density of the correlation for rice farmers versus non-rice farmers. Perfect insurance (whether self-insurance or otherwise) implies zero correlation. Almost all households have a positive correlation, meaning they consume less when their revenue falls.

But it does show that households can move in and out of these jobs easily. The second most variable category is miscellaneous activities, which are similarly easy to pick up and drop. Some examples include collecting "bamboo shoots to sell," "dress making," and "basket weaving."

Figure 10 shows the correlation between revenue and consumption, which is direct evidence that these extra activities do not perfectly insure against risk. I compute the correlation between monthly revenue and consumption expenditure for each household over however many months I observe it (72 months for the majority). If households are equally risk-averse, perfect insurance implies consumption should be uncorrelated with current revenue; indeed, consumption should be constant. [20] A household without perfect insurance cuts consumption when revenue falls, making the correlation positive. A higher correlation is evidence of less insurance. The figure plots the density of the

[20]Mazzocco and Saini (2012) show that when some households are less risk-averse, these households may insure the more risk-averse households. Then there would be a positive correlation between the consumption and revenue of the insurers.

correlation among rice farmers and non-rice farmers. Since zero is modal it appears many households do have near-perfect insurance, but many more do not. The distribution is heavily skewed towards less insurance with rice farmers particularly uninsured. [21] Some households have a negative correlation because of sampling error: the true correlation might be zero, but my estimate fluctuates around the truth and lands below zero for some households.

5 Risk and Under-Specialization Results

5.1 Main Results

Table 3 reports the results of the four tests derived in Section 2 and implemented as described in Section 3. Column 1 estimates (1) using the baseline specification. Column 2 controls for time dummies (one for each of the 73 months of data), and Column 3 additionally controls for a differential time trend for farmers expecting a harvest. Columns 4 and 5 estimate (4) first without and then with household fixed effects, and Columns 6 and 7 estimate (3) first without and then with the differential trend. Aside from the ordinary least squares regressions reported in Column 4 and 5, all regressions use the generated measure of volatility. I calculate the p-values and confidence intervals of these regressions using a two-stage bootstrap. [22] The bootstrap, which I describe in detail in the online appendix, corrects for the generated volatility

[21] The result may seem at odds with the high degree of insurance Townsend (1994) finds, but recall his result is that household consumption moves only with village-level and not household-level income. Figure 10 does not control for village-level shocks because a household cares only about having stable consumption, not where instability comes from. The shock I use for identification in Section 3 is a village-level shock: the international price of rice. It is precisely the village's inability to hedge against the price that drives households to under-specialize.

[22] It is not clear how to bootstrap the F-statistic on the excluded instruments or the Hansen's J Statistic. However, I can simply replace the generated volatility with $|P_{t-1} - P_{t-2}|$ in the first stage. Since this is perfectly collinear with the generated measure it produces algebraically identical IV coefficients, but since it is not generated the standard F and J statistics are valid.

measure and within-household correlation in the error term across time. [23]

The model's first test—Test 1—states that greater risk causes entry into more activities. The variable $[Rice\ Farmer] \times [Volatility]$ controls for any differential response of rice farmers to price volatility that is unrelated to their income (that is, the effect of volatility on farmers who expect no harvest). Thus the effect of income risk on activities is isolated by the coefficient on $[ExpectingHarvest] \times [Volatility]$ in Column 1 of Table 3, and as predicted it is positive and significant. The model also predicts in Test 2 that higher expected returns to the primary activity (rice farming) should cause a decrease in activities. The coefficient on $[ExpectingHarvest] \times [Mean]$ confirms that higher returns have a negative and significant effect on the number of activities. Columns 2 and 3 verify that both results hold when I control for time fixed-effects and a differential trend. Though the estimate of the effect of risk on activities becomes smaller, it remains positive and significant.

Test 3 states that if the extra activities cause (side) revenue to fall, then the failure to specialize is costly. I implement the test by running using the regressions in Columns 2 and 3 as a first-stage regression for Equation 3 using $[ExpectingHarvest] \times [Mean]$ and $[ExpectingHarvest] \times [Volatility]$ as instruments for the number of activities. Columns 6 and 7 of Table 3 show that the two-stage least squares coefficient on $[Activities]$ is negative and significant, confirming that under-specialization is costly. Columns 4 and 5 report the results of the simple ordinary least squares regression of revenue on number of activities, first without and then with household fixed effects (Equation 4). Test 4 states that the ordinary least squares coefficient on $[Activities]$ should be biased positively relative to the two-stage least squares coefficient because the farmers who pay lower costs for additional activities are exactly those who select into more activities. The coefficient is biased so strongly the sign flips, making under-specialization appear efficient.

[23] I cluster by household rather than village because there is variation in harvest times within villages.

Table 3
Main Results

	(1) Activities	(2) Activities	(3) Activities	(4) Revenue	(5) Revenue	(6) Revenue	(7) Revenue
Activities				1851.26***	2870.85***	-13883.30**	-23476.68*
				[0.000]	[0.000]	[0.035]	[0.062]
Mean	-0.00*						
	[0.096]						
Volatility	-0.08***						
	[0.000]						
Rice Farmer							
- × Mean	0.01***	0.00	0.00			-128.97	-107.16
	[0.010]	[0.359]	[0.314]			[0.135]	[0.286]
- × Volatility	-0.20***	-0.10***	-0.10***			-272.90	-800.18
	[0.002]	[0.009]	[0.010]			[0.618]	[0.344]
Expecting Harvest							
- Main	1.82***	1.89***	1.29***			4993.52	4725.61
	[0.006]	[0.000]	[0.008]			[0.310]	[0.442]
- × Mean	-0.02***	-0.02***	-0.01***			(Exc. Inst.)	(Exc. Inst.)
	[0.000]	[0.000]	[0.000]				
- × Volatility	0.18***	0.05*	0.06**			(Exc. Inst.)	(Exc. Inst.)
	[0.001]	[0.089]	[0.040]				
Recent Harvest							
- Main	-0.76	-0.57	-0.56			-34753.04**	-38092.75*
	[0.475]	[0.303]	[0.292]			[0.034]	[0.054]
- × Mean	-0.03***	-0.01***	-0.01***			300.00	191.14
	[0.000]	[0.002]	[0.006]			[0.129]	[0.350]
- × Volatility	0.41***	0.17***	0.17***			234.48	1683.32
	[0.002]	[0.004]	[0.008]			[0.819]	[0.258]
Household Fixed-Effects	Yes	Yes	Yes	No	Yes	Yes	Yes
Time Fixed-Effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Differential Trend	No	No	Yes	No	No	No	Yes
F-Stat Exc. Inst.						13.60	7.64
Hansen's J Stat.						0.12	0.63
Households	743	743	743	743	743	743	743
Observations	48329	48329	48329	48329	48329	48329	48329

Note: These regressions run the four tests of the theory of risky income (see Section 2): Test 1 (Risk): risk increases the number of activities; Test 2 (Returns): higher returns decrease the number of activities; Test 3 (Cost): more activities may cause side revneue to fall; Test 4 (OLS Bias): OLS is biased upwards. Column 1 estimates Equation 1, Column 2 adds time dummies, and Column 3 adds differential trends. Columns 4 and 5 estimate Equation 4 first without and then with household fixed effects, Column 6 estimates Equation 3, and Column 7 controls for differential trends. The bracketed values are p-values. I compute the p-values in Columns 4 and 5 using asymptotic standard errors that cluster by household. I compute the p-values in all other columns using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C). The value of the F-statistic on the excluded instruments from the first stage meets common standards for strength. The value of the J-statistic for overidentification is much too small to reject the null of exogenous instruments.

What do the sizes of these coefficients mean? Since the average price volatility for all available months is 8.8, the regression in Column 1 implies a 10 percent rise in volatility causes the farmer to enter $.18/(1/8.8) * 10/100 = .16$ additional activities. A similar calculation shows that the more conservative estimates in Column 2 implies a 10 percent rise in volatility causes the farmer to enter .05 activities. Recall from Table 1, however, that the international price of rice is not perfectly correlated with the actual price the farmer receives. This may be because government price supports give the farmer some insurance. Regardless of the cause, since in Table 1 the international price has a regression coefficient of roughly 1/3, a one unit rise in the volatility of the international rice price predicts a 1/3 unit rise in the volatility of the price the farmer receives. We can adjust the earlier numbers by dividing by 1/3, yielding estimates of .48 and .16 for the baseline and conservative estimates. The baseline estimate suggests the household enters an additional activity when local prices become 21 percent more volatile.

The two-stage least squares estimate in Column 6 implies that in taking on this activity the household will forego over 13 thousand baht, or over 60 percent of its average monthly revenue. According to the model in Section 2 this estimate is actually biased upward (towards zero), suggesting the true cost is even higher than implied. But recall that the average household has a little over four activities at once, making an additional activity a very large increase. Further, if the cost of an activity varies across households the estimate is not the average cost. If there is an upper bound on the number of activities a household can juggle, then the households with fewer activities are those most likely to respond to the instruments. These are also the households for whom an extra activity is most costly. Then the estimate, which is the continuous equivalent of the local average treatment effect, might be higher than the average cost of a side activity. Thus the point estimates in Columns 6 and 7 should be interpreted with caution.

The responses of households who had a recent harvest bear some explanation. First, the coefficient on $[RecentHarvest] \times [Mean]$ is negative. Since the expected price after the harvest is correlated with the price received at harvest, the negative coefficient confirms Figure 8 and the results of Adhvaryu, Kala, and Nyshadham (2013), both of which say that households increase their number of activities in response to bad shocks. Finally, the positive and significant coefficient on $[RecentHarvest] \times [Volatility]$ seems puzzling, as risk should not matter after a household has had its harvest. There are two explanations for this. First, since the current volatility is correlated with past volatility, this may just reflect that the household faced risk before the harvest and took on extra activities. Since the household cannot drop the extra activities immediately after harvest—temporary jobs must be finished and small businesses must be wound down—the household may still have more activities than usual after harvest. The second possibility is that a high current volatility implies the price has moved drastically in the recent past. Since the current expected price ([Mean]) does not perfectly capture the price at harvest, a high volatility means it is more likely the household had a low price at harvest. Since households take on activities to recover from low prices, the coefficient on post-harvest volatility may be picking up the response to negative income shocks.

Finally, it may seem surprising that the coefficient on $[RiceFarmer] \times [Volatility]$ is negative, meaning farmers without harvests reduce their number of activities. But recall that farmers of rice are also consumers of rice. As shown by Turnovsky, Shalit, and Schmitz (1980), except under extreme assumptions volatility will actually increase consumer welfare because the indirect utility function is quasi-convex in prices. Moreover, when prices become unusually low a consumer of rice can buy enough to last many months. The household need not work in so many activities because what income it has will go further. But for a rice farmer who expects a harvest, income is itself risky, forcing the farmer to keep hedging. Recent empirical work confirms that price volatility may help consumers; for example, Bellemare (2015) finds that higher price volatility actually reduces social unrest.

5.2 Robustness and Specification Tests

This section runs several checks to verify the results are both both robust and not spurious. For brevity I report only the coefficients of interest—the full tables are in Online Appendix D.1.

Table 4 reports several new specifications used to confirm the results are not driven by something other than volatility. The theory in Section 2 assumes the total labor supplied by the household is fixed, but in truth the household may work less when the returns to its labor grow riskier. Alternatively, the household might send some members to work abroad or in Bangkok. Columns 1 and 2 show that the effects of higher volatility and higher returns on the number of activities remain unchanged when I control for the household's total labor and the number of household members. The coefficient on $[ExpectingHarvest]\times[Volatility]$ remains positive and significant in both the specification that controls for just time dummies (top panel) and the specification that controls for both time dummies and differential trends (bottom panel). Likewise, Column 7 shows that the effect of additional activities on revenue remains negative and significant. This is not surprising, as I find in unreported regressions that the effect of volatility on total labor is tiny and statistically insignificant, suggesting households really are splitting a fixed supply of labor across more activities.

If volatility is a proxy for weather shocks the coefficient on $[RecentHarvest]\times[Volatility]$ may measure a response to weather rather than a response to risk. I argue in Section 3.1 that local weather shocks are unlikely to move the international rice price, and thus unlikely to be correlated with the volatility of prices. Column 3 confirms that controlling for rainfall shocks does not change the results. I define the monthly rainfall shock as the percentage deviation in rainfall from the mean for that month over the period from 1985 to 1998. I control for both the direct effect and the interaction with $[ExpectingHarvest]$. The coefficient of interest remains positive and significant. Column 8 confirms that the second-stage result also remains unchanged. [24]

Columns 4 and 9 of Table 4 both answer a simple concern: should we believe Thai rice farmers use a model of autoregressive conditional heteroskedasticity to decide how to spend their time? The model only formalizes a simple intuition: when prices fluctuate, they are risky. Columns 4 and 9 confirm that using simpler measures of the mean and volatility—the current price and the absolute value of the change in the price since last month—does not much change the results.

If the volatility of the price is just a proxy for unexpected decreases in the price, then what I assume is a response to risk may in truth be a response to changes in the household's permanent income. If this story is true, then the household should respond more strongly to simple changes in the price than to my measure of volatility, which is proportional to only the absolute value of the change. Column 5 of Table 4 runs a regression that replaces my measure of volatility with the simple change in price. Households expecting a harvest do not respond to simple changes in the price.

The reader may also worry whether the expected price and the volatility are valid instruments for side revenue. If the price of rice and the price of corn, say, are correlated then the expected price is no longer a valid instrument for rice farmers who also grow corn. Column 6 of Table 4 verifies that the second stage results hold when I use a measure of revenue that excludes earnings from crops.

I measure "specialization" with the number of economic activities, but this measure may seem arbitrary. Table 5 tests whether two alternate measures of specialization respond to risk. The first measure is an indicator for whether anyone in the household holds an unsteady job, which I define as a job that

[24] I measure rainfall by province, as this is the finest geographic identifier released by the Townsend Thai project.

Table 4
Robustness

	(1) Activities (Labor)	(2) Activities (HH Size)	(3) Activities (Rainfall)	(4) Activities (Non-Gen)	(5) Activities (Price Change)	(6) Revenue (Exc. Crop Rev.)	(7) Revenue (Labor & Size)	(8) Revenue (Rainfall)	(9) Revenue (Non-Gen)
Activities						-11232.10**	-17060.61***	-14071.44**	-9305.78*
						[0.031]	[0.009]	[0.028]	[0.064]
Expecting Harvest									
- × Mean	-0.01***	-0.01***	-0.02***	-0.02***	-0.01***	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]				
- × Volatility	0.05*	0.05*	0.07*	0.02***	-0.00	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.094]	[0.088]	[0.058]	[0.003]	[0.881]				
Household Labor	0.01***						275.42***		
	[0.000]						[0.000]		
Household Size		0.13***					950.03		
		[0.000]					[0.312]		
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Differential Trend	No	No	No	No	No	No	No	No	No
F-Stat Exc. Inst.						13.60	13.15	13.73	20.66
Hansen's J Stat.						0.02	0.27	0.43	0.37
Households	743	743	743	743	743	743	743	743	743
Observations	48160	48164	48329	48329	48329	48329	48160	48329	48329

	(1) Activities (Labor)	(2) Activities (HH Size)	(3) Activities (Rainfall)	(4) Activities (Non-Gen)	(5) Activities (Price Change)	(6) Revenue (Exc. Crop Rev.)	(7) Revenue (Labor & Size)	(8) Revenue (Rainfall)	(9) Revenue (Non-Gen)
Activities						-17835.42**	-30862.50**	-20164.36**	-12865.05**
						[0.042]	[0.032]	[0.048]	[0.038]
Expecting Harvest									
- × Mean	-0.01***	-0.01***	-0.01***	-0.02***	-0.01***	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.006]	[0.000]	[0.002]	[0.000]	[0.002]				
- × Volatility	0.06**	0.06*	0.09**	0.03***	-0.00	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.026]	[0.052]	[0.016]	[0.001]	[0.738]				
Household Labor	0.01***						354.60***		
	[0.000]						[0.002]		
Household Size		0.13***					2126.60		
		[0.000]					[0.166]		
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Differential Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-Stat Exc. Inst.						7.64	6.71	8.45	15.22
Hansen's J Stat.						1.01	0.80	3.44	2.39
Households	743	743	743	743	743	743	743	743	743
Observations	48160	48164	48329	48329	48329	48329	48160	48329	48329

Note: Columns 1 and 2 verify that controlling for household labor and household size do not change the effect that the mean and volatility of rice prices have on the number of activities. Column 7 confirms that controlling for both does not change the effect of activities on revenue. Column 3 shows that controlling for province-level rainfall shocks does not change the results. Columns 4 and 9 verify that a simpler measure of volatility–the absolute change in rice prices–does not qualitatively change the results for either number of activities or revenue. Column 5 verifies that the effect is caused by volatility rather than the (non-absolute) change in prices, as would be so if the effect were driven by a change in permanent income rather than volatility. Column 6 confirms that even excluding income from crop-plots the revenue results holds. The bracketed values are p-values. I compute the p-values in Columns 4, 5, and 9 using asymptotic standard errors clustered by household. I compute all other p-values using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C).

Table 5
Check: Other Measures

	Have Unsteady Job			Non-Crop Activities		
	(1)	(2)	(3)	(4)	(5)	(6)
Expecting Harvest						
- $\times$ Mean	-0.00**	-0.00	-0.00	-0.01***	-0.00***	-0.00
	[0.020]	[0.150]	[0.482]	[0.000]	[0.001]	[0.218]
- $\times$ Volatility	0.02***	0.01	0.01*	0.06***	0.02	0.03
	[0.009]	[0.111]	[0.080]	[0.010]	[0.278]	[0.122]
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	No	Yes	Yes	No	Yes	Yes
Differential Trend	No	No	Yes	No	No	Yes
Households	743	743	743	743	743	743
Observations	48329	48329	48329	48329	48329	48329

Note: I define an “unsteady job” as one held for less than five months. Columns 1–3 look for effects of volatility on an indicator for whether anyone in the household had an unsteady job. I define “non-crop activities” as the total number of activities minus the number of crop-plots farmed. Columns 4–6 look for effects on the number of non-crop activities. I present all coefficients with p-values calculated using the two-stage bootstrap.

lasts for five months or fewer. If volatile prices drive a household into casual labor it is another sign that risk causes inefficient under-specialization. The second measure is the number of non-crop economic activities. Since these households farm rice it is a sign of under-specialization if they expand their activities beyond the fields.

Columns 1 and 4 show that farmers who expect a harvest are more likely to get unsteady jobs and will take on more non-crop activities when rice prices grow more volatile. After controlling for month fixed-effects and differential trends the estimates are no longer significant in all specifications. This may be because month fixed-effects absorb much of the variation in $[Expecting\ Harvest] \times [Volatility]$. Since I must adjust the standard errors for generated regressors the reduced variation makes it hard to find effects. But the regressions provide some suggestive evidence to support the main results.

I identify my effects by comparing the response of households who expect a harvest during times of high volatility to those who expect harvests at other times. Could I be inadvertently comparing most farmers to those that reap many harvests or harvest at unusual times? All of the specifications reported in Table 6 drop households that have more than one harvest in the years from 1999 to 2003 (these are the only years for which I observe all twelve months). Columns 3 and 4 also control for a dummy for whether the household got its harvest in the months from October to December, which is when the vast majority of households get their harvest. By controlling for this dummy I control for households that choose to harvest at unusual times. [25] The coefficient is much the same as estimated in the main specifications of Table 3.

Though several of my specifications control for differential trends among farmers expecting a harvest, we may worry that a simple linear trend does not properly control for existing trends. If correlated with volatility these trends would bias the estimated effect of volatility on activities. But if any such trend

[25] Note that the household fixed-effects already control for households that always harvest at unusual times. This dummy additionally controls for households that choose an unusual harvest time in some years but not others.

Table 6
Check: Time of Harvest

	Activities (1)	Activities (2)	Activities (3)	Activities (4)
Expecting Harvest				
- $\times$ Mean	-0.02***	-0.02***	-0.02***	-0.01***
	[0.000]	[0.000]	[0.000]	[0.000]
- $\times$ Volatility	0.07*	0.07*	0.07*	0.08**
	[0.050]	[0.052]	[0.060]	[0.044]
Harvest Oct-Dec			0.23***	0.24***
			[0.000]	[0.000]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trends	No	Yes	No	Yes
Households	659	659	659	659
Observations	43257	43257	43257	43257

Note: These regressions drop all households that reported more than one harvest at any time in the years from 1999 to 2003 (the years for which I observe all twelve months). Columns 2 and 4 control for an indicator for whether the household had its harvest in the range from October to December, which is when most farmers harvest. This effectively controls for households that harvest at unusual times. I present all coefficients with p-values calculated using the two-stage bootstrap.

Table 7
Check: Placebo Test

	(1) Activities (3 Month Lead)	(2) Activities (3 Month Lead)	(3) Activities (12 Month Lead)	(4) Activities (12 Month Lead)
Expecting Harvest				
- × Mean	-0.01***	-0.01*	-0.01***	-0.01*
	[0.000]	[0.056]	[0.000]	[0.066]
- × Future Volatility	0.00	-0.00	-0.02	0.01
	[0.982]	[0.876]	[0.736]	[0.858]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trends	No	Yes	No	Yes
Households	735	735	715	715
Observations	45840	45840	39384	39384

Note: I confirm that using future price volatility—volatility in three or twelve months—has no impact on farmers expecting a harvest. This suggests farmers are responding to the volatility of the price of their harvest and not just following a differential trend.

exists, we would expect *future* levels of volatility to affect the *current* behavior of farmers expecting a harvest. Table 7 reruns the tests using future volatility in place of current volatility. I estimate the effect of both a 3 month and 12 month lead in the level of volatility. These placebo tests show that future volatility has no effect on current behavior. The estimated coefficients are close to zero, flip signs across specifications, and have p-values close to 1. Aside from easing concerns about differential trends, the non-effect of the 12-month lead confirms that there is no evidence of seasonality in volatility.

Is what I measure really a response to risk? To answer this question I examine whether households with better insurance make a smaller response to changes in the volatility. In poor countries a household often relies on family and friends for support in hard times. [26] Figure 11 shows that the rice farmers

[26]Rosenzweig (1988) found that households structure themselves to ease income sharing. Townsend (1994) and more recently Munshi and Rosenzweig (2009) find village and caste networks provide insurance in India. Yang and Choi (2007) show that rural Filipino households

Figure 11
Households Receive More Transfers when Prices are Low at Harvest

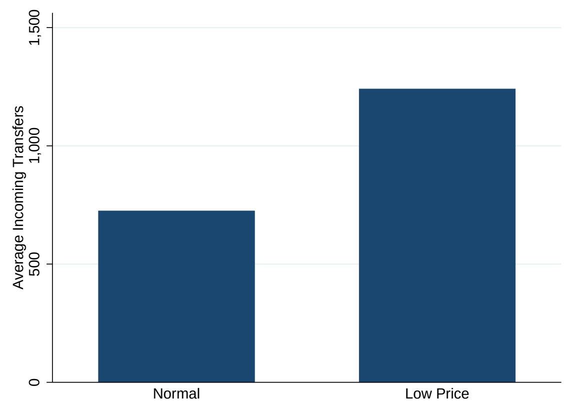

Note: The first bar depicts average incoming transfers for households harvesting rice when the international rice price is “normal”—above the bottom quartile of all prices I observe in the period covered by the monthly panel. The second bar depicts the average transfers when prices are “low”—in the bottom quartile. Rice farmers receive more money when the value of their harvest is low.

in my sample are no different. When the international price is low rice farmers tend to receive more transfers. I calculate for each household the monthly correlation between its net incoming transfers and its revenue, and call a household “insured” if that correlation is negative.

Columns 1 and 2 of Table 8 report the response of the uninsured, and Columns 3 and 4 the response of the insured. As expected, the response of the insured is smaller and insignificant. Since my measure of insurance is not exogenous I cannot rule out that households with insurance differ from uninsured households in ways that change how they might respond to volatility. Keeping that caveat in mind, the result is consistent with the theory of risky income.

Finally, I show several more robustness checks in the Online Appendix. One might worry that an aggregate shock to volatility might have general equilibrium effects on wages, which may be the true driver of behavior. In Ap-

who suffer bad rainfall shocks receive more remittances from overseas family.

Table 8
Check: Insurance

	(1) Activities (Uninsured)	(2) Activities (Uninsured)	(3) Activities (Insured)	(4) Activities (Insured)
Expecting Harvest				
- × Mean	-0.02***	-0.01***	-0.02***	-0.01**
	[0.000]	[0.002]	[0.000]	[0.036]
- × Volatility	0.09*	0.09*	0.04	0.05
	[0.067]	[0.072]	[0.277]	[0.130]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trends	No	Yes	No	Yes
Households	270	270	473	473
Observations	16933	16933	31396	31396

Note: I split the sample into households who receive transfers of income when their consumption is low ("insured") and those that do not ("uninsured"). I confirm that volatility has a larger effect on the number of activities among households that are uninsured.

pendix D.3 I show that controlling for wages does not change the results. Though general equilibrium effects no doubt influence the farmer's decision, this robustness check suggests the results are not driven by such effects. One might also worry that the triple-difference approach, by comparing rice farmers to non-rice farmers and farmers expecting a harvest to those that do not, is biased or misleading. Appendix D.2 shows that restricting the sample to rice farmers or to rice farmers expecting a harvest give similar results.

6 The Alternative Theory: Lumpy Investments

If "the poor cannot raise the capital they would need to run a business that would occupy them fully" (Banerjee and Duflo, 2007) then poor households cannot specialize. Suppose a man can learn to sew or bake but cannot can sew more than a few shirts or bake more than a few loaves unless he buys a sewing machine or an oven. Since he cannot afford either investment he cannot grow either business. To support his family he must sell both shirts and bread. This is the theory of lumpy investments. [27]

To test the theory I exploit a government program that produced quasi-experimental variation in the supply of credit. The theory predicts that households that get more credit should be better able to make the lumpy investments that let them specialize. The Million Baht Program gave one million baht to a fund for public lending set up in every village in my sample. Kaboski

[27] The inability to invest may create another source of under-specialization: the need to take on extra jobs because one may only work so long at any single task. Suppose labor and capital are complements, and make it simple with an extreme example: perfect complementarity. Suppose an activity m produces revenue with production function $y^m = A^m \min[L, K]$, with $m = T, B$ for tailoring or baking. Suppose $A^T > A^B$ for some household. If the household's labor endowment is $\bar{L}$, it will specialize in tailoring with $K^* = \bar{L}$. But suppose increasing capital beyond $\tilde{K} < K^*$ requires a lumpy investment the household cannot afford. If the household specialized, it would be left with $\bar{L} - \tilde{K}$ units of unused labor. In other words, it would be idle. The alternative is to spend its remaining time baking, so its total revenue is $A^T \tilde{K} + A^B(\bar{L} - \tilde{K}) < A^T \bar{L}$.

and Townsend (2009, 2011), who are the first to exploit the program, argue that the boundaries of villages in Thailand are set by bureaucratic fiat rather than economic logic. The sizes of villages are effectively random. Since every village got the same amount of credit, the per-household increase in credit is also random, with smaller villages exogenously given more credit. Kaboski and Townsend (2011) confirm in their first table that small and large villages have parallel trends.

To exploit the program I must know the number of household in each village, which is not recorded in the monthly data. Instead I use the annual data, which follows a nearly identical population and should yield comparable results. [28] The effect of the program is measured by a dummy for the year of implementation interacted with the average injection of credit per household. I use both the level and log of the per-household injection (1 million/number of households). The theory predicts that the coefficients should be statistically significant and negative.

According to Table 9 the coefficients, even when significant, are positive. If anything credit causes households to enter more activities. This result is inconsistent with the theory of lumpy investment, but might be consistent with the model from Section 2. If risk is really what drives under-specialization and some households want more activities but cannot afford to pay the fixed-cost, giving them credit might let them enter more activities.

Is it possible that the true effect is negative but I lack the power to detect it? The bottom rows of Table 9 use the 95 percent confidence interval of each estimate to compute the largest plausible decrease in the number of activities caused by credit. The estimates are rescaled to show how many activities the household would exit if the supply of credit rose by 21 percent. Once rescaled the estimates are directly comparable to the back-of-the-envelope calculation made in Section 5.1.

[28] Since the annual data cover more households than the monthly data, if anything I stack the odds towards finding statistically significant evidence for the theory of lumpy investment.

Table 9
Testing the Theory of Lumpy Investments

	(1)	(2)	(3)	(4)
	Activities	Activities	Activities	Activities
	b/se	b/se	b/se	b/se
2001 X Credit/HH	10.073		12.112	
	(9.05)		(8.23)	
2002 X Credit/HH	14.475		8.517	
	(9.82)		(10.72)	
2001 X Log Credit		0.106		0.124*
		(0.08)		(0.07)
2002 X Log Credit		0.128*		0.077
		(0.07)		(0.09)
Villages	64	64	64	64
Households	1228	1228	706	706
Observations	13745	13745	9884	9884
Largest effect: 2001	-0.014	-0.009	-0.008	-0.002
Largest effect: 2002	-0.009	-0.000	-0.023	-0.010

Note: The regressions test the lumpy investment theory using the Million Baht Program. The coefficient on the interaction of the average per-household credit injection (one million divided by number of households) with the year of implementation (2001) estimates the effect of relaxed credit on number of activities. The measure of number of activities is similar as possible to that in the risk regressions. The alternative specification uses the log of the injection (one million divided by number of households). The first two columns use the largest possible sample of households while the last two use a balanced panel. The rows labeled “Largest effect” show the largest reduction in activities ruled out by the 95 percent confidence interval. This effect is scaled to show the effect of a 21 percent increase in credit. These numbers can be compared to those from Section 5.1, where I computed that a 21 percent decrease in risk would cause the household to exit 1 activity. All inference uses asymptotic standard errors clustered at the village. These standard errors are reported in parentheses.

According to that calculation, a 21 percent decrease in risk would cause the household to exit either 1 activity (using the baseline estimates) or 1/3 of an activity (using the conservative estimates). By contrast, a 21 percent increase credit has a much smaller effect. Table 9 suggests I can with 95 percent confidence rule out a decrease of more than 0.023 activities. At least in this context I can rule out that credit is a more important cause of under-specialization than risk.

Though I cannot rule out that the result would have been different in another context, this result is supported by other work that studies other places. Bianchi and Bobba (2012), for example, find that the Progresa conditional cash transfer helped households start businesses by mitigating risk rather than alleviating a lack of funds. Likewise, Karlan, Knight, and Udry (2012) find that giving small African firms extra capital does not cause them to grow.

That said, both my study and these others test only a limited form of the theory. The smallest villages received per-household credit injections of half the median income. If households need sewing machines the credit injection could pay for them. Since most micro credit charities believe small entrepreurs need small loans, finding no effect from a small rise in the supply of credit has implications for those programs. But it is possible that an economy only transitions to specialized jobs when a few entrepreneurs build big firms that can offer everyone else a single salaried job. Such a transformation is too big to be financed by the Million Baht Program or any other credit intervention studied in this literature.

7 Summary

I show that Thai rice farmers expecting a harvest increase their number of economic activities when confronted with more volatile prices. My estimates suggest a 21 percent rise in volatility causes a household to enter 1 extra activity. I use this exogenous change in the number of activities to verify that under-specialization reduces revenue. Finally, I test an alternative theory of under-specialization—that the poor run many small businesses because they cannot afford the lumpy investments needed to grow any one—and find no supporting evidence.

The pin-maker wastes time when he switches from straightening wires to cutting them, and I find evidence of this waste in rural Thailand. My results do not measure the talent wasted when the poor forego expertise in a single trade. This kind of under-specialization, which changes the structure of an economy, is a long-run cost that requires a long-run study.

Finally, I study only two possible causes of under-specialization. Imperfect markets for factor inputs, distortions caused by government policy, and even social convention might also drive households to under-specialize. I leave these questions to future research.

References

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham. 2013. "Booms, Busts, and Household Enterprise: Evidence from Coffee Farmers in Tanzania." .

Bandyopadhyay, Sushenjit and Emmanuel Skoufias. 2012. "Rainfall Variability, Occupational Choice, and Welfare in Rural Bangladesh." *Review of Economics of the Household* :1–46.

Banerjee, A.V. and E. Duflo. 2007. "The Economic Lives of the Poor." *The Journal of Economic Perspectives: A Journal of the American Economic Association* 21 (1):141.

Bellemare, Marc F. 2015. "Rising Food Prices, Food Price Volatility, and Social Unrest." *American Journal of Agricultural Economics* 97 (1):1–21.

Bianchi, Milo and Matteo Bobba. 2012. "Liquidity, Risk, and Occupational Choices." *The Review of Economic Studies* :rds031.

Bliss, Christopher John and Nicholas Herbert Stern. 1982. *Palanpur: The Economy of an Indian Village*. Oxford University Press.

Cole, Shawn, Xavier Giné, and James Vickery. 2013. "How Does Risk Management Influence Production Decisions? Evidence from a Field Experiment."

Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H Dar. 2014. "Risk and the Modernization of Agriculture."

Engle, R.F. 1982. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica: Journal of the Econometric Society* :987–1007.

Gummert, M. and Joe Rickman. 2011. "When to Harvest." Tech. rep., International Rice Research Institute.

International Rice Research Institute. 2015. "Rice Production Manual." Tech. rep., International Rice Research Institute.

Jacoby, Hanan G and Emmanuel Skoufias. 1997. "Risk, Financial Markets, and

Human Capital in a Developing Country." *The Review of Economic Studies* 64 (3):311–335.

Kaboski, J. and R. Townsend. 2009. "The Impacts of Credit on Village Economies." *MIT Department of Economics Working Paper No. 09-13* .

Kaboski, J.P. and R.M. Townsend. 2011. "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative." *Econometrica* 79 (5):1357–1406.

Karlan, Dean, Ryan Knight, and Christopher Udry. 2012. "Hoping to Win, Expected to Lose: Theory and Lessons on Micro Enterprise Development."

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. "Agricultural Decisions after Relaxing Credit and Risk Constraints*." *Quarterly Journal of Economics* 129 (2).

Mazzocco, M. and Shiv Saini. 2012. "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences." *The American Economic Review* 102 (1):428–468.

Morduch, Jonathan. 1990. "Risk, Production and Saving: Theory and Evidence from Indian Households." *Harvard University, Manuscript* .

Munshi, Kaivan and Mark Rosenzweig. 2009. "Why Is Mobility in India So Low? Social Insurance, Inequality, and Growth." Tech. rep., National Bureau of Economic Research.

Murphy, K.M. and R.H. Topel. 2002. "Estimation and Inference in Two-Step Econometric Models." *Journal of Business & Economic Statistics* 20 (1):88–97.

Pagan, Adrian. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review* 25 (1):221–247.

Reardon, Thomas, Julio Berdegué, Christopher B Barrett, and Kostas Stamoulis. 2007. "Household Income Diversification into Rural Nonfarm Activities." *Transforming the rural nonfarm economy: opportunities and threats in the developing world* :115–141.

Rosenzweig, Mark R. 1988. "Risk, Implicit Contracts and the Family in Rural

Areas of Low-Income Countries." *The Economic Journal* 98 (393):1148–1170.

Rosenzweig, Mark R and Oded Stark. 1989. "Consumption Smoothing, Migration, and Marriage: Evidence from Rural India." *The Journal of Political Economy* :905–926.

Rosenzweig, M.R. and H.P. Binswanger. 1993. "Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments." *The Economic Journal* 103 (416):56–78.

Townsend, R., A. Paulson, S. Sakuntasathien, T. J. Lee, and M. Binford. 1997. "Questionnaire Design and Data Collection for NICHD Grant *Risk, Insurance and the Family* and NSF Grants." URL http://hdl.handle.net/1902.1/10672.

Townsend, R.M. 1994. "Risk and Insurance in Village India." *Econometrica: Journal of the Econometric Society* :539–591.

Townsend, Robert M. 2012. "Townsend Thai Project Monthly Survey (1-60) Initial Release." URL http://hdl.handle.net/1902.1/14795. Murray Research Archive [Distributor] V1 [Version].

Turnovsky, Stephen J, Haim Shalit, and Andrew Schmitz. 1980. "Consumer's Surplus, Price Instability, and Consumer Welfare." *Econometrica: Journal of the Econometric Society* :135–152.

Walker, Thomas S and James G Ryan. 1990. *Village and Household Economics in India's Semi-Arid Tropics*. Johns Hopkins University Press.

Yang, Dean and HwaJung Choi. 2007. "Are Remittances Insurance? Evidence from Rainfall Shocks in the Philippines." *The World Bank Economic Review* 21 (2):219–248.

A Proofs

A.1 Generalizing the Risk and Return Predictions

Letting $M \in \{1, 2, \dots\}$, the optimal labor allocation is

$$L_p = \frac{\bar{w}_+ M + \alpha\sigma_s^2}{\alpha\left(M\sigma_p^2 + \sigma_s^2\right)} \tag{5}$$

Consider the threshold fixed cost that separates households who choose M activities from those who choose $M+1$ activities:

$$-e^{-\alpha\bar{C}(M)+\frac{\alpha^2}{2}V(M)} = -e^{-\alpha\bar{C}(M+1)+\frac{\alpha^2}{2}V(M+1)}$$

The threshold is

$$\bar{F}_M = \frac{\sigma_s^2\left(\alpha\sigma_p^2 - \bar{w}_+\right)^2}{2\alpha\left(M\sigma_p^2+\sigma_s^2\right)\left((M+1)\sigma_p^2+\sigma_s^2\right)}.$$

The derivatives with respect to σ_p^2 and $\bar{w}_p$ are

$$\frac{\partial\bar{F}_M}{\partial\sigma_p^2} = \frac{\sigma_p\sigma_s^2\left(\alpha\sigma_p^2-\bar{w}_+\right)\left(\alpha\sigma_s^2\left((2M+1)\sigma_p^2+2\sigma_s^2\right)+\bar{w}_+\left(2M\left((M+1)\sigma_p^2+\sigma_s^2\right)+\sigma_s^2\right)\right)}{\alpha\left(M\sigma_p^2+\sigma_s^2\right)^2\left((M+1)\sigma_p^2+\sigma_s^2\right)^2}$$

$$\frac{\partial\bar{F}_M}{\partial\bar{w}_p} = -\frac{\sigma_s^2\left(\alpha\sigma_p^2-\bar{w}_+\right)}{\alpha\left(M\sigma_p^2+\sigma_s^2\right)\left((M+1)\sigma_p^2+\sigma_s^2\right)}$$

By assumption $\left(\alpha\sigma_p^2-\bar{w}_+\right) > 0$, implying that $\frac{\partial\bar{F}_M}{\partial\sigma_p^2} > 0$ and $\frac{\partial\bar{F}_M}{\partial\bar{w}_p} < 0$ for all M. Then a rise in the riskiness of the primary activity will cause all the thresholds to rise, meaning households will be willing to pay more for any number of activities. This will cause the average number of activities in the sample to rise. Likewise a rise in the average return decreases all thresholds and decreases the average number of activities.

QED

A.2 Verifying the Cost Prediction

Assume that σ_p^2 and $\bar{w}_p$ are independent (as they are in the framework of Section 2.2). Take the linear approximation of expected side revenue with respect to these two instruments:

$$\begin{aligned}
\mathbb{E}[y_s] &\approx \bar{w}_s - F\left(\frac{\partial M}{\partial \sigma_p^2}\sigma_p^2 + \frac{\partial M}{\partial \bar{w}_p}\bar{w}_p\right) + \bar{w}_s\left(-\frac{\partial L_p}{\partial M}\left[\frac{\partial M}{\partial \sigma_p^2}\sigma_p^2 + \frac{\partial M}{\partial \bar{w}_p}\bar{w}_p\right] - \frac{\partial L_p}{\partial \sigma_p^2}\sigma_p^2 - \frac{\partial L_p}{\partial \bar{w}_p}\bar{w}_p\right) \\
&= \bar{w}_s + \left(-F - \bar{w}_s\frac{\partial L_p}{\partial M}\right)\left[\frac{\partial M}{\partial \sigma_p^2}\sigma_p^2 + \frac{\partial M}{\partial \bar{w}_p}\bar{w}_p\right] + \bar{w}_s\left[-\frac{\partial L_p}{\partial \sigma_p^2}\sigma_p^2 - \frac{\partial L_p}{\partial \bar{w}_p}\bar{w}_p\right] \\
&= \bar{w}_s + \left(-F - \bar{w}_s\frac{\partial L_p}{\partial M}\right)\hat{M} + \bar{w}_s\left[-\frac{\partial L_p}{\partial \sigma_p^2}\sigma_p^2 - \frac{\partial L_p}{\partial \bar{w}_p}\bar{w}_p\right] \\
\rightarrow y_s &= \bar{w}_s + \left(-F - \bar{w}_s\frac{\partial L_p}{\partial M}\right)\hat{M} + \bar{w}_s\left[-\frac{\partial L_p}{\partial \sigma_p^2}\sigma_p^2 - \frac{\partial L_p}{\partial \bar{w}_p}\bar{w}_p\right] + \varepsilon
\end{aligned}$$

where $\hat{M}$ is the predicted number of activities from the first-stage regression, $\bar{w}_s\left[-\frac{\partial L_p}{\partial \sigma_p^2}\sigma_p^2 - \frac{\partial L_p}{\partial \bar{w}_p}\bar{w}_p\right]$ is the direct effect of labor reallocation from changes in the volatility and average returns to the primary activity, and ε is an approximation error. All partial derivatives are evaluated at the point $(\sigma_p^2, \bar{w}_p) = (0, 0)$ and thus uncorrelated with the instruments.

By the Omitted Variable Bias formula, the instrumental variables estimate is consistent for the value

$$\gamma_A = -F - \bar{w}_s\frac{\partial L_p}{\partial M} - \underbrace{\left[\bar{w}_s\frac{\partial L_p}{\partial \sigma_p^2}\frac{Cov(\hat{M}, \sigma_p^2)}{Var(\hat{M})} + \bar{w}_s\frac{\partial L_p}{\partial \bar{w}_p}\frac{Cov(\hat{M}, \bar{w}_p)}{Var(\hat{M})}\right]}_{\eta}$$

If $-\bar{w}_s\frac{\partial L_p}{\partial M} - \eta > 0$, then $\gamma_A > -F$, which implies that if $\gamma_A < 0$ then $-F < 0$ and thus under-specialization is costly.

First I show that $\frac{\partial L_p}{\partial M} < 0$. From the expression for L_p found in (5) in Appendix A.1 we have that a 1 unit increase in M will cause a rise in the numerator of $\bar{w}_+$ and a rise in the denominator of $\alpha\sigma_p^2$. Since by assumption $\alpha\sigma_p^2 > \bar{w}_+$, the denominator rises by more than the numerator and the total effect is negative. Thus, $\frac{\partial L_p}{\partial M} < 0$.

Now I show that $\eta < 0$. The expression equals

$$\eta = \frac{\bar{w}_s}{Var(\hat{M})}\left[\underbrace{\frac{\partial L_p}{\partial \sigma_p^2}}_{-}\underbrace{\frac{\partial M}{\sigma_p^2}}_{+}Var(\sigma_p^2) + \underbrace{\frac{\partial L_p}{\partial \bar{w}_p}}_{+}\underbrace{\frac{\partial M}{\bar{w}_p}}_{-}Var(\bar{w}_p)\right]$$

where I apply the definition of $\hat{M}$, the independence of σ_p^2 and $\bar{w}_p$, the predictions of the effects of risk and returns to the number of activities to get the signs of $\frac{\partial M}{\sigma_p^2}$ and $\frac{\partial M}{\bar{w}_p}$, and take the deriviatives of L_p found in (5) in Appendix A.1 with respect to σ_p^2 and $\bar{w}_p$. This proves that $\eta < 0$.

QED

B Detailed Data Appendix (For Online Publication)

B.1 Time Series Variables

- **Consumer Prices:** From Bank of Thailand monthly index, acquired from Global Financial Data database. Data were used with permission of Global Financial Data.

- **International Rice Price:** Acquired from IMF monthly commodity price data. Deflated using monthly consumer price index.

B.2 Panel Variables

- **Rice Harvest:** From module 7 (Crop Harvest) section of the monthly survey. Keep only un-milled rice (both sticky and non-sticky). Define "rice harvest soon" as a reported positive harvest of unmilled rice in the subsequent three months. Define "rice harvest past" as having had positive harvest of unmilled rice in the current or previous three months. Define "rice farmer" (or "rice harvest ever") as having had a positive rice harvest at any point in the survey span.

- **Crop-Plots:** From module 5 (Crop Activities) section of the monthly survey. Make the monthly aggregate of "value transacted" for each households sale of each crop. This is the revenue from crops. For number of crop plots, I use the "projected harvest" table, which asks farmers to predict revenue for each productive crop. Every entry corresponds to a different perceived revenue stream for the farmer, so I take number of crop-plots as simply the count of these for each household in each month.

- **Aquaculture:** From module 10 (Fish-Shrimp) of the monthly survey. For each household, make monthly aggregates of the value of fish and shrimp output; this is the revenue from aquaculture. I compute whether a household does aquaculture as whether it reports raising fish/shrimp or having shrimp ponds in a given month.

- **Large Businesses:** From module 12 (Household Business) of the monthly survey. For each household, make monthly aggregates of the cash and in-kind revenue plus the value of products/services consumed by the household; this is the revenue from large businesses. Compute the number of businesses for each household as the number of entries in the household report of revenues.

- **Small/Miscellaneous Businesses:** From module 24 (Income) of the monthly survey. For each household, make monthly aggregates of the cash and in-kind revenue for each "other" income source; this is the revenue from miscellaneous businesses. Compute the number of miscellaneous activities for each household as the number of entries in the household report of revenues.

- **Number of Jobs:** From module 11 (Activities-Occupation). For each person and each job number in any month, mark if it was worked the previous two and the following two months (note that jobs are not assigned job numbers in their first months, so technically I only check the previous one month as it must have been worked the month before to have an ID). If so, it is a "steady job." I count each households total number of jobs and steady jobs each month, then compute the number of unsteady jobs as the difference. For each job and each month, sum the cash and in-kind payments and aggregate by household-month. This is the monthly job revenue.

- **Number of Activities:** I define number of activities as simply the sum of the number of crop plots, the number of livestock activities, the indicator for practice of aquaculture, the number of large businesses, the number of jobs, and the number of miscellaneous activities.

- **Total Revenue, Consumption, and Transfers:** Total revenue is the sum of revenue from crop activities, livestock activities, aquaculture, large businesses, jobs, and miscellaneous activities. Total consumption is the sum of all domestic expenditures by both cash and credit plus consumption of home-produced goods. Expenditures reported at a weekly rather than monthly frequency (in module 23W, Weekly Expenditures Update) are aggregated by month for each household and added to those reported at a monthly frequency (in module 23M, Monthly Expenditures Update). Transfers are defined as the household's net incoming transfers. More precisely, I aggregate by household-month the transfers from people inside and outside the village and subtract similarly aggregated transfers to people inside and outside the village (all found in module 13 on Remittances). I use only transfers not earmarked for a specific event because these unplanned transfers are more like insurance.

C Inference: The Two-Stage Bootstrap (For Online Publication)

As the predicted mean and volatility are both generated regressors, I must adjust my inference to account for their presence. It is easy to see that under my assumptions the full estimators match the conditions for Murphy and Topel (2002). Directly applying their analytic expressions is inconvenient and also problematic because small sample bias in the time series estimates might produce an abnormal small sample distribution for the estimated parameters. But the asymptotic normality their propositions guarantee also ensures the validity of bootstrapped confidence intervals and hypothesis tests.

I implement the procedure as outlined in Figures 12-14. First, I prepare the time series of rice prices for resampling. I form "blocks" consisting of the contemporaneous price and however many lags I need to estimate the time series model. I then group every observation into one or more "blocks of blocks," contiguous interlocking sets of observations and their associated lags.

Next, I run the bootstrap replications. Each replication follows five intermediate steps. First, I sample with replacement the blocks of blocks of rice prices to construct a bootstrapped time series of equal length to the original time series. I estimate the parameters of the time series model on the bootstrapped data. I then resample with replacement households (together with all their monthly observations) from the panel to construct a bootstrapped panel with as many households as the original panel. Then I use the estimated time series model to predict the conditional mean and variance of the international rice price for each household-month observation. Finally, I estimate the panel specification and record the resulting coefficients. I run 1000 replications for the risk specification, 2000 replications for the IV specifications, and 500 replications for the robustness checks.

The final step is to compute confidence intervals and p-values. To construct confidence intervals, I use the dataset of estimated parameters from bootstrap replications to find the 2.5th and 97.5th percentiles. These are the boundaries of the 95% confidence interval. To construct p-values, I compute the absolute t-statistic centered around the original parameter estimate for each replication. The fraction of these absolute t-statistics that is greater than the original t-statistic is the p-value.

Figure 12

Appendix—Bootstrap, Step 1: Forming Blocks of Blocks

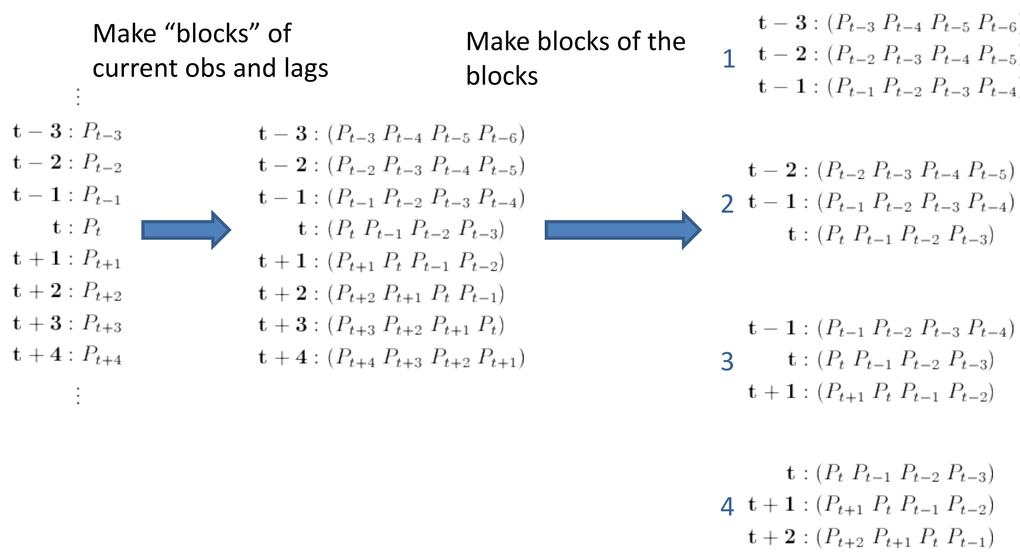

Note: First, I prepare the time series of rice prices for resampling. I form “blocks” consisting of the current price and however many lags are needed to estimate the time series model. I then group every observation into one or more “blocks of blocks,” adjacent interlocking sets of observations and their associated lags.

Figure 13
Appendix—Bootstrap, Step 2: Bootstrap Replications

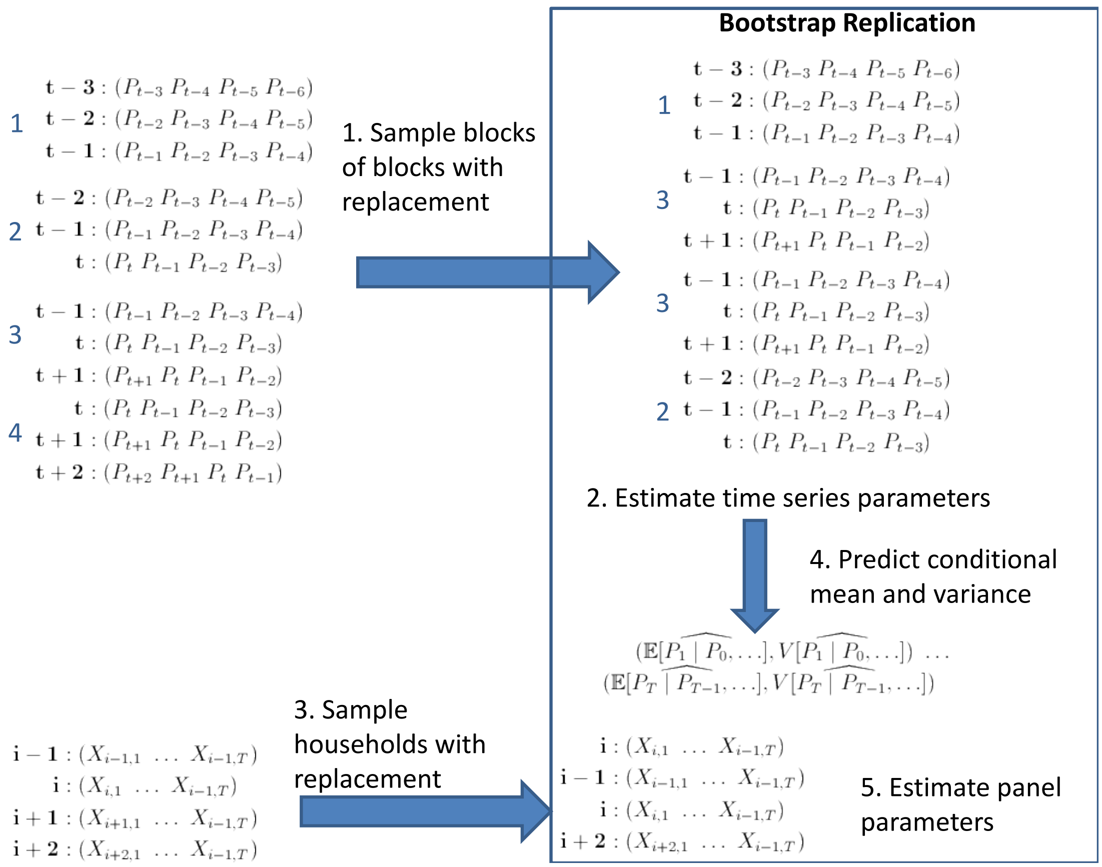

Note: Next I run the bootstrap replications. Each replication follows five steps. First I sample with replacement the blocks of blocks of rice prices to construct a bootstrapped time series of the same length as the original time series. Next I estimate the parameters of the time series model on the bootstrapped data. I then resample with replacement households (together with all their monthly observations) from the panel to construct a bootstrapped panel with as many households as the original panel. Then I use the estimated time series model to predict the conditional mean and variance of the international rice price for each household-month observation. Finally, I estimate the panel specification and record the resulting coefficients. I run 2000 replications for the risk specifications and 3000 replications for the IV specifications.

Figure 14
Appendix—Bootstrap, Step 3: Constructing Confidence Intervals and P-Values

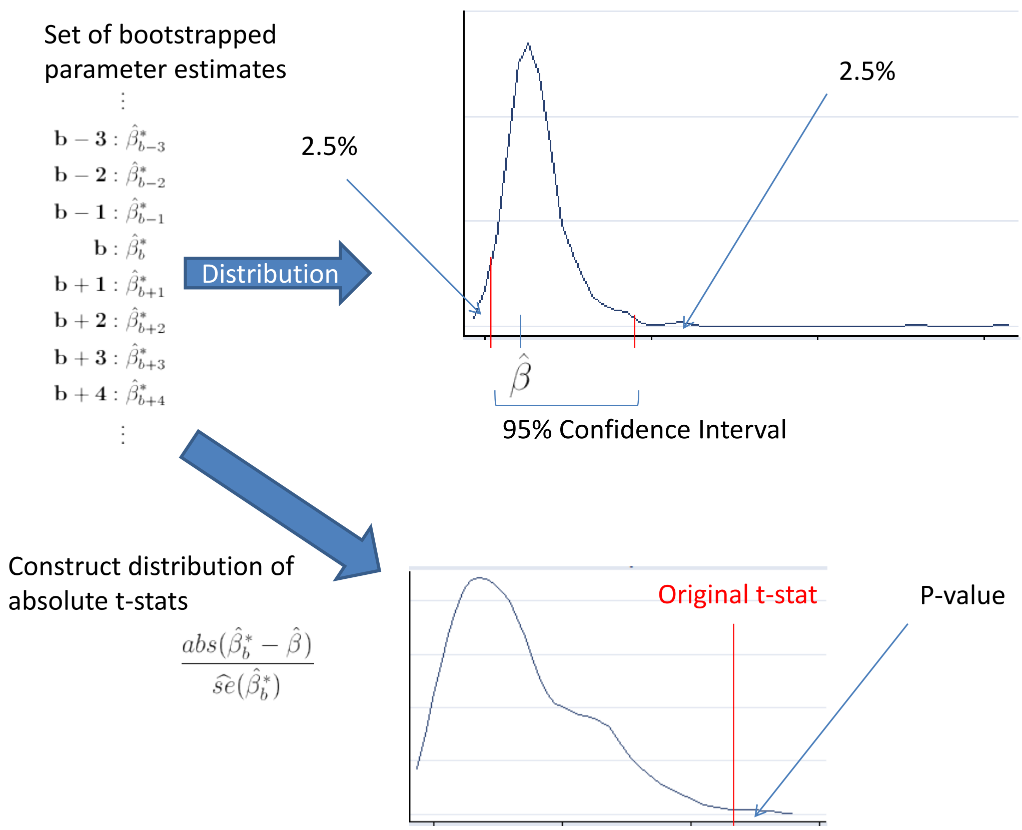

Note: I compute the absolute t-statistic centered around the original parameter estimate for each replication. The fraction of these absolute t-statistics that is greater than the original t-statistic is the p-value.

D Other Tests of Robustness (For Online Publication)

D.1 Full Tables of Robustness Checks from Section 5.2

Table 10
Appendix—Robustness

	(1) Activities (Labor)	(2) Activities (HH Size)	(3) Activities (Rainfall)	(4) Activities (Non-Gen)	(5) Activities (Price Change)	(6) Revenue (Exc. Crop Rev.)	(7) Revenue (Labor & Size)	(8) Revenue (Rainfall)	(9) Revenue (Non-Gen)
Activities						-11232.10**	-17060.61***	-14071.44**	-9305.78*
						[0.031]	[0.009]	[0.028]	[0.064]
Rice Farmer									
- × Mean	0.00	0.00	0.00	0.01***	0.00	-112.59	-91.88	-132.68	-21.25
	[0.249]	[0.441]	[0.352]	[0.001]	[0.854]	[0.178]	[0.278]	[0.156]	[0.795]
- × Volatility	-0.10***	-0.10**	-0.09***	-0.05***	0.02***	-493.25	-791.17	148.59	-397.12
	[0.006]	[0.013]	[0.010]	[0.000]	[0.000]	[0.378]	[0.265]	[0.792]	[0.120]
Expecting Harvest									
- Main	1.78***	1.86***	1.76***	2.58***	2.09***	7295.65	4897.58	5092.92	1537.60
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.118]	[0.345]	[0.292]	[0.725]
- × Mean	-0.01***	-0.01***	-0.02***	-0.02***	-0.01***	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]				
- × Volatility	0.05*	0.05*	0.07*	0.02***	-0.00	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.094]	[0.088]	[0.058]	[0.003]	[0.881]				
Recent Harvest									
- Main	-0.54	-0.61	-0.54	1.30***	0.24	-28550.62*	-34442.87**	-32556.34*	148.02
	[0.306]	[0.281]	[0.356]	[0.000]	[0.423]	[0.052]	[0.030]	[0.052]	[0.992]
- × Mean	-0.01***	-0.01***	-0.01***	-0.02***	-0.01*	167.55	288.46	300.70	-44.62
	[0.002]	[0.005]	[0.002]	[0.000]	[0.062]	[0.381]	[0.129]	[0.160]	[0.758]
- × Volatility	0.16***	0.17***	0.17***	0.05***	-0.04***	1018.61	161.63	-67.12	965.18**
	[0.003]	[0.002]	[0.002]	[0.000]	[0.000]	[0.390]	[0.878]	[0.956]	[0.012]
Household Labor	0.01***						275.42***		
	[0.000]						[0.000]		
Household Size		0.13***					950.03		
		[0.000]					[0.312]		
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Differential Trend	No	No	No	No	No	No	No	No	No
F-Stat Exc. Inst.						13.60	13.15	13.73	20.66
Hansen's J Stat.						0.02	0.27	0.43	0.37
Households	743	743	743	743	743	743	743	743	743
Observations	48160	48164	48329	48329	48329	48329	48160	48329	48329

Table 11
Appendix—Robustness

	(1) Activities (Labor)	(2) Activities (HH Size)	(3) Activities (Rainfall)	(4) Activities (Non-Gen)	(5) Activities (Price Change)	(6) Revenue (Exc. Crop Rev.)	(7) Revenue (Labor & Size)	(8) Revenue (Rainfall)	(9) Revenue (Non-Gen)
Activities						-17835.42**	-30862.50**	-20164.36**	-12865.05**
						[0.042]	[0.032]	[0.048]	[0.038]
Rice Farmer									
- × Mean	0.00	0.00	0.00	0.01***	0.00	-95.30	-54.00	-102.83	15.85
	[0.236]	[0.472]	[0.354]	[0.001]	[0.870]	[0.212]	[0.564]	[0.242]	[0.848]
- × Volatility	-0.10***	-0.10***	-0.09***	-0.05***	0.02***	-854.61	-1689.33	-159.24	-476.34*
	[0.010]	[0.008]	[0.010]	[0.000]	[0.000]	[0.234]	[0.122]	[0.816]	[0.090]
Expecting Harvest									
- Main	1.06**	1.20**	1.04**	2.10***	1.60***	6923.57	4527.99	3569.31	311.46
	[0.018]	[0.020]	[0.046]	[0.000]	[0.000]	[0.164]	[0.460]	[0.520]	[0.945]
- × Mean	-0.01***	-0.01***	-0.01***	-0.02***	-0.01***	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.006]	[0.000]	[0.002]	[0.000]	[0.002]				
- × Volatility	0.06**	0.06*	0.09**	0.03***	-0.00	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)	(Exc. Inst.)
	[0.026]	[0.052]	[0.016]	[0.001]	[0.738]				
Recent Harvest									
- Main	-0.53	-0.61	-0.53	1.29***	0.25	-30671.50*	-39079.16**	-33170.13*	5793.46
	[0.270]	[0.282]	[0.324]	[0.000]	[0.419]	[0.070]	[0.040]	[0.050]	[0.712]
- × Mean	-0.01***	-0.01***	-0.01***	-0.02***	-0.01*	90.86	140.53	217.35	-114.69
	[0.000]	[0.000]	[0.000]	[0.000]	[0.062]	[0.608]	[0.496]	[0.280]	[0.486]
- × Volatility	0.16***	0.17***	0.17***	0.05***	-0.04***	2016.33	2088.62	847.41	1085.79**
	[0.006]	[0.002]	[0.000]	[0.000]	[0.000]	[0.146]	[0.226]	[0.502]	[0.012]
Household Labor	0.01***						354.60***		
	[0.000]						[0.002]		
Household Size		0.13***					2126.60		
		[0.000]					[0.166]		
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Differential Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-Stat Exc. Inst.						7.64	6.71	8.45	15.22
Hansen's J Stat.						1.01	0.80	3.44	2.39
Households	743	743	743	743	743	743	743	743	743
Observations	48160	48164	48329	48329	48329	48329	48160	48329	48329

Table 12
Appendix—Check: Other Measures

	Have Unsteady Job			Non-Crop Activities		
	(1)	(2)	(3)	(4)	(5)	(6)
Mean	0.00***			-0.00*		
	[0.000]			[0.076]		
Volatility	-0.01***			-0.01		
	[0.007]			[0.302]		
Rice Farmer						
- × Mean	-0.00*	-0.00***	-0.00***	-0.00	-0.00	-0.00*
	[0.058]	[0.000]	[0.002]	[0.336]	[0.113]	[0.090]
- × Volatility	0.01	0.01**	0.01**	-0.02	0.01	0.00
	[0.244]	[0.040]	[0.022]	[0.173]	[0.726]	[0.758]
Expecting Harvest						
- Main	-0.07	-0.03	-0.10	0.11	0.29	-0.16
	[0.369]	[0.661]	[0.234]	[0.630]	[0.113]	[0.518]
- × Mean	-0.00**	-0.00	-0.00	-0.01***	-0.00***	-0.00
	[0.020]	[0.150]	[0.482]	[0.000]	[0.001]	[0.218]
- × Volatility	0.02***	0.01	0.01*	0.06***	0.02	0.03
	[0.009]	[0.111]	[0.080]	[0.010]	[0.278]	[0.122]
Recent Harvest						
- Main	0.22***	0.10	0.10	0.80***	0.64***	0.65***
	[0.001]	[0.196]	[0.158]	[0.000]	[0.002]	[0.000]
- × Mean	-0.00***	-0.00	-0.00	-0.01***	-0.01**	-0.01***
	[0.000]	[0.834]	[0.810]	[0.000]	[0.011]	[0.006]
- × Volatility	0.01*	-0.01	-0.01	0.03*	-0.01	-0.01
	[0.060]	[0.134]	[0.114]	[0.072]	[0.479]	[0.476]
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	No	Yes	Yes	No	Yes	Yes
Differential Trend	No	No	Yes	No	No	Yes
Households	743	743	743	743	743	743
Observations	48329	48329	48329	48329	48329	48329

Table 13
Appendix—Check: Time of Harvest

	Activities (1)	Activities (2)	Activities (3)	Activities (4)
Rice Farmer				
- × Mean	0.00*	0.00	0.01**	0.01*
	[0.075]	[0.126]	[0.040]	[0.062]
- × Volatility	-0.10***	-0.10**	-0.09***	-0.09***
	[0.010]	[0.012]	[0.005]	[0.010]
Expecting Harvest				
- Main	2.13***	1.87***	1.88***	1.52***
	[0.000]	[0.002]	[0.000]	[0.006]
- × Mean	-0.02***	-0.02***	-0.02***	-0.01***
	[0.000]	[0.000]	[0.000]	[0.000]
- × Volatility	0.07*	0.07*	0.07*	0.08**
	[0.050]	[0.052]	[0.060]	[0.044]
Recent Harvest				
- Main	0.02	0.02	0.08	0.08
	[0.975]	[0.974]	[0.930]	[0.924]
- × Mean	-0.02***	-0.02***	-0.02***	-0.02***
	[0.000]	[0.000]	[0.000]	[0.000]
- × Volatility	0.22***	0.22***	0.25***	0.25***
	[0.005]	[0.006]	[0.000]	[0.006]
Harvest Oct-Dec			0.23***	0.24***
			[0.000]	[0.000]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trends	No	Yes	No	Yes
Households	659	659	659	659
Observations	43257	43257	43257	43257

Table 14
Appendix—Check: Placebo Test

	(1) Activities (3 Month Lead)	(2) Activities (3 Month Lead)	(3) Activities (12 Month Lead)	(4) Activities (12 Month Lead)
Rice Farmer				
- × Mean	-0.01	-0.01*	-0.00	-0.00
	[0.100]	[0.084]	[0.322]	[0.296]
- × Future Volatility	0.03	0.03	-0.10**	-0.10**
	[0.312]	[0.332]	[0.038]	[0.018]
Expecting Harvest				
- Main	1.91***	1.33***	2.02***	1.19*
	[0.000]	[0.002]	[0.000]	[0.090]
- × Mean	-0.01***	-0.01*	-0.01***	-0.01*
	[0.000]	[0.056]	[0.000]	[0.066]
- × Future Volatility	0.00	-0.00	-0.02	0.01
	[0.982]	[0.876]	[0.736]	[0.858]
Recent Harvest				
- Main	0.28	0.28	-1.69**	-1.68**
	[0.598]	[0.548]	[0.022]	[0.032]
- × Mean	0.00	0.00	-0.00	-0.00
	[0.200]	[0.174]	[0.468]	[0.484]
- × Future Volatility	-0.12**	-0.13**	0.21**	0.21**
	[0.030]	[0.042]	[0.020]	[0.032]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trends	No	Yes	No	Yes
Households	735	735	715	715
Observations	45840	45840	39384	39384

Table 15
Appendix—Check: Insurance

	(1) Activities (Uninsured)	(2) Activities (Uninsured)	(3) Activities (Insured)	(4) Activities (Insured)
Rice Farmer				
- $\times$ Mean	0.01**	0.01**	-0.00	-0.00
	[0.027]	[0.036]	[0.922]	[0.936]
- $\times$ Volatility	-0.09*	-0.09*	-0.10**	-0.10**
	[0.069]	[0.068]	[0.013]	[0.014]
Expecting Harvest				
- Main	1.65***	1.52*	1.99***	1.15*
	[0.001]	[0.078]	[0.000]	[0.052]
- $\times$ Mean	-0.02***	-0.01***	-0.02***	-0.01**
	[0.000]	[0.002]	[0.000]	[0.036]
- $\times$ Volatility	0.09*	0.09*	0.04	0.05
	[0.067]	[0.072]	[0.277]	[0.130]
Recent Harvest				
- Main	-0.77	-0.77	-0.39	-0.38
	[0.186]	[0.174]	[0.557]	[0.560]
- $\times$ Mean	-0.01	-0.01	-0.01***	-0.01***
	[0.221]	[0.226]	[0.006]	[0.002]
- $\times$ Volatility	0.14**	0.14**	0.19***	0.19***
	[0.025]	[0.022]	[0.003]	[0.002]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trends	No	Yes	No	Yes
Households	270	270	473	473
Observations	16933	16933	31396	31396

D.2 Regressions Using Only Farmers

Here I run several simpler specifications that use only rice farmers or rice farmers expecting a harvest. Tables 16 and 17 restrict the sample to only rice farmers. (The interactions of mean and volatility with $[Rice\ Farmer]$ are excluded because they are now colinear with the main effect of the mean and volatility.) Table 18 further restricts the sample to months in which each rice farmer is expecting a harvest. Since there is only one group the interactions are no longer useful; the regression simply compares how each rice farmer expecting a harvest reacts to low versus high volatility. The first-stage F-statistic in these regressions is so low that I do not attempt the second-stage analysis.

D.3 Controlling for Wages

Suppose wages are correlated with volatility in the rice price (for example, because of general equilibrium effects). Then the extra jobs the household takes up may not be a response to risk but rather a response to better earnings in side activities. Tables 19 and 20 show that controlling for median village wages does not change the coefficients of interest. I use median wages at the village level rather than individual wages because households might be willing to take lower paying jobs to hedge against risk, which would introduce an artificial correlation between wages and the regressor of interest.

E Alternative Model: Minimum Labor Inputs (For Online Publication)

Is it plausible that the kinds of activities a rice farmer can enter three months before his harvest would, as my model assumes, have a lumpy fixed cost? Finding casual labor or growing cassava may be easy if the farmer has already done so every time prices turned volatile in the past. In this appendix I build a model

Table 16
Appendix—Main Regression with Only Rice Farmers

	Activities	Activities	Activities	Activities
Mean	0.00*	0.01***	0.01***	0.01***
	[0.050]	[0.000]	[0.000]	[0.000]
Volatility	-0.28***	-0.25***	-0.26***	-0.25***
	[0.000]	[0.006]	[0.000]	[0.002]
Expecting Harvest				
- Main	1.82***	1.72***	0.92	0.63
	[0.004]	[0.006]	[0.142]	[0.326]
- × Mean	-0.02***	-0.02***	-0.02***	-0.02***
	[0.000]	[0.000]	[0.000]	[0.000]
- × Volatility	0.18***	0.19***	0.19***	0.22***
	[0.003]	[0.004]	[0.002]	[0.002]
Recent Harvest				
- Main	-0.76	-0.94	-0.98	-0.88
	[0.482]	[0.338]	[0.348]	[0.386]
- × Mean	-0.03***	-0.03***	-0.03***	-0.03***
	[0.000]	[0.000]	[0.000]	[0.000]
- × Volatility	0.41***	0.39***	0.40***	0.39***
	[0.000]	[0.006]	[0.000]	[0.002]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trend	No	No	Yes	Yes
Rainfall Shocks	No	Yes	No	Yes
Households	354	354	354	354
Observations	23613	23613	23613	23613

Note: I confirm the results hold when I exclude non-farmers. The bracketed values are p-values. I compute the p-values using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C).

Table 17
Appendix—Main Regression with Only Rice Farmers (Second Stage)

	Revenue	Revenue	Revenue	Revenue	Revenue	Revenue
Activities	1073.95***	2087.20***	-7548.36*	-6890.26*	-6700.01*	-4131.07
	[0.000]	[0.000]	[0.056]	[0.098]	[0.096]	[0.322]
Mean			-77.37*	-72.43	-37.88	-32.93
			[0.094]	[0.142]	[0.308]	[0.382]
Volatility			-1915.27**	-1564.96	-1743.26*	-1037.83
			[0.046]	[0.110]	[0.074]	[0.250]
Expecting Harvest						
- Main			6961.02	6207.05	3454.75	2030.71
			[0.122]	[0.178]	[0.374]	[0.526]
- × Mean			(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)
			(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)
- × Volatility			(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)
			(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)	(Excl. Inst.)
Recent Harvest						
- Main			-21727.29*	-19250.84	-17979.06	-14326.56
			[0.092]	[0.152]	[0.146]	[0.194]
- × Mean			119.87	130.39	108.60	171.49
			[0.218]	[0.238]	[0.288]	[0.146]
- × Volatility			1339.76	943.43	1077.55	-6.61
			[0.144]	[0.296]	[0.300]	[0.996]
Household Fixed-Effects	No	Yes	Yes	Yes	Yes	Yes
Differential Trend	No	No	No	No	Yes	Yes
Rainfall Shocks	No	No	No	Yes	No	Yes
F-Stat Exc. Inst.			33.83	34.20	30.54	32.20
Hansen's J Stat.			1.67	3.06	7.47	10.28
Households	354	354	354	354	354	354
Observations	23613	23613	23613	23613	23613	23613

Note: I confirm the second stage results hold when I exclude non-farmers. The bracketed values are p-values. I compute the p-values using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C).

Table 18
Appendix—Main Regression with Only Rice Farmers Currently Expecting a Harvest

	Activities	Activities	Activities
Mean	-0.01***	0.00	0.01
	[0.010]	[0.982]	[0.238]
Volatility	0.05	0.06*	0.11**
	[0.112]	[0.054]	[0.028]
Household Fixed-Effects	Yes	Yes	Yes
Trend	Yes	Yes	Yes
Rainfall Shocks	Yes	Yes	Yes
Harvest Fixed-Effects	No	Yes	Yes
Year Fixed-Effects	No	No	Yes
Households	352	352	352
Observations	5539	5539	5539

Note: I confirm the results hold when I restrict the sample to months in which farmers expect a harvest. By definition this excludes non-farmers. The bracketed values are p-values. I compute the p-values using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C).

Table 19
Appendix—Robustness: Controlling for Wages
Does Not Change the Results

	Activities (1)	Activities (2)	Activities (3)	Activities (4)	Activities (5)
Rice Farmer					
- × Mean	0.01***	0.00	0.00	0.00	0.00
	[0.010]	[0.380]	[0.388]	[0.388]	[0.358]
- × Volatility	-0.20***	-0.10**	-0.10**	-0.10***	-0.09**
	[0.002]	[0.012]	[0.012]	[0.008]	[0.012]
Expecting Harvest					
- Main	1.89***	1.92***	1.77***	1.17**	0.88*
	[0.002]	[0.000]	[0.000]	[0.020]	[0.078]
- × Mean	-0.02***	-0.02***	-0.02***	-0.01***	-0.01***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.002]
- × Volatility	0.17***	0.05	0.07**	0.06*	0.09**
	[0.002]	[0.106]	[0.044]	[0.062]	[0.030]
Recent Harvest					
- Main	-0.70	-0.54	-0.51	-0.54	-0.50
	[0.504]	[0.304]	[0.328]	[0.332]	[0.342]
- × Mean	-0.03***	-0.01***	-0.01***	-0.01***	-0.01***
	[0.000]	[0.004]	[0.000]	[0.002]	[0.004]
- × Volatility	0.40***	0.17***	0.17***	0.17***	0.17***
	[0.000]	[0.000]	[0.006]	[0.004]	[0.006]
Wage	-0.01***	-0.01***	-0.01***	-0.01***	-0.01***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Household Fixed-Effects	Yes	Yes	Yes	Yes	Yes
Time Fixed-Effects	No	Yes	Yes	Yes	Yes
Rainfall Shocks	No	No	Yes	No	Yes
Differential Trend	No	No	No	Yes	Yes
Households	743	743	743	743	743
Observations	48329	48329	48329	48329	48329

Note: The bracketed values are p-values. I compute the p-values using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C).

Table 20
Appendix—Robustness: Controlling for Wages
Does Not Change the Results (cont.)

	(1) Revenue	(2) Revenue	(3) Revenue	(4) Revenue
Activities	-13807.27**	-14039.89***	-24118.94**	-19734.11**
	[0.030]	[0.010]	[0.036]	[0.038]
Rice Farmer				
- × Mean	-129.43	-133.60	-106.72	-100.05
	[0.146]	[0.110]	[0.270]	[0.240]
- × Volatility	-276.18	148.38	-862.67	-140.11
	[0.622]	[0.786]	[0.340]	[0.844]
Expecting Harvest				
- Main	4896.29	5023.28	4344.98	2946.65
	[0.360]	[0.318]	[0.436]	[0.566]
Recent Harvest				
- Main	-34607.68**	-32427.87**	-37882.50*	-32329.78*
	[0.032]	[0.034]	[0.058]	[0.056]
- × Mean	300.49	301.23	182.48	218.31
	[0.132]	[0.102]	[0.384]	[0.264]
- × Volatility	213.17	-88.67	1744.51	746.90
	[0.832]	[0.928]	[0.262]	[0.598]
Wage	-33.44	-37.88	-121.87	-89.01
	[0.498]	[0.442]	[0.194]	[0.242]
Household Fixed-Effects	Yes	Yes	Yes	No
Time Fixed-Effects	Yes	Yes	Yes	Yes
Rainfall Shocks	No	Yes	No	Yes
Differential Trend	No	No	Yes	Yes
F-Stat Exc. Inst.	13.82	13.96	7.07	8.28
Hansen's J Stat.	0.14	0.42	0.69	3.73
Households	743	743	743	743
Observations	48329	48329	48329	48329

Note: The bracketed values are p-values. I compute the p-values using a two-stage bootstrap that corrects for generated regressors and clusters by household (see Appendix C).

without fixed costs where risk still causes under-specialization. The prediction's robustness is why I emphasize that my model of risk and under-specialization is not *the* model, but just a convenient tool to formalize the intuition.

Let the household's utility function be as before and for simplicity consider the case of choosing between perfect specialization and one side activity. The household can costlessly enter a side activity but must allocate it at least $\underline{L} > 0$ units of labor. The lower-bound on labor choice captures the idea that it is not worth an employer's time to hire a worker for only a few hours per week, so even work that does not require paying a fixed cost does require a lumpy investment of time. I need the lumpiness to make specialization optimal for some degree of riskiness. Otherwise the household always has a side activity and only varies how much it works on the side activity instead of whether it has one at all. I also assume the average return to the side activity is strictly less than the average return to the primary activity—that is, $\bar{w}^p - \bar{w}^s = w^+ > 0$. The household faces the trade-off

	$\boldsymbol{M = 0}$	$\boldsymbol{M = 1}$
$\bar{\boldsymbol{C}}$	$\bar{w}^p$	$\bar{w}^p - w^+(1 - L^p)$
$\boldsymbol{V}$	σ_p^2	$(L^p)^2\sigma_p^2 + (1 - L^p)^2\sigma_s^2$

The opportunity cost of the side activity is $w^+(1-L^p)$, and since it is no less than $w^+\underline{L} > 0$ the household still loses a discrete chunk of expected revenue when it diversifies. Although it does not literally pay a fixed cost the household's trade-off between the mean and variance of consumption is similar to the one it faced in the original model. They are not identical—for example, the cost of diversification is now uncertain—but similar enough for risk to cause under-specialization.

Figure 15 gives the intuition. With perfect specialization the household's expected utility is maximized when the primary activity's returns have zero variance, but expected utility falls steeply as the variance rises. The household can flatten the utility-variance relationship by moving some labor from

Figure 15
Appendix—Intuition of the Alternative Model

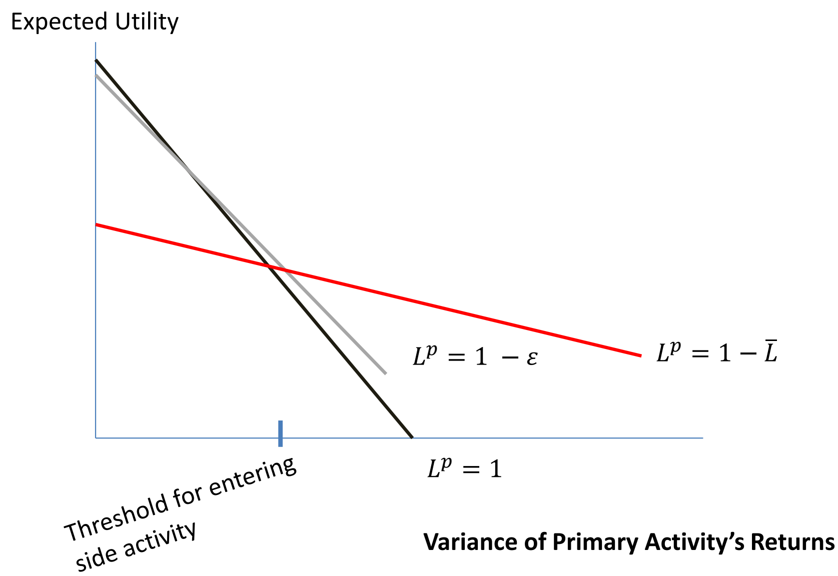

the primary activity to the side activity. Without a lower bound on labor devoted to the side activity, the household would always move ε units of labor to the side activity and be happier without perfect specialization. But with a lower bound the household must accept a discretely lower and flatter utility-variance relation. If the variance of the side activity is low, the household prefers specialization. But when the variance exceeds a critical threshold the household prefers to diversify. If w^+ has a nondegenerate distribution the average number of activities will rise continuously with the variance. Then the lower bound model makes the same prediction $\frac{d\mathbb{E}[M]}{d\sigma_p^2} > 0$ as the fixed cost model from the main text.

F Other Tables Referenced in the Main Text (For Online Publication)

Table 21
Appendix—Modeling the Rice Price as a Random Walk

	(1)
	P_t
P_{t-1}	0.995***
	(0.00)
N	389
R^2	0.995

Note: The random walk specification describes the data well. It models the current price of rice as the previous month's price plus a random innovation: $P_t = P_{t-1} + \varepsilon_t$.

Table 22
Appendix—Robustness: Main Results Excluding Pre-Harvest Rice Sales

	(1) Activities	(2) Activities	(3) Revenue	(4) Revenue
Activities			-14195.18**	-20750.62*
			[0.027]	[0.058]
Rice Farmer				
- × Mean	0.00	0.00	-86.06	-68.73
	[0.395]	[0.410]	[0.307]	[0.448]
- × Volatility	-0.09**	-0.09**	-368.73	-739.82
	[0.011]	[0.016]	[0.531]	[0.360]
Expecting Harvest				
- Main	1.37***	0.96**	-238.27	-1304.60
	[0.000]	[0.038]	[0.953]	[0.746]
- × Mean	-0.01***	-0.01***	(Exc. Inst.)	(Exc. Inst.)
	[0.000]	[0.000]		
- × Volatility	0.05*	0.06*	(Exc. Inst.)	(Exc. Inst.)
	[0.091]	[0.060]		
Recent Harvest				
- Main	-0.63	-0.62	-27329.39*	-29798.14*
	[0.263]	[0.240]	[0.077]	[0.090]
- × Mean	-0.01***	-0.01***	147.78	74.07
	[0.003]	[0.002]	[0.390]	[0.654]
- × Volatility	0.17***	0.17***	1095.82	2078.94
	[0.003]	[0.006]	[0.370]	[0.188]
Household Fixed-Effects	Yes	Yes	Yes	Yes
Time Fixed-Effects	Yes	Yes	Yes	Yes
Differential Trend	No	Yes	No	Yes
F-Stat Exc. Inst.			10.05	6.04
Hansen's J Stat.			0.01	0.74
Households	743	743	743	743
Observations	47395	47395	47395	47395

Note: I exclude observations when households claim they sold rice while still expected their harvest. Volatility still causes households to enter more activities (Column 1) and the extra activities are costly (Column 2).