Review of Economic Studies (2021) **88**, 2275–2307 doi:10.1093/restud/rdaa084

Advance access publication 24 December 2020

The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity

KEVIN DONOVAN

Yale School of Management

First version received January 2017*; Editorial decision July* 2020*; Accepted November* 2020 (*Eds.*)

I consider the aggregate impact of low intermediate input intensity in the agricultural sector of developing countries. In a dynamic general equilibrium model with idiosyncratic shocks, incomplete markets, and subsistence requirements, farmers in developing countries use fewer intermediate inputs because it limits their exposure to uninsurable shocks. The calibrated model implies that Indian agricultural productivity would increase by 16% if markets were complete, driven by quantitatively important increases in both the average real intermediate share and measured TFP through lower misallocation. I then extend the results to consider the importance of risk in other contexts. First, the introduction of insurance decreases cross-country differences in agricultural labour productivity by 14%. Second, scaling the introduction of improved seeds to decrease downside risk reduces inequality by reallocating resources from rich to poor farmers via equilibrium effects. This reallocation substantially increases aggregate productivity relative to what would be expected from extrapolating the partial equilibrium impact.

Key words: Agriculture, General equilibrium, Intermediate inputs, Misallocation, Productivity, Risk.

JEL Classifications Codes: O11, O41.

1. INTRODUCTION

Differences in agricultural labour productivity between the richest and poorest countries are twice as large as differences in aggregate labour productivity. In spite of this, the least developed countries in the world employ over 70% of their population in the agricultural sector, suggesting that agricultural productivity is critical for understanding aggregate income differences.[1]

One possible cause of low agricultural productivity is low intermediate intensity (*e.g.* fertilizer, pesticide), as the intermediate input share (the value of intermediates as a share of harvest) differs

1. This argument has been made in various forms starting with Restuccia *et al*. (2008) and Caselli (2005) and this cross-country gap is also is the focus of this article. Relatedly, Vollrath (2009) and Gollin *et al*. (2014) focus on nominal productivity gaps between the agricultural and non-agricultural sectors of the same country, and show that the magnitude of the gap is decreasing in a country's income.

The editor in charge of this paper was Michele Tertilt.

by a factor of ten between the poorest and richest countries in the world. Moreover, I document in Section 2 that this positive cross-country correlation does not exist in other sectors, suggesting that it may be an important margin for understanding why the agricultural sector exhibits significantly lower labour productivity than the nonagricultural sector in developing countries.

The goal of this article is to study the relationship between this low intermediate intensity and agricultural risk, and in particular, the importance of such a channel when taking into account general equilibrium effects. I develop and quantify a model in which agricultural risk generates inefficiently low intermediate use, then use the model to study the impact of risk on agricultural and aggregate productivity both within and across countries. I further apply the results to policy, showing how the taking equilibrium effects into account plays a critical role in understanding the gains from implementing policy at-scale.

To do so, I begin by developing a model in which farmers produce agricultural output using intermediate inputs and labour, and are subject to incomplete markets and random fluctuations in farm productivity. Consistent with the temporal sequencing of agricultural decisions, intermediates are chosen before the realization of shocks. Together, these features require that farmers internalize the impact of this *ex ante* intermediate choice on their *ex post* consumption. This link matters for optimal intermediate expenditures—higher expenditures expose farmers to lower consumption in the event of a bad shock, and they therefore optimally scale back expenditures to limit the risk of low consumption. More formally, it implies a risk-generated gap between the expected marginal revenue and marginal cost of intermediates, which I show is isomorphic to a reduced-form tax on intermediate inputs in a complete markets model (in a Hsieh and Klenow, 2009, sense), despite the lack of any direct frictions in the intermediate input market.

I then turn to quantifying the implications of this model. I calibrate the baseline model to India, using a mix of aggregate and micro-level data. The nominal intermediate share in agriculture is 10%, substantially lower than in richer countries. I match harvest-level volatilities to track shocks in the economy, while also allowing smoothing through *ex post* labour (*e.g.* Kochar, 1999) and savings. Finally, I include a number of commonly modelled relevant sector-neutral and agriculture-specific features of the economy. I test a number of model predictions against household-level panel data on Indian farmers, including the relationship between savings, intermediate intensity, and consumption volatility. The model matches the broad patterns observed in the micro data.

I therefore turn to the quantitative implications of the model, beginning with the cost of risk in India. In the baseline model without insurance, I find that the nominal intermediate share in India is 0.31, compared to the calibrated Cobb–Douglas exponent of 0.40. This implies a sizable risk-induced wedge in India, though cannot capture the full gap (roughly, the model captures about one-third of the gap in nominal intermediate shares between India and the U.S.). I measure the cost of risk by comparing this model to an identical Indian economy with complete markets. Doing so increases agricultural productivity in India by 16% while increasing the nominal intermediate share to (the Cobb–Douglas exponent) of 0.40. Thus, risk plays a quantitatively important role in understanding Indian agricultural productivity.

There are two key forces underlying this aggregate result. The first is that risk decreases the average household's willingness to use intermediates. This manifests as a lower average real intermediate share and naturally lowers labour productivity. The second is that risk induces misallocation, as it more severely affects poorer households. This lowers aggregate TFP, further amplifying the cost of risk. I decompose the importance of these two channels, and find that 36% of the total impact comes from lower misallocation, while the remainder comes from the increase in the average real intermediate share. Thus, changes in both the first and second moment of the real intermediate share play an important quantitative role.

I then extend the analysis across countries, with the goal of studying whether the model with uninsured risk can amplify cross-country productivity differences compared to a complete markets model. Intuitively, one might suspect that risk is more damaging in poor countries—their relative poverty (here, generated by low TFP) implies that a bad shock realization should have more dire consequence for consumption. This type of intuition, however, is incomplete in general equilibrium, and I show theoretically that it relies critically on the inclusion of subsistence requirements in the utility function. The key model mechanism here is that, while the poor are indeed more concerned about risk, the equilibrium price responds to compensate them for it. I show that without subsistence, the price in poor countries exactly offsets the extra burden of risk they face relative to their rich country counterparts. Thus, while risk depresses productivity *within* an economy (as I show in India), the quantitative magnitude of this cost is identical *across* countries without subsistence.

A subsistence requirement breaks this result: as the price rises to compensate farmers for risk-taking, it simultaneously increases the cost of satisfying subsistence, which has the countervailing effect of increasing farmer risk aversion. Specifically, it decreases their available income net of required subsistence payments, effectively making them poorer and therefore more risk averse. The results show that the same non-homothenticity generating income effects in structural transformation (*e.g.* Kongsamut *et al.*, 2001; Herrendorf *et al.*, 2013, 2014) simultaneously modifies household risk aversion, which allows the model to qualitatively replicate the patterns we observe in the data: a positive cross-country correlation between nominal intermediate shares and agricultural productivity.

To quantify this cross-country channel, I introduce a second economy calibrated to the U.S. In the baseline no-insurance model, the U.S. nominal intermediate share is 0.40 (again, compared to the Cobb–Douglas exponent of 0.40). That is, there is essentially no risk-generated wedge. As expected, introducing insurance has no quantitative impact on aggregate moments. Thus, when combined with the Indian results above, the risk-induced cross-country productivity gap is substantially larger than would otherwise be implied by the complete markets model. The introduction of complete insurance in both countries decreases the U.S.–India agricultural productivity gap from a factor of 35.9 to 30.9 (a 14% decline), while aggregate GDP per worker differences decline from a factor of 8.2 to 7.4 (a 10% decline).[2]

Finally, motivated by the importance of general equilibrium price movements in both the Indian and cross-country results, I use the model to highlight their importance in understanding the returns from a concrete policy proposal.[3] Specifically, I study the impact of scaling a promising randomized controlled trial on the introduction of improved seeds that reduce downside risk in India (Emerick *et al.*, 2016). In partial equilibrium, the model naturally predicts that such an intervention would increase intermediate expenditures and yield. Moreover, it predicts a decline in savings, as such seeds reduce the need to hold buffer stock savings to smooth consumption. When I run the experiment in the calibrated Indian model, the resulting treatment effects along these dimensions are consistent in both sign and magnitude with the empirical results.

I then use the model to simulate the introduction of such an intervention at scale. I find that equilibrium effects provide an additional channel through which this intervention increases productivity, stemming from the channels highlighted in the previous results. As the agricultural

2. A useful implication from this result is that risk generates a complementary amplification channel for other agricultural distortions considered in a complete markets framework, many of which are considered in the papers cited below. I take this up in the Supplementary Appendix.

3. Rural agricultural markets are poorly integrated into the world food market in developing countries, suggesting that prices are unlikely to be determined internationally. See, for instance, Adamopoulos (2011, Adamopoulos) and Sotelo (Sotelo) among others for this lack of integration.

price falls in response to more productive farms, it simultaneously has two effects. It both lowers the expected return to intermediates while making farmers more risk tolerant (by lowering the cost of subsistence). However, the relative intensity of these two channels differ across farmers. While rich farmers in India—essentially unconstrained by risk—lower intermediate intensity due to the lower expected return, the poor increase intermediate intensity due to the outsized role played by their decline in risk aversion. Put differently, equilibrium effects effectively redistribute farm income from rich to poor, generating a convergence in real intermediate shares. This decline in misallocation (and, more broadly, consumption inequality) provides an important increase in agricultural productivity that is un-captured by the partial equilibrium results and is critical to understanding the effect of such a policy at scale.

1.1. *Related literature*

This article contributes to a number of different literatures. Most closely, it joins a recent macro literature on the role of agriculture in understanding cross-country income differences, including Lagakos and Waugh (2013), Herrendorf and Schoellman (2015), Tombe (2015), and Caunedo and Keller (2020). Moreover, the model provides an otherwise absent complementary amplification channel for distortions considered in a complete markets framework, including transportation costs (Adamopoulos, 2011; Gollin and Rogerson, 2014), distortions that vary farm size (Adamopoulos and Restuccia, 2014) or technology choice (Yang and Zhu, 2013), and more general distortions that affect input markets (Gollin *et al.*, 2004). Most closely related is the work of Restuccia *et al.* (2008), who also focus on the role of intermediate inputs. Building off their work, this article contributes a micro-founded rationale for distortions in the intermediate input market by linking intermediate choices with consumption risk.

I show that the distortions emphasized in these papers induce misallocation across agricultural production units when combined with consumption risk. In this sense, the paper more broadly relates to the literature relating establishment-level distortions to aggregate productivity (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009). Recent work has focused on the role of financial development (Buera *et al.*, 2011; Midrigan and Xu, 2014; Moll, 2014), and in agriculture specifically, land (Restuccia and Santaeulàlia-Llopis, 2017; Gottlieb and Grobovšek, 2019) and capital markets (Adamopoulos *et al.*, 2019). The difference between such theories and the one developed here is that these theories all explicitly distort the relevant input market. In contrast, the theory developed here shows how distortions in other parts of the economy (namely, the consumption market) can be observationally similar to such input market distortions. I discuss the links between such models in more detail in the Supplementary Appendix.

Finally, this article highlights the importance of equilibrium consequences for understanding the full impact of development policy at scale. See Buera *et al.* (2020), Bergquist *et al.* (2019), Greenwood *et al.* (2019), and Fujimoto *et al.* (2019) for broadly similar implications in response to a varied set of policies. This adds important context for experimental insurance interventions (Mobarak and Rosenzweig, 2014; Karlan *et al.*, 2014; Cai *et al.*, 2015; Cole *et al.*, 2017) or other work highlighting the individual-level importance of risk (Dercon and Christiaensen, 2011; Emerick *et al.*, 2016; Brooks and Donovan, 2020).

2. MOTIVATING EVIDENCE

Before turning to the model, this section documents facts about the intermediate share across countries. First, I confirm the results of Restuccia *et al.* (2008) and show that the same results extend to nominally priced intermediate shares. I then show that the same relationship does not hold in other sectors of the economy, suggesting that this relationship is a relevant margin for understanding the agricultural productivity gap.

Define the real intermediate input share as $X^{j*} := p_x^* X^j / p_a^* Y_a^j$, where X^j is intermediate input consumption in agriculture of country j, Y_a^j is agricultural output, and p_x^* and p_a^* are international prices of intermediates and agricultural output. The nominal intermediate share is given by $\widehat{X}^j := p_x^j X^j / p_a^j Y_a^j$, where the only difference is that intermediates and output are valued at nominal country-specific prices p_x^j and p_a^j. As discussed in Section 1, the influential work of Restuccia *et al.* (2008) finds that real intermediate shares differ substantially across countries. Using data from Prasada Rao (1993), which is constructed from Food and Agricultural Organization (FAO) statistics and underlies the Restuccia *et al.* (2008) analysis, Figure 1a reproduces their finding of a strong positive correlation between GDP per capita and the real intermediate share.[4]

One possibility is that the relationship between the real intermediate share and income is driven exclusively by the fact that farmers in poor countries pay higher prices for intermediates. Hsieh and Klenow (2007), for example, find that higher prices in poor countries are sufficient to explain cross-country variation in aggregate investment rates, and Restuccia *et al.* (2008) find an important role for price variation in agriculture specifically. However, this does not seem to be the totality of the story here. Figure 1b uses the same data and plots the nominal share. There remains strong positive correlation of 0.65. The tenth percentile country, as ranked by GDP per capita, has a nominal intermediate share that is one-fourth of the intermediate share in the U.S.[5,6]

This suggests a potentially important role for distortions that do not operate simply by changing the price that farmers pay for intermediate inputs. Of course, such a distortion is relevant for productivity only to the extent that it impacts real variables. This article provides such a theory, based on the notion of a risk-generated wedge between the marginal revenue and marginal cost of intermediate inputs, and investigates its importance for real outcomes across countries.

2.1. *Comparison to non-agriculture*

As a last step before the model, I compare intermediate input shares across agriculture and non-agriculture using the data from the 2005 World Input–Output Database (Timmer *et al.*, 2015). Figures 2a and b plot the cross-country nominal intermediate shares in agriculture and non-agriculture. Only the agricultural sector has a positive slope. A further benefit of the WIOD is that it allows me to decompose the production sector of intermediates consumed in agriculture. Figure 2c shows that there is no relationship between income and intermediates that are produced, then consumed, in the agricultural sector. Instead, the positive relationship in Figure 2a is

4. Chen (2020) finds a qualitatively similar pattern in capital shares across countries. In the Supplementary Appendix, I show that cross-country variation in intermediates is larger than capital. I emphasize, however, that this result should not be taken to suggest that capital is unimportant in the development process, only that intermediate variation is substantial. For example, agricultural capital quality, much like any capital stock, is notoriously difficult to measure. Caunedo and Keller (2020) study this in detail and point to important cross-country implications of variation in capital quality.

5. As further robustness of the result, I construct aggregate nominal intermediate shares using micro data from the Living Standard Measurement Studies (LSMS) released by the World Bank in the Supplementary Appendix. The same pattern emerges.

6. The positive correlation between nominally and real priced shares imply that the ration p_x/p_a does not vary systematically with development. Figure 1c confirms this, but note that there is still substantial variation in the price ratio across countries.

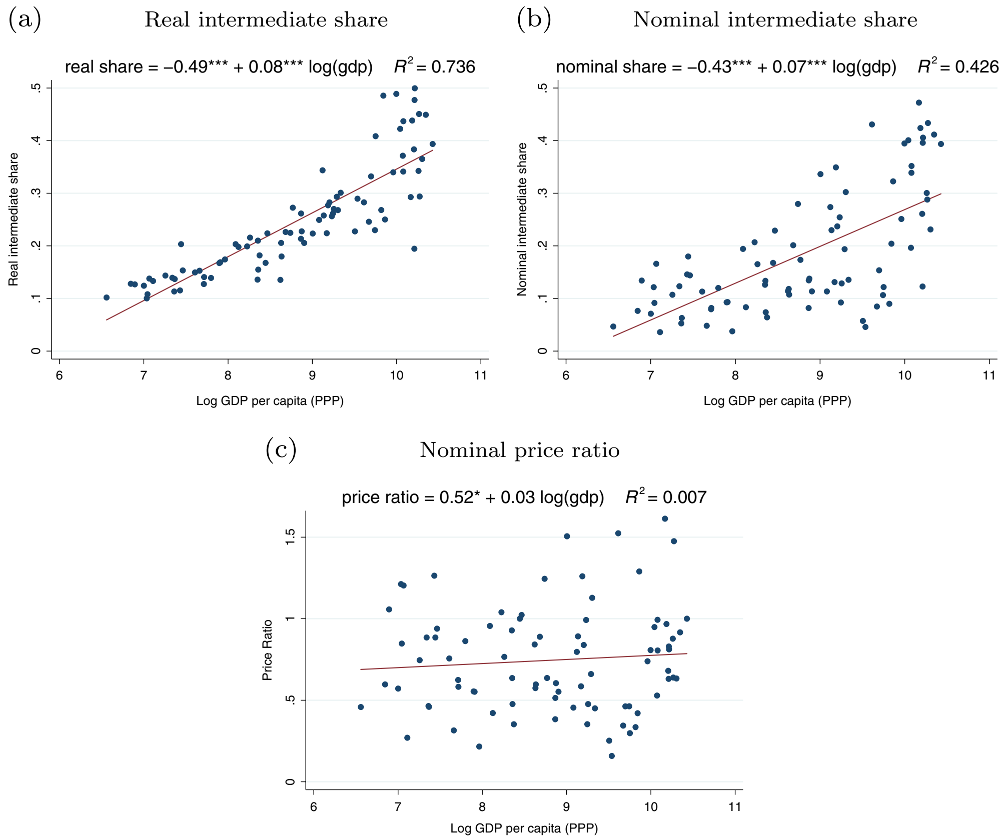

FIGURE 1

Cross-country intermediate shares (1985)

Notes: This figure plots the real and nominal intermediate shares, along with the relative price ratio of intermediates to agricultural output. They are derived from Prasada Rao (1993). Significance of coefficient estimates in the linear regression of each moment on log GDP per capita are denoted by ***, **, and *.

entirely accounted for by variation in intermediates produced in the non-agricultural sector (Figure 2d).[7,8]

The rest of this article is devoted to developing and quantifying a model to understand the causes and consequences of the correlation in agriculture.

3. MODEL

Time is discrete, and a model period is one year. There are two sectors, sector a for agriculture and sector m for manufacturing, which includes all non-agriculture. The manufacturing good is the numeraire, so its output price is normalized to $p_{mt} = 1$ for all t. Within an economy, decisions are made by a measure one of infinitely lived households.

7. See Bartelme and Gorodnichenko (2015) and Fadinger, Ghiglino, and Teteryatnikova (Fadinger *et al.*) for more detailed studies of variation within non-agricultural sectors.

8. In the Supplementary Appendix, I show the same results hold in U.N. data, which includes broader country coverage, but does not allow me to distinguish the sector producing intermediates.

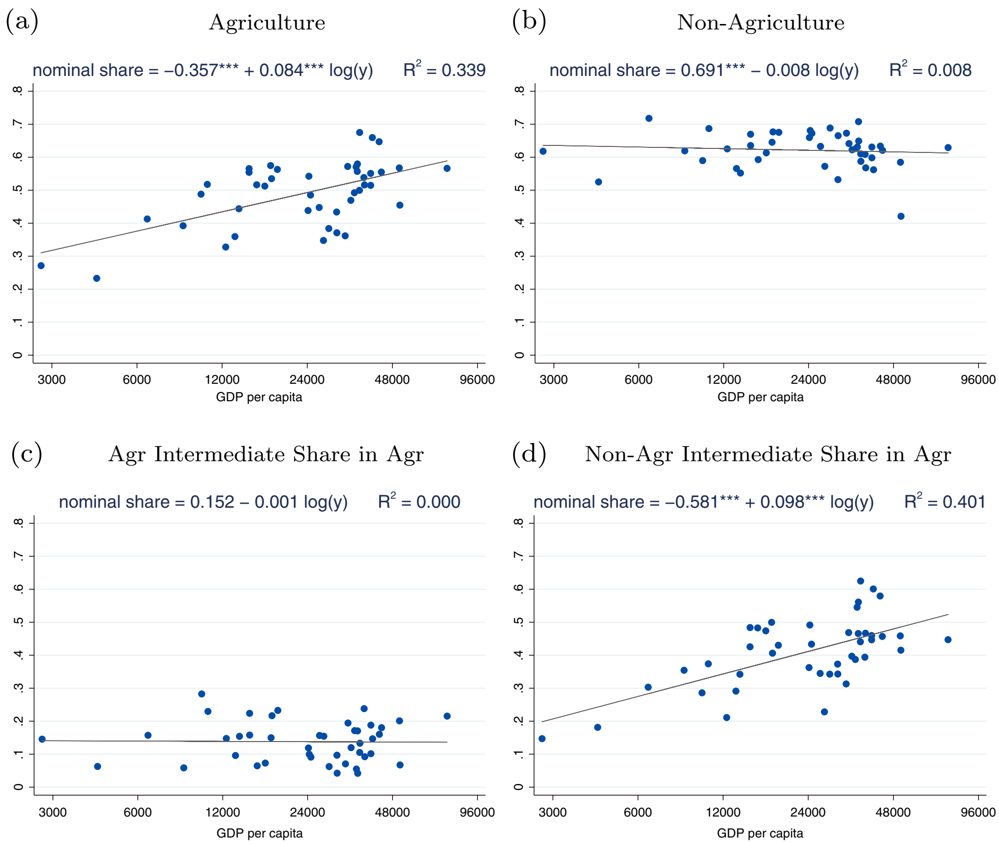

FIGURE 2

Sectoral nominal intermediate shares (2005)

Notes: This figure plots nominal intermediate input shares derived from World Input–Output Database. Non-agriculture is computed as the entire economy net of the agriculture sector. Also included are regression equations for each share regressed against log GDP per capita. Significance of coefficient estimates at 0.01, 0.05, 0.1 levels denoted by ***, **, and *.

3.1. Technology

3.1.1. Manufacturing. The manufacturing output good can be used as either consumption or intermediate inputs in agricultural production. Production is characterized by a stand-in firm which uses only labour services N_{mt} to produce output according to the constant returns to scale production function $Y_{mt} = AN_{mt}$, where A is a sector neutral TFP parameter. The parameter A is country-specific, and is a measure of the overall productivity of the economy. The firm maximizes profits at each date t, so that N_{mt} is the solution to

$$\max_{N_{mt} \geq 0} AN_{mt} - w_t N_{mt} \tag{3.1}$$

where w_t is the wage paid per unit of N_{mt}. In a competitive equilibrium, $w_t = A$ for all t.

3.1.2. Agriculture. Each household is endowed with one farm that requires intermediate inputs x and labour n_a. Production occurs according to the decreasing returns to scale production function $y_{at} = z_t A x_t^{\psi} n_{at}^{\eta}$, where $\psi + \eta < 1$ and A is, again, sector neutral TFP.[9] The shock z_t is a household-specific productivity shock drawn from a time-invariant distribution with cumulative distribution function $Q(z)$ and support on $[\underline{z}, \overline{z}]$.[10,11] The realization of z_t is i.i.d. with respect to both households and time. I assume the law of large numbers holds, so that the distribution of shocks across households is certain. Intermediate inputs are purchased from the manufacturing sector, at the price $p_x \geq 1$, which varies across countries.[12]

3.2. *Household utility and decisions*

A household values consumption from both sectors a and m, and maximizes expected utility

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t u(c_{at}, c_{mt})\right]$$

with discount factor $\beta \in (0,1)$. The period t utility flow takes the form $u(c_{at}, c_{mt}) = \alpha \log(c_{at} - \bar{a}) + (1-\alpha)\log(c_{mt})$, where c_{jt} is consumption from sector $j \in \{a, m\}$ and $\bar{a} > 0$ is the subsistence requirement of agricultural consumption. The utility function is consistent with the structural transformation paths when households consume gross output (Herrendorf *et al*., 2013).

Households do not have access to insurance markets, so that the consumption can only be insured through self-insurance. Savings is denominated in the agricultural good and depreciates at a country-specific rate δ to capture differences in agricultural savings technologies across countries (*e.g.* high spoilage rates of crop storage).

3.2.1. Decision timing. At time $t-1$, households save b_t units of the agricultural good. A fraction δ depreciates, and the household enters time t with $(1-\delta)b_t$ units of savings. The period t decision problem is broken down into two stages denoted *ordering* and *production*, which are separated by the realization of the shock z.

In the ordering stage, each household chooses intermediates x_t to use in their farm. After ordering, z_t is realized. All production and consumption occurs in the production stage. First, a household chooses how to allocate labour between the agricultural sector, where they can work on the household farm, and in the manufacturing sector, where they can work for wage w_t which is taxed at rate $\tau \geq 0$.[13] Note that this allows labour to be used to smooth consumption across shock realizations. After labour is decided, all production takes place. There is a centralized market for buying and selling goods, implying a unique equilibrium price p_a. Profits are made, all factors of production are paid, and consumption and savings choices $(c_{at}, c_{mt}, b_{t+1})$ take place.[14]

9. This assumption follows from the fact that land is in fixed supply (*e.g.* Restuccia and Santaeulàlia-Llopis, 2017). Note that decreasing returns implies that any force lowering agricultural employment will therefore directly increase labour productivity.

10. Throughout, it is assumed TFP A is high enough to guarantee subsistence can be satisfied for all economies in some set $\mathcal{A} \subset \mathbb{R}_+$. The results should be interpreted as holding for economies with TFP in that set.

11. Note that increased intermediate intensity does not decrease the variance of shocks. This is supported by micro evidence in both developed and developing countries (Just and Pope, 1979; Traxler *et al*., 1995).

12. The implicit assumption made is that there exists a technology to turn one unit of manufacturing output into $1/p_x$ units of intermediate input. This is a simple way to capture the fact that intermediate inputs are more expensive in developing countries.

13. The tax revenue is used to buy manufacturing goods that are destroyed.

14. I abstract from manufacturing risk here because manufacturing households are significantly richer. As I show in Section 6, the impact of risk is small at sufficiently high levels of savings.

3.3. *Recursive problem*

The timing described above implies that the household state variable is savings b, and the aggregate state is the distribution of savings across all households, $\mu(b)$. I suppress the dependence of the decision problem on $\mu(b)$ as I focus on the stationary equilibrium.

At the production stage, once the choice of x is made and z realized, the value of entering time t with $(1-\delta)b$ savings is

$$v^p(x,b,z) = \max_{c_a,c_m,n_a,b'} \alpha\log(c_a-\bar{a})+(1-\alpha)\log(c_m)+\beta v^o(b') \tag{3.2}$$

subject to constraint set

$$\begin{aligned} &p_a c_a + c_m + p_a b' = p_a z A x^{\psi} n_a^{\eta} - p_x x + (1-\tau)w(1-n_a) + p_a(1-\delta)b \\ &b' \geq 0, \quad c_a \geq \bar{a}, \quad c_m \geq 0, \end{aligned}$$

where v^o is the value of entering the ordering stage at $t+1$ with b' units of savings in the stationary equilibrium. The production problem in (3.2) defines decision rules as a function of the intermediate choice. Working backwards, the ordering stage value of entering time t with b savings is

$$v^o(b) = \max_{x\geq 0} \int_z v^p(x,b,z)dQ(z). \tag{3.3}$$

This defines the decision rule for intermediate inputs $x(b)$ and therefore the production stage decision rules $c_a(b,z)$, $c_m(b,z)$, $n_a(b,z)$, and $b'(b,z)$.

3.4. *Stationary equilibrium*

The stationary competitive equilibrium of this economy is defined by an invariant distribution $\mu=\mu^*$, a value function v^o, decision rules x, n_a, b', c_a, c_m, labour choice N_m, and prices p_a and w, such that (i) the value function v^o solves the households's problem given by (3.2) and (3.3) with the associated decision rules, (ii) N_m solves the sector m firm problem (3.1), (iii) the law of motion for μ, $\Lambda(\mu)$, implies $\Lambda(\mu^*)=\mu^*$, and μ^* is consistent with $Q(z)$ and decision rules, and (iv) markets clear:

(a) Labour market:

$$N_m = 1 - \int_z\int_b n_a(b,z)d\mu dQ(z)$$

(b) Agricultural goods market:

$$\int_b\int_z c_a(b,z)dQ(z)d\mu + \delta\int_b b d\mu = \int_b\int_z zAx(b)^{\psi} n_a(b,z)^{\eta} dQ(z)d\mu$$

(c) Manufacturing good market:

$$\int_b\int_z c_m(b,z)dQ(z)d\mu + p_x\int_b x(b)d\mu + \tau A\int_b (1-n_a(b,z))dQ(z)d\mu = AN_m$$

3.5. *Brief discussion on modelling choices*

Before turning to the results, I briefly discuss some modelling choices, all of which are discussed in more detail in the Supplementary Appendix.

First, the household problem in (3.2) allows intermediates to be purchased out of *ex post* farm revenue, implying no direct frictions that limit borrowing for intermediates. This decision is made for two reasons. First, it isolates the role of the missing consumption insurance market, contrasting with recent work that directly distort the input market via borrowing constraints in otherwise similar models to the one laid out here (Buera *et al.*, 2011; Midrigan and Xu, 2014; Moll, 2014). I take up this comparison in the Supplementary Appendix, showing that they both micro-found the same distortion, but have different implications for the relationship between consumption volatility and intermediate intensity. Second, and more practically, the assumption is consistent with the expansion of borrowing opportunities for farmers in India during this time. In ICRISAT micro data (discussed in more detail in the calibration), 74% of all outstanding debt value and 94% of all outstanding formal debt value (co-ops, commercial banks, financial companies, and microfinance banks) is for agricultural inputs. Thus, the model captures the key margin of farmer borrowing.

Second, the model has no margin of *ex ante* occupational choice. I similarly take this up in the Supplementary Appendix and show that the results are similar in this case. I exclude it here in the interest of analytical clarity that allows for a more simple elucidation of the model mechanics.

Third, the model assumes i.i.d. shocks, which potentially play a important role in quantitative studies of misallocation (*e.g.* Buera and Shin, 2011). I show in the Supplementary Appendix that the auto-correlation of harvest realizations is low in Indian micro-data, and contrast this to studies in other countries and sectors.

Finally, the model assumes that the underlying shock distributions are identical across countries. In the Supplementary Appendix, I show that available cross-country measures of underlying risk do not vary with income, echoing recent work by Adamopoulos and Restuccia (2018) on the relatively small importance of geographical differences in contributing to cross-country income differences. Moreover, I show that even without this empirical result, increasing variance in poor countries is unlikely to play a major role, due to the particulars of how the equilibrium price responds to higher shock variance.

4. CHARACTERIZATION AND ANALYTIC RESULTS

I begin by clarifying the mechanics of the model, with the goal of highlighting the role of risk. I start by showing how risk affects a single country under different market structures, then extend those same results to a cross-country context. I assume throughout that $p_x=1$ and $\tau=0$ in all economies to focus on the role of TFP, though none of the results change if they are included. To further make the results are sharp as possible, I consider the static version of the model (identically, $\delta=1$ for all economies). All proofs are relegated to the Supplementary Appendix.

4.1. *Formalizing the impact of risk within an economy*

I isolate the impact of risk by comparing the model developed in the previous section (denoted by superscript *IM* for incomplete markets) with a complete markets version of the economy (denoted by superscript *CM* for complete markets).[15] How market structure affects the intermediate input

15. The complete markets economy is identical in every respect, except that households can trade a full set of state-contingent assets before the realization of z. These trades allow farmers to maximize *ex ante* expected profit. It does not allow them to attain the profits that would be achieved by choosing intermediates after the shock is known.

choice can be seen by comparing the first order conditions in the *IM* and *CM* economies. Farmers maximize expected profit with complete markets, so combining the first order conditions for n_a and x yields

$$Ap_a^{1/(1-\eta)}F'(x)\int_Z z^{1/(1-\eta)}dQ(z)=1, \tag{4.1}$$

where $F(x)=x^{\psi/(1-\eta)}\left(\eta^{\eta/(1-\eta)}-\eta^{1/(1-\eta)}\right)$ and $F'(\cdot)$ is the derivative with respect to x. Without the ability to trade these claims, the same optimality condition in the *IM* economy yields

$$Ap_a^{1/(1-\eta)}F'(x)\int_Z z^{1/(1-\eta)}\left(\frac{\widetilde{u}'(C(x,z))}{\mathbb{E}_z[\widetilde{u}'(C(x,z))]}\right)dQ(z)=1, \tag{4.2}$$

where $\widetilde{u}'$ is the derivative of utility flow with respect to consumption expenditures $C:=p_ac_a+c_m$, given by

$$\widetilde{u}(C):=u(c_a(C),c_m(C))=\Omega-\alpha\log(p_a)+\log(C-p_a\bar{a}) \quad \text{for constant } \Omega. \tag{4.3}$$

The two optimality conditions in (4.1) and (4.2) are identical except for the addition of marginal utility to the integrand of (4.2). These risk neutral probabilities formalize the idea that households internalize the impact intermediate expenditures have on consumption, and shift weight toward outcomes with high marginal utility. Therefore, economy *IM* tilts the weight assigned by every household toward low shock realizations. Some algebra on (4.2) yields an implicit equation for the aggregate nominal intermediate share in the economy,

$$\frac{X}{p_aY_a}=\psi\left[\frac{\int_Z z^{1/(1-\eta)}\left(\frac{\widetilde{u}'(C(x,z))}{\mathbb{E}_z[\widetilde{u}'(C(x,z))]}\right)dQ(z)}{\int_Z z^{1/(1-\eta)}dQ(z)}\right]<\psi, \tag{4.4}$$

where the inequality follows from Jensen's inequality. Throughout the remainder of this paper, I refer to the term in the brackets of (4.4) as the "risk-induced wedge." This wedge is the focus of both this paper and the broad motivation behind the micro literature on rainfall insurance cited in the introduction.[16] When combined with decreasing returns in agricultural production, (4.4) immediately implies that missing insurance markets depress agricultural labour productivity (*i.e.* Y_a/N_a) within a given economy (*i.e.* fixing TFP A). This is summarized in Proposition 1.[17]

Proposition 1 *For any level of TFP A,*

(a) ***Risk-induced wedge:*** *The nominal intermediate share $X/(p_aY_a)$ is higher in the CM than IM economy. That is,*

$$\frac{X^{IM}}{p_a^{IM}Y_a^{IM}}<\psi=\frac{X^{CM}}{p_a^{CM}Y_a^{CM}}$$

16. Moreover, note that the Cobb–Douglas assumption implies that the magnitude of this wedge is given by the ratio of the realized nominal intermediate share (the left-hand side of 4.4) and the implied nominal intermediate share with complete markets (equal to ψ). Thus, the nominal intermediate share—while not directly relevant for productivity—provides an important moment summarizing the cost of risk.

17. The formal proof follows almost exactly the results derived in Restuccia *et al*. (2008), once one realizes that risk is a micro-foundation of the reduced form distortion studied in that paper (but does not show up in the price of intermediates). See the Supplementary Appendix for details on the isomorphism between the incomplete markets model studied here and a reduced-form distortion in a complete markets model.

(b) ***Labour productivity gains from insurance:*** *Agricultural labour productivity, Y_a/N_a, is higher in the CM than IM economy. That is,*

$$\frac{Y_a^{CM}/N_a^{CM}}{Y_a^{IM}/N_a^{IM}} > 1$$

Thus, as expected, risk makes any economy less productive.

4.2. *Risk and productivity across countries*

A second question is the extent to which the gains from insurance highlighted in Proposition 1 vary across countries of different income levels. A seemingly reasonable intuition is that because of low TFP, farmers in poor countries should bare a larger cost of risk from higher marginal utility. It turns out the result is more subtle that this when the agricultural price can vary across countries. Proposition 2 summarizes, with discussion following.

Proposition 2 *In the model with uninsurable shocks (Economy IM), the following results hold in the competitive equilibrium for any two economies defined by TFP levels $A^1 < A^2$:*

(a) ***Risk-induced wedge:*** *The nominal intermediate share is such that*

$$\frac{X^{1IM}}{p_a^{1IM} Y_a^{1IM}} \leq \frac{X^{2IM}}{p_a^{2IM} Y_a^{2IM}} \text{ with equality only if } \bar{a} = 0$$

(b) ***Labour productivity gains from insurance:*** *The gains in agricultural labour productivity when insurance is introduced are such that*

$$\frac{Y_a^{1CM}/N_a^{1CM}}{Y_a^{1IM}/N_a^{1IM}} \geq \frac{Y_a^{2CM}/N_a^{2CM}}{Y_a^{2IM}/N_a^{2IM}} \text{ with equality only if } \bar{a} = 0.$$

While Proposition 1 shows that risk both generates a wedge in intermediates and lowers labour productivity, Proposition 2 shows that translating those results into a cross-country context relies critically on the interaction of TFP differences with subsistence requirements. Specifically, when $\bar{a} = 0$, both the size of the wedge and the gains from insurance highlighted in Proposition 1 are identical across countries despite differences in TFP.

While the formal proof is in the Supplementary Appendix, the intuition relies on how TFP changes differently affect farmer risk aversion. In this model, relative risk aversion for the utility function (4.3) is given by

$$R(C;\bar{a},p_a) = \frac{C}{C - p_a\bar{a}}. \tag{4.5}$$

The key insight here is that the same non-homothenticity that generates income effects in structural transformation (Kongsamut *et al.*, 2001; Herrendorf *et al.*, 2013, 2014) simultaneously changes the properties of household risk aversion, as $\bar{a} > 0$ implies household exhibit decreasing relative risk aversion. Not only does $\bar{a} > 0$ imply DRRA, but it also allows the price p_a to directly influence the level of risk aversion by influencing the cost of subsistence $p_a\bar{a}$. This role of $\bar{a}$ is what I exploit to generate the results.

To see how this operates intuitively, first consider the world without subsistence, $\bar{a} = 0$. As emphasized in the first order condition (4.2), risk re-weights shock realizations toward those that generate high marginal utility. In partial equilibrium, this is enough to qualitatively replicate the motivating cross-country relationship between agricultural productivity and the nominal intermediate share with standard reasoning: curvature in the utility function implies that poor countries overweight low shock realizations, and therefore use fewer intermediates than their rich counterparts. In general equilibrium, however, this argument is incomplete. In response to low input expenditures, farmers are incentivized to produce in equilibrium through a higher return on investment, driving up the price p_a. In the Supplementary Appendix, I construct this $\bar{a}=0$ equilibrium and show that the higher price *perfectly* offsets increased risk aversion in poor countries. This means that while risk lowers productivity in any given country, the cost of risk is identical across countries when $\bar{a}=0$.

Once subsistence is introduced, the same equilibrium channel is operational, but with one important change. Now, that higher price has a second, countervailing effect: it directly increases risk aversion by increasing the cost of subsistence $p_a\bar{a}$. This second effect eliminates the ability of the price to perfectly offset the partial equilibrium increase in risk aversion generated by low TFP in poor countries, and thus guarantees the cost of risk is larger in poor countries than rich.

Overall, this section shows two things. First, risk depresses agricultural productivity. Second, risk differentially affects poor countries only when combined with subsistence requirements, highlighting the importance of general equilibrium effects when extrapolating risk to a broader cross-country context. Of course, these results are derived in a simplified model in the pursuit of analytical clarity. To fully quantify the impact of risk, I turn back to the full dynamic model laid out in Section 3 and study the impact of this risk-generated wedge on critical real variables – intermediate inputs, labour, and output.

5. CALIBRATION AND TESTING MODEL PREDICTIONS

To study the quantitative implications of risk in general equilibrium, I first calibrate the baseline model to India using a combination of micro and aggregate statistics. Because I will eventually study the cross-country results, I also construct a second economy designed to match features of the United States. This model economy differs only in its aggregate productivity A, along with a set of agriculture-specific distortions detailed below, while holding the remaining calibration fixed. All quantitative results follow in Section 6.

The calibration works as follows: I choose parameters in the Indian and U.S. economies, solve them both, and continually update until the moments are matched. The reason I am required to loop over both countries is because I directly target the gap in real GDP per capita across the two countries. This is discussed in detail below.

I calibrate the baseline Indian model in part with the ICRISAT Village Level Studies (VLS) panel survey from India, covering 2009–14. The data include a panel of households across 18 villages in 5 states, with detailed information on inputs and outputs collected at a high frequency, and prices.

5.1. *Calibration*

There are 12 parameters both the Indian and U.S. economies. Four of these vary between the U.S. and India economies, while the other nine are held fixed across countries.

The four parameters that vary are sector-neutral TFP A, storage depreciation δ, the earnings tax τ, and the price of intermediates p_x. In the U.S., these parameters are set to create an undistorted benchmark, implying $\tau^{\text{U.S.}}=0$, $p_x^{\text{U.S.}}=1$, $\delta=0$, and a normalization of $A^{\text{U.S.}}=1$.[18]

18. This choice of $\delta=0$ is irrelevant for the quantitative results. The U.S. model economy is sufficiently rich that the savings distribution at any negative interest rate is nearly degenerate at zero.

I then choose the twelve parameters of the Indian economy to match a number of aggregate and micro-level moments. Of these twelve, six are chosen jointly, while the other six are set exogenously. These steps are discussed in turn, with a full list of moments and parameter values in Table 1.

5.1.1. Indian parameters set exogenously. There are six parameters set exogenously. First, in terms of the utility function, I set $\beta=0.96$ because the model period is a year and $\alpha=0.005$ following Restuccia *et al.* (2008) and Lagakos and Waugh (2013).[19] The latter implies that in the long run, agricultural consumption accounts for 0.5% of consumption.

Second, I exogenously set some parts of the production technology, in particular, the technological exponents of the Cobb–Douglas production function. These cannot be set to match factor shares in the Indian economy, as the realized shares combine both the technological parameters and the distortions I seek to investigate. Instead, I assume the technologies are the same across U.S. and India and choose $\eta=0.42$ and $\psi=0.40$ to be consistent with U.S. estimates (Restuccia *et al.*, 2008). As I will show below, this is possible because the incompleteness of markets has no quantitative effect on the U.S. economy. Thus, these can be set as if matched directly to the nominal shares in the U.S..

The remaining exogenously chosen parameters are two of the four Indian distortions. I set $p_x^{\text{India}}=2.77$ from Restuccia *et al.* (2008). The last parameter set exogenously is the tax on manufacturing labour income, τ^{India}. I calibrate this using ICRISAT data. ICRISAT includes household earnings and days worked in both agricultural and nonagricultural wages over the 6 year panel. I run the regressions

$$\log(w_{isvt})=\alpha X_{ivt}+\theta_{vt}+\gamma\,\mathbb{1}_{isvt}[s=m]+\varepsilon_{isvt} \quad \text{for } j\in\{a,m\},$$

where w_{isvt} is the wage earned by household i in village v at year t in sector $s\in\{a,m\}$. X controls for differences in household composition, education, and age that are not included in the model by may affect wages, while θ_{vt} is a village-year fixed effect. The key term here is the coefficient γ, which estimates the change in wages if the wage is earned in the manufacturing sector. This is given by the indicator $\mathbb{1}[s=m]$.

The results imply $\hat{\gamma}=0.32$, which I use as my measure of τ.

5.1.2. Indian parameters chosen jointly. With the six parameters in hand, I am left with six parameters to target. They are the subsistence requirement $\overline{a}$, sector-neutral TFP A, the depreciation on savings δ, and three features of the shock distribution. I assume the shock distribution is a mean zero truncated log-normal distribution, which requires bounds $(\underline{z},\overline{z})$ and standard deviation σ_z. While chosen jointly, each has an intuitive match among the target moments. I describe each target with its model counterpart individually, with the caveat that all are jointly chosen to match moments in equilibrium.

First, I set $\overline{a}$ to match a 44% agricultural employment share in India, consistent with World Bank (2018). I choose sector-neutral TFP $A^{\text{India}}=0.20$ to match a real GDP per worker ratio of 8.2 between the U.S. and India. The remaining 4 moments are selected to match moments from ICRISAT. The targeted moments are the top and bottom 1% realizations of the 2014 cross-sectional harvest distribution (matched to $\underline{z}$ and $\overline{z}$), the average household-level standard deviation of harvest (σ_z), and the value of savings relative to harvest (δ). These are discussed in turn.

19. The discount factor is arbitrary here given the exogeneity of the interest rate. I recalibrate the model in the Supplementary Appendix using lower discount rate of $\beta=0.90$ and show that the results are nearly identical.

TABLE 1
Calibration targets and fit

Model parameter	Description	Parameter Value	Target moment	Source	Target Value	Model Value
Common parameters in both countries						
α	Utility weight on agriculture	0.005	Long run agr. consumption share	common literature value	0.005	0.005
β	Discount factor	0.96	Standard value	common literature value	0.96	0.96
$\bar{a}$	Subsistence consumption†	0.03	Agricultural employment share	World Bank (2018)	0.44	0.44
ψ	Exponent on intermediates	0.40	Agr. intermediate share in U.S.	Restuccia *et al*. (2008)	0.40	0.40
η	Exponent on labour	0.42	Agr. labour share in U.S.	Restuccia *et al*. (2008)	0.42	0.42
σ_z	Standard deviation of shock†	0.32	Avg. household harvest variance over time	Computed from ICRISAT	0.67	0.72
$\underline{z}$	Minimum shock realization†	0.001	1st percentile/Mean harvest realization	Computed from ICRISAT	0.03	0.04
$\bar{z}$	Maximum shock realization†	5.00	99th percentile/Mean harvest realization	Computed from ICRISAT	6.64	6.33
India-specific parameters						
τ	Tax on manufacturing wages	0.32	Relative agr/non-agr wage	Computed from ICRISAT	0.32	0.32
δ	Storage depreciation rate†	0.10	Avg. savings / Harvest	Computed from ICRISAT	0.77	0.77
p_x	Relative price of intermediates	2.77	Intermediate price	Restuccia *et al*. (2008)	2.77	2.77
A	Sector-neutral TFP†	0.20	GDP per worker ratio between U.S. and India (2014)	World Bank (2018)	8.23	8.21
U.S.-specific parameters						
τ	Tax on manufacturing wages	0	Distortion-less benchmark	n.a.	0	0
δ	Storage depreciation rate	0	Distortion-less benchmark	n.a.	0	0
p_x	Relative price of intermediates	1	Distortion-less benchmark	n.a.	1	1
A	Sector-neutral TFP	1	Normalization	n.a.	1	1

Notes: Unless otherwise specified, all targets refer to moments from the Indian economy. Descriptions marked with a † are jointly calibrated, while those without are set exogenously.

The depreciation rate of savings δ determines the availability of smoothing by running down savings. This requires defining a notion of savings in the empirics. I measure savings as the total value of all net savings, livestock, and crop storage relative to consumption expenditures.[20] Recall from Section 3 that the model already includes any borrowing for farm inputs. Since ICRISAT provides the reason for borrowing, I compute the proper measure of net savings as the total value of all savings net of any outstanding borrowing for consumption.[21] The total value of savings relative to harvest value averages 77% and I set $\delta^{\text{India}}=0.10$ to match this fact. This is consistent with large costs to save in developing countries, a point similarly highlighted in Lagakos *et al.* (2020).[22]

I next compute the individual-level harvest variation. This uses the six years panel in ICRISAT. However, like with the previously estimated wages, the data include variation due to heterogeneity in household size, education, and village-level variation that are not modelled here. To the extent that these are predictable, directly using variance in the data would attribute them to unanticipated shocks. Instead I follow Kaboski and Townsend (2011) and others and purge the data of these factors with the regressions

$$\log(Y_{ivt})=\alpha+\beta X_{ivt}+(1+g_v)t+\varepsilon_{ivt} \tag{5.1}$$

where g_v is a village-level time trend, and X_{ivt} is a set of controls for household i at year t that include number of men, women, and children, and age, gender, and education of the household head. The R^2 on this regression is 0.32, so that these features account for about one-third of the total variation in harvest. I match the standard deviation of the error term ε_{ivt} to harvest variation in the stationary equilibrium. This implies $\sigma_z=0.32$.[23]

Finally, I measure the top and bottom 1% of harvest realizations normalized by the mean. Conditional on the standard deviation, this implies bounds for the shock distribution of $\underline{z}=0.001$ and $\overline{z}=5.00$ to match the fact that the first percentile harvest realization is 3% of the mean, while the the ninety-ninth percentile realization is 6.64 times the mean.

While only the household-averages of harvest variance and the savings rate are targeted, I also check whether the model provides a reasonable approximation of the full density. Figure 3 plots the results.[24] The model matches the savings distribution well (Figure 3b). Harvest volatility (Figure 3a) is a slightly worse fit. The model has difficulty matching the thickness in the right tail, and thus compensates by shifting the mode of the distribution slightly to the right of its empirical counterpart.

20. All three of these stocks have been documented as an important smoothing vehicle in India. See Rosenzweig and Wolpin (1993) on the role of livestock and Udry (1995), Fafchamps et al. (1998), and Kazianga and Udry (2006) on smoothing via stored crops.

21. As discussed in Section 3, most borrowing is for farm inputs. Moreover, when measured this way, only 3% of ICRISAT households have negative asset holdings. Thus, the constraint of positive asset holdings is a reasonable one.

22. While this article does not take a stand on the underlying microeconomic causes of the high cost of savings, there are a variety of potential explanations. In addition to the high storage depreciation rates for crops, Jakiela and Ozier (2016) and Goldberg (2017) highlight the cost of social and family pressure, while Anagol *et al.* (2017) discuss the value of holding livestock despite the potentially negative interest rate. See Karlan *et al.* (2014) for a review on the high cost of saving in the developing world.

23. Alternatively, I could calibrate the shock distribution to the United States or some other rich country. However, because poor countries are much closer to subsistence, they respond more strongly to changes in the shock distribution. As previously discussed, rich countries respond little to changes in the variance of shocks, as they act quite close to profit maximizing households. Therefore, it is more important for the quantitative results to properly match the shock distribution of the poor country, while matching the variance of shocks in rich countries is of little quantitative importance. I therefore assume that they are the same to minimize differences across economies.

24. I compute the implied volatility in the model by simulating 500,000 individuals for 6 years each, consistent with the ICRISAT panel length.

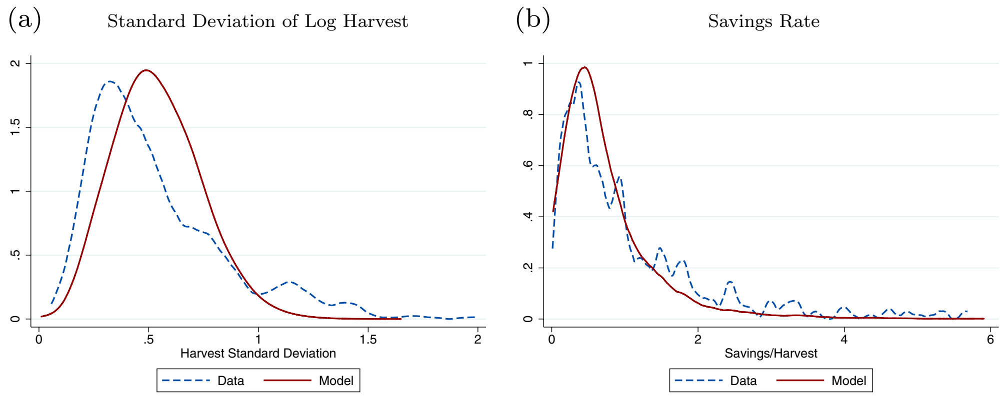

FIGURE 3

Fit of model distributions

Notes: Savings rate is measured as the value of savings as a share of the value of harvest. The dashed lines are computed from ICRISAT. The solid lines are the implied model counterparts, which are computed by simulating 500,000 individuals in the Indian model economy for 6 periods.

Table 1 summarizes value of all parameters and how they match the targeted moments. Overall, the model matches the data well on targeted moments. One moment the model misses is consumption volatility, which the model underestimates. Thus, if anything, the model is providing too many smoothing opportunities for households.[25]

5.2. *Model versus data: comparing micro moments in India*

Before turning to the results, I assess the model's predictions on untargeted moments of the Indian data, covering the relationship between consumption volatility, savings, and intermediate intensity. To compute the regressions in the model, I simulate 500,000 households in the Indian stationary equilibrium. The point estimates from this procedure measure the true model relationships (net of the small error induced by "only" simulating 500,000 households), and thus do not require standard errors. However, for comparison's sake, I compute them by bootstrapping 1000 samples of populations equal to the ICRISAT sample size.

The results are broken into two sections. The first covers to the relationship between savings and intermediate use. The second discusses the relationship between consumption volatility and intermediate use. In both cases, the model matches the data well.

5.2.1. Savings and intermediate intensity. I consider the relationship between savings and agricultural outcomes with the regression

$$outcome_{i,v,t} = \alpha + \beta b_{i,v,t-1} + \theta_{v,t} + \varepsilon_{i,v,t}, \tag{5.2}$$

where the village-year fixed effect θ_{vt} allows me to consider the variation within a village-year (and is thus the correct counterpart to the stationary model). Savings is lagged to limit any risk of observing the empirical counterpart of the *ex post* savings response $b'(b,z)$ instead of the *ex ante*

25. Previous versions of this paper targeted this moment directly, by rebating some of the tax τ to households. The current calibration, with the under-prediction of consumption volatility, would push that transfer to zero.

TABLE 2
Savings and agricultural production (model versus data)

	Expenditures		Nominal share		Harvest	
	Data	Model	Data	Model	Data	Model
Savings	0.414 (0.021)***	0.426 (0.016)***	−0.018 (0.020)	0.069 (0.010)***	0.382 (0.019)***	0.301 (0.024)***
Obs	2389	n.a.	2367	n.a.	2389	n.a.
R^2	0.335	0.714	0.085	0.033	0.334	0.119

Notes: Significance at 0.01, 0.05, 0.1 levels denoted by ***, **, and *. The model standard errors are bootstrapped using 1000 samples of 205 individuals. Dependent and independent variables are normalized by sample mean.

state b. I consider intermediate expenditures ($p_x x$ in model notation), the nominal intermediate share ($(p_x x)/(p_a y_a)$), and harvest value ($p_a y_a$). In all regressions, both dependent and independent variables are normalized by their respective means. The results are presented in Table 2.

First, Table 2 shows that the model and data match well when considering intermediate expenditures, where $\widehat{\beta}^{\text{data}}=0.41$ and $\widehat{\beta}^{\text{model}}=0.43$. Both are significant at 1%. In contrast, the realized intermediate shares are partially driven by unexpected volatility in output. Correspondingly, the R^2 decreases from 0.71 to 0.03 in the model and from 0.34 to 0.09 in the data. Moreover, despite the strong relationship between savings and expenditures in the model, the magnitude of relationship get weaker when I consider the share. The model prediction is somewhat stronger than the data, with $\widehat{\beta}^{\text{model}}=0.07$.[26] Lastly, the final two columns consider the relationship between harvest value and lagged savings. Unlike the intermediate share regression, the relationship between harvest and savings is significant at 1% in both model ($\widehat{\beta}^{\text{model}}=0.30$) and data ($\widehat{\beta}^{\text{data}}=0.38$), though the model slightly under-predicts the quantitative relationship.[27]

5.2.2. Consumption volatility and intermediate intensity.

I next turn to the relationship between consumption volatility (measured as the coefficient of variation in consumption) and intermediate use during the 6-year panel. To compute consumption, I follow a similar procedure to harvest and compute

$$\log(C_{ivt})=\alpha+\beta X_{ivt}+(1+g_v)t+\varepsilon_{ivt}. \tag{5.3}$$

I then measure consumption as

$$\widehat{\log(C_{ivt})}=\widehat{\alpha}+\widehat{\beta X}+\widehat{\varepsilon}_{ivt},$$

26. This is to be expected, given the model's simple design to highlight risk. Adding alternative channels and aspects of rural life could potentially match this insignificant effect. Hence, I focus on the difference between the expenditure value and the share.

27. Despite the fact that both intermediate share and harvest are *ex post* realizations, the relationship with savings is much stronger for harvest. The model provides a simple rationale for this result. Solving for the realized nominal intermediate share and harvest value implies

$$\frac{p_x x}{p_a y_a(x,z)} \propto x^{\frac{1-\eta-\psi}{1-\eta}},$$

$$p_a y_a(x,z) \propto x^{\frac{\psi}{1-\eta}}.$$

The calibrated model implies $(1-\eta-\psi)/(1-\eta)=0.33$ and $\psi/(1-\eta)=0.67$, so that variation in intermediate expenditures—highly correlated variation in savings—is a larger component of variation of harvest than the intermediate share.

TABLE 3
Consumption volatility and intermediate inputs (model versus data)

	Model	Data (estimated)	Data (direct)
c.v. of consumption	0.027	0.121	0.103
	(0.031)	(0.042)*	(0.039)*
Obs	n.a.	478	478
R^2	0.002	0.008	0.301

Notes: Significance at 0.01, 0.05, 0.1 levels denoted by ***, **, and *. The model standard errors are bootstrapped using 1000 samples of 205 individuals. Column 2 computes consumption as the resulting estimates from regression (5.3). Column 3 computes household consumption directly from ICRISAT data, deflating by the Indian CPI. Dependent and independent variables are normalized by sample mean.

where $\overline{X}$ is the average value of the controls. I then compare volatility of consumption with intermediate intensity. Specifically, denoting $\widehat{x}_{it}$ as the realized nominal intermediate share of household i at time t, I run the regression

$$\frac{\sum_{t=2009}^{t=2014}\widehat{x}_{it}}{6}=\alpha+\beta\left(\frac{\sigma^C}{\mu^C}\right)_i+\varepsilon_i, \tag{5.4}$$

where $(\sigma^C/\mu^C)_i$ is the coefficient of variation in consumption for household i over the 6 years of data. I compute the same regression estimates in the model by simulating 500,000 households for 6 years. The preferred data comparison is the consumption estimates from regression (5.3), in which $C_{it}=\exp(\widehat{\log(C_{ivt})})$, for the reasons discussed in Section 5.1.2. For completeness, I also report results the household consumption directly from ICRISAT (deflated by the Indian CPI). The results are in Table 3. The model is consistent with the positive relationship between the consumption coefficient of variation and the average intermediate share found in the data, in which the model predicts $\widehat{\beta}^{\text{model}}=0.03$, compared to $\widehat{\beta}^{\text{est}-\text{data}}=0.12$ using the estimated consumption data. The point estimate is lower in the model, but within the 95% confidence interval of either empirical result.[28] Finally, note that the noise in both the realized intermediate shares and consumption process predicted by the model is partially driven by shocks. Thus, over a short time horizon the model predictions are insignificant when drawing from the relatively small bootstrap sample.

6. QUANTITATIVE RESULTS

I now use the model to understand the equilibrium impact of risk. This proceeds in three steps. Section 6.1 begins by focusing on India. I compare two Indian economies, one without and one with insurance, and use the results to highlight the channels through which risk dampens productivity in a relatively poor country. Section 6.2 then links these results to cross-country income differences via a comparison to the U.S. model economy, a country in which the mechanisms highlighted in India have little quantitative bite. Finally, motivated by the mechanisms

28. The results in Table 3 are derived from the link between consumption and input decisions in the model. If I replace the distorted consumption market with an explicit input market distortion (*e.g.* through collateral constraints), the model can similarly depress aggregate intermediate intensity, but predicts a strong negative relationship between consumption volatility and intermediate intensity at the household level. Thus, at least within the context of such Aiyagari (1994)-style models, this moment can be used to distinguish between these two possible microfoundations of low intermediate intensity. See the Supplementary Appendix for more details.

TABLE 4
Aggregate Indian moments

	Labour productivity			
Economy	Agriculture	GDP p.c.	p_xX/p_aY_a	N_a (%)
Incomplete markets	1.00	1.00	0.31	44.3
Complete markets	1.16	1.11	0.40	35.9
%Δ	16.3	11.4	28.5	−18.8

Notes: Incomplete markets productivity measures are normalized to one. GDP p.c. stands for real GDP per capita, my measure of aggregate labour productivity in the model.

highlighted in the first two sets of results, Section 6.3 returns to India to highlight the importance of general equilibrium effects when considering policy responses in agriculture. Specifically, I study the impact of scaling a promising policy response to risk—the introduction of improved seed designed to limit downside risk. I compare the partial equilibrium model predictions to a randomized controlled trial in India, then assess the model's predictions for introducing such a technology at scale, focusing on the importance of the general equilibrium effects highlighted in the previous sections.

6.1. *The impact of risk in the Indian economy*

Table 4 reports aggregate moments from the Indian economy under the baseline incomplete markets model and a counterfactual complete markets model. On aggregate, the introduction of insurance causes the nominal intermediate share to increase from 0.31 to 0.40, an increase of 28%. This risk-induced distortion has important productivity and employment effects. The employment share in agriculture decreases from 44.3% to 35.9%, a 19% decline, while labour productivity increases both in agriculture and at the aggregate.[29] Agricultural labour productivity increases by 16.3%, while GDP per capita increases by 11.4%.

What forces increase Indian productivity when insurance is introduced? There are two key channels: insurance affects both the first and second moments of the real intermediate share distribution across households. First, insurance increases the average real intermediate share. Naturally, this increases agricultural labour productivity. The second channel is that eliminating risk increases allocative efficiency across farmers, even conditional on the mean realization. This decrease in misallocation also contributes to increased agricultural productivity by increasing what would be measured agricultural TFP.

The relative importance of these two channels is a quantitative question. I isolate risk-induced misallocation by asking how the baseline model compares to a counterfactual model with an identical average real intermediate share, but no variation across households. That is, it isolates the importance of cross-household variation. In practice, this requires a complete markets model (to eliminate variance) with an intermediate price that is 21% higher than baseline, increasing from $p_x = 2.77$ to $p_x = 3.36$ (to equalize average real intermediate shares). I refer to this economy as the "counterfactual" complete markets economy.

29. I measure real Indian GDP using U.S.-model prices. Explicitly, letting superscript I denote India and U denote the U.S., GDP per capita is measured as

$$y^I = p_a^U Y_a^I - p_x^U X^I + Y_m^I.$$

Agricultural productivity is measured as Y_a^I/N_a^I, and when used, real agricultural value added per worker is measured as $(p_a^U Y_a^I - X^I)/N_a^I$.

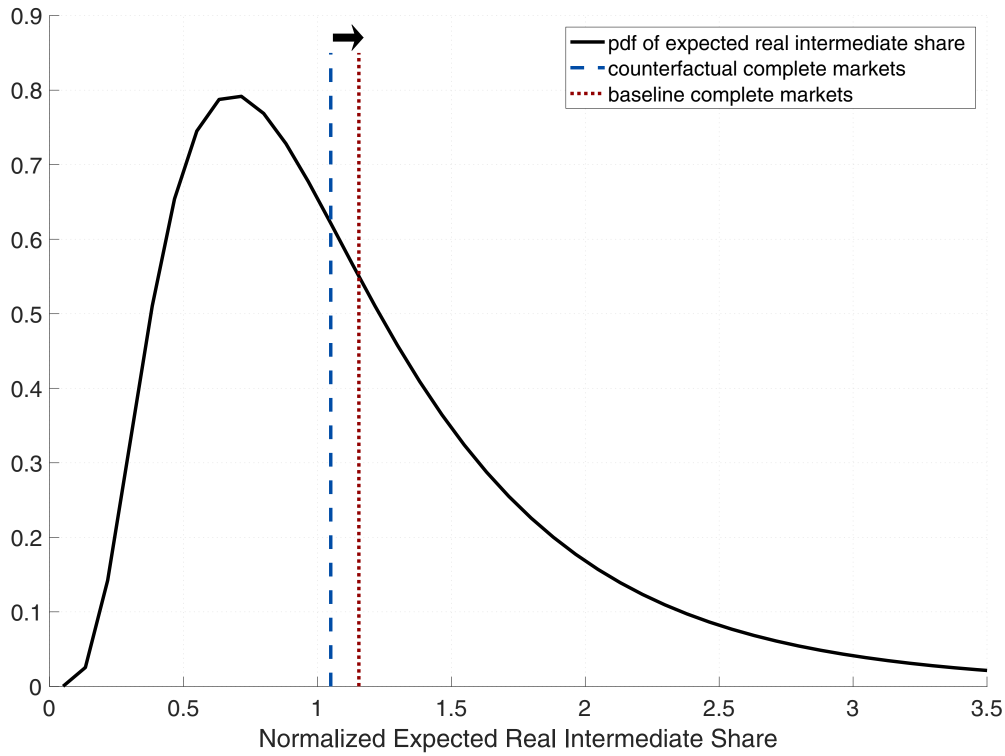

FIGURE 4

Distribution of expected real intermediate shares in India

Notes: The solid curve is the probability density function of expected real intermediate shares in the baseline Indian incomplete markets model, measured as deviations from the mean household-level realization (*i.e.* 0.5 equals half of the mean realization). The vertical lines are aggregate real intermediate shares in the two complete markets model for India—the baseline complete markets model ($p_x = 2.77$) and the counterfactual model ($p_x = 3.36$).

Figure 4 shows the procedure graphically. It plots the p.d.f. of the expected real intermediate share in the baseline economy, along with those of the baseline complete markets model ($p_x = 2.77$) and the counterfactual complete markets model ($p_x = 3.36$).[30] The latter two naturally have degenerate distributions given the model setup, thus eliminating any cross-household variation. The incomplete markets model, on the other hand, implies substantial variation in the expected real intermediate share of Indian farmers, with a coefficient of variation of 0.18. Moreover, note that variation in the real intermediate share is equal to variation in the nominal intermediate share, as farmers in the same country face the same prices. This links the variation in real shares, and thus productivity, to the risk-induced nominal wedge between the marginal revenue and price of intermediate inputs.

Table 5 shows the equilibrium results from these various model economies. Rows one and two have the same real intermediate share. Yet despite that, the second economy utilizes 11% less employment in agriculture. This is the importance of misallocation—even conditional on identical average real intermediate share, the risk-induced variation across households lowers productivity. Eliminating only this variation, while holding the average real share fixed, increases

30. The expected real intermediate share of a household is given by $x/\mathbb{E}_z[y_a(z,x)]$, while the nominal share is $(p_x x)/(p_a \mathbb{E}_z[y_a(z,x)])$.

TABLE 5
Decomposition of aggregate effects in India

	Exogenously varied	Equilibrium response			
	p_x	X/Y_a	N_a	Y_a/N_a	p_a
(1) Baseline incomplete markets	1.00	1.00	1.00	1.00	1.00
(2) Counterfactual complete markets: equalized real intermediate share	1.21	1.00	0.89	1.06	0.94
(3) Baseline complete markets: baseline calibration	1.00	1.10	0.81	1.16	0.86

Notes: This table decomposes the amplification and misallocation effects in the Indian model economy relative to the baseline incomplete markets model. The *Baseline Incomplete Markets* row is normalized to one.

agricultural productivity by 6%, and accounts for 36 of the total impact on agricultural labour productivity. The remaining 64% is due to the 10% increase in the average real intermediate share induced by insurance. Both margins therefore play an important role in generating the aggregate effects of Table 4.

The relative price of agriculture is 16% (= 1/0.86) higher in the baseline incomplete markets model (row one) than the baseline complete markets model (row three), compared to only 10% (= 0.94/0.86) in the counterfactual complete markets model. As discussed in Section 4, this higher price is critical to maintain the misallocation effect. It increases the cost of subsistence and thus risk aversion, which generates the relatively high variance across households. How quantitatively reasonable is this effect? The model predicts that $(p_x^{\text{India}}/p_a^{\text{India}})/(p_x^{\text{U.S.}}/p_a^{\text{U.S.}}) = 0.58$. The same statistic implied by Restuccia *et al.* (2008) is 0.55, so the amplification predicted by the model is well within reason of variation in the data.

Finally, note that this misallocation effect has important quantitative consequences for existing theories of agricultural productivity conducted in complete markets models, Gollin *et al.* (2004), Restuccia *et al.* (2008), and Gollin and Rogerson (2014) among many others. The assumption of complete markets eliminates the possibility of this misallocation effect, and likely understates the true cost of such distortions. I return to this issue in the Supplementary Appendix.

6.2. *Implications for cross-country productivity differences*

Since the quantitative magnitudes of the channels discussed above rely on the link between poverty and willingness to take risk, they are likely to impact rich countries differently. This intuition is formalized in Section 4, and I now turn to the quantitative magnitudes. I do so by comparing the Indian economy to the calibrated U.S. economy described in the previous section. I isolate the cross-country impact of risk by comparing India to the U.S. under two scenarios: one in which markets are incomplete (in both countries) and risk therefore plays a role, and then in a counterfactual economy in which markets are complete. The latter identifies the cross-country productivity gap induced directly by country-specific differences in agricultural distortions (p_x, τ, δ) and sector-neutral TFP A (see Table 1 for values).

The main results on productivity and input mix are in Table 4. Indeed, the relative wealth of U.S. farmers implies that insurance has little impact. Both the nominal intermediate share and employment vary little with the introduction of complete markets. Figure 5 plots the risk-generated wedge across various levels of savings in the U.S. and India, showing that while it can

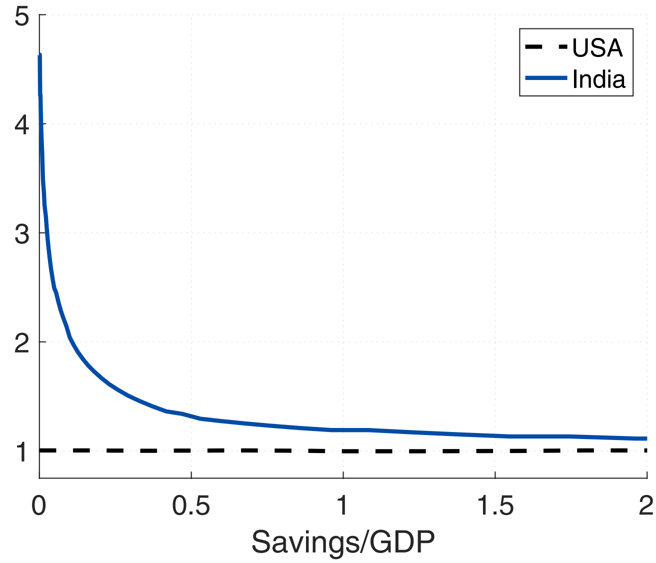

FIGURE 5

Risk-generated distortion ($\phi^{\text{risk}}(b)$)

Notes: Shows the implied household-level distortion generated by risk in both the Indian and U.S. model economies. As discussed in the text, it is the percentage difference between the expected marginal revenue and marginal cost of intermediate use, and is derived formally in the Supplementary Appendix. A value of one implies no distortion. The horizontal axis is the value of savings normalized by that economy's GDP.

be quite large at low levels of savings in India, it is nearly uniformly non-existent among richer U.S. farmers.[31]

Thus, when combined with the results of Section 6.1, both cross-country input and productivity gaps are substantially larger with incomplete markets. The results are in Table 6. The model captures $(0.40-0.31)/(0.40-0.10)=30\%$ of the nominal intermediate share difference between the U.S. and India.[32][33]

The agricultural productivity gap between the two countries then declines, as insurance in both countries primarily benefits Indian productivity. The agricultural labour productivity gap declines by 14% from 35.9 to 30.9 when insurance is introduced. Relatedly, the ratio of real agricultural value added per worker declines from 46.8 to 45.9. As one would expect, this gap is smaller than that measured in output per worker, as the depression in Indian intermediate use weakens the change in value added terms. Here, that difference is only 2%, driven by the large increase in the real intermediate share. At the aggregate level, these sectoral changes imply that the cross-country gap in real GDP per worker declines by 10.2%, from 8.2 to 7.4.

31. This is the wedge between expected marginal revenue and the price of intermediate inputs, defined as $\phi^{risk}(b) := \frac{\psi}{(p_x x(b))/(p_a \mathbb{E}[y_a(b)])}$. See the Supplementary Appendix for a formal derivation of this relationship as a reduced-form tax wedge in a complete markets model.

32. Restuccia *et al.* (2008) report a nominal intermediate share of 0.09 from FAO data and the World Input-Output Database (Timmer *et al.*, 2015) report 0.13, similar to the $0.10\times$ used here.

33. An alternative exercise is to lower exogenous agricultural productivity until the predicted real intermediate share differences matches the data. Targeting this moment directly implies a productivity too low to satisfy subsistence among the households who receive the lowest shock realization while holding no savings.

TABLE 6
Quantitative cross-country results

Economy	Labour productivity gap (U.S./India)		p_xX/p_aY_a		N_a (%)	
	Agriculture	GDP p.c.	U.S.	India	U.S.	India
Data	51.6	8.2	0.40	0.10	1.4	44.3
Model with						
Incomplete markets	35.9	8.2	0.40	0.31	1.4	44.3
Complete markets	30.9	7.4	0.40	0.40	1.4	35.9

Overall, the introduction of risk provides a substantial amplification of cross-country productivity differences at both the agricultural and aggregate level relative to a complete markets model.

6.3. *The equilibrium impact of improved seeds in India*

The results so far highlight two things: risk plays a quantitatively important role in understanding productivity, and moreover, equilibrium changes play a key role in generating the overall cost of risk. While useful for highlighting these mechanisms, comparing the uninsured and completed insured Indian economies is perhaps not a directly policy-relevant comparison, and will certainly overstate the gains from any reasonable policy proposal. I take this issue up here. I study the importance of the channels highlighted in Sections 6.1 and 6.2 for understanding the gains from a plausible and concrete policy change—the introduction of flood-tolerant seeds. I validate the model's predictions against results from a randomized controlled trial in India (Emerick *et al.*, 2016), and the assess the model's predictions for implementation at scale, taking the equilibrium effects of the last section into account.

6.3.1. Basics of the experiment. Emerick *et al.* (2016) provide flood-tolerant seeds—Swarma-Sub1—to Indian farmers in an area in which flooding affects 19% of plots annually. The randomized controlled trial delivered free Swarma-Sub1 with a short explanatory information sheet, but farmers were otherwise not provided with managerial skills or information, such as how to plant the field or how much fertilizer to use. Excluding its flood tolerant properties, these seeds are identical to standard seeds (Swarma) and leave plots unaffected in non-flood years given a fixed input bundle. Thus, from an *ex ante* perspective, these seeds both increase the mean and decrease the variance of a weather shock.

Treatments were randomized in a two-step procedure. Villages were randomized into treatment and control. In the treatment villages, five farmers were randomly selected to receive the new seeds, while ten were randomly selected to serve as part of the control group. In each control village, five farmers were selected as control as well.[34] This randomization allows for a causal interpretation of the estimates. Moreover, usage conditional on treatment status was high, which sidesteps issues related to interpreting the difference between intent-to-treat and treatment-on-the-treated effects.[35]

34. Note that the small number of farmers treated in each village makes it unlikely that one would expect village-level general equilibrium effects from the experiment.

35. This issue arises in a number of papers that introduce formal rainfall insurance, which have proven difficult to implement for a variety of reasons outside the scope of this paper. See discussions in Cole *et al.* (2013) and Karlan *et al.* (2014).

6.3.2. Running the experiment in the model. In the model, this experiment corresponds to a change in the probability distribution faced by Indian farmers randomly selected to use Swarma-Sub1, while holding the remaining parameters fixed at the values described in Section 5. I construct the new shock distribution faced by treatment farmers by shifting weight away from "flood" realizations toward "non-flood" realizations. For all $z \leq \widehat{z}$, where $G(\widehat{z}) = 0.19$, I decrease the probability of occurrence by a fixed factor (denoted λ), then uniformly distribute that probability to all $z > \widehat{z}$. Formally, after discretizing the shock distribution, this new probability distribution P_1 is constructed as

$$P_1(z) = \begin{cases} \lambda P_0(z) & \text{if } z \leq \widehat{z} \\ P_0(z) + \frac{(1-\lambda)\sum_{z \leq \widehat{z}} P_0(z)}{\sum_{z > \widehat{z}} \mathbb{1}_z} & \text{if } z > \widehat{z}. \end{cases}$$

I run the experiment in the model and compute the short-run partial equilibrium results, iterating over λ until the treatment effect for farm yields match those generated in Emerick *et al.* (2016). That procedure works as follows. I randomly allocate treatment status to 500,000 treatment and control farmers, each of whom starts from a randomly selected point in the baseline Indian stationary equilibrium. Treatment farmers are then subjected to an unexpected shift in their probability distribution from P_0 to P_1. For this exercise, I assume treatment farmers treat the introduction of new seeds as a one-time shock, and believe the program will end after one period. I compare these outcomes in a simple cross-sectional regression in the period immediately after the change in P,

$$outcome_i = \alpha + \beta T_i + \varepsilon_i, \tag{6.1}$$

where $T_i = 1$ if household i is treated.[36] I compute the estimate $\widehat{\beta}$ for yields, and update the guess of λ until the estimate matches Emerick *et al.* (2016). This iterative procedure implies $\lambda^* = 0.863$.

6.3.3. Short run partial equilibrium results. First, I consider other untargeted outcomes to validate the model. Specifically, I run regression (6.1) in the empirical and model-created datasets, testing the partial equilibrium impact on fertilizer expenditures and the share of savings held as a buffer stock. I run the same regressions as in Emerick *et al.* (2016), and normalize their results for easier comparison. To compute standard errors in the model, I draw 1,000 samples with observations equal to the sample size reported in Emerick *et al.* (2016). The results are in Table 7.

The model matches the yield treatment effect by construction but also matches the fact that treatment households both increase fertilizer use and decrease savings. While the quantitative magnitude both responses are muted relative to the empirics when measured by the point estimates, both model predictions fall within the 95% (and, more strictly, the 90%) confidence interval of the relevant empirical estimate. Thus, both the sign and magnitudes of the model response seem to be within reason of the data.[37]

36. Specifically the timing works as follows. A farmer enters the period with b savings then learns her shock distribution (P_0 or P_1) before making any decisions. The rest of the timing proceeds as in the model: farmers make fertilizer choices, their shock z is realized, then *ex post* decisions on consumption, savings, and labour are made. The regression specification and timing are consistent with the timing of impact evaluation in Emerick *et al.* (2016).

37. The harvest and fertilizer treatment effects are naturally related through the production function. Some algebra on the production function implies that it is possible to write the treatment effects as $T_y = \frac{1}{1-\eta}\Delta\bar{z} + \frac{\psi}{1-\eta}T_x$ where T_y and T_x are treatment effects on (log) yield and intermediates, and $\Delta\bar{z}$ is the imposed change in average shock realization between treatment and control. This bounds T_x as a function of T_y, given by $T_x \in \left(0, \frac{(1-\eta)T_y^*}{\psi}\right)$ where T_y^* is the targeted value for yield treatment. Note, however, that these treatment effects depend on the elasticities of input choices with respect to λ.

TABLE 7
Randomized controlled trial results (model versus data)

	Yield		Fertilizer expenditures		Savings/harvest	
	Empirical	Model	Empirical	Model	Empirical	Model
Treatment	0.101 (0.028)***	0.101 (0.027)**	0.111 (0.050)**	0.058 (0.026)**	−0.071 (0.024)***	−0.048 (0.032)
Obs	4,573	n.a.	1,237	n.a.	1,167	n.a.
R^2	0.160	0.003	0.619	0.001	0.113	0.002

Notes: Significance at 0.01, 0.05, 0.1 levels denoted by ***, **, and *. The model standard errors are bootstrapped using 1000 samples of number of observations reported for each outcome. Note that the yield results are carried out at the plot level in Emerick *et al.* (2016), hence the larger sample size for this outcome. All outcomes are normalized by the control group mean.

TABLE 8
Aggregate impact of new seed

	X	N_a	Y_a	$\frac{p_x X}{p_a Y_a}$	$\frac{X}{Y_a}$	$\frac{Y_a}{N_a}$	Real GDP p.c.	p_a
Baseline economy	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
New seed (partial equilibrium)	1.06	1.10	1.10	0.96	0.96	1.00	0.95	1.00
New seed (general equilibrium)	0.92	0.92	0.97	0.99	0.95	1.05	1.06	0.95

Notes: Each row are moments from the steady state of the listed economy. Row 1 is normalized. Row 2 fixes prices as the baseline economy level, and changes only the shock distribution. Row 3 allows prices to adjust as well. Real GDP is measured at baseline U.S. model prices.

While the rationale for higher fertilizer expenditure is clear, perhaps less obvious is that savings declines. Intuitively, treatment farmers are less concerned about risk, and therefore have less need to hold a buffer stock. Therefore, they repurpose those savings to consumption and intermediate spending. Since this *ex post* savings realization is affected by the individual-level shock realization, the P-value ($P=0.15$) falls above standard cutoffs for statistical significance, though the standard interpretation of statistical significance is a bit strained here, as this is indeed the true model effect.

Overall, the model matches well the predictions of the RCT, and I therefore turn to the impact of delivering such an innovation at scale, taking seriously the equilibrium implications of such a roll out.

6.3.4. Implications at scale and the importance of GE. To test the impact of rolling out Swarma-Sub1 at scale across India, I compute the stationary equilibrium under the new probability distribution under two economies. The first holds the agricultural price fixed at its baseline level, while the second allows it to adjust in response to the new distribution. Table 8 provides the aggregate moments from stationary equilibria of these partial and general equilibrium cases.

To begin, note that the total amount of fertilizer in the economy increases by 6% when prices do not change, as implied by the treatment effects of Table 7. This increases the marginal return to labour. Households therefore shift labour toward agriculture and increase production. However, this influx of labour implies that agricultural productivity remains constant.[38] Yields increase similarly, by 10%.[39]

38. This is straightforward to show in partial equilibrium. For any two economies with the same price p_a,

$$\frac{Y_a^1/N_a^1}{Y_a^2/N_a^2} = \frac{(1-\tau^1)A^1}{(1-\tau^2)A^2}.$$

39. Note that real GDP per capita declines in partial equilibrium. The relatively high out-of-equilibrium price in India generates an increase in intermediates, which once deflated to U.S. levels generates the decline. I focus primarily on agricultural productivity here, as it is more easily compared across the various models.

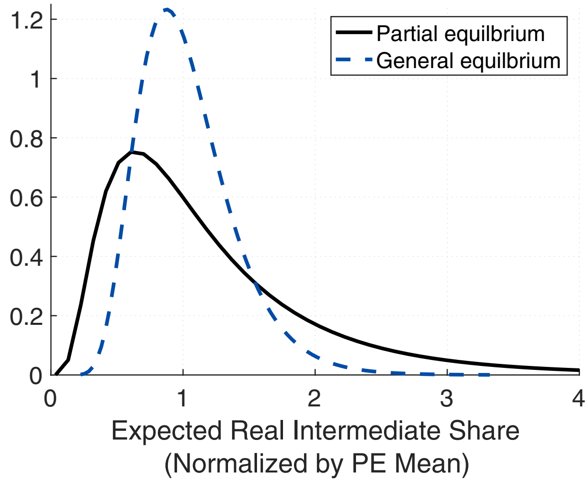

FIGURE 6

Density of normalized expected real intermediate shares after scaling improved seeds

Notes: Both curves are normalized by the mean real intermediate share in the partial equilibrium model.

One perhaps surprising result is that the nominal real intermediate share actually declines slightly in partial equilibrium. The intuition for this result comes from farmers' risk-neutral probabilities, which imply high weight on low shock realizations. Thus, while these new seeds shift some weight from the probability distribution from low to high realizations, the *effective* weights from the farmer's perspective shift much less. This implies a lower elasticity of intermediates relative to harvest for small, partial equilibrium changes like the one generated here.[40]

The general equilibrium decline in the price changes this, and the nominal intermediate share increases by 3% relative to the partial equilibrium case. This follows from the discussion in Section 6.1—the declining agricultural price decreases the cost of subsistence, and with it, relative risk aversion. Figure 6 plots the distribution of expected real intermediate shares across households in both the partial and general equilibrium models, both normalized by the mean partial equilibrium realization. One can see the decline in micro-level variance.

Interestingly, however, the declining variance is not simply the poor catching up to the rich, as Figure 6 shows that the result is coming from both the left and right tails of the distribution. On aggregate, the real intermediate share actually decreases by 5%. The key intuition for this seemingly divergent set of results is to understand to how the impact of the price change varies along the savings distribution. Since rich farmers approximately maximize profits, the lower price does little to their (already quite small) risk-induced wedge. Instead, it simply lowers the marginal return to intermediates, and they lower expenditures in response. This generates the contraction in the right tail of Figure 6. The poor, on the other hand, benefit substantially from this lower price. While the marginal return declines for them as well, it is more than made up for by the decrease in the risk distortion.[41] This generates the change in the left tail of Figure 6.

40. This is not a generic result. Depending on the scale of the exogenously-imposed change, the model can generate a larger treatment effect for intermediates than harvests. This is only to say that small changes in the shock distribution may not generate large changes in the intermediate share given the way in which farmers weight shock realizations.

41. Recall, relative risk aversion for a given level of consumption expenditures C is $C/(C-p_a\overline{a})$. Thus, the poor are the most risk averse households in the economy.

Taken together, the induced price change is, in effect, a redistributive equilibrium response that decreases misallocation and generates a more equitable distribution of resources in the economy.

The net effect of these two channels in equilibrium is positive, as labour productivity increases in both agriculture (5%) and on aggregate (6%). Thus, this misallocation effect can potentially be quite large, even overcoming the corresponding decline in intermediate input intensity by generating a more efficient allocation across production units.

These results make two important policy-relevant points. First, the partial equilibrium results show that in experiments changing both the first and second moments of the shock distribution, policy makers need to exercise caution in defining success. In particular, without a corresponding change in the price, the impact on the nominal share can be muted or even negative. Second, general equilibrium effects play a critical role in understanding the gains from such a policy by lowering misallocation. These features are critical to understanding the full gains from such a policy change or product introduction, and rely on properly accounting for equilibrium effects.

7. DISCUSSION

Before concluding, I briefly discuss a number of additional results that frame the those presented in the main text. All of these exercises are available in the Supplementary Appendix.

7.1. *Credit constraints*

The model used here builds off a class of Aiyagari (1994)-style models with limited borrowing and uninsured shocks that has been used to study misallocation in input markets in a number of recent papers (*e.g.* Buera *et al.*, 2011, among many others). The difference here is that I distort the consumption market instead of the input market, though as I show in the Supplementary Appendix, they can both micro-found the same reduced form input-market distortion. To what extent are those missing markets distinguishable within this model?

To study this, I replace the consumption market distortion with a direct intermediate input market distortion via a collateral constraint. Unlike the main text model, now farmers are required to fund intermediates out of own savings. The comparison highlights an important distinction—the collateral constraints model implies that households with higher intermediate shares have lower consumption volatility, unlike what is observed in the data. I formalize these results and discuss the underlying intuition for this difference in the Supplementary Appendix.[42]

7.2. *Ex ante occupational choice and labour demand*

The model developed here abstracts from *ex ante* occupational choice and labour demand. Of course, in reality, most non-agricultural work is provided by households that specialize in the non-agricultural sector. Does abstracting from this *ex ante* selection into sectors matter quantitatively? To study this, I extend the model in two ways. I first include an *ex ante* selection between sectors, and then also require some agricultural labour to be chosen before the shock realization. Both extensions increase the gains from insurance relative to the baseline model, implying that the results presented here are conservative along these dimensions.

42. This is not an argument that credit constraints are immaterial for intermediate use. One could certainly construct a model that simultaneously involves credit constraints and rationalizes this result. However, the point here is to say that simply adjusting the market failure faced by farmers in this class of models is insufficient to rationalize the relationship between intermediate intensity and consumption variation.

In terms of occupational selection, if the rich select into manufacturing while the poor select into agriculture (as is true empirically) this extended model predicts that the gains from insurance are in fact larger than the baseline model. The rationale is straightforward: if *ex ante* selection removes the most undistorted farmers, the production burden falls on the poorest and most risk-averse remaining farmers. This has the direct effect of making the average farmer more distorted (relative to the baseline model) but also the indirect general equilibrium effect of increasing the agricultural price, which further distorts farmers through risk aversion. On *ex ante* labour demand, shift more inputs pre-risk limits farmers' ability to respond to shocks, thus increasing the distortion. I quantify both margins in the Supplementary Appendix.

7.3. *Different shocks or different responses?*

Another alternative is that poor countries simply face more uncertainty than rich countries, which drives the differences across countries. I show that this is unlikely to play a key role for two reasons. First, using spatial rainfall data, I show that there is little cross-country correlation between rainfall variability and GDP per capita. This echoes recent work by Adamopoulos and Restuccia (2018) who show that natural disadvantages (soil and land quality, for example) are not responsible for low agricultural productivity in poor countries. Instead, these results suggest that the response to risk is key for understanding the relationship between risk and intermediate inputs, not variation in the exogenous shocks themselves.

As a second test, I counterfactually vary shock variance and show it has little effect on aggregate outcomes. The rationale for this result is due to the interaction of shock variance and general equilibrium price movements in a model with subsistence. Details of both tests are available in the Supplementary Appendix.

7.4. *Mismeasurement*

In a quantitative study of risk, mismeasurement is potentially an important concern. There are two ways in which this potentially matters. There are comparisons between model and data. In this context, mis-measurement in harvest and consumption—assuming mean zero noise—will have the usual effects. Harvest is included only as a dependent variable (see regression (5.2) and results in Table 2), and thus suggests the standard error is biased upward. Consumption is used as an independent variable in regression (5.4). Noise in this variable will bias the relationship toward zero. Thus, the positive relationship would be stronger than suggested by the results in the main text.

The second issue is how mismeasurement impacts model construction and quantitative results. Specifically, it affects how to choose parameters to match moments. Given that the model under-predicts consumption volatility, noise in this dimension is perhaps less of a concern. More first order here is harvest noise. Since the model is targeted to match the variation in harvest, this would imply the shock distribution requires a lower variance than the calculated value $\sigma_z = 0.32$.

7.5. *Permanent productivity difference*

If the poor are sufficiently and persistently low skilled, they may optimally have lower real intermediate shares. Buera and Shin (2011), Moll (2014), and Midrigan and Xu (2014) highlight how a greater role for permanent productivity lowers the impact of misallocation across production units in the presence of financial constraints. I take this issue up in the Supplementary Appendix. I show that the autocorrelation in household-level harvest over time is relatively low, which implies a small role for permanent productivity differences. The ICRISAT data imply one and three year autocorrelations 0.75 and 0.67. For comparison, Midrigan and Xu (2014) find the same moments to be 0.90 and 0.87 among Korean manufacturing firms, and use this to deduce that permanent productivity differences play a critical role in overall cross-sectional productivity variance across plants. Here, the model requires that only 28% of cross-sectional variation in productivity comes from permanent differences, inducing a small quantitative effect—agricultural productivity rises by about 1% in this model compared to the baseline model.

Sectoral differences in output persistence and its relationship to sectoral productivity is an interesting avenue for future work.

7.6. *Other results*

Finally, the Supplementary Appendix includes a number of other results from the model. I vary the share of shocks that are insurable, provide additional results varying subsets of parameters that differ between the Indian and U.S. model economies, show how the model amplifies other shocks, and provide quantitative results under other DRRA utility functions.

8. CONCLUSION

This article quantifies the role of production risk in accounting for agricultural and aggregate productivity in general equilibrium. In poor countries, farmers use fewer intermediate inputs, driving down agricultural productivity. The model provides a risk-based foundation for misallocation across agricultural production units, but the ability to generate relatively larger distortions in poor economies depends critically on the inclusion of a subsistence requirement. Quantitatively, the model captures about one-third of the difference in intermediate input shares between the India and the U.S. This has important quantitative implications for productivity.

Moreover, the model implies an important role for general equilibrium in the context of real-world policies such as insurance or seeds that decrease downside-risk. The quantitative implications for this are large. Building off RCT results in Emerick *et al.* (2016), I compute the new equilibrium of introducing flood-resistant seeds nation-wide. The general equilibrium model implies a 6% increate in agricultural productivity and a 5% increase in GDP per worker. The partial equilibrium model implies only small gains along both of these margins.

Finally, I note that the model cannot capture the entirety of intermediate input gap between the U.S. and India. Thus, the model demands additional distortions that affect this margin. Many have been proposed in the literature—credit, information, behavioural biases, savings constraints, and trade costs to name a few. While the goal of this article was to study the impact of risk directly, an important avenue for future work is to understand the relationship between such distortions. As I show in the Supplementary Appendix, certain types of distortions are complementary to risk, in that the quantitative importance of those distortions is amplified by its presence. The converse is of course true as well—savings constraints, for example, limit the ability of farmers to self-insure and magnify the impact of risk. Understanding these relationships in the context of general equilibrium will provide a broader view of the proper policy responses to low input use in the developing world.

Acknowledgments. This article previously circulated as "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences." I thank the editor, Michèle Tertilt, and five anonymous referees for many suggestions that greatly improved this article. I also thank conference and seminar participants at Arizona State, ITAM, Notre Dame, Rochester, the St. Louis Fed, Virginia, The World Bank, the CASEE Macro Reunion Conference, the Econometric Society NASM, and the Society for Economic Dynamics, including comments and insights from Alex Bick, Berthold Herrendorf, Chris Herrington, Joe Kaboski, David Lagakos, Ed Prescott, B. Ravikumar, Todd Schoellman, and especially Richard Rogerson for a detailed discussion of the article. Will Jianyu Lu provided excellent research assistance.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online. And the replication packages are available at https://doi.org/10.5281/zenodo.4139612.

Data Availability Statement

The data underlying this article are available in the Zenodo.org repository, at https://doi.org/10.5281/zenodo.4139612

REFERENCES

ADAMOPOULOS, T. (2011), "Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences", *International Economic Review,* **52**, 489–521.

ADAMOPOULOS, T., "Spatial Integration and Agricultural Productivity: Quantifying the Impact of New Roads" (Working Paper).

ADAMOPOULOS, T., BRANDT, L., LEIGHT, J. et al. (2019), "Misallocation, Selection and Productivity: A Quantitative Analysis with Panel Data from China" (NBER Working Paper 23039).

ADAMOPOULOS, T. and RESTUCCIA, D. (2014), "The Size Distribution of Farms and International Productivity Differences", *American Economic Review,* **104**, 1667–1697.

ADAMOPOULOS, T. and RESTUCCIA, D. (2018, April), "Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data" (NBER Working Paper 24532).

AIYAGARI, S. R. (1994), "Uninsured idiosyncratic risk and aggregate saving", *Quarterly Journal of Economics,* **109**, 659–684.

ANAGOL, S., ETANG, A. and KARLAN, D. (2017), "Continued Existence of Cows Disproved Central Tenets of Capitalism?", *Economic Development and Cultural Change,* **65**, 583–618.

BARTELME, D. and GORODNICHENKO, Y. (2015), "Linkages and Economic Development" (NBER Working Paper 21251).

BERGQUIST, L., FABER, B., FALLY, T. et al. (2019), "Scaling Up Agricultural Policy Interventions: Theory and Evidence from Uganda" (Working Paper).

BROOKS, W. and DONOVAN, K. (2020), "Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua", *Econometrica*, **88**, 1965–1997.

BUERA, F., KABOSKI, J. and SHIN, Y. (2020), "The Macroeconomics of Microfinance", *Review of Economic Studies*, forthcoming.

BUERA, F. J., KABOSKI, J. P. and SHIN, Y. (2011), "Finance and Development: A Tale of Two Sectors", *American Economic Review*, **101**, 1964–2002.

BUERA, F. J. and SHIN, Y. (2011), "Self-insurance vs. Self-financing: A Welfare Analysis of the Persistence of Shocks", *Journal of Economic Theory*, **146**, 845–862.

CAI, H., CHEN, Y., FANG, H. et al. (2015), "The Effect of Microinsurance on Economic Activities: Evidence from a Randomized Field Experiment", *Review of Economics and Statistics*, **97**, 287–300.

CASELLI, F. (2005), "Accounting for Cross-Country Income Differences", in Aghion, P. and Durlauf, S. (eds.), *Handbook of Economic Growth* (Elsevier), 679–741, https://www.sciencedirect.com/science/article/pii/S1574068405010099.

CAUNEDO, J. and KELLER, E. (2020), "Capital Obsolescence and Agricultural Productivity", *Quarterly Journal of Economics*, forthcoming.

CHEN, C. (2020), "Technology Adoption, Capital Deepening, and International Productivity Differences", *Journal of Development Economics*, **143**.

Cole, S., Giné, X., Tobacman, J. et al. (2013), "Barriers to Household Risk Management: Evidence from India", *American Economic Journal: Applied Economics*, **5**, 104–135.

COLE, S., GINÉ, X. and VICKERY, J. (2017), "How Does Risk Management Influence Production Decisions? Evidence from a Field Experiment", *Review of Financial Studies*, **30**, 1935–1970.

DERCON, S. and CHRISTIAENSEN, L. (2011), "Consumption Risk, Technology Adoption and Poverty Traps: Evidence from Ethiopia", *Journal of Development Economics*, **96**, 159–173.

EMERICK, K., DE JANVRY, A., SADOLET, E. et al. (2016), "Technological Innovations, Downside Risk, and the Modernization of Agriculture", *American Economic Review*, **106**, 1537–1561.

FADINGER, H., GHIGLINO, C. and TETERYATNIKOVA, M. (), "Income Differences, Productivity and Input-Output Networks" (Working Paper, Mannheim).

FAFCHAMPS, M., UDRY, C. and CZUKAS, K. (1998), "Drought and Saving in West Africa: Are Livestock a Buffer Stock?" *Journal of Development Economics*, **55**, 273–305.

FUJIMOTO, J., LAGAKOS, D. and VANVUREN, M. (2019), "The Aggregate Effects of "Free" Secondary Schooling in the Developing World" (Working Paper).

GOLDBERG, J. (2017), "The Effect of Social Pressure on Expenditures in Malawi", *Journal of Economic Behavior and Organization*, **143**, 173–185.

GOLLIN, D., LAGAKOS, D. and WAUGH, M. E. (2014), "The Agricultural Productivity Gap", *Quarterly Journal of Economics*, **129**, 939–993.

GOLLIN, D., PARENTE, S. L. and ROGERSON, R. (2004), "Farm Work, Home Work and International Productivity Differences", *Review of Economic Dynamics*, **7**, 827–850.

GOLLIN, D. and ROGERSON, R. (2014), "Productivity, Transport Costs, and Subsistence Agriculture", *Journal of Development Economics*, **107**, 38–48.

GOTTLIEB, C. and GROBOVŠEK, J. (2019), "Communal Land and Agricultural Productivity", *Journal of Development Economics*, **138**, 135–152.

GREENWOOD, J., KIRCHER, P., SANTOS, C. et al. (2019), "An Equilibrium Model of the African HIV/AID Epidemic", *Econometrica*, **87**, 1081–1113.

HERRENDORF, B., ROGERSON, R. and VALENTINYI, A. (2013), "Two Perspectives on Preferences and Structural Transformation", *American Economic Review* (North Holland: Imprint of Elsevier), **103**, 2752–2789, https://www.sciencedirect.com/science/article/pii/B9780444535405000069.

HERRENDORF, B., ROGERSON, R. and VALENTINYI, A. (2014), "Growth and Structural Transformation", in Aghion, P. and Durlauf, S. (eds) *Handbook of Economic Growth* 855–941.

HERRENDORF, B. and SCHOELLMAN, T. (2015), "Why is Measured Productivity so Low in Agriculture?", *Review of Economic Dynamics*, **18**, 1003–1022.

HSIEH, C. and KLENOW, P. (2007), "Relative Prices and Relative Prosperity", *American Economic Review*, **97**, 562–585.

HSIEH, C. and KLENOW, P. J. (2009), "Misallocation and manufacturing tfp in china and india", *Quarterly Journal of Economics*, **124**, 1403–1448.

JAKIELA, P. and OZIER, O. (2016), "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economics", *Review of Economic Studies*, **83**, 231–268.

JUST, R. E. and POPE, R. D. (1979), "Production Function Estimation and Related Risk Considerations", *American Journal of Agricultural Economics*, **61**, 276–284.

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

KARLAN, D., OSEI, R. D., OSEI-AKOTO, I. et al. (2014), "Agricultural Decisions After Relaxing Credit and Risk Constraints", *Quarterly Journal of Economics*, **129**, 597–652.

KARLAN, D., RATAN, A. and ZINMAN, J. (2014), "Savings By and For the Poor: A Research Review and Agenda", *The Review of Income and Wealth*, **60**, 36–78.

KAZIANGA, H. and UDRY, C. (2006), "Consumption Smoothing? Livestock, Insurance and Drought in Rural Burkina Faso", *Journal of Development Economics*, **79**, 413–446.

KOCHAR, A. (1999), "Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India", *Review of Economics and Statistics*, **81**, 50–61.

KONGSAMUT, P., REBELO, S. and XIE, D. (2001), "Beyond balanced growth", *Review of Economic Studies*, **68**, 869–882.

LAGAKOS, D., MOBARAK, A. and WAUGH, M. (2020), "The Welfare Effects of Encouraging Rural-Urban Migration", Working Paper.

LAGAKOS, D. and WAUGH, M. E. (2013), "Selection, Agriculture, and Cross-Country Productivity Differences", *American Economic Review*, **103**, 948–980.

MIDRIGAN, V. and XU, D. (2014), "Finance and Misallocation: Evidence from Plant-Level Data", *American Economic Review*, **104**, 422–458.

MOBARAK, A. M. and ROSENZWEIG, M. (2014), "Selling Formal Insurance to the Informally Insured" (Yale Economics Growth Center Discussion Paper No. 1007).

MOLL, B. (2014), "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?", *American Economic Review*, **104**, 3186–3221.

PRASADA RAO, D. (1993), "Intercountry Comparisons of Agricultural Output and Productivity" (FAO Economic and Social Development Paper No. 112).

RESTUCCIA, D. and ROGERSON, R. (2008), "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments", *Review of Economic Dynamics*, **11**, 707–720.

RESTUCCIA, D. and SANTAEULÀLIA-LLOPIS, R. (2017), "Land Misallocation and Productivity" (Working Paper, University of Toronto).

RESTUCCIA, D., YANG, D. T. and ZHU, X. (2008), "Agriculture and Aggregate Productivity: A Quantitative Cross-country Analysis", *Journal of Monetary Economics*, **55**, 234–250.

ROSENZWEIG, M. R. and WOLPIN, K. I. (1993), "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India", *Journal of Political Economy*, **101**, 223–244.

SOTELO, S. (), "Domestic Trade Frictions and Agriculture", *Journal of Political Economy*, **128**, 2690–2738.

TIMMER, M., DIETZENBACHER, E., LOS, B. et al. (2015), "An Illustrated User Guide to the World Input-Output Database: The Case of Global Automotive Production", *Review of International Economics*, **23**, 575–605.

TOMBE, T. (2015), "The Missing Food Problem: Trade, Agriculture, and International Productivity Differences", *American Economic Journal: Macroeconomics*, **7**, 226–258.

TRAXLER, G., FALCK-ZEPEDA, J., ORTIZ-MONASTERIO R. and SAYRE, K. (1995), "Production Risk and the Evolution of Varietal Technology", *American Journal of Agricultural Economics*, **77**, 1–7.

UDRY, C. (1995), "Risk and Saving in Northern Nigeria", *American Economic Review*, **85**, 1287–1300.

UNITED NATIONS (2015), *National Account Statistics: Main Aggregates and Detailed Tables* (Table 2.3).

VOLLRATH, D. (2009), "How Important Are Dual Economy Effects for Aggregate Productivity?", *Journal of Development Economics*, **88**, 325–334.

WORLD BANK (2018, April), *World Development Indicators* (The World Bank).

YANG, D. and ZHU, X. (2013), "Modernization of Agriculture and Long-Term Growth", *Journal of Monetary Economics*, **60**, 367–382.