

Essays in Growth and Development

by

Kevin Donovan

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2013 by the
Graduate Supervisory Committee:

Edward C. Prescott, Chair

Berthold Herrendorf

David Lagakos

Todd Schoellman

ARIZONA STATE UNIVERSITY

May 2013

ABSTRACT

This dissertation consists of three essays that broadly deal with the growth and development of economies across time and space.

Chapter one is motivated by the fact that agricultural labor productivity is key for understanding aggregate cross-country income differences. One important proximate cause of low agricultural productivity is the low use of intermediate inputs, such as fertilizers, in developing countries. This paper argues that farmers in poor countries rationally choose to use fewer intermediate inputs because it limits their exposure to large uninsurable risks. I formalize the idea in a dynamic general equilibrium model with incomplete markets, subsistence requirements, and idiosyncratic productivity shocks. Quantitatively, the model accounts for two-thirds of the difference in intermediate input shares between the richest and poorest countries. This has important implications for cross-country productivity. Relative to an identical model with no productivity shocks, the addition of agricultural shocks amplifies per capita GDP differences between the richest and poorest countries by nearly eighty percent.

Chapter two deals with the changes in college completion in the United States over time. In particular, this paper develop a dynamic lifecycle model to study the increases in college completion and average IQ of college students in cohorts born from 1900 to 1972. I discipline the model by constructing historical data on real college costs from printed government reports covering this time period. The main finding is that that increases in college completion of 1900 to 1950 birth cohorts are due primarily to changes in college costs, which generate a large endogenous increase in college enrollment. Additionally, evidence is found that supports cohorts born after 1950 underpredicted sharp increases in the college earnings premium they eventually received. Combined with increasing college costs during this time period, this generates a slowdown in college completion, consistent with empirical evidence for cohorts born after 1950. Lastly, the rise in average college stu-

dent IQ cannot be accounted for without a decrease in the variance of ability signals. This is attributed the increased precision of ability signals primarily to the rise of standardized testing.

Chapter three again deals with cross-country income differences. In particular, it is concerned with the fact that cross-country income differences are primarily accounted for by total factor productivity (TFP) differences. Motivated by cross-country empirical evidence, this paper investigates the importance individuals who operate their own firms because of a lack of other job opportunities (need-based entrepreneurs). I develop a dynamic general equilibrium labor search model with with entrepreneurship to rationalize this misallocation across occupations and assess its role for understanding cross-country income differences. Developing countries are assumed to have tighter collateral constraints on entrepreneurs and lower unemployment benefits. Because these need-based entrepreneurs actually have a comparative advantage as workers, they operate smaller and less productive firms, lowering aggregate TFP in developing countries.

DEDICATION

To my parents and Peggy, all of whom have had to put up with me for the past six years. Without their patience this would not have been possible.

ACKNOWLEDGEMENTS

I owe a huge debt to Berthold Herrendorf, David Lagakos, Ed Prescott, and Todd Schoellman for their constant support and guidance throughout this project. The chapters included here have also benefited tremendously from comments and insights by participants at Arizona State, Exeter, Georgia State, Miami, Notre Dame, Rochester, the St. Louis Fed, SUNY Buffalo, Virginia, Washington State, the CASEE Macro Reunion Conference at ASU, the 2011 CEA Annual Meetings, the 2012 Econometric Society NASM, the 2012 Midwest Macro Meeting, and the 2011 SED Annual Meeting, especially those from Alex Bick, Chris Herrington, Dan Lawver, Richard Rogerson, Don Schlagenhauf, and Alan Taylor. I thank all of them, without holding them responsible for any errors that remain. The second chapter is joint work with Christopher Herrington.

TABLE OF CONTENTS

	Page
LIST OF TABLES	viii
LIST OF FIGURES	ix
CHAPTER	
1 AGRICULTURAL RISK, INTERMEDIATE INPUTS, AND CROSS-COUNTRY	
PRODUCTIVITY DIFFERENCES	1
1.1 Introduction	1
1.2 Motivating Evidence	5
Intermediate Input Shares Across Countries	6
Comparison to Manufacturing and Services	7
1.3 Model	9
Technology	9
Village Utility and Decisions	10
Recursive Problem	12
Stationary Equilibrium	13
Discussion	14
1.4 Characterization and Analytic Results	16
Subsistence Requirements and Risk	17
Characterizing Intermediate Input Shares	18
1.5 Quantitative Exercise and Calibration	23
Common Parameters	24
Economy Specific Parameters	27
1.6 Quantitative Results	29
Impact of Agricultural Risk	30
The Role of Intermediate Input Price	32
Changes in Savings Technologies	34
1.7 Conclusion	35

CHAPTER	Page
2 FACTORS AFFECTING COLLEGE COMPLETION AND STUDENT ABILITY IN THE U.S. SINCE 1900	37
2.1 Introduction	37
2.2 Model	40
Timing and Recursive Problem	42
2.3 Calibration	44
Historical Time Series Data	44
Life-Cycle Wage Profiles	45
Exogenous Parameters	46
Calibrated Parameters	48
2.4 Results	50
Benchmark Model Fit	50
Discussion of Benchmark Results	53
Counterfactual Experiments	56
2.5 Robustness	59
Correlation of Ability and Initial Assets	60
College Costs Including Room and Board	61
2.6 Conclusion	62
3 NEED-BASED ENTREPRENEURSHIP AND AGGREGATE PRODUCTIVITY ACROSS COUNTRIES	64
3.1 Introduction	64
3.2 Empirical Motivation	66
3.3 Model	67
Capital and Labor Markets	68
Timing and Recursive Formulation	69
Stationary Equilibrium	71
Need-Based Entrepreneurship in the Model	72

CHAPTER	Page
3.4 Conclusion	72
REFERENCES	73
APPENDIX A	79
3.5 Labor Share Parameter, η	80
3.6 Changes in the Shock Distribution	80
APPENDIX B	82
APPENDIX C	85
3.7 Proof of Proposition 1	86
3.8 An Additional Lemma for the Proof of Proposition 2	86
3.9 Proof of Proposition 2	88
3.10 Proof of Proposition 3	89
APPENDIX D	91
3.11 Construction of Inputs	92
3.12 Output	93
3.13 Decomposition of Residuals	93
APPENDIX E	94

LIST OF TABLES

Table	Page
1.1 Relationship between Intermediate Input Share and Log GDP per Capita (PPP), by Sector	8
1.2 Post Harvest Weight Loss (%) for Selected Countries and Crops for 2007 . . .	29
1.3 Parameter Values for Two Economies	30
1.4 Impact of Productivity Shocks on Labor Productivity	31
1.5 Decomposition of Quantitative Results	33
2.1 Measures of Fit for Various Model Specifications	52
3.1 Model Results for Different η	80
3.2 Model Results for Different σ_z	81

LIST OF FIGURES

Figure	Page
1.1 Intermediate Share in Agriculture	6
1.2 Intermediate Shares in Three Sectors	8
1.3 Stochastic Discounting for Different Subsistence Levels	22
1.4 Empirical Distribution of Shocks	28
1.5 Poor economy model implications for different p_x	33
1.6 Poor economy model implications for different depreciation rates δ	35
2.1 College Completion and Average Student Ability in the U.S. since 1900	38
2.2 Benchmark Model Results	51
2.3 College Enrollment Conditional on High School Graduation	53
2.4 College Pass Rate	54
2.5 College Completion if Enrollment Rates and Pass Rates were Constant	55
2.6 Education Premia Implied by Estimated Life-Cycle Wage Profiles	56
2.7 Results with Imperfect Foresight of Education Earnings Premia	57
2.8 College Costs	58
2.9 Results with Alternative College Costs	59
2.10 Results with Positive Correlation between Ability and Initial Assets	60
2.11 Results for College Costs including Tuition, Fees, Room, and Board	62
3.1 Fraction of Entrepreneurs that are Self-Employed	67

Chapter 1

AGRICULTURAL RISK, INTERMEDIATE INPUTS, AND CROSS-COUNTRY PRODUCTIVITY DIFFERENCES

1.1 Introduction

Differences in agricultural labor productivity between the richest and poorest countries are significantly larger than differences in aggregate labor productivity. This point has been made recently by Caselli (2005) and Restuccia, Yang, and Zhu (2008). In spite of this, the least developed countries in the world employ over eighty percent of their population in the agricultural sector. Since these countries employ such a large fraction of their population in a particularly unproductive sector, basic accounting suggests that understanding agricultural productivity differences are key in understanding aggregate differences.

One possible cause of agricultural productivity differences is that farmers in developing countries use fewer intermediate inputs, such as fertilizer. Empirically, the intermediate input share in agriculture is positively correlated with per capita income, and ranges from a low of 0.04 in Uganda to 0.40 in the United States. Moreover, I document in Section 1.2 that this positive cross-country correlation between the sectoral intermediate input share and income does not exist in other sectors, suggesting that it may be an important margin for understanding why the agricultural sector exhibits significantly lower labor productivity than the nonagricultural sector in developing countries. The goal of this paper is to provide a theory to understand the cross-country correlation between the agricultural intermediate input share and per capita income, and in turn, quantitatively assess its role for cross-country productivity differences.

The basic idea put forth here is that the relatively low intermediate input intensity in developing countries is a rational response to the risk generated by agricultural productivity shocks. Because intermediate decisions are made before the realization of shocks (e.g. weather), the absence of insurance markets requires farmers to internalize the impact

this choice will have on ex-post consumption. In particular, purchasing a large amount of intermediate inputs, then getting hit with a bad shock (e.g. drought) leads to extremely low consumption. The extent to which this consideration impacts the ex-ante intermediate choice depends critically on the income level of farmers. Low shock realizations are certainly bad for everyone, but they are particularly disastrous for farmers in extremely poor countries, since consumption moves dangerously close to subsistence. These farmers are less willing to take on the risk associated with intermediate inputs usage, thus driving down labor productivity in developing countries.

I formalize this idea in a dynamic general equilibrium model in which farm inputs are chosen jointly with consumption. Farmers therefore maximize expected utility instead of expected profit, consistent with a large empirical literature reviewed by Mor Duch (1995). I further assume that farmers in all countries face idiosyncratic productivity shocks, incomplete markets, and subsistence requirements. These features imply that each possible shock realization is weighted not only by the probability of the shock, but also the farmer's realized marginal utility. As total factor productivity (TFP) decreases in poor countries, a farmer's incomes move closer to subsistence. Marginal utility at low shock realizations increases relative to farmers in rich countries, and therefore implies that poor farmers put relatively more weight on bad potential outcomes. This extra weight on low realizations shows up as a wedge between the profit-maximizing marginal value and price of intermediate inputs, and causes farmers to decrease their ex-ante intermediate input choice. Aggregating over all farmers therefore generates a positive correlation between per capita income and the intermediate input share, even though all underlying farm technologies are Cobb-Douglas. As developing countries use fewer intermediate inputs, agricultural labor productivity decreases.

Quantifying the implications of the model for cross-country productivity requires taking a stand on two key features of the economy. The first is the distribution of productivity shocks in agriculture. The distribution of shocks controls the probability of getting a low

realization, which in turn controls the probability of consumption being near subsistence. I estimate the distribution of productivity shocks with plot-level data drawn from eight Indian villages from the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT) Village Level Surveys. This data set contains daily diaries of input and output usage from over 8000 plots, and has the benefit of recording a vast array of inputs and output quantities and values. The estimated shock distribution implies the probability of getting a shock less than forty percent of the mean is approximately ten percent, suggesting that farmers face quite a bit of risk.

Second, as pointed out by Aiyagari (1994) and Deaton (1991), a countervailing force is that individuals can limit the distortionary impact of risk through savings. I assume that savings is limited to storage of agricultural output, which is consistent with observed responses to shocks found by Fafchamps, Udry, and Czukas (1998) and Kazianga and Udry (2006). I discipline the savings technology by utilizing a new set of storage depreciation rates from Africa, constructed by the African Postharvest Loss Information System. Depreciation ranges from ten to thirty percent, suggesting that while in principle farmers can save their way away from this risk, in practice it is extremely costly.

The main quantitative exercise then compares the stationary equilibria of a rich and a poor economy, which are calibrated to capture the relevant differences between the richest and poorest countries in the world. The rich economy is calibrated to match key sectoral features of the United States, including the intermediate input share and employment share in agriculture. The poor economy differs along three dimensions. It has lower economy-wide TFP, intermediate inputs are more costly, and the depreciation of agricultural storage is higher. Naturally, these exogenous differences generate lower labor productivity in the poor economy. To isolate the impact of the theory developed here, I ask how much *larger* productivity differences are in the model with shocks, relative to the identical model with no productivity shocks in the agricultural sector.

The quantitative results imply that the seemingly sub-optimal intermediate input choices in agriculture are in fact optimal responses by small-scale farmers to large risk, and a key factor affecting labor productivity across countries. The calibrated model predicts that the poor economy has an intermediate input share of 0.22, compared to the U.S. intermediate share of 0.40. This is fifty eight percent of the difference found in the data, in which the poorest countries have intermediate input shares that average 0.09. By virtue of the assumed Cobb-Douglas production technologies, any predicted difference in the intermediate input shares is due entirely to the addition of agricultural productivity shocks. These predicted differences in intermediate input shares then amplify the output per worker differences generated by exogenous factors. Relative to an identical model without productivity shocks, the addition of shocks in the agricultural sector amplifies agricultural output per worker differences by slightly over forty percent and GDP per worker differences by seventy five percent.

This paper joins recent work by Jones (2011), Koren and Tenreyro (2012), and the aforementioned Restuccia, Yang, and Zhu (2008) by emphasizing the role of intermediate inputs for economic development. I differ by focusing on the role of productivity shocks for depressing intermediate input intensity in developing countries, and rely on the interaction of these shocks with incomplete consumption insurance. This Bewley (1986)-Aiyagari (1994) framework has been exploited by a number of recent papers to study cross-country productivity, including Buera, Kaboski, and Shin (2011a) and Buera and Shin (2013), though these papers generate input misallocation through financial distortions. I generate similar deviations from undistorted profit-maximization by relying on the inability of firm owner-operators to insure their own consumption, more along the lines of Midrigan and Xu (2012). Their focus on the manufacturing sector yields a much smaller quantitative impact. Due to the importance of the agricultural sector in developing countries, and in particular small-scale farming, the quantitative results here suggest that misallocation due to uninsurable risk is an important margin for understanding labor productivity.

As it specifically relates to intermediate input usage in agriculture, this theory stands in contrast to the recent work of Duflo, Kremer, and Robinson (2011). Backed by experimental evidence from western Kenya, they develop a theory in which under-investment in fertilizer is a result of present-biased farmers who do not necessarily understand their preferences. These farmers delay fertilizer purchases until the last moment, then due to their lack of knowledge of their own preferences, may not buy fertilizer at all. Here, under-investment is a rational response to large uninsurable risks in the agricultural sector, though both theories suggest a role for properly executed subsidies to intermediate purchases. The importance of uninsurable risk is supported by recent empirical work including Karlan et al. (2012) and Zerfu and Larson (2010).

This paper is not, however, the first to utilize the agricultural sector to better understand aggregate income differences. Other explanations include work by Adamopoulos and Restuccia (2011) on distortions limiting farm size, Lagakos and Waugh (2012) on occupational selection, and the possibility of mismeasurement due to home production by Gollin, Parente, and Rogerson (2004) or under-estimation of value added by Herrendorf and Schoellman (2012). None consider the role of productivity shocks for understanding input decisions. A number of these theories can be easily embedded in the framework developed here, which would serve only to magnify the quantitative results.

The rest of the paper proceeds as follows. Section 1.2 provides motivating evidence of intermediate input share differences in a cross section of countries. Section 1.3 describes the model. Section 1.4 shows how the interaction of productivity shocks and subsistence requirements theoretically generate differences in the intermediate input share across countries. Turning to the quantitative work, Section 1.5 details the calibration and Section 1.6 presents the quantitative results of the model. Finally, Section 1.7 concludes.

1.2 Motivating Evidence

In this section, I first show that the intermediate input share in agriculture is positively correlated with per capita GDP in a cross section of countries. In Section 1.2, I show that this positive correlation is limited to the agricultural sector. The manufacturing and service sectors exhibit no such relationship between intermediate input shares and income, supporting the hypothesis that differences in intermediate input shares can help explain why agriculture is particularly unproductive in developing countries.

Intermediate Input Shares Across Countries

The intermediate input share in agriculture of country j is

$$\hat{X}^j := \frac{p_x^j X^j}{p_a^j Y_a^j} \quad (1.2.1)$$

where X is the quantity of nonagricultural intermediate inputs, such as fertilizer, and Y_a is the quantity of agricultural output. The prices faced by the farmer are denoted p_x^j and p_a^j , and are denominated in local currency units. The price p_x^j takes into account any sector-specific distortions that increase the intermediate input price, such as transportation costs. Since I am interested in the decisions of farmers, these are the relevant prices. I construct this share with data from Prasada Rao (1993), which covers 84 countries in 1985 and is derived from Food and Agricultural Organization (FAO) statistics. Figure 1.1 plots the intermediate input share in agriculture with log PPP GDP per capita on the horizontal axis, from the Penn World Tables version 7.0, Heston, Summers, and Aten (2011).

There is a clear positive relationship between income level and the intermediate share in agriculture, with a correlation of 0.65. Not only is there a positive correlation, but the level difference between rich and poor countries is large. The lowest intermediate share in the sample belongs to Uganda, and is one-tenth that of the United States. The tenth percentile country, as ranked by GDP per capita, has an intermediate input share that is one-fourth of the intermediate share in the United States.¹

¹If rich countries are producing different crops than developing countries, one might suspect that the





Figure 1.1: Intermediate Share in Agriculture

Comparison to Manufacturing and Services


Since this paper is concerned with understanding why agriculture is so much less productive than nonagriculture, an important question is whether other sectors exhibit the same relationship between the intermediate share and per capita GDP. I therefore turn to statistics from the United Nations *System of National Accounts* (SNA) to construct intermediate input shares for (a) “Agriculture, hunting, forestry; fishing,” (b) “Manufacturing,” and (c) “Education; health and social work; other community, social and personal services.” These intermediate shares are plotted for in Figure 1.2 for the 49 countries with data available for all sectors.

Figure 1.2a confirms the relationship between the agricultural intermediate input share and per capita GDP present in the FAO statistics. Figures 1.2b and 1.2c, however, show the intermediate input shares in manufacturing and services exhibit no such relationship. This result relates to the results of Hsieh and Klenow (2007). That is, if intermediate input intensity in manufacturing or services differs across countries, it seems to be driven by distortions that manifest themselves in higher intermediate input prices. Once these prices result is driven by different production techniques for these different types of output. While I cannot directly test this, I do group countries by latitude to control for the type of agricultural production, and compare within-group variation. The same correlation holds within groups.



(a) Agriculture, hunting, forestry; fishing

(b) Manufacturing



(c) Education; health and social work; other community, social and personal services

Figure 1.2: Intermediate Shares in Three Sectors

differences are accounted for by denominating intermediate shares in domestic currency, there is no difference across countries. By contrast, even after accounting for intermediate price differences across countries, the agricultural intermediate share still exhibits a positive correlation with income. The distortions driving differences in intermediate share differences in agriculture are therefore not driven exclusively by cross-country price differences. To summarize, Table 1.1 presents the results of a simple linear regression of the sectoral intermediate share on log PPP GDP per capita.

Only agriculture has a slope significantly different from zero, implying that the positive relationship between the intermediate input share and per capita income is unique to the agricultural sector. This result suggests that this positive relationship may be an important factor for understanding income differences across countries. The rest of this paper is devoted to developing and quantifying a model to understand the cause of this

Table 1.1: Relationship between Intermediate Input Share and Log GDP per Capita (PPP), by Sector

	Agriculture	Manufacturing	Services
Constant	-0.43*** (0.11)	0.59*** (0.06)	0.21* (0.11)
Log GDP per capita (PPP)	0.10*** (0.01)	0.01 (0.01)	0.02 (0.01)
R^2	0.52	0.03	0.03

Note: Standard errors are in parentheses. Significance at 0.01, 0.05, 0.1 levels denoted by ***, **, and *

correlation in agriculture and assess its impact on cross-country productivity differences.

1.3 Model

This section lays out a multi-sector dynamic general equilibrium model in the spirit of Bewley (1986) and Aiyagari (1994). The key features are owner-operated farms, incomplete markets, idiosyncratic productivity shocks, and subsistence requirements.

The model period is a year, and time is discrete and runs $t = 0, 1, 2, \dots$. There are two sectors, sector a for agriculture and sector m for manufacturing, which includes all nonagriculture. 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 villages. While the distinction of a “village” is irrelevant to the theory, it is important for the quantitative results. Intuitively, this distinction is important because, as pointed out by Townsend (1994), individuals are relatively well insured against purely idiosyncratic shocks. Assuming incomplete markets at the individual level would therefore overstate my results. Covariate risk, such as weather, is more difficult to insure against, implying villages are the decision making units subject to incomplete markets.

Technology

The manufacturing output good can be used as either consumption or as intermediate inputs in agricultural production. Production is characterized by a stand-in firm which uses only labor 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

$$\max_{N_{mt} \geq 0} AN_{mt} - w_t N_{mt} \quad (1.3.1)$$

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

Turning to agriculture, each village is endowed with one farm that requires intermediate inputs x and labor 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. Land is a fixed factor, and normalized to one.² The shock z_t is a village-specific productivity shock drawn from a time-invariant distribution with cumulative distribution function $Q(z)$ and support on $[\underline{z}, \bar{z}]$. The realization of z_t is i.i.d. with respect to both villages and time. I assume the law of large numbers holds, so that the distribution of shocks across villages is certain.

Intermediate inputs are purchased from the manufacturing sector, at the price $p_x \geq 1$. This price is allowed to differ across countries but not time, with the implicit assumption being that there exists a technology that turns one unit of manufacturing output into $1/p_x$

²See Adamopoulos and Restuccia (2011) for the impact of farm size differences on aggregate productivity. Adding these differences here would serve only to magnify the quantitative impact, without changing the theoretical results.

units of intermediate input. This is a simple way to capture the fact that intermediate inputs are more expensive in developing countries.

Village Utility and Decisions

A village values consumption from both sectors a and m , and maximizes total expected village utility given by

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

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 subsistence requirement of the agricultural good. Herrendorf, Rogerson, and Valentinyi (2012) find that this utility function is consistent with the structural transformation path of the U.S..

Villages do not have access to insurance markets, so that the shock can only be insured against through self-insurance. To this end, villages can save only by storing agricultural output. This storage depreciates at a country-specific rate δ to capture differences in agricultural storage technologies across countries. This assumption is discussed further in Section 1.3.

Decision Timing

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

In the ordering stage, each village chooses intermediates x_t to use in their farm. After ordering, the shock z_t is realized. After the shock, all production and consumption occurs in the production stage. First, a village chooses how to allocate a village decides how to allocate labor between the agricultural sector, where they can work on the village

farm, and in the manufacturing sector, where they can work for wage w_t which is taxed at rate $\tau \geq 0$. I assume that tax revenue is rebated as a lump-sum transfer $T(b, z)$. After labor is decided, all production takes place. There is a centralized market for buying and selling goods, implying that there is 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.

The distortion τ is designed to capture the fact that the marginal value of labor is lower in agriculture than in manufacturing, a fact discussed in Gollin, Lagakos, and Waugh (2011). Tax revenue is rebated lump-sum as to not change the total income available to the village. This distortion is not required to generate a correlation between A and the intermediate input share, but still has a major quantitative impact.

Recursive Problem

The timing described above implies that the village state variable is savings b_t , and the aggregate state is the distribution of savings across all villages, denoted $\mu_t(b)$. Since I will be studying the stationary equilibrium, I suppress the dependence of the decision problem on the aggregate state $\mu_t(b)$.

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

$$V^P(x, b, z) = \max_{c_a, c_m, n_a, b'} \alpha \log(c_a - \bar{a}) + (1 - \alpha) \log(c_m) + \beta V^O(b') \quad (1.3.2)$$

subject to constraint set

$$\begin{aligned} p_a c_a + c_m + p_a b' &= p_a z A x^\psi n_a^\eta - p_x x + (1 - \tau) w (1 - n_a) + p_a (1 - \delta) b + T(b, z) \\ b' &\geq 0 \\ c_a &\geq \bar{a}, \quad c_m \geq 0 \end{aligned}$$

where V^O is the value of entering the ordering stage at $t + 1$ with b' units of savings in the stationary equilibrium. The first constraint is the village budget constraint, and the

second captures the inability to borrow. The harvesting problem in (1.3.2) defines decision rules $c_a(x, b, z)$, $c_m(x, b, z)$, $n_a(x, b, z)$ and $b'(x, b, z)$. Working backwards, the ordering stage value of entering time t with b savings is

$$V^O(b) = \max_{x \geq 0} \int_z V^P(x, b, z) dQ(z). \quad (1.3.3)$$

This defines the decision rule for intermediate inputs $x(b)$. For future use, aggregate variables will be denoted by capital letters

$$\begin{aligned} N_{at} &= \int_b \left[\int_z n_a(b, z) dQ(z) \right] d\mu_t \\ X_t &= \int_b x(b) d\mu_t \\ Y_{at} &= \int_b \left[\int_z z A x(b)^\psi n_a(b, z)^\eta dQ(z) \right] d\mu_t \end{aligned}$$

so that the intermediate input share in agriculture can be written as

$$\widehat{X}_t = \frac{p_x X_t}{p_a Y_{at}}. \quad (1.3.4)$$

Stationary Equilibrium

I will study the stationary competitive equilibrium of this economy. This is defined by an invariant distribution $\mu = \mu^*$, a value function V^O , decision rules x, n_a, b', c_a, c_m , labor choice N_m , prices p_a and w , and a transfer function $T(b, z)$ such that

1. The value function V^O solves the villages's problem given by (1.3.2) and (1.3.3) with the associated decision rules
2. N_m solves the sector m firm problem (1.3.1)
3. Markets clear

(a) Manufacturing labor market:

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

(b) Agricultural consumption market:

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

(c) Manufacturing consumption market:

$$\int_b \int_z c_m(b, z) dQ(z) d\mu + p_x \int_b x(b) d\mu = AN_m$$

4. The state contingent transfer balances for all (b, z)

$$T(b, z) = \tau w(1 - n_a(b, z))$$

5. The law of motion for μ , denoted $\mu'(\mu)$, is such that $\mu'(\mu^*) = \mu^*$, and μ^* is consistent with $Q(z)$ and decision rules

Discussion

Before turning to the theoretical and quantitative results, I briefly digress to discuss the assumptions related to timing of input choices, the shock distribution, and the savings technology.

Timing

While intermediate inputs are chosen before the realization of the shock, labor is chosen after. This captures the fact that off-farm labor is an important form of insurance for farmers, which is discussed in Kochar (1999). Assuming that labor is also chosen before the shock would amplify the quantitative results, while depriving villages of an empirically relevant form of insurance. Moreover, it implies a constant labor share across countries, consistent with the work of Gollin (2002). Note, however, that this timing implies that villages order intermediate inputs before the shock is realized, and then requires them to take delivery and pay for the intermediates regardless of the shock realization. This setup allows me to parsimoniously capture the risky intermediate input decision, while remaining equivalent to more complicated timing arrangements in which intermediates are produced before they are purchased.

Productivity Shocks

There are three main details of the shock distribution that are worth discussing. First, I assume that the productivity shock is i.i.d. over time. In Bewley models such as this, the ability to self-insure decreases as the persistence of the shock increases. In this sense, I am giving the village the best possible chance to self-insure by assuming z_t is i.i.d.. The quantitative results still suggest that risk is an important margin for understanding intermediate input choices.

Second, the shock realization is independent of the level of intermediate input usage. In particular, using intermediate inputs cannot decrease the variance of shocks. It is well established that fertilizer does not decrease the variance of farm yields. The classic reference is Just and Pope (1979), who find that nitrous fertilizers actually increase yield variance. Traxler et al. (1995) find no effect of fertilizer on yield variance in Mexico. A component of intermediate inputs that does decrease variance is water used for irrigation. However, according to statistics from the World Bank Development Indicators, only 3% of arable cropland in Sub-Saharan Africa is irrigated. If pasture land is included, only 0.4% is irrigated. Uganda for example, irrigates 0.1% of its arable cropland. The poorest ten percent of countries almost exclusively irrigate less than one percent of arable cropland. Because irrigation is used so sparingly, total intermediate inputs should not have a variance-decreasing effect.

Lastly, the distribution of z must have a lower bound sufficiently far from zero. Otherwise, an equilibrium may not exist, since a village may not be able to satisfy subsistence. This turns out not to be an issue when the model is calibrated.

Savings

I assume that the only savings technology available is costly storage of the agricultural good, and insurance is not available. The assumption on the lack of insurance markets is

certainly not controversial in developing countries, and numerous empirical studies have pointed to this as a key feature limiting fertilizer use, including Karlan et al. (2012) and Zerfu and Larson (2010). However, there are many ways to save around risk, and self-insurance has been shown to be effective in limiting the impact of risk in Bewley models. Here, I discuss a few possibilities for savings, and appeal to a vast empirical literature that supports my assumption.

One option is savings banks. In addition to paying no interest, Dupas and Robinson (2011) find that rural savings banks in Kenya actually charge both a start-up fee and a variable fee for every transaction, making savings accounts an expensive way to save for the extremely poor. In twelve of thirteen developing countries considered, Banerjee and Duflo (2007) find that less than 14% of all people living on under \$1 a day have savings accounts. Moreover, saving in cash subjects individuals to depreciation through inflation.

Another option is that villages can simply save in intermediate inputs, and not be subject to storage depreciation. Duflo, Kremer, and Robinson (2011) point out that re-selling fertilizer in western Kenya typically involves a twenty percent discount in addition to the costs of finding a seller. Interestingly, this cost is roughly equal to the calibrated depreciation of agricultural storage. Therefore, the addition of this second savings technology would have no impact on equilibrium outcomes.

A last possibility is savings through capital, which Udry (1996) finds is primarily accounted for by livestock. Empirical studies in Africa however, including Fafchamps, Udry, and Czukas (1998) and Kazianga and Udry (2006), have found little response of livestock sales to negative shocks (e.g. drought). In fact, the latter finds that self-insurance through grain storage is the key method through which farmers attempt to smooth consumption. During a prolonged drought in Burkina Faso, Reardon, Matlon, and Delgado (1988) find that although there was little change in livestock holdings, cereal stocks were almost

completed depleted.³ This leads me to the assumption that self-insurance is available only through storage of the agricultural output good. However, as pointed out in Kazianga and Udry (2006), this method of self-insurance still yields little consumption smoothing in the data. This same outcome is generated by the model developed here, due to the fact that storing agricultural output is incredibly costly in African countries. Using a new set of storage depreciation rates, I find that storage losses of twenty percent are common. In Zimbabwe, for example, almost 30% of maize produced is lost in storage. This is further detailed in the calibration of Section 1.5.

1.4 Characterization and Analytic Results

This section provides some analytic results to clarify the mechanics of the model. Section 1.4 begins by discussing the importance of subsistence requirements in a model with productivity shocks. Subsistence requirements generate decreasing relative risk aversion, which has important implications for how villages make risky intermediate input decisions. The implications of this result are discussed in Section 1.4, which shows that TFP differences generate intermediate share differences if and only if the economy includes incomplete markets, idiosyncratic shocks, and subsistence requirements.

Subsistence Requirements and Risk

This paper exploits a feature of the utility function that has yet to be explored in a cross-country framework. Namely, subsistence requirements change the relative risk aversion of a standard constant relative risk aversion (CRRA) utility function to decreasing relative risk aversion (DRRA). It is not immediately clear how to define relative risk aversion, however, because villages value two types of consumption. It turns out that the utility function can be rewritten as a function of total income, which allows relative risk aversion to be directly defined over income levels. To see this, first define y as the total income at the production

³On an aggregate level, Lagakos and Waugh (2012) find that capital per worker differences account for similar percentages of output per worker differences in agriculture and non-agriculture. While capital per worker differences are important for understanding aggregate output per worker differences, they are not responsible for the fact that agriculture is significantly *less* productive than non-agriculture.

stage, given savings b , intermediate choice x , shock z , and the optimal savings decision rule b'

$$y(x, b, z) = p_a z A x^\psi n_a^\eta - p_x x + (1 - \tau) w (1 - n_a) + p_a (1 - \delta) b + T(b, z) - p_a b'.$$

With income y , a village purchases enough agricultural consumption to satisfy subsistence \bar{a} , then splits the rest of their income between the two sectors based on the relative weights assigned by the price p_a and utility parameter α .

$$\begin{aligned} c_a(y) &= \bar{a} + \frac{\alpha}{p_a} (y - p_a \bar{a}), \\ c_m(y) &= (1 - \alpha) (y - p_a \bar{a}). \end{aligned}$$

Using these decision rules, the utility flow can be rewritten as a function of total income y ,

$$\tilde{u}(y) := u(c_a(y), c_m(y)) = \Omega - \alpha \log(p_a) + \log(y - p_a \bar{a}) \quad (1.4.1)$$

where $\Omega = \alpha \log(\alpha) + (1 - \alpha) \log(1 - \alpha)$. Because utility \tilde{u} is only a function of income y , relative risk aversion with respect to total income y , given \bar{a} and price p_a , can be defined as

$$R(y|\bar{a}, p_a) = \frac{y}{y - p_a \bar{a}}.$$

If $\bar{a} = 0$, this is a standard log CRRA utility function. However if $\bar{a} > 0$, the utility function instead exhibits decreasing relative risk aversion (DRRA), consistent with the household evidence of Ogaki and Zhang (2001) from India and Pakistan.

With this form of the period utility function, the production stage utility can be written

$$V^P(x, b, z) = \Omega - \alpha \log(p_a) + \log(y(x, b, z) - p_a \bar{a}) + \beta V^O(b'(x, b, z)). \quad (1.4.2)$$

The indirect utility function at the ordering stage is then

$$V^O(b) = \Omega - \alpha \log(p_a) + \max_{x \geq 0} \int_z \left[\log(y(x, b, z) - p_a \bar{a}) + \beta V^O(b'(x, b, z)) \right] dQ(z). \quad (1.4.3)$$

Equations (1.4.2) and (1.4.3) illustrate the key tension between expected income and expected utility. While profits drive production stage utility by increasing y , the ordering stage choice of x maximizes expected utility, of which income is only one component. The other is the risk associated with the choice of x . While farm profit increases utility, higher x implies large exposure to risk. To limit this exposure and decrease the variation in production utility, the village must decrease its ex-ante choice of x . Thus, the optimal intermediate input choice must balance the desire for both high income and low exposure to risk. This balancing act allows subsistence requirements to play an important role. Since $\bar{a} > 0$ implies DRRA, the inclusion of subsistence requirements alters the way farmers undertake risky investments for different levels of TFP. Section 1.4 shows that this the inclusion of subsistence requirements interacts with TFP differences and uninsurable shocks to generate differences in the intermediate share across economies.

Characterizing Intermediate Input Shares

In this section, I show that TFP differences generate differences in intermediate input shares if and only if the economy includes incomplete markets, idiosyncratic shocks, and subsistence requirements. This qualitatively replicates the positive correlation between intermediate input shares and income detailed in Section 1.2. To make these results as sharp as possible, I consider the static version of the model (identically, $\delta = 1$ for all economies). Furthermore, because the two exogenous distortions p_x and τ are not required to generate this positive correlation between the intermediate input share and productivity A , I fix $\tau = 0$ and $p_x = 1$ in all economies. This leaves sector-neutral TFP A as the only difference between any two model economies. All proofs are relegated to Appendix C.

The assumed productivity shocks must translate into consumption risk for subsistence requirements to play a role. Therefore, to assess the role of productivity shocks and incomplete markets, I compare the model developed above (denoted by superscript I for incomplete markets) with a complete markets version (denoted by superscript C for com-

plete markets). The complete markets version is identical, except that villages are allowed to trade a full set of state contingent assets before the realization of z . How this affects intermediate input choices can be seen by comparing the first order conditions with respect to x in the I and C economies. Farmers maximize expected profit with complete markets, because consumption is fully insured. The ordering stage first order condition in the complete markets economy is therefore

$$Ap_a^{1/(1-\eta)} F'(x) \int_Z z^{1/(1-\eta)} dQ(z) = 1 \quad (1.4.4)$$

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 I economy), the first order condition of equation (1.4.3) yields

$$Ap_a^{1/(1-\eta)} F'(x) \int_Z z^{1/(1-\eta)} \left(\frac{\tilde{u}'(y(x,z))}{\mathbb{E}_z[\tilde{u}'(y(x,z))]} \right) dQ(z) = 1 \quad (1.4.5)$$

where \tilde{u} is defined as in equation (1.4.1), and \tilde{u}' is the derivative with respect to income y . The two first order conditions are exactly the same except for the addition of marginal utility to the integrand of equation (1.4.5). Instead of just weighting each shock realization by the probability that it occurs, incomplete markets imply that villages weight by their risk-neutral probabilities. This captures the fact that villages internalize the impact their intermediate choice has on consumption in the absence of insurance markets. The addition of marginal utility to the weight assigned by the village implies that those realizations of z that imply a higher than average marginal utility are weighted relatively more heavily by a village that faces uninsurable risk. Similarly, those realizations of z that imply a lower than average marginal utility are assigned less weight. Thus, the inclusion of incomplete markets tilts the weight assigned by every village toward “bad” outcomes relative to a profit-maximizing village. This leads naturally to Proposition 1.

Proposition 1. *In the competitive equilibrium, the intermediate share is lower in the incomplete markets economy (I) than the complete markets economy (C) for a given TFP A.*

That is,

$$\frac{X^I}{p_a^I Y_a^I} < \frac{X^C}{p_a^C Y_a^C} = \psi$$

Graphically, this result can be seen in Figure 1.3. In the complete markets economy (C), the stochastic discount factor is equal to one at every realization of z , and is shown in Figure 1.3 as the horizontal dotted line at one. In the incomplete markets economy (I) this changes. Stochastic discounting at low z realizations increases, which can be seen in the solid line of Figure 1.3.

Proposition 1 only considers the role of incomplete markets within one economy. The more interesting issue is how the intermediate input share reacts to changes *across* economies, because the empirical evidence presented in Section 1.2 suggests a positive correlation between the intermediate input share and TFP A . First, with $\bar{a} = 0$, the interaction of incomplete markets and agricultural productivity shocks is irrelevant in accounting for the fraction of the labor force in agriculture, the intermediate share, or agricultural productivity differences.

Proposition 2. *In the model with uninsurable shocks (I economy) and $\bar{a} = 0$, the following results hold in the competitive equilibrium:*

1. $n_a(z)$ is independent of A
2. The intermediate share $X/(p_a Y_a)$ is independent of A
3. For two economies with TFP levels A^1 and A^2 , agricultural output per worker differences in the I economy are the same as in the C economy. That is,

$$\frac{Y_a^{1C}/N_a^{1C}}{Y_a^{2C}/N_a^{2C}} = \frac{Y_a^{1I}/N_a^{1I}}{Y_a^{2I}/N_a^{2I}}$$

While Proposition 1 shows that the equilibrium intermediate input share is lower with incomplete markets, Proposition 2 shows that when $\bar{a} = 0$, it does not differ *across*

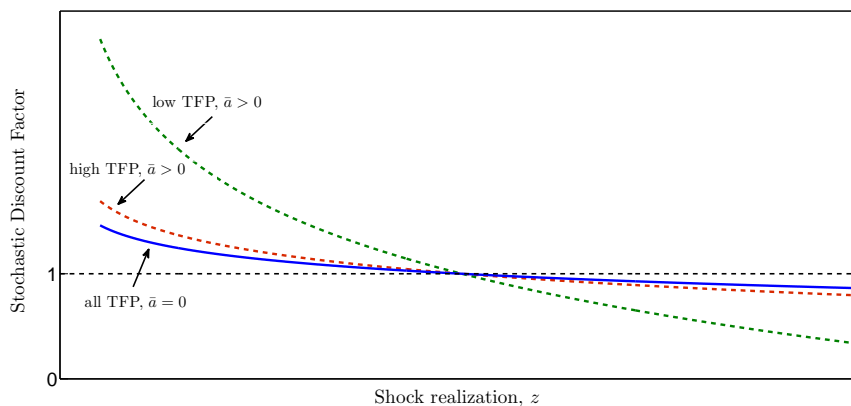
economies. This is due to the fact that $\bar{a} = 0$ implies that the period utility function exhibits CRRA. Intuitively, A can be thought of as a decrease in the mean income realization. With CRRA, the general equilibrium price of agricultural output increases to incentivize poor villages to take on more risk. In fact, it exactly offsets the decrease in mean income, and makes the stochastic discount factor independent of A . This can be seen in the solid line of Figure 1.3, which shows that stochastic discounting for any realization of z is identical for all levels of TFP. Moreover, the third result in Proposition 2 shows that, in the absence of subsistence requirements, the lack of insurance markets plays no role in understanding labor productivity differences across countries. That is, the predicted agricultural productivity differences in the model with uninsurable risk and no subsistence requirements are exactly the same as they would be with complete markets. The inclusion of subsistence requirements breaks this result. When $\bar{a} > 0$, the period utility function now exhibits DRRA, causing the stochastic discount factor to depend on the level of TFP. Proposition 3 shows that the interaction of productivity shocks and subsistence requirements can qualitatively replicate the empirical correlation between the intermediate share and TFP from Section 1.2.

Proposition 3. *In the competitive equilibrium, the intermediate share is increasing in A if and only if $\bar{a} > 0$.*

In an economy with incomplete markets, idiosyncratic shocks, and subsistence requirements, TFP differences are able to generate differences in the intermediate input share that are qualitatively consistent with the evidence provided in Figure 1.1, while leaving out any one of these features implies a constant intermediate share across countries. Technically, this result is driven by the interaction of two features implied by subsistence requirements: DRRA utility and an income elasticity that is less than one with respect to the agricultural good. The intuition, however, is as follows. Poor farmers have relatively less income than their rich counterparts for all realizations of z . With subsistence requirements,

this difference increases as z decreases. Since farmers weigh each realization of z by their marginal utility at that realization, farmers in poor economies put relatively more weight on low z than their rich counterparts, as can be seen in Figure 1.3. This causes the intermediate good share to decrease in economies with low A .⁴

Figure 1.3: Stochastic Discounting for Different Subsistence Levels



While Proposition 3 is consistent with the aggregate evidence presented in Section 1.2, the model achieves this result by predicting that villages leave potential profits unrealized due to their own risk aversion. This naturally leads to the obvious concern of empirical support for this prediction. In fact, there is significant evidence of under-investment (relative to the profit-maximizing choice) in fertilizer. Duflo, Kremer, and Robinson (2008), for example, use a randomized trial in Kenya to show that increases in fertilizer use can dramatically increase farm yields. Moreover, mounting evidence supports the model’s prediction that this under-investment is driven by the interaction of individual risk aversion and uninsurable shocks. This evidence is derived from both randomized trials by Karlan et al. (2012) and panel surveys by Zerfu and Larson (2010). Given that the model is theoretically consistent with both household and aggregate evidence on intermediate input choices, I

⁴This result is not generically true in the dynamic model, due to the presence of borrowing constraints. However, in quantitative simulations, the result holds. Intuitively, when $\bar{a} = 0$ the precautionary savings motive is quite low, in part due to the negative real return of agricultural storage. The stationary equilibrium therefore has all individuals at zero savings, and is identical to the static equilibrium for which the results are proved.

now turn back to the full dynamic model to assess the quantitative impact for cross-country labor productivity.

1.5 Quantitative Exercise and Calibration

The goal is now to quantitatively assess the importance of agricultural productivity shocks for cross-country productivity and intermediate input shares. I do so by comparing the model's predictions for a rich and poor country. The rich country is designed to capture the relevant features of the U.S. economy. I normalize $A = 1$ and calibrate the model with no distortions ($p_x = 1$ and $\tau = 0$) so that the stationary equilibrium matches a number of features of the U.S. economy, including the intermediate input share in agricultural and the sectoral composition of employment. The poor economy differs in its level of TFP A , the depreciation rate of storage δ , the intermediate input price p_x , and the tax rate τ . These are all chosen to match the relevant features of the tenth percentile country as ranked by per capita GDP. To construct the tenth percentile country, I take the average values from the bottom fifteen to five percent of countries. This averages out some of the variation in intermediate input shares and intermediate input prices. See Appendix B for more details.

I then proceed to consider two quantitative experiments. The first experiment considers the impact on intermediate input shares and labor productivity. Because the poor economy differs along a number of dimensions, some differences in labor productivity will be exogenously fed into the model. Recall, however, that the model with no productivity shocks generates no differences in intermediate input shares. Therefore, to isolate the impact of intermediate input share differences, I ask how much *larger* productivity differences are in the model with shocks, relative to the identical model with no shocks.

The second exercise is to vary the exogenous parameters p_x , τ , and δ in the poor model economy while holding all other parameters fixed. The goal of this exercise is to understand (1) their role in account for intermediate input share differences and (2)

the complementarity of these different distortions for understanding the results.⁵ Section 1.5 presents the parameters that are the same across economies. Section 1.5 details the differences between the two economies in the baseline calibration. Table 1.3 lists all the parameters chosen.

Common Parameters

The parameters that are the same in both economies are the agricultural production technology parameters (except for TFP), utility parameters, and the shock distribution. These are discussed in turn.

The farm production parameters are the shares of intermediates, ψ , and labor, η . These are chosen to match the aggregate intermediate input share and labor share in agriculture in the United States in 1985. The exponent on intermediates, ψ , is set slightly above 0.40 to match an intermediate input share of 0.40 in the rich economy. This is consistent with the statistics from Prasada Rao (1993) presented in Figure 1.1. Since labor is chosen after the realization of all uncertainty, the parameter η is exactly equal to the payments to labor as a share of gross agricultural output. However, this share is difficult to define in the absence of capital. I therefore choose $\eta = 0.40$, which is consistent with the labor share in Restuccia, Yang, and Zhu (2008). Because estimates of this parameter vary widely, Appendix A considers the sensitivity of the results to this parameter.

Next I calibrate the utility function parameters. Since the model period is a year, I set $\beta = 0.96$. The remaining parameters are the weight on agricultural consumption, α , and subsistence \bar{a} . The parameter α controls the share of agricultural output in GDP in the long run as TFP approaches infinity. I set $\alpha = 0.005$, following Restuccia, Yang, and Zhu (2008) and Lagakos and Waugh (2012). The parameter \bar{a} is chosen so that the rich

⁵The goal of this paper is not to explain these distortions but, given that they exist, to understand their interaction with productivity shocks in the agricultural sector. See Adamopoulos (2011) for the role of the transportation sector and Estevadeordal and Taylor (2013) for the role of import tariffs in accounting for high intermediate input prices. Theories of sectoral differences in the marginal value of labor have been proposed by Caselli and Coleman (2001) and ?, while Gollin, Lagakos, and Waugh (2011) quantitatively investigate the causes using detailed micro data.

economy has an equilibrium agricultural employment share of 2.84%, consistent with the U.S. in 1985. This implies $\bar{a} = 0.048$.

The distribution of shocks to agricultural production is identical in the two economies. While these distributions may be quite different in reality, the distinction is quantitatively irrelevant. U.S. villages act similar to profit maximizers because they are so far from subsistence. When A decreases however, villages become much more sensitive to the shape of this distribution because they are (ex-ante) closer to subsistence. Therefore the distribution is chosen to match the poor economy, and I make the innocuous assumption that the distribution is the same in the rich economy. This, of course, requires data from a developing country.

This is not the only issue confronting the construction of the shock distribution, as estimating a distribution that is consistent with the model presents a number of other challenges. First, inputs and outputs in the model are aggregated. Shocks are to “agricultural production,” for example, not just maize production. Construction of aggregated quantity indices require both prices and quantities for a variety of inputs and outputs. Second, as pointed out by Townsend (1994), not all risk is uninsurable. While purely idiosyncratic risk can be insured relatively well through informal arrangements within villages, covariate risk that hits entire villages simultaneously is not. The quantitative results will be overstated if I include insurable output variation as risk.

With these issues in mind, I turn to the International Crops Research Institute for the Semi-Arid Tropics Village Level Surveys (ICRISAT VLS). This data set contains daily diaries of plot-level activities, including input and output usage, from ten different Indian villages. It also contains both quantities and valuations of these goods, so that I can impute prices and construct the aggregate quantity indices required by the model. Just as importantly, the notion of a village in the data and a village in my model match. Ogaki and Zhang (2001) use the same data and reject risk sharing across these villages, suggesting

that village-level shocks are uninsurable. This data set is therefore able to overcome the main issues required in constructing the shock distribution.

I provide an overview of the procedure used to construct the shock distribution here, while a more detailed explanation is given in Appendix D. The data covers three seasons per village-year. Since the model includes only one production stage per year, the first goal is to construct village-year-season quantity indices of agricultural output (Y), agricultural intermediates (I), nonagricultural intermediates (X), labor (N), and capital (K). I refer to these as “categories.” While capital is not in the model, I do not want to include variation in capital and land quality as risk, so I must include it here. The production function from the model then provides me with residuals, which can be used to construct village-level shocks.

One minor issue is prices are not directly reported. Instead, they are imputed from quantities and values. These are only recorded when the input is used or output produced, so I do not have a complete time series of prices for all inputs and outputs. To combat this, I first construct the total village-year-season value of each category (I , X , N , K , and Y) used. This simply requires summing up the value of all goods in each category across all farms in that village. The next step is to construct the quantity of the aggregated indices, which requires the construction of a price index and therefore prices. While missing some prices, every village has at least one (and usually exactly one) good in every category that is used every period. I choose this good as my numeraire, and use the price time series as my village-category price deflator. This gives me the needed village-year-season quantity indices for the required inputs and outputs.

Plugging the quantities into the production function from the model generates Solow residuals for village i over year-season time periods t

$$z_{i,t}^* = \frac{Y_{i,t}^{data} - I_{i,t}^{data}}{(X_{i,t}^{data})^\psi (N_{i,t}^{data})^\eta (K_{i,t}^{data})^{1-\psi-\eta}}$$

I assume that these residuals $z_{i,t}^*$ are comprised of exponential time trend, and village and

seasonal fixed effects, and the random shock z . That is

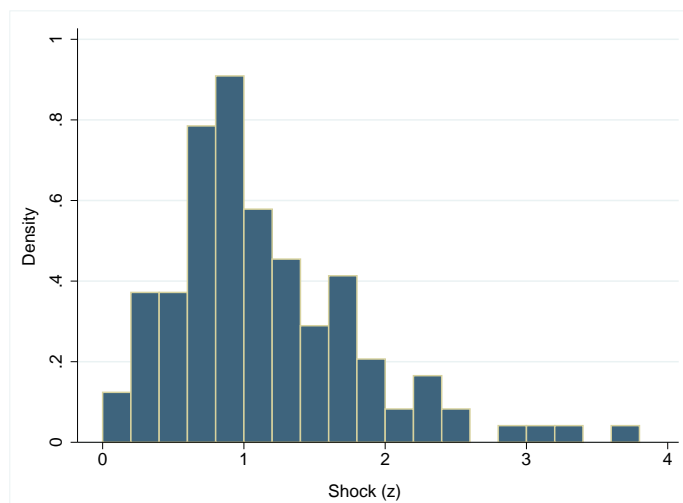
$$z_{i,t}^* = \mathbf{vs}(1 + g)^t z_{i,t}.$$

To isolate $z_{i,t}$, I take logs and run the following regression

$$\log(z_{i,t}^*) = \log(\mathbf{v}) + \log(\mathbf{s}) + t \log(1 + g).$$

The error from the regression, $\varepsilon_{i,t}$, is equal to $\log(z_{i,t})$. Therefore, $z_{i,t}$ is drawn from a log-normal distribution, and the shock $z_{i,t} = \exp(\varepsilon_{i,t})$. The empirical probability distribution of $z_{i,t}$ is displayed in Figure 1.4. The underlying error term ε has a standard deviation of $\sigma_\varepsilon = 0.59$, and a mean of zero. To match in the model, I assume that $\log(z)$ is drawn from a truncated normal distribution with standard deviation 0.59. This distribution is assumed to have support on $[\log(0.10), \log(4)]$. Consistent with the error term from the regression, all shocks therefore fall in the interval $[0.10, 4]$. The continuous distribution is approximated by a twenty point discrete distribution.

Figure 1.4: Empirical Distribution of Shocks



Economy Specific Parameters

The two economies differ along four dimensions: TFP A , the tax rate τ , intermediate input price p_x , and lastly the depreciation of stored goods δ ,

For the U.S. economy, TFP is normalized to $A = 1$. I discipline the TFP in the poor country by manufacturing labor productivity. Since manufacturing labor productivity is equal to A , I set $A = 0.25$ in the poor economy, which is roughly consistent with nonagricultural labor productivity differences between the richest and poorest countries.

Since the rich model economy is assumed frictionless, $\tau = 0$ and $p_x = 1$. For the poor model economy, I choose $\tau = 0.40$. This is roughly consistent with differences in the marginal value of labor across sectors found in Vollrath (2009). The intermediate price in the poor economy is $p_x = 3$. This is taken from Restuccia, Yang, and Zhu (2008), who use FAO data to show that there is a strong correlation between per capita income and intermediate input prices across countries.

Agricultural storage technologies differ between rich and poor countries. For one, the abundance of silos in developed countries storage provide prima facie evidence of difference in depreciation rates between rich and poor countries. This depreciation rate plays an important role in this analysis, since it controls the ability of villages to save their way away from subsistence. New statistics from the African Post Harvest Loss Information System (APHLIS) allow me to discipline this storage technology. APHLIS is a network of local experts that aggregates statistics on weight loss into comparable measures across African countries and crops. Table 1.2 presents the estimated weight loss data for a number of crops in a selection of African countries.⁶

While these weight losses already paint a dire picture of storage in Africa, a distinction must be made between weight and quality losses. Since the model contains no notion of quality, the exact empirical counterpart would be depreciation of the *value* of agricultural output. However, quality losses are notoriously difficult to measure, and certainly do not change one-to-one with weight losses. With this caveat in mind, I conservatively set

⁶See Hodges et al. (2010) for a more complete review of APHLIS. Considering more countries only emphasizes the results. The large weight losses presented in Table 1.2 are present in almost all countries in the data set.

Table 1.2: Post Harvest Weight Loss (%) for Selected Countries and Crops for 2007

	Maize	Wheat	Sorghum	Millet	Rice
Eritrea	17.9	12.9	12.2	10.9	–
Ethiopia	16.4	12.4	12.4	12.1	11.3
Kenya	21.1	12.9	12.7	11.9	13.2
Malawi	19.6	13.4	13.0	12.9	11.6
Mozambique	21.0	–	12.8	12.6	11.4
Rwanda	17.5	14.5	12.5	–	11.3
Sudan	18.0	12.9	12.2	10.7	–
Tanzania	22.0	14.4	12.5	12.3	11.2
Median	19.6	12.9	12.5	12.1	11.4

Note: Data from APHLIS

$\delta = 0.15$ in the poor economy. I set $\delta = 0.03$ in the rich economy. Since the rich model economy has little need for precautionary savings, changing this value does not influence the results.

1.6 Quantitative Results

Section 1.6 considers the calibrated model's ability to predict differences in intermediate input shares and labor productivity. I find that the model is consistent with the fact that developing countries have lower intermediate input shares and higher employment in agriculture. This generates significantly lower labor productivity in the agricultural sector. In Section 1.6, I investigate the implications of changing the exogenous distortions p_x , τ , and δ .

Impact of Agricultural Risk

The main model results are presented in Table 1.4. Columns two and three, under the heading *Labor Productivity Gap*, present the agricultural and aggregate labor productivity differences between the two economies. These are measured as the rich-to-poor ratio in

Table 1.3: Parameter Values for Two Economies

Parameter	Rich	Poor
<i>Specific</i>		
A	1.00	0.25
τ	0.00	0.40
p_x	1.00	3.00
δ	0.03	0.15
<i>Common</i>		
ψ	0.40	0.40
η	0.40	0.40
α	0.005	0.005
\bar{a}	0.048	0.048
$z, Q(z)$	<i>(see Figure 1.4)</i>	

agricultural output per worker and GDP per worker. The latter is measured as GDP at the rich economy model price, since the total labor force is normalized to one. To understand how these labor productivity gaps are generated, columns four and five are the intermediate input shares in both economies, while columns six and seven are the employment share in agriculture. Table 1.4 shows that the addition of agricultural productivity shocks to the model generates significant amplification of labor productivity differences. Agricultural productivity differences are amplified 41% from 23.8 to 33.5, or fifteen percentage points closer to the ratio of 63.7 found in the data. The increase is even larger at the aggregate level, in which the addition of shocks amplifies GDP per worker differences by 75% from 4.3 to 7.5. This is fourteen percentage points closer to the ratio of 23.1 found in the data. The model with agricultural productivity shocks gets significantly closer to the data along both productivity dimensions, implying that agricultural productivity shocks are a key component of aggregate income differences across countries.

Table 1.4: Impact of Productivity Shocks on Labor Productivity

Economy	<i>Labor Productivity Gap</i>		$p_x X / p_a Y_a$		N_a (%)	
	Agriculture	Aggregate	Rich	Poor	Rich	Poor
<i>Data</i>	63.7	23.1	0.40	0.09	2.8	82.0
<i>Model with</i>						
Shocks	33.5	7.5	0.40	0.22	2.8	56.8
No shocks	23.8	4.3	0.40	0.40	2.8	46.6

The amplification occurs due to a change in input decisions. By virtue of the Cobb-Douglas production function, the model with no shocks predicts no change in the intermediate input share across countries. Once agricultural shocks are included, the intermediate share prediction for the poor economy decreases from 0.40 to 0.22. This decrease captures 58% of the actual difference between rich and poor countries. The lack of intermediate inputs used in the poor country forces villages to substitute more labor to reach subsistence consumption. The prediction of the agricultural labor force increases from 46.6 to 56.8 percent of the population, an increase of 22%. Just as with the intermediate input shares, the model with shocks is better aligned with the data along this dimension.

Decomposition of Results

This amplification generated by productivity shocks depends on a number of exogenous distortions in the model, including agriculture-specific distortions p_x and τ , and the storage depreciation rate δ . I isolate the impact of each of these by computing a series of counterfactual poor model economies. Each one assumes that one of these features is equal to the calibrated U.S. level, instead of the higher level originally calibrated. The results are presented in Table 1.5. The first two columns are agricultural and aggregate output per worker in the rich economy relative to the poor economy. The last two columns are the agricultural employment share and intermediate input share in the poor economy (the inputs in the rich economy obviously do not change). The first row is the calibrated model with all the differences already discussed. Then I turn off differences in depreciation δ , agricultural

distortions (p_x, τ) , and then both simultaneously. Lowering the depreciation rate decreases agricultural productivity differences by 13% and GDP per worker differences by 11%. This difference, though somewhat small, is driven by the fact that lowering the depreciation rate allows villages to save away from subsistence. This makes them more willing to take on risk and use more intermediate inputs, and as such, the intermediate input share increases by 32%. Lower price distortions work somewhat similarly, though with a much larger magnitude. Agricultural labor productivity decreases by 64% when distortions are lowered to the U.S. level. This is partially driven by the fact that the intermediate input share increases by 45%, freeing up labor to move into the manufacturing sector. The intuition is similar to that of lower depreciation rates. Villages move further from subsistence with lower distortions, and therefore are willing to take on more risk. The last line shows that the impact of lowering the depreciation rate is muffled in the presence of low agricultural distortions. If distortions are at the U.S. level already, a change in the depreciation rate lowers agricultural productivity differences by only 4%, and increases the intermediate input share by only 9%. As villages move away from subsistence in response to lower distortions, they have little need for precautionary savings, especially with a negative real return. Although the impact of agricultural distortions dwarf that of differences in storage depreciation in terms of accounting for agricultural labor productivity, the complementarity between the two is key for understanding the full impact of agricultural productivity shocks in developing countries.

Table 1.5: Decomposition of Quantitative Results

	<i>Labor Productivity Gap</i>		<i>Poor Economy Inputs</i>	
	Agriculture	Aggregate	N_a (%)	$p_x X / p_a Y_a$
<i>Baseline</i>	33.5	7.5	56.8	0.22
<i>No cross-country difference in</i>				
δ	29.2	6.7	49.8	0.29
(p_x, τ)	12.2	5.1	20.8	0.32
δ or (p_x, τ)	11.6	5.1	19.1	0.35

The Role of Intermediate Input Price

The relatively high price of intermediate inputs in developing countries has been pointed to as an important factor limiting growth, with implications ranging from agricultural productivity to trade policy. I further investigate the role played by the intermediate input price for understanding labor productivity and the intermediate input share in the poor model economy. Figure 1.5 plots the response of agricultural labor productivity and the intermediate input share to changes in the intermediate input price for two cases: with shocks (the solid line) and without shocks (the dotted line).

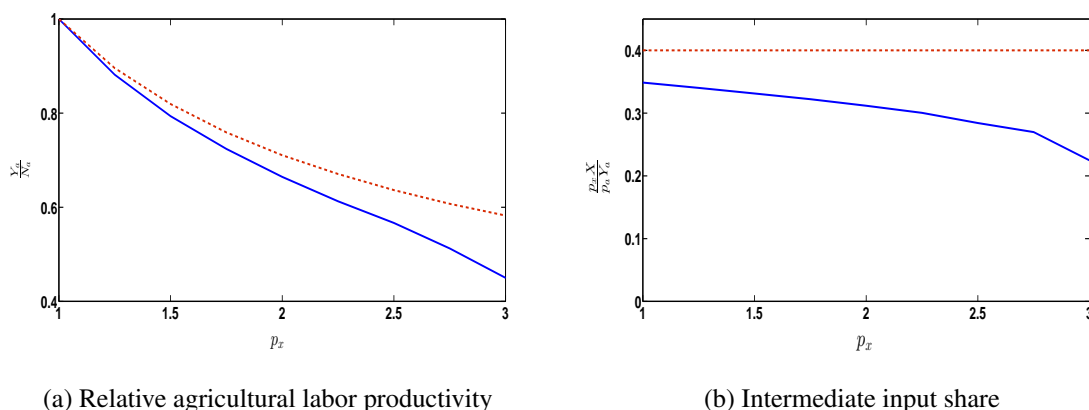


Figure 1.5: Poor economy model implications for different p_x

The first thing to point out is that a higher intermediate input price lowers labor productivity substantially. As intermediate inputs become more expensive, farmers substitute relatively cheaper labor services. This “substitution effect” lowers labor productivity regardless of the presence of shocks. The addition of shocks, however, predicts a larger impact of increasing the price. This is driven by an “income effect” present in this model, but not in the model without shocks. As intermediate inputs become more expensive, farmers become poorer and less willing to take on risk, and therefore reduce their intermediate input usage even further beyond what is generated by the substitution effect. This can be

seen in Figure 1.5b, which shows that the intermediate input share decreases in response to higher prices. This does not occur in the absence of these shocks, due to the absence of this income effect. As the intermediate input share gets lower, the predicted decrease in labor productivity becomes larger.

Through the combined force of these two effects, this model predicts significantly larger losses in productivity from a change in intermediate input prices. The model predicts that an increase of the price from one to three decreases labor productivity by about forty percent without shocks, and by fifty five percent with shocks. Therefore, this model presents a new margin through which intermediate input prices can affect labor productivity, and implies that the gains from lowering these distortions are larger than predicted in models without risk. If the intermediate input share was lowered to U.S. levels, agricultural labor productivity would increase by approximately 120%, and the intermediate input share would increase by 55% due to the decreased importance of this income effect. The model without shocks predicts an increase in labor productivity of 71%, and no change in the intermediate input share.

Changes in Savings Technologies

Table 1.5 shows that the ability to save has a significant impact if it can remove villages from very close to subsistence, but otherwise seems to have little impact. In this section, I evaluate this further by comparing agricultural productivity for a variety of depreciation rates. For a variety of depreciation rates, Figure 1.6 plots the poor economy response of the intermediate input share and agricultural productivity for two cases: the baseline calibration of $(p_x, \tau) = (3, 0.40)$ (solid line) and the case where $(p_x, \tau) = (1, 0)$ (dashed line).

As in Table 1.5, Figure 1.6 shows that lowering the depreciation rate increases productivity, but depends on the level of distortions. With no distortions, a decrease in depreciation from $\delta = 0.15$ to $\delta = 0.03$ is required to increase agricultural labor productivity by 5%. With the calibrated distortions, depreciation only needs to decrease from $\delta = 0.15$

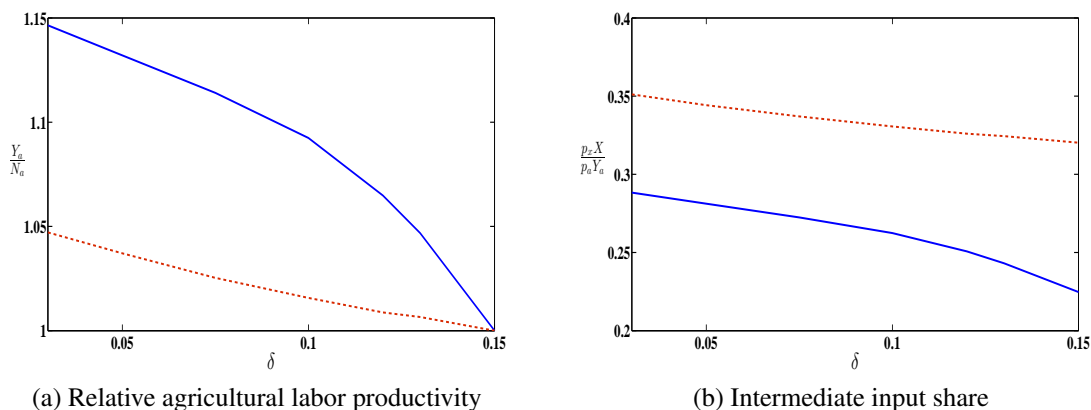


Figure 1.6: Poor economy model implications for different depreciation rates δ

to $\delta = 0.13$ to generate an identical increase in productivity. This is due to the fact that a change in depreciation has little impact on input choices once villages are sufficiently far from subsistence. In fact, in the absence of agricultural distortions, a decrease in depreciation from $\delta = 0.15$ to $\delta = 0.03$ only increases the intermediate input share by 10%. In the presence of these distortions, a decrease from $\delta = 0.15$ to $\delta = 0.13$ provides a similar percentage increase, just as was the case for productivity. Therefore, there is scope for δ to impact productivity, but it depends almost entirely on the economy being extremely poor.

1.7 Conclusion

This paper quantifies the role of idiosyncratic production risk in accounting for sectoral output per worker differences in a two sector general equilibrium model. In poor countries, farmers use fewer intermediate inputs, driving down agricultural productivity. The model captures about sixty percent of the difference in intermediate input shares between the richest and poorest countries, even though underlying farm technologies are Cobb-Douglas. Technically, this result is due to the interaction of uninsurable risk with DRRA preferences generated by subsistence requirements. This has important quantitative implications for productivity across countries. Relative to an identical model with no productivity shocks, agricultural productivity differences are amplified by about forty percent, while aggregate

productivity differences are amplified by seventy five.

The model also provides a new channel through which sector-specific distortions can impact labor productivity. Since distortions decrease income, farmers become more risk averse and therefore choose to use even fewer intermediate inputs. Quantitatively, these distortions are key to understanding the complete impact of agricultural risk. Counterfactual experiments show that lowering these distortions facilitates increased self-insurance on the part of farmers, decreasing the impact of agricultural risk on intermediate input shares and productivity. The model predicts that decreasing these distortions to U.S. levels increases the intermediate input share by 45% in the poorest countries in the world.

Chapter 2

FACTORS AFFECTING COLLEGE COMPLETION AND STUDENT ABILITY IN THE U.S. SINCE 1900

2.1 Introduction

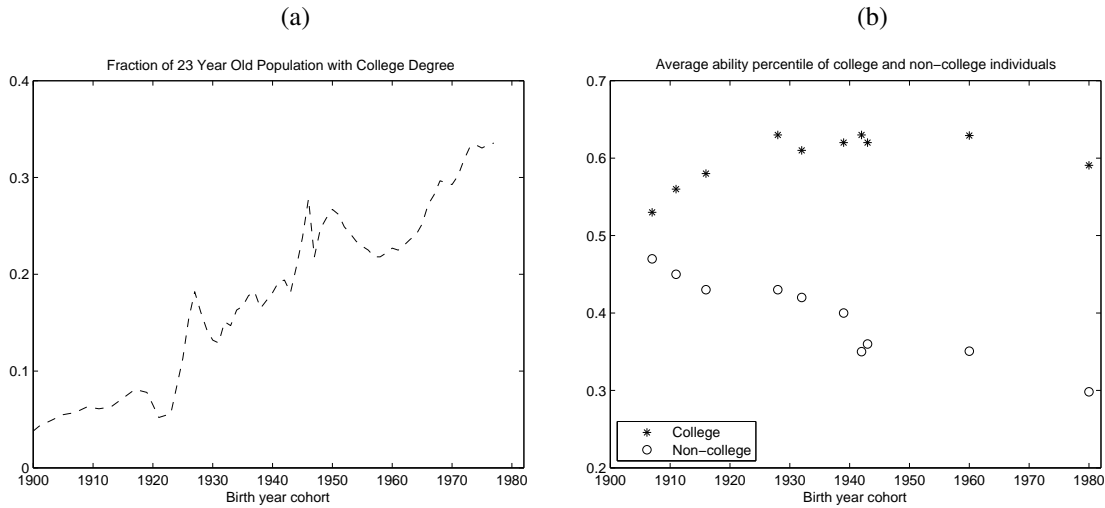
The twentieth century saw a dramatic expansion of higher education in the United States. Among those in the 1900 birth cohort, less than 4% held a bachelor's or first professional degree at age 23, but by the 1970 birth cohort this share had risen to more than 30%. Panel (a) of Figure 2.1 plots this series for all cohorts from 1900 through 1977.¹ Concurrent with the increase in college attendance, the ability gap widened substantially between college students and those individuals with a high school degree and no college experience, i.e., “non-college” individuals. This pattern is seen in Panel (b) of Figure 2.1, which plots the average IQ percentile (the proxy for “ability”) of college and non-college individuals.² For example the average college student born in 1907 had an IQ in the 53rd percentile, very close to the average non-college individual whose IQ was in the 47th percentile. Yet over the next several decades, the average IQ percentile increased among college enrollees and decreased among those with only a high school degree. Most intriguing is that this trend of increased ability sorting occurred even as the share of students attempting college grew steadily larger.

The goal of this paper is to understand the causes of these two empirical trends. However, this task is complicated by the vast number of changes in both the aggregate economy and education sector over this time period. I combat this by developing an overlapping generations lifecycle model populated by high school graduates who are heterogeneous in both ability and financial assets. An important feature of the model is that

¹The 1977 cohort was 23 years old in 2000 when this data series ends. Data for cohorts born up to 1967 are taken from Snyder (1993), and from 1968 through 1977 are the authors' calculation.

²These two data trends have also been documented by other authors, including Hendricks and Schoellman (2012). In panel (b), data points for cohorts prior to 1950 are from Taubman and Wales (1972). The 1960 data point is from the NLSY79, as calculated by Hendricks and Schoellman (2012). The 1980 data point is my calculation based on data from the NLSY97.

Figure 2.1: College Completion and Average Student Ability in the U.S. since 1900



individuals only see a noisy signal of their true ability when making risky decisions about college enrollment. I incorporate newly constructed data on college costs obtained from historical printed government sources. Additionally, I estimate life-cycle wage profiles for men and women in each birth-year cohort in order to accurately model the opportunity costs of wages foregone by college attendees and the education earnings premia realized by those who either complete some college or successfully graduate from college.

I calibrate parameters of the model to match the U.S. data and then conduct a series of experiments in order to understand changes in college completion and ability sorting over time. First, I find that the secular increase in high school completion is responsible for less than half of the increase in college completion over the entire time period. The remainder is due to changes in college enrollment and completion rates *conditional on high school graduation*. Interestingly, however, the key features of the model allowing us to match the data depend critically on the time period considered. For cohorts born from 1930 to 1950, I find that changes in college costs are key for generating the increase in college completion, as they generate a large endogenous increase in college enrollment. Endogenous changes in the average ability of college students also affects college completion rates, but the impact

is quantitatively much smaller. For cohorts born after 1950, the benchmark model significantly overpredicts college completion rates in the data. I show that this is likely due to a sharp increase in the growth rate of the college earnings premium. While the college earnings premium was roughly flat for cohorts born between 1900 and 1950, the growth rate increased sharply for cohorts born after 1950. I find that modifying the model to allow for imperfect forecasting of the college wage premium improves substantially the predictions for college completion for cohorts born after 1950, while leaving the results for cohorts born before 1950 largely unaffected.

In terms of capturing increased ability sorting over time, I consistently find that changes in economic factors (i.e., earnings premia, college costs, opportunity costs, and asset endowments) have little impact. Instead, the key feature in the model that accounts for this is uncertainty about ability. I show that a decrease in the variance of ability signals can generate an increase in ability sorting similar to that in the data. I attribute this change to the increases in standardized testing which improved students knowledge of their own ability relative to other students in their cohort, as discussed in Hoxby (2009).

This paper is related to a large literature on the joint determination of enrollment changes and ability sorting, but previous work focuses almost exclusively on the post-World War II period. Lochner and Monge-Naranjo (2011) look at the role of student loan policies with limited commitment, and shows that this can generate ability sorting. My focus on an earlier time period excludes the student loan innovations they consider, so I instead investigate other factors that may be relevant in understanding ability sorting. Garriga and Keightley (2007) consider the impact of different education subsidies for enrollment and time-to-degree decisions, in a model with borrowing constraints and risky education investment. Hendricks and Leukhina (2011) consider the role of borrowing constraints and learning in understanding the evolution of educational earnings premia. Like this paper, Altonji (1993) and Manski (1989) assume that high school students do not perfectly know their own ability, and they use this feature to investigate the role of preferences, ability, and

earnings premia for enrollment and dropout. Cunha, Heckman, and Navarro (2005) extend the model developed in Willis and Rosen (1979) to include uncertain ability, and find that roughly sixty percent of the variability in returns to schooling is forecastable.

Hendricks and Schoellman (2012) study the same time period as I do, but they take data on college completion and student ability as given in order to understand changes in the college earnings premium in a complete markets model. By contrast, I seek to understand the economic factors that affected college completion and average student ability for cohorts since 1900. Perhaps most related to this paper is Castro and Coen-Pirani (2012), who ask whether educational attainment over time can be explained by earnings premia in a complete markets model. They find that it cannot. My model, with limited borrowing and uncertainty about ability, matches college attainment well for early cohorts, but shares the problem that the model overpredicts attainment after 1950 due to the increase in the earnings premia for these cohorts. In both, disregarding individuals' ability to perfectly forecast future earnings premia helps the model fit, but not entirely.

My work also relates to a number of empirical papers on the impact of different economic forces on historical post-secondary completion, including college costs and income by Campbell and Siegel (1967), work on student ability by Taubman and Wales (1972), academic quality by Kohn, Manski, and Mundel (1976) and borrowing constraints by Hansen and Weisbrod (1969).

2.2 Model

In this section, I develop an overlapping generations model to investigate the causes of increased college completion and increased ability sorting.³ The relevant features include borrowing limits, uncertain ability, and risky completion of college education.

³The counterpart to ability in the data is IQ.

Demographics and Preferences

Time in the model is discrete, and a model period is one year. Each period, N_{mt} males and N_{ft} females are born, each of whom lives for a total of T periods. Let $a = 1, 2, \dots, T$ denote age. Each individual maximizes expected lifetime consumption

$$\mathbb{E}_0 \sum_{a=1}^T \beta^{a-1} \left(\frac{c_a^{1-\sigma} - 1}{1-\sigma} \right)$$

Endowments and Signals

Individuals are *ex-ante* heterogeneous along three dimensions: their sex, m or f , initial asset endowment k_0 , and ability to complete college, denoted α . The probability that any individual completes his or her current year of college is given by $\pi(\alpha)$, where $\pi' > 0$. Log initial assets, $\log(k_0)$, and ability α are drawn from a joint normal distribution with correlation ρ_t , means $\mu_{\alpha,t}$ and $\mu_{k,t}$, and standard deviations $\sigma_{\alpha,t}$ and $\sigma_{k,t}$. Note that the parameters on the joint distribution for $\{\alpha, k_0\}$ are potentially time-varying.

While sex and asset endowments are perfectly observable, ability α is not. Instead, each individual receives a signal $\theta = \alpha + \varepsilon$ at the beginning of life. The error term is $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. Note that because assets and ability are jointly distributed, individuals actually receive two pieces of information about ability – the signal θ and asset endowment k_0 . Let $v = (k_0, \theta)$ be the information an individual has about his true ability. After the initial college enrollment decision, ability α becomes publicly observable.

Education Decisions

The population I am considering consists of high school graduates, so that birth in this model translates to a high school graduation in the real world. At birth, every individual decides whether or not to enroll in college, given sex, asset endowment k_0 , and signal θ . This is the only time this decision can be made. Once enrolled in college, individuals can only exit college by graduating or failing out with annual probability $\pi(\alpha)$. After failure, individuals enter the labor force and may not re-enroll, consistent with the finality

of dropout decisions discussed in Card and Lemieux (2001). Graduating college requires C years of full-time education at a cost of λ_t per year. If an individual decides to not enter college, he or she immediately enters the labor market and begins to work.

Labor Market

I adopt the common assumption that individuals of different ages, a , sex s , and education, e , are different inputs into a constant returns to scale production function that requires only labor. Therefore, wages depend on age, sex, education level, and the year. I write wages as $w_{a,t}(e, s)$ for $s \in \{f, m\}$ and $e \in \{0, 1, \dots, C\}$. While ability α has no direct effect on realized wages, it does affect expected wages because higher ability students are more likely to graduate college and earn higher wages.

Savings Market

Each individual can borrow and save at an exogenous interest rate r_t . I assume individuals must die with zeros assets, so $k_{T+1} = 0$. Borrowing is constrained to be a fraction $\gamma \in [0, 1]$ of expected discounted future earnings. Therefore, individuals must keep assets k_t each period above some threshold \bar{k} , where

$$\bar{k} = -\gamma \cdot \mathbb{E} \sum_{n=a}^{n=T} \frac{w_{n,t}}{1+r_t}$$

Note that both the expectations operator and wage can depend on a number of factors, including ability α , age a , year t , education e , and sex s . Therefore, the borrowing constraint will be written as the function $\bar{k}(\alpha, a, t, e, s)$. In a slight abuse of notation, I will write $\bar{k}(a, t, e, s)$ when the borrowing constraint does not depend on ability α , as is the case once an individual finishes college.

Timing and Recursive Problem

At the beginning of year t , N_{mt} men and N_{ft} women are born at age $a = 1$. Again, each individual is initially endowed with assets k_0 , sex s , ability α , and a signal θ of true ability. Immediately, each individual decides whether or not to enroll in college. If he or she enrolls

in college, true ability is immediately realized, and the individual proceeds through college. In the case of failure (due to $\pi(\alpha)$) or graduation, he or she proceeds to the labor market and works for the remainder of his or her life. Individuals who do not enroll in college proceed directly to the labor market, where they receive the wage associated with age a , education $e = 0$, and sex s .

Recursive Problem for Worker

For individuals currently not enrolled in college, their ability is irrelevant for their decision problem. Therefore, the value of entering year t at age a with assets k , years of college education e , and sex $s \in \{f, m\}$ is:

$$\begin{aligned} V_{a,t}^w(k, e, s) &= u(c) + \beta V_{a+1,t+1}^w(k', e, s) \\ s.t. \quad c + k' &= (1+r)k + w_{a,t}(e, s) \\ k' &\geq \bar{k}(a, t, e, s) \\ k_{T+1} &= 0 \end{aligned}$$

Recursive Problem for College Student

If instead an individual is currently enrolled in college, he has already completed e years of his education and must pay λ_t in college costs for the current year. The probability that he passes and remains enrolled the next year, however, depends on his ability α . Recall that α is known with certainty as soon as the education decision is made, so there is no uncertainty about ability.

The value of being enrolled in college at year t at age a , with assets k , ability α , e years of education completed, and sex $s \in \{f, m\}$ is:

$$\begin{aligned} V_{a,t}^c(k, \alpha, e, s) &= u(c) + \beta \left[\pi(\alpha) V_{a+1,t+1}^c(k', \alpha, e+1, s) + (1 - \pi(\alpha)) V_{a+1,t+1}^w(k', \alpha, e, s) \right] \\ s.t. \quad c + k' - \lambda_t &= (1+r)k \\ k' &\geq \bar{k}(\alpha, a, t, e, s) \\ \pi(\alpha) &= 0 \quad \text{if } a = C \quad \forall \alpha \end{aligned}$$

The last restriction simply states that if $a = C$, that individual is graduating college and cannot acquire any more years of college education.

The College Enrollment Decision

Given the value of being enrolled in college and working, it is possible then to define the educational decision rule at the beginning of life. Recall that at this point, α is unknown, but each individual receives a signal $v = (k_0, \theta)$. Each individual then constructs beliefs over possible ability levels by using Bayes' Rule.

Let $F(\alpha; k_0, \theta)$ be the cumulative distribution function of beliefs (as defined by Bayes' Rule) over ability levels. Given all this, an individual born in year t of sex s with assets k_0 and signal θ enters college if and only if the expected value of entering college is higher than the (certain) value of entering the workforce. This is given by the inequality

$$\int_{\alpha} V_{1,t}^c(k_0, \alpha, 1, s) F(d\alpha; k_0, \theta) \geq V_{1,t}^w(k_0, 0, s) \quad (2.2.1)$$

2.3 Calibration

The goal of this paper is to assess the role played by a number of features of the economy in understanding ability sorting and college enrollment over time. I therefore take a multi-faceted approach to parameterizing the model. First, I construct historical data series for N_{mt} , N_{ft} , and λ_t , which are incorporated directly into the model. Second, I estimate life-cycle wage profiles $w_{a,t}(e, s)$, which are taken as given by model individuals solving their dynamic problem. Third, I exogenously choose values for T , C , r_t , β , ρ_t , $\mu_{\alpha,t}$, $\mu_{k,t}$, $\sigma_{\alpha,t}$, $\sigma_{k,t}$, and $\pi(\alpha)$. Finally, I calibrate $\sigma_{\varepsilon,t}$, and γ in order to match important features of the time series data. Each of these are discussed in more detail below.

Historical Time Series Data

As previously mentioned, N_{mt} males and N_{ft} females are “born” into the model each year, meaning they graduate high school and enter the model eligible to make college enrollment decisions. I take high school completion, and thus the population of potential college en-

rollees, as exogenous. The series for N_{mt} and N_{ft} are taken directly from the U.S. Statistical Abstract Historical Statistics, and I use linear interpolation to supply missing values.

Annual college costs per student, λ_t , are calculated as the average tuition and fee expenses paid out-of-pocket by students each year.⁴ Note that because I measure average *out-of-pocket* costs in the data, λ_t accounts for changes over time in the average amount of financial aid received by students in the form of public and private scholarships and grants. Full details of the data construction are relegated to Appendix E. Briefly, however, I compute λ_t each period as the total revenues from student tuition and fees received by all institutions of higher education divided by the total number of students enrolled in those institutions. The complete time series is constructed by splicing together data from historical print sources including the *Biennial Surveys of Education* (1900 to 1958) and the *Digests of Education Statistics* (since 1962).

Life-Cycle Wage Profiles

Life-cycle wage profiles $w_{a,t}(e,s)$ are estimated using decennial U.S. Census data from 1940 through 2000, along with American Community Survey (ACS) data from 2006-2010. Each ACS data set is a 1% sample of the U.S. population, so that when combined they constitute a 5% of the U.S. population, similar to a decennial census. The data are collected from the Integrated Public-Use Microdata Series (IPUMS), Ruggles et al. (2010), and include wage and salary income, educational attainment, age, and sex. From age and education data I compute potential labor market experience, x , as age minus years of education minus six. I assume that wages can be drawn from one of three education categories - high school, some college, or college. These correspond to $e = 0$, $e \in [1, C - 1]$ and $e = C$ in the model. For each education category, I estimate wage profiles for the non-institutionalized population between ages 17 and 65 who report being in the labor force using the following

⁴Additional student expenses, such as room and board, could also be included, and in fact I do consider these costs as a robustness exercise in Section 2.5. I choose to leave these out of the benchmark specification because such costs are usually more accurately classified as consumption rather than education expenses, and must be paid regardless of college enrollment status.

regression:

$$\log(w_{i,t}) = \delta_{i,t}^b + \sum_{j=1}^4 \beta_j^s x_{i,t}^j \quad (2.3.1)$$

where i denotes individuals, b is birth-year cohort, s is sex, and x is potential labor market experience. In words, I regress log wages on a full set of birth year dummies plus sex specific quartics in experience.

Exogenous Parameters

Parameters set exogenously prior to solving the model are: T , C , r_t , β , ρ_t , $\mu_{\alpha,t}$, $\mu_{k,t}$, $\sigma_{\alpha,t}$, $\sigma_{k,t}$, and $\pi(\alpha)$. I set the length of working life at $T = 48$, implying that individuals born into the model at age 18 would retire at age 65. The number of periods required to complete college is $C = 4$, so that all individuals in the model have post-secondary education $e \in \{0, 1, 2, 3, 4\}$.⁵ The real interest rate is set to $r_t = 0.04$ in all periods, and the discount rate is $\beta = 0.96$, a standard value in models with annual periods.

I now turn to the parameters for the joint normal distribution over $\{\alpha, k\}$. Recall from Section 2.2 that α only affects an individual's probability of passing college. Furthermore, my interest in "ability" is limited to understanding changes over time in the average ability of college versus non-college students *within cohorts*. In other words, I only care here about the relative ability of students within the same birth year, as in the data from Figure 2.1b, not across birth years. As this is the objective, I do not have to worry about trends in average student ability (such as the so-called "Flynn effect") and can normalize the ability distribution for each birth cohort. For this reason, I set $\mu_{\alpha,t} = 0$ and $\sigma_{\alpha,t} = 1$, for all t , so the distribution for α is a standard normal, conditional on k_0 .

Unlike with ability, I am certainly concerned about changes over time in the mean and variance of the initial asset distribution. I interpret k_0 as a reduced-form way of capturing all of the personal financial resources available to a new high school graduate, including but not limited to parental gifts and bequests, and the individual's own income and savings.

⁵I am not presently concerned with educational attainment beyond the bachelor's degree level, so I do not model post-graduate education in this paper.

Additionally, since the model does not allow for individuals to work while in college, I interpret initial assets to also include the present value of income earned while enrolled. With this in mind, I require that the mean and standard deviation of initial assets in the model to track the mean and standard deviation of income in U.S. data. To this end, I start with $\mu_{k,t}$ equal to the annual mean real income per person, as in Piketty and Saez (2006) so that the average real asset endowment in the model equals the actual real mean income in the U.S. each year. Then, in order to account for the fact that $\mu_{k,t}$ includes the individuals' own earnings while in college, I adjust it upward for men and downward for women so that the difference between mean asset endowments for men and women matches the gender earnings gap in the estimated wage profiles during college years.

Piketty and Saez (2006) also provide historical data on the share of income received by the top ten percent of individuals, as well as the cut-off income level for the 90th percentile. Assuming that the U.S. income distribution is log-normal as predicted by Gibrat's law, I can use these data to back out the implied standard deviation of the U.S. income distribution each year. The procedure is as follows. Let real income in year t , denoted Y_t , be a random variable with realization y_t such that $Y_t \sim \ln \mathcal{N}(\mu_t, \sigma_t^2)$ and the associated cumulative distribution function is $F_Y(y_t; \mu_t, \sigma_t^2)$. Observed data are the real mean income in the U.S. in year t , denoted \bar{y}_t , and the 90th percentile of real income in year t , denoted $y_{90,t}$. A standard property of the log-normal distribution is that $\mathbb{E}[Y_t] = \exp(\mu_t + \frac{\sigma_t^2}{2})$. Since $\mathbb{E}[Y_t] = \bar{y}_t$ is observed, I can guess a value $\tilde{\sigma}_t^2$ and solve for the associated mean of the distribution:

$$\tilde{\mu} = \ln(\bar{y}_t) - \frac{\tilde{\sigma}_t^2}{2}$$

Next, I compute $1 - F_Y(y_{90,t}; \tilde{\mu}, \tilde{\sigma}_t^2)$, which would be the fraction of total income received by those with income above the threshold value $y_{90,t}$ if the mean and variance of the income distribution were actually $\tilde{\mu}$ and $\tilde{\sigma}_t^2$. This process continues iteratively until I find a value σ_t^2 , and associated μ_t such that the fraction of income received by the top ten percent equals that observed in the data. I then set $\sigma_{k,t} = \sigma_t$.

The last parameter related to the stochastic endowment process that I need to determine is ρ_t , the correlation between ability and initial asset endowments. Lacking the rich historical data that would be required to properly identify this parameter, I will assume for the benchmark parameterization that $\rho_t = 0$ for all t , so that ability and assets are independent random variables. Intuitively, though, one would expect some positive correlation between a student's financial resources and his or her probability of completing college. It is well known, for example, that parental income is positively related with student test scores and performance. This is discussed in Black, Devereux, and Salvanes (2005) and Cameron and Heckman (1998). Moreover, this correlation also implies a more precise signal of ability. Thus, I later examine in Section 2.5 how the results may change as I allow ρ to increase.

Finally, I need to set the annual probability of passing college, $\pi(\alpha)$. Note that $\pi(\alpha)$ is a reduced form way to capture college non-completion for any reason, including failure and voluntary drop-out. I employ the simple assumption that an individual's *cumulative* probability of completing college equals her percentile rank in the ability distribution. For example, an individual whose ability is higher than 75% of the peers in her birth-year cohort will complete college with probability 0.75, conditional on enrollment. With the length of college set to $C = 4$, there are 3 independent opportunities for failure - after the first, second, and third years of school. Thus, the *annual* probability $\pi(\alpha)$ is simply the cumulative probability raised to the power one-third.

Calibrated Parameters

Finally, I choose the borrowing constraint, γ , and the variance of the noise on the ability signal, $\sigma_{\epsilon,t}$, to replicate the two main data series of interest – college completion and the average ability of college relative to non-college individuals. The borrowing constraint is set to $\gamma = 0.025$ in order to match the time series of college completion. Intuitively, this means that in any given period an individual can borrow up to 2.5% of his expected lifetime

income. Post-schooling, this amount is known with certainty because the wage profiles are given, but during college the expected lifetime income is conditional on the probability of passing college.

Unfortunately, I do not have direct evidence on the precision with which individuals in a given cohort know their own ability relative to their peers. At a qualitative level, it is likely that this precision has increased – i.e., $\sigma_{\epsilon,t}$ has likely decreased – over time. In the early part of the 20th century, no standardized exams existed to compare students within cohorts across schools. Those college admissions exams that did exist were generally school-specific, so there was little scope for comparison of students across schools. During World War I, the U.S. military began testing recruits using the Army Alpha and Army Beta aptitude tests. By World War II, these tests were replaced by the Army General Classification Test (AGCT), a precursor to the Armed Forces Qualification Test (AFQT). On the civilian side, the introduction of the Scholastic Aptitude Test (SAT) in 1926 started a trend toward more widespread use of standardized exams as a college admissions criteria. As standardized testing became more common, students obtained more and more precise signals of their own ability relative to peers. In the modern era, virtually every student contemplating college takes either (or both) of the SAT or the ACT (American College Testing) exams. Even those who do not take these college admissions exams still have quite precise information about their relative ability because other standardized exams are mandated at public schools.

With this historical background in mind, I make the following assumptions on the time series structure of $\sigma_{\epsilon,t}$. For cohorts making college decisions prior to World War II, i.e., those born 1900 through 1923 and graduating high school from 1918 through 1941, I assume that $\sigma_{\epsilon,t}$ decreases linearly from $\sigma_{\epsilon,1900} = 2$ to $\sigma_{\epsilon,1923} = 0.2$. For cohorts born after 1923, $\sigma_{\epsilon,t}$ remains constant at 0.2. This is an admittedly ad hoc construction, but in a simple way it captures the trend of each subsequent cohort getting slightly better information than the previous cohort as aptitude and ability tests became more common in the time between

the world wars. By the completion of World War II, such tests were in widespread use and students likely had quite precise signals about their own ability relative to peers.

2.4 Results


The main computational exercise consists of first simulating the model for U.S. birth cohorts from 1900 through 1972 (i.e., students who graduated high school from 1918 through 1990), verifying that the model replicates important features of the historical data, and then running counterfactual simulations to quantify the impact of changes in direct college costs, education earnings premia, and opportunity costs of college (foregone wages) on college completion and average student ability. Having discussed the benchmark model parameterization, I now examine how well the simulated model matches U.S. data.

Benchmark Model Fit

Figure 2.2 depicts the model predictions along with historical U.S. data for college completion and average student ability. The measure of college completion that I choose to match is the share of 23-year-olds with a college degree. While educational attainment is often measured later in life to capture those who complete college at older ages, I prefer this series for a couple of reasons. First, to my knowledge it is the only measure of college completion with consistent time series data for birth cohorts back to 1900. Second, the model is not constructed to evaluate college enrollment decisions of older students who: (i) are generally less financially-dependent upon parents when paying for education; (ii) face different opportunity costs of school after having been in the workforce for some time; and (iii) may anticipate different return on investment in education due to later-life completion.

Panel (a) of Figure 2.2 shows that, overall, the model replicates well the trends in U.S. college completion over much of the 20th century, with one notable exception. The model does not capture the initial decline and subsequent increase in college completion for cohorts born in the 1950s and 1960s. This deviation is due primarily to the modeling assumption that individuals know their lifetime wage profile with certainty, implying that

Figure 2.2: Benchmark Model Results



they can perfectly forecast changes in the education earnings premium. Later I consider alternative assumptions, and find that the model can generate more accurate predictions over this time period.

Panel (b) of Figure 2.2 plots the average ability percentile of students who attempt college (even if they do not complete), and those who have only high school education. While I only have a few reliable data points to match, those I do have show a clear pattern of increased sorting by ability over time. For cohorts born at the beginning of the 20th century, college and non-college students had similar ability on average, but the ability gap widened throughout the century. This general pattern is also predicted by the model.

In order to facilitate quantitative comparison with alternative specifications, I also provide measures of model fit over various time periods in Table 2.1. The measure of fit I report is the sum of squared deviations between model and data. The columns labeled "Fraction of 23-year-olds with college degree" refer to the series in Panel (a) of Figure 2.2. For this series, I compute the fit over all cohorts 1900-1972, and three subsamples: 1900-1925, 1926-1950, and 1951-1972. As seen in the "Benchmark" model specification

Table 2.1: Measures of Fit for Various Model Specifications

Model \ Cohorts	Fraction of 23-year-olds with College Degree					Average Ability Difference
	1900-1972	1900-1925	1926-1950	1951-1972	1900-1972	
Benchmark	0.158	0.015	0.009	0.133	0.034	
Imperfect foresight	0.055	0.020	0.013	0.022	0.028	
Constant costs rel. to income	0.134	0.024	0.011	0.099	0.034	
Corr(α, k_0) = 0.30	0.183	0.023	0.017	0.143	0.038	
Include room and board	0.159	0.019	0.017	0.123	0.035	

in Panel (a) of Figure 2.2, the model matches the data very closely for cohorts born pre-1950, but does less well for cohorts born after 1950. The column labeled “average ability difference” measures how well the model matches the difference between the average ability percentile of college and non-college individuals. I only report the full sample for this statistic because there are so few data points to match within the sub-sample periods.

Discussion of Benchmark Results

The measure of college completion – the fraction of twenty-three year olds with a college degree – can be decomposed as

$$\frac{P^{grad}}{P^{23}} = \left(\frac{P^{HS}}{P^{23}} \right) \left(\frac{P^{enroll}}{P^{HS}} \right) \left(\frac{P^{grad}}{P^{enroll}} \right) \quad (2.4.1)$$

where P^{HS} , P^{enroll} , and P^{grad} are the number of people that complete high school, enroll in college, and graduate college. The model’s predictions for college completion can be decomposed into the three terms on the right hand side of equation (2.4.1). While the first is exogenous, the second and third terms are endogenous to the model. In this section, I use this decomposition to understand what drives the change in college completion predicted by the model.

Figure 2.3: College Enrollment Conditional on High School Graduation





Figure 2.4: College Pass Rate



First, Figure 2.3 plots the share of high school graduates that enroll in college, as predicted by the model. In the language of equation (2.4.1), this is P^{enroll} / P^{HS} . Figure 2.3 shows that for cohorts born between 1900 and 1920, college enrollment rates conditional on high school graduation were between 30 and 50 percent, albeit with a lot of noise. This rate increased for cohorts born in the 1920s and generally remained between 50 and 60 percent for cohorts through 1950, after which the rate again increased substantially.

The third term in equation (2.4.1) is the share of college enrollees that graduate by age twenty-three. This is given by the ratio P^{grad} / P^{enroll} and is plotted in Figure 2.4. While Figure 2.4 shows that the college pass rate has a fair amount of year-to-year noise, the hump-shaped trend is still evident. From the 1900 through 1930 birth cohorts, the college pass rate increased from about 51% to nearly 61%. After the 1930 cohort, however, this trend reverses, and the pass rate steadily declines back down to around 53%. This result is consistent with evidence from Bound, Lovenheim, and Turner (2010), who compare the high school class of 1972 (roughly birth cohort 1954) to that of 1992 (birth cohort 1974) and find a significant decrease in college completion conditional on enrollment. In my model, this pattern is due entirely to the ability composition of college students. Recall from Panel (b) of Figure 2.2 that the average ability of college enrollees was generally increasing

through the 1930 cohort, then decreasing in the following cohorts. Unfortunately, I have found no reliable historical data to compare with the model’s predicted pass rates. However, the National Center for Education Statistics (NCES) does provide more recent data I can use for a rough comparison. For the cohort beginning college in 1996 (assuming they are around 18 years old on average, this would be approximately the 1976 birth cohort), the share completing college within five years was 50.2%.⁶ The last birth cohort in the model is 1972, so the comparison is not perfect, but the model pass rate of 53.1% for that cohort is quite close.

I now isolate the effects of the college enrollment and college pass rates through two counterfactual experiments. I ask two questions. First, how does college completion change relative to the benchmark if there were no endogenous increase in the college enrollment rate, as in Figure 2.3? Second, how does college completion change if there were no endogenous changes in the college pass rate, as in Figure 2.4? Results from these two experiments are plotted in Figure 2.5, along with the benchmark prediction for college completion.

Figure 2.5: College Completion if Enrollment Rates and Pass Rates were Constant

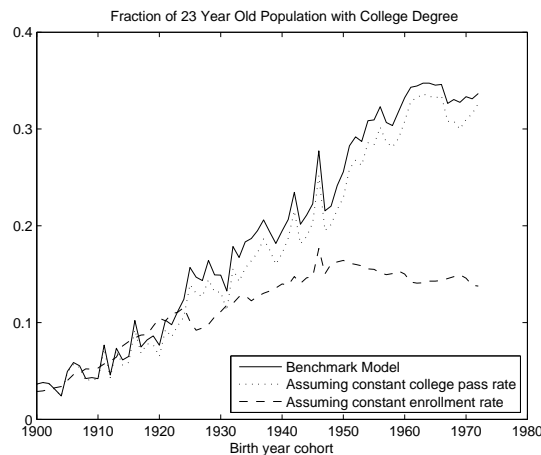


Figure 2.5 shows that if the college enrollment rate had remained constant instead

⁶See Table 341 in the 2010 *Digest of Education Statistics*.

of rising after the 1920 cohort, the model would have under-predicted college completion rates by more than half by the end of the time series.⁷ Similarly, if the college pass rate had instead remained constant at the 1900 value of 51.5%, then college completion would have been several percentage points lower than in the benchmark model. It is clear, however, that the quantitative effects of changes in college enrollment are much larger than those due to changing college completion rates.

Counterfactual Experiments

What if individuals do not have perfect foresight of education earnings premia?

Figure 2.6 shows that for cohorts born in the U.S. prior to 1950, the education premia implied by the estimated life-cycle wage profiles exhibit some year to year variation, but essentially no trend. Beginning around the 1950 cohort, however, the college earnings premia began increasing steadily. I now examine how the model predictions for college completion and average student ability would differ if, instead of predicting changes in the education premium exactly, model individuals expected an historical average education earnings premia to prevail in the future as well.

For this exercise, I assume that the high school wage for each cohort is observable, but the earnings premia for individuals who complete college or some college are not observable. Rather, individuals observe a moving average of the earnings premia earned by previous cohorts and assume their own cohort's earnings premia will be the same. Thus, as the true college earnings premium begins rising, newly born cohorts will predict the increase imperfectly and with several years lag.

Figure 2.7 shows the model predictions under this counterfactual experiment, assuming a 25-year moving average. Relative to the benchmark model results, notice that the model now comes much closer to the actual college completion rate in the data for cohorts born after 1950. The model still does not capture all of the decline for the cohorts in the

⁷In Figure 2.5, I assume that the college enrollment rate conditional on high school graduation is constant at 36.9%, which is the average enrollment rate for cohorts 1900 through 1920.

Figure 2.6: Education Premia Implied by Estimated Life-Cycle Wage Profiles

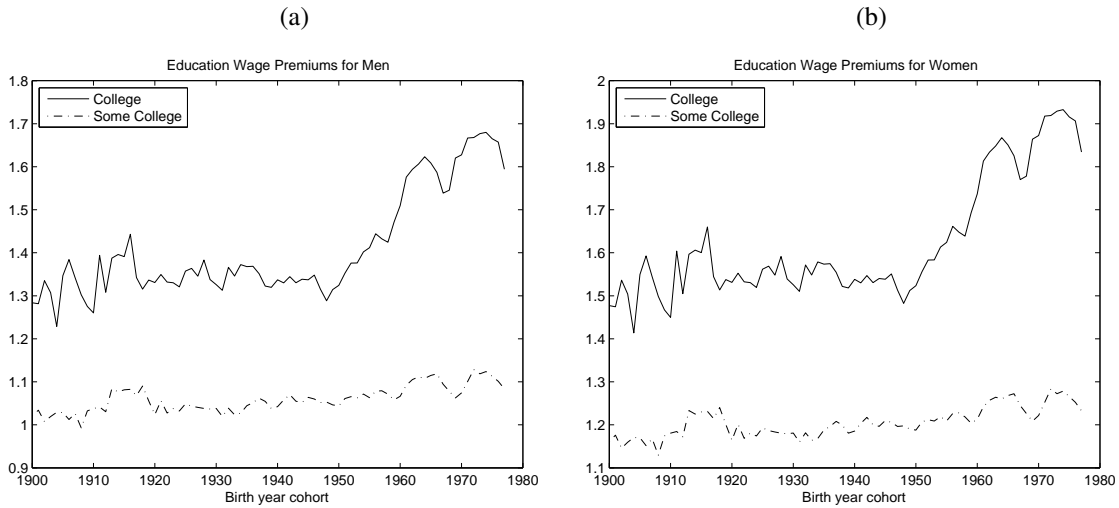



Figure 2.7: Results with Imperfect Foresight of Education Earnings Premia



1950s, but as Table 2.1 clearly shows, this specification fits the data much better than the benchmark assumption that individuals perfectly forecast changes in the education premia. Over the entire time period, the sum of squared deviations declines by almost two-thirds from the benchmark value of 0.158 to 0.055. All of this gain is due to the 1951-1972 cohorts, where the sum of squared deviations changes from 0.133 to 0.022, a decrease


of more than 83%. Additionally, the model's ability to match changes in average ability of college and non-college students also improves under this specification. According to the last column of Table 2.1, the sum of squared deviations declines from 0.034 to 0.028. These improvements strongly suggest that perfect foresight of education earnings premia is a problematic assumption. Accurately modeling students' expectations about the returns to education is crucial for understanding college enrollment decisions, particularly during periods of time when education premia are changing rapidly.

What if real college costs increased proportional to real disposable incomes?

I now ask how college completion rates and average student ability would have differed over the time period in question if real college costs were constant with respect to real average income. Figure 2.8 depicts the actual time series data for real college costs that I use in the benchmark model (solid line), along with a hypothetical series for college costs which are a constant fraction of annual real average income (dashed line). From 1920 to around 1940, the actual series exceeds the hypothetical series due to the fact that per student tuition and fees spiked relative to income during the Great Depression. Then from the early 1940s until about 1990, the hypothetical series is above the actual series. Holding all else constant, I would expect that individuals in the counterfactual model facing the hypothetical college costs should attend college in greater numbers for the cohorts born from about 1900 to 1920 (those in school from around 1920 to 1940), and fewer of those born after 1920 would attend college.

Figure 2.9 largely confirms these predictions. Relative to the data, the model predicts too many people attending college for those cohorts born between about 1910 and 1925. For the cohorts from 1925 through 1950, the model does predict slightly fewer college graduates, but still matches the data quite closely. And finally, for the cohorts born after 1950, the model still predicts more college graduates than in the data. However, as can be seen in Table 2.1, the model fit improves over this period since the sum of squared

Figure 2.8: College Costs



deviations fall from 0.133 to 0.099, a decrease of more than 25%. Turning to Panel (b) of Figure 2.9, there are hardly any discernible differences in average ability of college and non-college students relative to the benchmark model. This can also be confirmed by noting that sum of squared deviations for the average ability difference in Table 2.1 is unchanged from the benchmark value of 0.034. I conclude that the fluctuations in real college costs relative to real income are not a major factor in accounting for the increased ability sorting over time.


2.5 Robustness

Having discussed the benchmark model results and counterfactual experiments, I now make a few remarks about the robustness of some modeling assumptions. In particular, I made the strong assumption that ability and initial assets were uncorrelated. I also assumed that room and board were excluded from college costs. I now relax these assumptions and see how they affect the results.

Correlation of Ability and Initial Assets


In the benchmark specification, I assumed that the random endowments for ability and assets were uncorrelated. However, there is evidence to suggest that these may be positively

Figure 2.9: Results with Alternative College Costs



correlated, and I want to understand how this affects the results. I maintain the assumption that α and $\log(k_0)$ share a bivariate normal distribution, only now I set $\rho = 0.3$. All other parameters are maintained as in the benchmark specification. Figure 2.10 shows the model predictions for college completion and ability sorting between college and non-college individuals.

Figure 2.10: Results with Positive Correlation between Ability and Initial Assets



Relative to the benchmark model results, two things are notable. The positive correlation between ability and assets increases college completion minimally throughout the time period, and it increases the difference in ability between college and non-college students during earliest birth cohorts. Both of these effects reduce the model fit slightly, as seen in Table 2.1. The increase in completion is simply due to the fact that higher ability students are now more likely to have greater financial resources as well, thus making them more likely to attend college. The effect on average ability is also quite intuitive. Recall that individuals receive information $v = (k_0, \theta)$, where $\theta = \alpha + \varepsilon$ is the noisy signal of true ability α . As ρ increases k_0 becomes more informative about α , so individuals with high initial assets will infer that they have higher ability, and thus be more likely to enroll in college. This increases the average ability of individuals who attempt college, while simultaneously decreasing the average ability of non-college individuals. The effect is largest for earlier birth cohorts because later birth cohorts received more accurate signals about their true ability.

College Costs Including Room and Board

College costs in the benchmark model were restricted only to tuition and fees. Now, I take a broader view of college costs and examine whether or not the results are sensitive to the inclusion of room and board expenses. Like the earlier time series data on college tuition and fees, I construct this data from printed historical government documents. The details are found in Appendix E. For this experiment, all calibrated values are maintained just as in the benchmark economy, with the exception of the borrowing constraint, γ . I need to adjust γ because students now face additional college expenses, so college completion rates would be too low if I held γ constant at the benchmark value. The new borrowing constraint which allows us to match the time series of college completion is $\gamma = 0.04$.

Figure 2.11 shows the model predictions for college completion and average student ability when room and board costs are included. Relative to the benchmark results in

Figure 2.11: Results for College Costs including Tuition, Fees, Room, and Board

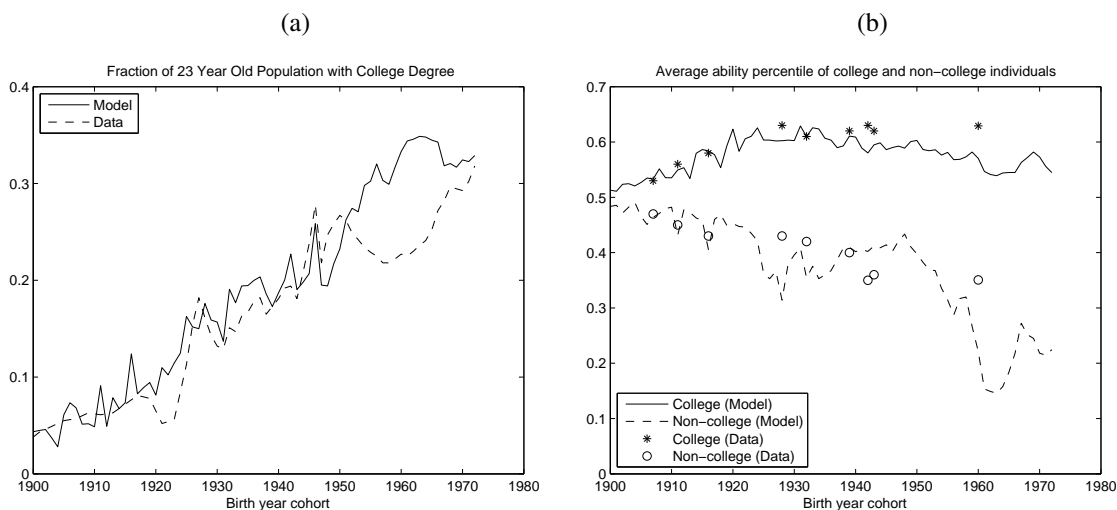


Figure 2.2, very little has changed. The model still predicts college completion rates in line with the data up until the 1950s and 1960s cohorts, when model and data diverge. Additionally, average ability of college and non-college students diverges over time just as in the benchmark model. Referring to Table 2.1, it is clear that while the model fits college completion slightly worse than the benchmark model pre-1950, it does slightly better post-1950. On the whole, this model fits almost exactly as well as the benchmark model for both college completion and average ability difference.

2.6 Conclusion

I develop an overlapping generations model with unobservable ability and borrowing constraints to investigate post-secondary completion and ability sorting in the birth cohorts of 1900–1972. To discipline the model, I digitize and utilize historical data series including statistics on college costs and high school graduation rates. I find that the share of high school graduates enrolling in college and the subsequent college pass rate are both key for understanding increased college graduation rates. However, I find no evidence that economic factors – including real college costs, opportunity costs, education wage premia, or asset endowments – have a major impact on increasing ability sorting over time. I do find,

however, that a decrease in the variance of ability signals can properly match this fact, a trend which I attribute to increases over time in standardized testing.

An important deviation between the benchmark model and historical data is that the model does not properly match college completion after the 1950 birth cohort. I show that this could be due to individuals having imperfect foresight about the college earnings premium. If individuals observe a moving average of the earnings premia from previous cohorts and use this to estimate the future earnings premium, then changes in the earnings premium are taken into account only with a lag. I build this into the model and find that it significantly improves the model's fit. I therefore view this as evidence of backward looking wage estimation when making college enrollment decisions.

An interesting use of this framework would be an extension to multiple countries. Evidence by Hanushek and Kimko (2000) suggests that ability is strongly related to growth, but Bils and Klenow (2000) find that the causality from formal schooling to economic growth is somewhat tenuous. If developing countries have very little ability sorting between education levels, as was the case in the early U.S., there may be a weak correlation between education level and labor efficiency. In a cross-country context, this could arise due to tighter borrowing constraints or less precise signals about true ability. I will explore this link in future research.

Chapter 3

NEED-BASED ENTREPRENEURSHIP AND AGGREGATE PRODUCTIVITY ACROSS COUNTRIES

3.1 Introduction

While income per capita differences across countries are large, human and physical capital can account for only a small share of these income differences. Instead, Hall and Jones (1999) and Klenow and Rodriguez-Clare (1997) find that income differences are primarily accounted for by TFP.

One possible source of these TFP differences is that a large share of business owners in developing countries operate firms not due to sufficient entrepreneurial ability or wealth, but because they have no other employment opportunities. In fact, while over ninety percent of business owners in the United States claim they are taking advantage of a business opportunity, only sixty percent of business owners in places like Brazil, India, Peru, and China claim the same. The other forty percent of business owners in these countries operate businesses explicitly because they have no other employment opportunities.¹ I refer to these individuals as *need-based* entrepreneurs. Put somewhat differently, need-based entrepreneurs would be willing to accept a job offer but simply to not have access to one, which suggests a significant amount of occupational misallocation between entrepreneurship and employment in developing countries. Motivated by this evidence, the goal of this paper is to provide a theory to help rationalize cross-country differences in need-based entrepreneurship, and in turn, quantitatively assess its role for TFP differences across countries.

To do so, I develop a general equilibrium model that embeds a costly search framework into a Bewley (1986)-Aiyagari (1994) incomplete markets model, further extended to include entrepreneurship and collateral constraints. The basic idea put forth here is that

¹These statistics are derived from the Global Entrepreneurship Monitor Surveys, and are detailed further in Section 3.2.

unemployed individuals in developing countries turn to entrepreneurship as a replacement for the lack of government-provided unemployment benefits. Without unemployment insurance, workers who lose their jobs are forced to turn to entrepreneurship to generate some income, rather than simply live for a short time on unemployment benefits. However, entrepreneurship takes time (e.g. set up costs, monitoring employees, making business decisions, etc.) that could otherwise be used to search for new work. In developed countries, these unemployment benefits afford the opportunity to search more intensely for work, allowing them to quickly leave unemployment and re-enter the workforce. So while the unemployed in developed countries quickly return to their preferred occupation, those in developing countries get “stuck” in entrepreneurship, even though their comparative advantage may be as a worker.

The quantitative results suggest that this framework is able to generate need-based entrepreneurship, but has very little quantitative impact. Future research will address this shortcoming.

This paper joins the rapidly expanding literature on misallocation of inputs, including Hsieh and Klenow (2009). Buera, Kaboski, and Shin (2011b), Midrigan and Xu (2010), and Moll (2012) emphasize financial distortions to generate misallocation of physical capital across firms, and find varying degrees of quantitative importance, while Greenwood, Sanchez, and Wang (2010) show that financial intermediation costs across countries can have a large quantitative impact. This paper utilizes similar financial distortions to also generate misallocation of ability across occupations. However, while the lack of search frictions in these models imply efficient occupational choice, the frictions that distort this choice are the focus of this paper. More similar to this paper, Poschke (2010) develops a model in which skill-biased technology generates two entrepreneurship cutoffs, so that a group of low-skilled individuals choose to be entrepreneurs. While this generates a mass of smaller and less productive firms, it shares the feature of the aforementioned papers that occupational choice is undistorted. This paper is also related to recent work in which labor

market frictions are built into growth models with incomplete markets, including Alvarez and Veracierto (2001), Krusell et al. (2008), Krusell, Mukoyama, and Şahin (2010), and Krusell et al. (2011). As in Acemoglu and Shimer (1999), I find that a certain level of unemployment insurance can increase aggregate productivity. Due to a difference in focus, none of these papers include entrepreneurship.

3.2 Empirical Motivation

This paper is motivated by the fact that a significant fraction of the population works as entrepreneurs because they cannot find work. This implies that a number of firms are run by low-skilled entrepreneurs. For evidence of this, I turn to the Global Entrepreneurship Monitor Surveys (GEM).² The GEM are harmonized national surveys of entrepreneurship in over fifty countries, from 2001 to 2010. Specifically helpful for this paper, they include a number of developing countries, the poorest of which is Uganda. Because the questions and results are standardized across countries, the GEM are useful for cross-country comparisons. Furthermore, the surveys are designed to capture all residual claimants, meaning that the survey is not limited to formal sector firms or firms with a sufficiently large workforce.

To capture the idea of need-based entrepreneurship, the survey asks “Why do you operate your business?” The possible responses are

1. To take advantage of a business opportunity
2. Because I had no other options
3. Some combination of the two
4. Do not know/refuse to answer

²Although the dataset is relatively new, this is certainly not the first paper to utilize it. See, for example, Poschke (2010) for the use of this dataset in an aggregate cross-country study.

I take those that answer “Because I had no other options,” and “Some combination of the two” to be my measure of self-employment.³

To construct my sample from the dataset, I take the most recent year with sufficient data for every country available. As a measure of a income level, I use real GDP per capita from the same year. This is from the Penn World Table version 7.0. The resulting dataset gives a cross-sectional view of entrepreneurship across countries. Further details are given in Appendix . Figure 3.1 shows a strong negative correlation between income level and the fraction of entrepreneurs that are need-based.⁴




Figure 3.1: Fraction of Entrepreneurs that are Self-Employed

Figure 3.1 shows that while only thirteen percent of U.S. entrepreneurs are need-based, that number rises to over sixty percent in Uganda. Even in other developing countries, the fraction of need-based entrepreneurs is well above thirty percent.

³Other definitions yield similar results.

⁴This result is certainly not unique to this paper. Poschke (2010), for example, also documents this fact.

3.3 Model

Time is discrete and infinite, running $t = 0, 1, 2, \dots$. There is a measure one of infinitely lived individuals, who maximize total lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [\log(c_t) + u(1 - h_t)]$$

in which $\beta < 1$ is the discount factor, c_t is consumption, and h_t is total effort in period t .

The function for disutility of effort is given by

$$u(1 - h_t) = \gamma \frac{(1 - h_t)^{1-\sigma} - 1}{1 - \sigma}.$$

Each individual is endowed with one unit of time and ability vector $z_t = (z_{et}, z_{wt})$ where z_{et} is entrepreneurial ability and z_{wt} is ability as a worker. This ability evolves jointly according to the autoregressive process

$$\log(z_{t+1}) = \rho \log(z_t) + \Gamma$$

Let $Q(z_{t+1}, z_t)$ be the transition function for z . Individuals accumulate assets a_t , and are constrained to hold non-negative assets.

There are three mutually exclusive occupations: worker w , entrepreneur e , or unemployed u . An entrepreneur operates a technology that employs workers. It is assumed that being a worker or entrepreneur requires a fixed amount of effort, $\bar{h} > 0$.

Capital and Labor Markets

The labor market is somewhat similar to that proposed in Alvarez and Veracierto (2001), although I make a number of simplifying assumptions. There are two islands in the economy: a “working” island, that includes all entrepreneurs and workers, and a “unemployed” island, where the unemployed live. The working island is a competitive labor market, which implies that entrepreneurs pay a wage w per unit of labor services. This wage clears the market.

Only the unemployed can search for work. If an unemployed individual exerts effort h searching, he finds the working island with probability h^η , where $\eta < 1$. Once employed, a worker is on the working island, and even if his employer disbands his firm, he can immediately find a new job on the island. However, all workers are exogenously separated from the working island with probability λ . If this occurs, the individual can either become unemployed or become an entrepreneur, but he cannot work.

Unemployed individuals also receive a payment b per period of being unemployed. This is financed by a government tax τ on entrepreneur wage payments. The government is subject to a balanced budget.

In terms of the capital market, there is a competitive financial intermediary who takes in all assets at the beginning of the period t and lends capital to entrepreneurs at the rental rate r . At the end of period t , the intermediary pays back $(1 + R)a$ to an individual who deposited a assets. The zero profit condition for the intermediary implies $r = R + \delta$, where δ is the depreciation rate of capital.

Contracting between the intermediary and entrepreneurs is subject to limited enforcement. If entrepreneur with assets a rents capital k , he can renege on the contract and steal $(1/\Delta)k$. The punishment is that the intermediary takes the deposited assets a . This leads to a simple collateral constraint $k \leq \Delta a$, as in Buera and Shin (2010) and Moll (2012). The parameter $\Delta \in [1, +\infty)$ indexes the financial development of the country, and spans between complete self-financing ($\Delta = 1$) and perfect contracting markets as Δ approaches infinity.

Timing and Recursive Formulation

With the labor and capital market assumptions, the aggregate state of this economy is the three-dimensional distribution across assets a , ability z , and occupations o , defined as $\mu(z, a, o)$. Because I investigate the stationary competitive equilibrium, I suppress the dependence of the value functions on μ .

A worker enters time t with ability z and assets a . His income is equal to wages w . He chooses to consume c and save $a' \geq 0$. At the start of $t + 1$, he becomes unemployed with probability λ . If he is separated, he must choose between being unemployed or an entrepreneur at $t + 1$. If he is not separated, he can choose to quit his job and become unemployed. Given all this, the value of being a worker at time t can be written as

$$\begin{aligned}
v^w(z, a) &= \max_{a', c} \log(c) + u(1 - \bar{h}) + \beta \int_z \left[\lambda \max\{v^u(z', a'), v^e(z', a')\} \right. \\
&\quad \left. + (1 - \lambda) \max\{v^u(z', a'), v^w(z', a'), v^e(z', a')\} \right] Q(dz', z) \\
s.t. \quad &c + a' = (1 + R)a + w \\
&a' \geq 0
\end{aligned}$$

where v^u is the value of being unemployed.

An entrepreneur enters at time t with ability z and assets a . He chooses inputs capital, k , and labor services, n , to produce output according to the production function

$$y(z_e) = z_e^\psi k^\theta n^{1-\psi-\theta}$$

As discussed above, contracting is subject to limited enforcement, implying the additional constraint $k \leq \Delta a$. Entrepreneurs are also taxed at rate τ to finance unemployment benefits. His profit is then

$$\pi(z, a) = z_e^\psi k^\theta n^{1-\psi-\theta} - rk - (1 + \tau)wn$$

Entrepreneurship requires the same time commitment as working, so his effort is \bar{h} . At the end of period t , the entrepreneur realizes tomorrow's productivity z' . Once he realizes z' , he can choose to remain an entrepreneur or become unemployed. The value

function for a current entrepreneur with individual state (z, a) is therefore

$$\begin{aligned}
v^e(z, a) &= \max_{k, n, c} \log(c) + u(1 - \bar{h}) + \beta \int_{z'} \max\{v^e(z', a'), v^u(z', a')\} Q(dz', z) \\
s.t. \quad c + a' &= (1 + R)a + z_e^\psi k^\theta n^{1-\psi-\theta} - rk - wn \\
k &\leq \Delta a \\
a' &\geq 0
\end{aligned}$$

The first constraint is the budget equation, and the second is a financial friction.

An unemployed individual receives income only through interest on his assets a and unemployment benefits bs . He can also choose search intensity h . Unlike entrepreneurs or workers, the unemployed have no other time obligations, so that total effort is equal to h . Therefore, the value of being unemployed with ability z and assets a is

$$\begin{aligned}
v^u(z, a) &= \max_{h, c, a'} \log(c) + u(1 - h) + \beta \int_{z'} \left[(1 - h^\eta) \max\{v^e(z', a'), v^u(z', a')\} \right. \\
&\quad \left. + h^\eta \max\{v^e(z', a'), v^w(z', a'), v^u(z', a')\} \right] Q(dz', z) \\
s.t. \quad c + a' &= (1 + R)a + b \\
a' &\geq 0 \\
h &\in [0, 1]
\end{aligned}$$

Stationary Equilibrium

A stationary equilibrium in this economy is a distribution $\mu(z, a, o)$ such that $\mu(z, a, o) = \Lambda(\mu(z, a, o))$, value functions v^e , v^w , and v^u , decision rules k , n , h , ϕ , and prices r , R , and w such that

1. Given prices, v^e , v^w , and v^u solve the individual's problem with the associated decision rules
2. Intermediaries make zero profit: $r = R + \delta$

3. μ is consistent with the decision rules and $Q(z, z')$

4. The government budget balances

$$b \int_{o=u} d\mu = \tau_w \int_{o=e} n(z, a, o) d\mu$$

5. Markets clear:

(a) Labor market:

$$\int_{o=e} n(z, a, o) d\mu = \int_{o=w} z_w d\mu$$

(b) Capital market

$$\int a d\mu = \int_{o=e} k(z, a, o) d\mu$$

(c) Consumption market

$$\int c(z, a, o) d\mu + \delta \int_{o=e} k(z, a, o) d\mu = \int_{o=e} z_e^\psi k(z, a, o)^\theta n(z, a, o)^{1-\psi-\theta} d\mu$$

Need-Based Entrepreneurship in the Model

The GEM surveys break entrepreneurs into two groups. There are those who are entrepreneurs to take advantage of a business opportunity, and those who are entrepreneurs because they have no other options. I want to map these groups into “opportunity-based” and “need-based” entrepreneurs. Luckily, the model allows a simple mapping between the empirical statistics and model individuals’ responses. If I were to ask individuals in the model the same GEM question, those who prefer entrepreneurship to working would answer “to take advantage of a business opportunity.” These are individuals in the model that are entrepreneurs with $v^e(z, a) \geq v^w(z, a)$. Those that would answer “because I had no other option” are those that are employed as entrepreneurs, but would prefer to be workers. They therefore have $v^e(z, a) < v^w(z, a)$. Therefore, the total share of the population that are need-based entrepreneurs in the stationary equilibrium is

$$NE = \int_{o=e} \mathbf{1}[v^w(z, a) > v^e(z, a)] d\mu \quad (3.3.1)$$

where $\mathbf{1}[\cdot]$ is the indicator function. The population of opportunity-based entrepreneurs is

$$OE = \int_{o=e} \mathbf{1}[v^w(z, a) \leq v^e(z, a)] d\mu \quad (3.3.2)$$

3.4 Conclusion

This chapter lays out a model that may have the potential to account for differences in need-based entrepreneurship across country. Future work will expand on quantifying the model.

REFERENCES

- Acemoglu, Daron and Robert Shimer. 1999. "Efficient Unemployment Insurance." *Journal of Political Economy* 107 (5):893–928.
- Adamopoulos, Tasso. 2011. "Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences." *International Economic Review* 52 (2):489–521.
- Adamopoulos, Tasso and Diego Restuccia. 2011. "The Size Distribution of Farms and International Productivity Differences." Working Paper.
- Aiyagari, S. Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics* 109 (3):659–684.
- Altonji, Joseph G. 1993. "The Demand for and Return to Education When Education Outcomes are Uncertain." *Journal of Labor Economics* 11 (1):48–83.
- Alvarez, Fernando. and Marcelo Veracierto. 2001. "Severance payments in an economy with frictions." *Journal of Monetary Economics* 47 (3):477–498.
- Banerjee, Abhijit V. and Esther Duflo. 2007. "The economic lives of the poor." *Journal of Economic Perspectives* 21 (1):141.
- Bewley, Truman F. 1986. "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers." In *Contributions to Mathematical Economics in Honor of Gerard Debreu*, edited by Werner Hildenbrand and Andreu Mas-Colell. 79–102.
- Bils, Mark and Peter J. Klenow. 2000. "Does Schooling Cause Growth?" *American Economic Review* 90 (5):1160–1183.
- Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. 2005. "Why the Apple Doesn't Fall Far: Understanding Intergenerational Transmission of Human Capital." *American Economic Review* 95 (1):437–449.
- Bound, John, Michael F. Lovenheim, and Sarah Turner. 2010. "Why have college completion rates declined? An analysis of changing student preparation and collegiate resources." *American Economic Journal: Applied Economics* 2 (3):129.
- Buera, Francisco and Yongseok Shin. 2010. "Financial Frictions and the Persistence of History: A Quantitative Exploration." Mimeo.
- Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin. 2011a. "Finance and Development: A Tale of Two Sectors." *American Economic Review* 101 (5):1964–2002.
- . 2011b. "Finance and Development: A Tale of Two Sectors." *American Economic Review* 101 (5):1964–2002.
- Buera, Francisco J. and Yongseok Shin. 2013. "Financial Frictions and the Persistence of History: A Quantitative Exploration." *Journal of Political Economy*. Forthcoming.

- Cameron, Stephen V. and James J. Heckman. 1998. "Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males." *Journal of Political Economy* 106 (2):262–333.
- Campbell, Robert and Barry N. Siegel. 1967. "The Demand for Higher Education in the United States, 1919–1964." *American Economic Review* 57 (3):482–494.
- Card, David and Thomas Lemieux. 2001. "Dropout and enrollment trends in the postwar period: What went wrong in the 1970s?" In *Risky behavior among youths: An economic analysis*. University of Chicago Press, 439–482.
- Caselli, Francesco. 2005. "Accounting for cross-country income differences." In *Handbook of Economic Growth*, edited by P. Aghion and S.N. Durlauf. 679–741.
- Caselli, Francesco and W.J. Coleman. 2001. "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation." *Journal of Political Economy* 109 (3):584–616.
- Castro, Rui and Daniele Coen-Pirani. 2012. "Do Returns to Skill Explain the Evolution of Educational Attainment in the U.S.?" Mimeo. Université de Montréal.
- Cunha, Flavio, James J. Heckman, and Salvador Navarro. 2005. "Separating uncertainty from heterogeneity in life cycle earnings." *Oxford Economic Papers* 57 (2):191–261.
- Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5):1221–1248.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2008. "How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya." *American Economic Review Papers and Proceedings* 98 (2):482–488.
- . 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6):2350–90.
- Dupas, Pascaline and Jonathan Robinson. 2011. "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya." NBER Working Paper 14693.
- Estevadeordal, Antoni and Alan M. Taylor. 2013. "Is the Washington Consensus Dead? Growth, Openness, and the Great Liberalization, 1970s - 2000s." *Review of Economics and Statistics* Forthcoming.
- Fafchamps, Marcel, Christopher Udry, and Katherine Czukas. 1998. "Drought and saving in West Africa: are livestock a buffer stock?" *Journal of Development Economics* 55 (2):273–305.
- Garriga, Carlos and Mark P. Keightley. 2007. "A General Equilibrium Theory of College with Education Subsidies, In-School Labor Supply, and Borrowing Constraints." Federal Reserve Bank of St. Louis Working Paper.

- Gollin, Douglas. 2002. "Getting Income Shares Right." *Journal of Political Economy* 110 (2):458–474.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2011. "The Agricultural Productivity Gap in Developing Countries." Working Paper.
- Gollin, Douglas, Stephen L. Parente, and Richard Rogerson. 2004. "Farm work, home work and international productivity differences." *Review of Economic Dynamics* 7 (4):827–850.
- Greenwood, Jeremy, Juan M. Sanchez, and Cheng Wang. 2010. "Quantifying the Impact of Financial Development on Economic Development." Mimeo.
- Hall, Robert E. and Charles I. Jones. 1999. "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quarterly Journal of Economics* 114 (1):83–116.
- Hansen, W.L. and B.A. Weisbrod. 1969. "The distribution of costs and direct benefits of public higher education: the case of California." *Journal of Human Resources* :176–191.
- Hanushek, Eric A. and Dennis D. Kimko. 2000. "Schooling, Labor-Force Quality, and the Growth of Nations." *American Economic Review* 90 (5):1184–1208.
- Hendricks, Lutz and Oksana Leukhina. 2011. "The Return to College: Selection Bias and Dropout." Mimeo. University of North Carolina.
- Hendricks, Lutz and Todd Schoellman. 2012. "Student Abilities During the Expansion of U.S. Education, 1950-2000." Mimeo. University of North Carolina.
- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi. 2012. "Two Perspectives on Preferences and Structural Transformation." *American Economic Review*. Forthcoming.
- Herrendorf, Berthold and Todd Schoellman. 2012. "Why is Measured Productivity so Low in Agriculture?" Working Paper.
- Heston, Alan, Robert Summers, and Bettina Aten. 2011. *Penn World Table Version 7.0*. Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania. URL <http://pwt.econ.upenn.edu/>.
- Hodges, R.J., M. Bernard, H. Knipschild, and F. Rembold. 2010. "African Postharvest Losses Information System - a network for the estimation of cereal weight losses." In *Proceedings of the 10th International Working Conference On Stored Products Protection*, edited by M.O. Carvalho. 958–964.
- Hoxby, Caroline M. 2009. "The Changing Selectivity of American Colleges." *Journal of Economic Perspectives* 23 (4):95–118.
- Hsieh, Chang-Tai and Peter J. Klenow. 2007. "Relative prices and relative prosperity." *American Economic Review* 97 (3):562–585.

- . 2009. “Misallocation and manufacturing TFP in China and India.” *The Quarterly Journal of Economics* 124 (4):1403–1448.
- Jones, Charles I. 2011. “Intermediate Goods and Weak Links in the Theory of Economic Development.” *American Economic Journal: Macroeconomics* 3 (2):1–28.
- Just, Richard E. and Rulon D. Pope. 1979. “Production function estimation and related risk considerations.” *American Journal of Agricultural Economics* 61 (2):276–284.
- Karlan, Dean, Robert Darko Osei, Isaac Osei-Akoto, and Christopher Udry. 2012. “Agricultural Decisions After Relaxing Credit and Risk Constraints.” NBER Working Paper 18463.
- Kazianga, Harounan and Christopher Udry. 2006. “Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso.” *Journal of Development Economics* 79 (2):413–446.
- Klenow, Peter and Andrés Rodríguez-Clare. 1997. “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?” In *NBER Macroeconomics Annual 1997, Volume 12*, NBER Chapters. National Bureau of Economic Research, Inc, 73–114.
- Kochar, Anjini. 1999. “Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India.” *Review of Economics and Statistics* 81 (1):50–61.
- Kohn, Meir G., Charles F. Manski, and David S. Mundel. 1976. “An Empirical Investigation of Factors Which Influence College-Going Behavior.” In *Annals of Economic And Social Measurement, Volume 5, number 4*, edited by P. Aghion and S.N. Durlauf. National Bureau of Economic Research, 391–419.
- Koren, Miklós and Silvana Tenreyro. 2012. “Technological Diversification.” *American Economic Review* Forthcoming.
- Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2008. “Aggregate implications of indivisible labor, incomplete markets, and labor market frictions.” *Journal of Monetary Economics* 55 (5):961–979.
- . 2011. “A three state model of worker flows in general equilibrium.” *Journal of Economic Theory* 146 (3):1107–1133.
- Krusell, Per, Toshihiko Mukoyama, and Ayşegül Şahin. 2010. “Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations.” *Review of Economic Studies* 77 (4):1477–1507.
- Lagakos, David and Michael E. Waugh. 2012. “Selection, Agriculture, and Cross-Country Productivity Differences.” *American Economic Review*. Forthcoming.
- Lochner, Lance J. and Alexander Monge-Naranjo. 2011. “The Nature of Credit Constraints and Human Capital.” *American Economic Review* 101 (6):2487–2529.

- Manski, Charles F. 1989. "Schooling as experimentation: a reappraisal of the postsecondary dropout phenomenon." *Economics of Education Review* 8 (4):305–312.
- Midrigan, Virgiliu and Daniel Xu. 2010. "Finance and Misallocation: Evidence from Plant-Level Data." Mimeo.
- . 2012. "Finance and Misallocation: Evidence from Plant-Level Data." Working Paper.
- Moll, Benjamin. 2012. "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" Working Paper.
- Morduch, Jonathan. 1995. "Income Smoothing and Consumption Smoothing." *Journal of Economic Perspectives* 9 (3):103–114.
- Ogaki, Masao and Qiang Zhang. 2001. "Decreasing relative risk aversion and tests of risk sharing." *Econometrica* 69 (2):515–526.
- Piketty, Thomas and Emmanuel Saez. 2006. "The evolution of top incomes: A historical and international perspective." *American Economic Review* 96 (2):200–205.
- Poschke, Markus. 2010. "Skill-Biased Change in Entrepreneurial Technology." Mimeo, McGill University.
- Prasada Rao, D.S. 1993. "Intercountry comparisons of agricultural output and productivity." FAO Economic and Social Development Paper No. 112.
- Reardon, Thomas, Peter Matlon, and Christopher Delgado. 1988. "Coping with household-level food insecurity in drought-affected areas of Burkina Faso." *World Development* 16 (9):1065–1074.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu. 2008. "Agriculture and aggregate productivity: A quantitative cross-country analysis." *Journal of Monetary Economics* 55 (2):234–250.
- Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. 2010. *Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]*. Minneapolis: University of Minnesota. URL <http://www.ipums.org/>.
- Snyder, Thomas D., editor. 1993. *120 Years of American Education: A Statistical Portrait*. United States Department of Education.
- Taubman, Paul and Terence Wales. 1972. *Mental Ability and Higher Educational Attainment in the 20th Century*. National Bureau of Economic Research.
- Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62 (3):539–591.

- Traxler, Greg, Jose Falck-Zepeda, J.I. Ortiz-Monasterio R., and Ken Sayre. 1995. "Production risk and the evolution of varietal technology." *American Journal of Agricultural Economics* 77 (1):1–7.
- Udry, Christopher. 1996. "Efficiency and Market Structure: Testing for Profit Maximization in African Agriculture." Working Paper.
- United Nations. 1993. *National Account Statistics: Main Aggregates and Detailed Tables*. Table 2.3. URL http://data.un.org/Data.aspx?q=Table+2.3&d=SNA&f=group_code%3a203.
- Vollrath, Dietrich. 2009. "How important are dual economy effects for aggregate productivity?" *Journal of Development Economics* 88 (2):325–334.
- Willis, Robert J. and Sherwin Rosen. 1979. "Education and Self-Selection." *Journal of Political Economy* 87 (5):7–36.
- Zerfu, Daniel and Donald F. Larson. 2010. "Incomplete markets and fertilizer use: Evidence from Ethiopia." World Bank Policy Research Paper 5235.

APPENDIX A

ROBUSTNESS

3.5 Labor Share Parameter, η

Estimates of the labor share vary substantially. In this section, I consider the how changes in the agricultural labor share parameter, η , impact the predictions of the model. As a share of value added, payments to labor are generally estimated to be below 0.50.⁵ This implies that as a share of total output, the labor share is almost certainly below 0.40. I therefore use this as the upper bound, and vary $\eta \in \{0.2, 0.3, 0.4\}$ while holding the rest of the calibration fixed. Table 3.1 lists the results.

Table 3.1: Model Results for Different η

Economy	<i>Labor Productivity Gap</i>		$p_x X / p_a Y_a$		N_a (%)	
	Agriculture	Aggregate	Rich	Poor	Rich	Poor
<i>Data</i>	63.7	23.1	0.40	0.09	2.8	82.0
<i>Model with</i>						
$\eta = 0.40$	33.5	7.5	0.40	0.22	2.8	56.8
$\eta = 0.30$	47.3	6.5	0.40	0.25	1.0	41.2
$\eta = 0.20$	76.2	5.6	0.40	0.30	0.01	27.0

Increasing η causes agricultural productivity differences to decrease, while aggregate productivity differences increase. This is due to two forces that work in opposite directions. First, higher η causes agricultural productivity to decrease in *both* economies. Because the rate of decrease is higher in the rich country, agricultural output per worker differences decrease. At the same time however, higher η causes the employment share in agriculture to increase. The rate of this increase is higher in the poor economy. Because aggregate productivity is an employment-weighted average of sectoral productivity, aggregate productivity differences tilt towards agricultural productivity differences. This increases aggregate productivity differences as η increases. Higher η also implies that the intermediate input share decreases in the poor economy. Because the parameter on intermediate inputs, ψ , is held fixed, higher η decreases the span of control of the production function, which decreases expected income to villages. To limit exposure to risk, villages decrease investment in intermediate inputs.

Overall, the model's basic predictions stand up to varying the labor share parameter in the production function.

3.6 Changes in the Shock Distribution

I investigate the importance of the shock distribution. I do so by varying the standard deviation of the underlying normal distribution σ_z , while holding the support \underline{z} and \bar{z} fixed. The results are presented in Table 3.2.

Interestingly, the higher the standard deviation of the shock distribution, the smaller productivity differences get, and the higher the intermediate input share in the poor econ-

⁵For further discussion, see Herrendorf and Schoellman (2012).

Table 3.2: Model Results for Different σ_z

Economy	<i>Labor Productivity Gap</i>		$p_x X / p_a Y_a$		N_a (%)	
	Agriculture	Aggregate	Rich	Poor	Rich	Poor
<i>Data</i>	63.7	23.1	0.40	0.09	2.8	82.0
<i>Model with</i>						
$\sigma_z = 0.59$	33.5	7.5	0.40	0.22	2.8	56.8
$\sigma_z = 0.75$	29.8	6.3	0.40	0.27	1.8	44.6
$\sigma_z = 0.85$	29.4	5.9	0.40	0.38	1.7	40.9
$\sigma_z = 0.95$	28.9	5.6	0.40	0.29	1.5	37.8

omy. This is due to the interaction of the low utility weight on agricultural consumption, α , and subsistence requirements \bar{a} . Intuitively, because α is so low, total agricultural output needs to be roughly \bar{a} . When σ_z is low, the price p_a must increase to incentivize people to produce with risky intermediate inputs. As σ_z increases, a larger and larger number of villages “luck” into a good shock, and are able to produce \bar{a} and the equilibrium price remains low.

APPENDIX B

DATA SOURCES FOR AGRICULTURAL STATISTICS

I make use of the publicly available data from Restuccia, Yang, and Zhu (2008) for statistics on aggregate productivity, agricultural productivity, labor, and intermediate input prices. This is augmented with purchasing power parities (PPP) for agricultural output and nonagricultural intermediate inputs from Prasada Rao (1993). The resulting dataset contains 84 countries, which are:

Algeria, Angola, Argentina, Australia, Austria, Bangladesh, Belgium, Bolivia, Brazil, Burkina Faso, Burundi, Cameroon, Canada, Chad, Chile, Columbia, Costa Rica, Côte d'Ivoire, Democratic Republic of the Congo, Denmark, Dominican Republic, Ecuador, Egypt, El Salvador, Ethiopia, Finland, France, Germany, Ghana, Greece, Guatemala, Guinea, Haiti, Honduras, Hungary, India, Indonesia, Iran, Iraq, Ireland, Israel, Italy, Japan, Kenya, Korea, Madagascar, Malawi, Malaysia, Mali, Mexico, Morocco, Mozambique, Nepal, Netherlands, New Zealand, Nicaragua, Niger, Nigeria, Norway, Pakistan, Papua New Guinea, Paraguay, Peru, Philippines, Portugal, Rwanda, Senegal, Somalia, South Africa, Spain, Sri Lanka, Sudan, Sweden, Switzerland, Syria, Thailand, Tunisia, Turkey, U.K. U.S.A. Uganda, Uruguay, Venezuela, and Zimbabwe.

I am interested a measure of the ninetieth percentile country relative to the tenth percentile country, similar to that used in Caselli (2005). As a measure of the rich country, I take average of the top ten percent of countries. Listed from largest to smallest income, they are USA, Canada, Switzerland, Australia, Norway, Netherlands, Belgium, and Germany. As a measure of the “tenth” percentile, I take an average of the countries that make up the bottom fifteen to five percent of countries, as ranked by PPP GDP per capita. They are Somalia, Rwanda, Mozambique, Uganda, Malawi, Chad, Zaire and Niger.

The productivity statistics are taken from Restuccia, Yang, and Zhu (2008). They are derived from PWT and FAO data. These averages imply a factor of 63.66 difference in agricultural output per worker and 23.18 difference in aggregate output per worker. On average, 82% of the population in the poor countries work in agriculture.

As in the text, the domestic intermediate share in agriculture of country j is

$$\hat{X}^j := \frac{p_x^j X^j}{p_a^j Y_a^j} \quad (3.6.1)$$

This measure is not directly reported in Prasada Rao (1993). He does however, report the real intermediate share in agriculture, defined as

$$\hat{X}^{j*} := \frac{p_x^* X^j}{p_a^* Y_a^j} \quad (3.6.2)$$

where p_x^* and p_a^* are international prices of intermediate inputs and agricultural output. Combining equations (3.6.1) and (3.6.2), it is possible to write the domestic intermediate share as

$$\hat{X}^j = \hat{X}^{j*} \left(\frac{p_x^j / p_x^*}{p_a^j / p_a^*} \right) \quad (3.6.3)$$

The price ratio in equation (3.6.3) can be calculated from reported purchasing power parities

$$PPP_a^j = \frac{p_a^j}{p_a^*}$$

$$PPP_x^j = \frac{p_x^j}{p_x^*}$$

where p_a^* and p_x^* are international (unreported) prices and (p_a^j, p_x^j) are (unreported) domestic prices for country j . The purchasing power parities are normalized to one in a baseline country, which in Prasada Rao (1993) is the USA. Therefore, $PPP_a^{US} = PPP_x^{US} = 1$, implying $\hat{X}^{US} = \hat{X}^{US*}$. Therefore, calculating the domestically priced intermediate share of all other countries reduces to

$$\hat{X}^j = \hat{X}^{j*} \left(\frac{PPP_x^j}{PPP_a^j} \right) \quad (3.6.4)$$

As mentioned, the real intermediate share and the ratio of PPPs are both reported, so this is sufficient to define the domestically priced intermediate input share. The poor group group of countries has, on average, a domestically priced intermediate input share of 0.09 and a real intermediate input share of 0.13. The right hand side of equation (3.6.4) is the statistic reported in Figure 1.1. The horizontal axis, GDP per capita, is real GDP per capita for 1985, variable *cgdp* from the Penn World Tables version 7.0 (PWT).

For the comparison of agriculture to manufacturing and services, I use a set of 49 countries from the UN SNA. The 49 countries with sufficient data for all three sectors Austria, Benin, Bolivia, Botswana, Burundi, Cameroon, Canada, Cape Verde, Chile, Hong Kong, Colombia, Cyprus, Denmark, Ecuador, El Salvador, Fiji, Finland, France, Gambia, Germany, Ghana, Hungary, Iceland, Italy, Jamaica, Japan, Jordan, Republic of Korea, Luxembourg, Mauritius, Mexico, Netherlands, New Zealand, Nigeria, Norway, Peru, Portugal, Rwanda, Seychelles, Sierra Leone, Spain, Sri Lanka, Swaziland, Sweden, Syrian Arab Republic, United Kingdom, Uruguay, Venezuela, and Zimbabwe.

From the UN SNA, I use “Output, at basic prices” and “Intermediate consumption, at purchaser’s prices” for the year 1985 for each of the three sectors. Dividing them gives the domestically priced intermediate input share by sector. Figure 1.2 plots this, along with variable *cgdp* for 1985 from PWT on the horizontal axis. Note that the intermediate share in agriculture derived from the UN statistics and the FAO statistics may differ. This is due to the fact that the UN statistics includes intermediate inputs produced in the agricultural sector, while the FAO statistics only consider nonagricultural intermediate inputs.

APPENDIX C

PROOFS

3.7 Proof of Proposition 1

Proof. First, the profit maximizing first order condition implies

$$\frac{x^C}{p_a^C Y_a^C} = \psi \quad (3.7.1)$$

Define x^* be the optimal choice for a farmer facing p_a^I , but with complete markets. Then the first order condition implies,

$$\frac{x^*}{p_a^I Y_a^*} = \psi \quad (3.7.2)$$

Comparing (3.7.1) and (3.7.2), the proposition is equivalent to proving that at the price p_a^I , $x^I < x^*$. This can be seen from the first order conditions. The first order condition from the ordering problem is

$$A(p_a^I)^{1/(1-\eta)} F'(x^I) \int_Z z^{1/(1-\eta)} \left(\frac{\tilde{u}'(y(x^I, z))}{\mathbb{E}_z[\tilde{u}'(y(x^I, z))]} \right) dQ(z) = 1 \quad (3.7.3)$$

where $F(\cdot)$ and \tilde{u} are defined as in the text. Note that $F(\cdot)$ is concave because $\psi + \eta < 1$. Now, consider the profit maximizing problem. The first order condition is

$$A(p_a^I)^{1/(1-\eta)} F'(x^*) \int_Z z^{1/(1-\eta)} dQ(z) = 1 \quad (3.7.4)$$

Since $u(\cdot)$ is concave,

$$\int_Z z^{1/(1-\eta)} \left(\frac{\tilde{u}'(y(x^*, z))}{\mathbb{E}_z[\tilde{u}'(y(x^*, z))]} \right) dQ(z) < \int_Z z^{1/(1-\eta)} dQ(z) \quad (3.7.5)$$

Since $F(\cdot)$ is concave, it follows that $x^I < x^*$. ■

3.8 An Additional Lemma for the Proof of Proposition 2

To prove the result, I first characterize the the equilibrium of an I economy with TFP A^2 and $\bar{a} = 0$ in terms of an economy with TFP A^1 and $\bar{a} = 0$. This is done in Lemma 1 below.

Lemma 1. *Consider two I economies characterized by TFP levels A^1 and A^2 , both with $\bar{a} = 0$. Denote the equilibrium for economy 1 as $(x^1, n_a^1(z), p_a^1)$. Then the equilibrium for economy 2, $(x^2, n_a^2(z), p_a^2)$ can be characterized as*

$$\begin{aligned} n_a^2(z) &= n_a^1(z) \\ x^2 &= \left(\frac{A^2}{A^1} \right) x^1 \\ p_a^2 &= \left(\frac{A^1}{A^2} \right)^\psi p_a^1 \end{aligned}$$

Proof. Two things must be checked for the proposed allocation to be a competitive equilibrium. First, the proposed equilibrium must satisfy the village optimization problem. That is, if $(p_a^1, x^1, n_a^1(z))$ is an equilibrium in economy 1, then $(p_a^2, x^2, n_a^2(z))$ satisfies the farmer's optimization problem in economy 2. Second, markets must clear. These are considered in turn.

The first thing to check is that the labor choice is identical between the two. Using the decision rules, I can check this using the first order conditions for $n_a^1(z)$ and $n_a^2(z)$.

$$\frac{n_a^1(z)}{n_a^2(z)} = \left(\frac{p_a^1 A^1 (x^1)^\psi}{p_a^2 A^2 (x^2)^\psi} \right)^{1/(1-\eta)}$$

Plugging in (p_a^2, x^2) implies

$$\frac{n_a^1(z)}{n_a^2(z)} = 1$$

For simplicity, I drop the superscript on $n_a(z)$, with the understanding that they are identical in both economies.

Next up is to check if x^2 satisfy the required first order conditions, given that x^1 satisfies the first order condition in Economy One. Note that when $\bar{a} = 0$, the production utility for a given income y can be written as

$$\begin{aligned} V^P(y) &= \alpha \log(c_a^1) + (1 - \alpha) \log(c_m^1) \\ &= \Omega - \alpha \log(p_a^1) + \log(y) \end{aligned} \quad (3.8.1)$$

where $\Omega = \alpha \log(\alpha) + (1 - \alpha) \log(1 - \alpha)$. Denote the income of a farmer who chooses intermediates x and gets hit with shock z in economy $j = 1, 2$ as

$$y^j(x, z) = p_a^j A^j z x^\psi n_a(z)^\eta - x + (1 - n_a(z)) A^j$$

Plugging in the proposed equilibrium yields the following relationship

$$y^2(x^2, z) = \left(\frac{A^2}{A^1} \right) y^1(x^1, z) \quad (3.8.2)$$

Equation (3.8.1) implies that

$$x^j = \arg \max_x \int_{\underline{z}} \log(y^j(x, z)) dQ(z)$$

After plugging in the optimal values for $n_a(z)$, the first order condition for this problem can be written as

$$\int_{\underline{z}} \left(\frac{\psi p_a^j A^j x^j \psi n_a(z)^\eta - 1}{y^j(x, z)} \right) = 0$$

Plugging in the proposed equilibrium yields a relationship between economies one and two

$$\int_{\underline{z}} \left(\frac{\psi p_a^2 A^2 x^2 \psi n_a(z)^\eta - 1}{y^2(x, z)} \right) = \left(\frac{A^1}{A^2} \right) \int_{\underline{z}} \left(\frac{\psi p_a^1 A^1 x^1 \psi n_a(z)^\eta - 1}{y^1(x^1, z)} \right)$$

Since an equilibrium is assumed in economy one, it follows then that

$$\int_{\underline{z}}^{\bar{z}} \left(\frac{\Psi p_a^2 z A^2 x^{2\Psi} n_a(z)^\eta - 1}{y^2(x, z)} \right) = 0$$

Therefore, the proposed economy two equilibrium satisfies a village's optimization problem.

To check market clearing, first note that aggregate sector a output for economy $j = 1, 2$ is

$$Y_a^j = A x^j \Psi \mathbb{E}_z(z n_a(z)^\eta)$$

Thus,

$$\frac{Y_a^1}{Y_a^2} = \left(\frac{A^1}{A^2} \right) \left(\frac{x^1}{x^2} \right)^\Psi \quad (3.8.3)$$

Therefore, at the proposed equilibrium,

$$\frac{Y_a^1}{Y_a^2} = \left(\frac{A^1}{A^2} \right)^{1+\Psi} \quad (3.8.4)$$

For any $\bar{a} \geq 0$, the total demand for sector a consumption is given by

$$D_a^j = (1 - \alpha) \bar{a} + \frac{\alpha}{p_a^j} \mathbb{E}_z[y^j(X^j, z)] \quad (3.8.5)$$

Using equation (3.8.2),

$$\frac{\mathbb{E}_z[y^1(x^1, z)]}{\mathbb{E}_z[y^2(x^2, z)]} = \frac{A^1}{A^2} \quad (3.8.6)$$

Since $\bar{a} = 0$, equations (3.8.5) and (3.8.6) and the prices p_a^1 and p_a^2 imply that

$$\frac{D_a^1}{D_a^2} = \left(\frac{A^1}{A^2} \right)^{1+\Psi} \quad (3.8.7)$$

Since the proof assumes an equilibrium in economy 1, equations (3.8.4) and (3.8.7) imply $Y_a^2 = D_a^2$ so that the agricultural output market clears in economy two. Since the labor market in sector m clears trivially, Walras' Law implies that the sector m output market also clears. ■

3.9 Proof of Proposition 2

Proof. With Lemma 1 in hand, the three claims of the proposition follow quickly.

$n_a(z)$ is independent of A

This follows directly from Lemma 1.

The intermediate input share is independent of A

Denote \hat{X}^j as the intermediate good share in economy $j = 1, 2$, so that \hat{X}^j is defined as

$$\hat{X}^j = \frac{x^j}{p_a^j Y_a^j} \quad (3.9.1)$$

First, note that total agricultural output in economy j is given as

$$Y_a^j = A^j (x^j)^\psi \mathbb{E}_z (z n_a^j(z)^\eta) \quad (3.9.2)$$

Using the fact that $n_a^1(z) = n_a^2(z)$ and plugging (3.9.2) into (3.9.1) gives

$$\frac{\hat{X}^1}{\hat{X}^2} = \left(\frac{x^1}{x^2} \right)^{1-\psi} \left(\frac{p_a^2}{p_a^1} \right) \left(\frac{A^2}{A^1} \right)$$

Plugging in the equilibrium found in Lemma 1, this gives

$$\begin{aligned} \frac{\hat{X}^1}{\hat{X}^2} &= \left(\frac{A^1}{A^2} \right)^{1-\psi} \left(\frac{A^1}{A^2} \right)^\psi \left(\frac{A^2}{A^1} \right) \\ &= 1 \end{aligned}$$

Since A^1 and A^2 are arbitrary, this completes the proof.

No increase in productivity relative to C economy

For any two economies characterized by TFP A^1 and A^2 and complete markets (the C economy), it is easy to show that in equilibrium,

$$\begin{aligned} n_a^1 &= n_a^2 \\ x^2 &= \left(\frac{A^2}{A^1} \right) x^1 \end{aligned}$$

Since this is the same as in the incomplete markets model (the I economy), relative agricultural labor productivity between the two economies is equal in both. ■

3.10 Proof of Proposition 3

Proof. Consider the equilibrium for economy 1 with TFP equal to A^1 . Denote this equilibrium $(p_a^1, x^1, n_a^1(z))$. Suppose that the intermediate good share is $\hat{X}^1 < \psi$, where the inequality follows from Proposition 1. Define x^{1C} to be the optimal choice of the farmer who faces p_a^1 but with complete markets. I know that the intermediate good share is $\hat{X}^{1C} = \psi$. Therefore, the ratio is

$$\frac{\hat{X}^1}{\hat{X}^{1C}} = \frac{\hat{X}^1}{\psi} = \left(\frac{x^1}{x^{1C}} \right)^{(1-\eta-\psi)/(1-\eta)}$$

Thus, I can write \hat{X}^1 as

$$\hat{X}^1 = \psi \left(\frac{x^1}{x^{1C}} \right)^{(1-\eta-\psi)/(1-\eta)}$$

Similarly, it follows that in Economy 2,

$$\hat{X}^2 = \psi \left(\frac{x^2}{x^{2C}} \right)^{(1-\eta-\psi)/(1-\eta)}$$

These equations show that the intermediate good share is directly related to how “far” the optimal choice of x is from the choice x^C . What’s left to show is that when $\bar{a} > 0$ and $A^1 > A^2$,

$$\frac{x^1}{x^{1C}} > \frac{x^2}{x^{2C}}$$

This follows from the fact that, when $\bar{a} > 0$, relative income net of subsistence,

$$\frac{y^1(z) - p_a^1 \bar{a}}{y^2(z) - p_a^2 \bar{a}}$$

is decreasing in z . ■

APPENDIX D

SHOCK DISTRIBUTION CONSTRUCTION

The data used was collected by ICRISAT. I use the version that was released by Stefan Dercon, via the Oxford University website. It is publicly available at <http://www.economics.ox.ac.uk/members/stefan.dercon/icrisat/ICRISAT/oldvls.html>.

The ICRISAT VLS is an unbalanced panel set covering 10 villages in India. The data covers the time period 1975 - 1984. The goal is to calculate the value of the following inputs at the village-year-season level: capital K , agricultural intermediates I , nonagricultural intermediates X , human labor hours N_a , and land L . Allowing for some abuse of notation, let these letters also denote the set of all inputs of that type, so K is the set of all capital goods in the economy, for example.

3.11 Construction of Inputs

The data includes five inputs: capital, human labor, non-agricultural intermediates, and agricultural intermediates. Each input category includes a number of different inputs, with quantities measured in different denominations. The first step is to put all inputs in a given category into units of a numeraire good using relative prices. The second step is then to construct a quantity index using the time series of the numeraire good. These are discussed in detail below.

A time period here is a season, and each year contains three seasons. So time t here should be read as a season-year pair.

Capital includes class code E , farm equipment and implements and class code M , major farm machinery, class code R , production capital assets, and the total value of land. Class code E includes basic farm equipment such as plows and hoes. Class code M includes major machinery such as tractors and electric pumps. Capital also includes bullock labor hours at the plot level, from both owned and rented bullocks. The value of an hour of an owned bullock is imputed from the rental rates of hired bullock hours, so they are valued equally.

The total value of capital at time t in village v can be computed simply by summing over the capital goods, and adding in the value of bullock labor V^B and land V^L .

$$V_{p,f,v,t}^{K1} = \left(\sum_{k \in K} V_{k,f,p,v,t} \right) + V_{f,p,v,t}^B + V_{f,p,v,t}^L \quad \forall (p, f, v, t)$$

The most commonly used capital good is the electric pump, code MK , so I compute the price $p_{v,t}^{MK} = V_{p,f,v,t}^{MK} / Q_{p,f,v,t}^{MK}$. From there, I put capital in terms of units of MK by simply dividing through by $p_{v,t}^{MK}$

$$V_{p,f,v,t}^{K2} = \frac{V_{p,f,v,t}^{K1}}{p_{v,t}^{MK}} \quad \forall (p, f, v, t).$$

The last step is to construct the quantity of capital, by using the time series $p_{v,t}^{MK}$ to get rid of price changes. Therefore, the total quantity of capital used on plot p on farm f in

village v at time t can be written as

$$Q_{p,f,v,t}^K = \frac{V_{p,f,v,t}^{K2}}{P_{v,0}^{MK}}$$

From there, I simply sum up and get the quantity of capital for village v at time t

$$Q_{v,t}^K = \sum_{p,f} Q_{p,f,v,t}^K$$

The rest of the inputs are constructed similarly. For labor, the Y files give hours of male, female, and child labor in the data. Since I calibrate to match the fraction of the population over 15 years of age, I include only male and female labor. Child labor is a small component with the lowest price (i.e. not as productive as an adult laborer). Including it makes no discernible difference. Similar to bullock labor hours, the Y files include disaggregated data on both family and hired workers. Once again though, the value of family labor is imputed from market value, so they are valued equally. Nonagricultural intermediates include pesticides, which are input codes 1A – 9A, and fertilizer (input codes A – Z). Agricultural intermediates can be included by using the Y files. Organic manure in the data are inputs 1 – 7. Seed is denoted as inputs CA – ZK. The quantity and values are in the Y files.

3.12 Output

Total value of output is given by summing over output values in the Y files by plot level. I include both actual production and by-products produced by farming. The procedure is identical to the one discussed above for inputs. The only difference is that each village produces a different (but not mutually exclusive) set of goods at output, and therefore each village gets a different numeraire output good.

3.13 Decomposition of Residuals

Now armed with input vector $(K, L, N_a, I, X)_{v,t}$ and output $Y_{v,t}$, I can calculate the residual

$$z_{v,t}^* = \frac{Y_{v,t}^a - I_{v,t}}{X_{vt}^\psi N_{a,vt}^\eta (K_{vt} + L_{vt})^{1-\psi-\eta}}$$

where η and ψ are taken from the calibration in the main text. The rest is explained in the main text.

APPENDIX E

EDUCATION DATA SOURCES

I take several historical data series as exogenous to the model, and this section details the construction of those series. Data are taken from several sources in order to construct a consistent series since 1900. From 1900 to 1958, most data were collected every two years and published in the Biennial Survey of Education (BSE). Since 1962, the Digest of Education Statistics (DES) has been published annually. Other publications including the annual U.S. Statistical Abstract, the Bicentennial Edition “Historical Statistics of the United States: Colonial Times to 1970”, and “120 Years of American Education: A Statistical Portrait” help in bridging breaks between series, as well as verifying continuity of series that may have changed names from year to year. Also, many data were revised in later publications, so I take the most recent published estimates where available.

First, let c_t be the total annual cost of college per student. I assume that the total cost for educating *all* students in the U.S. in a given year equals the total revenues received in the current period by all institutions of higher education. Dividing this by the total enrollment each year yields the total annual cost per student. Alternatively, one could use the total current expenditures rather than revenues as the measure of total cost, but this makes little difference quantitatively because revenues and expenditures track each other quite closely. In addition, the revenue data is preferable because it allows us to determine how much of costs are paid out-of-pocket by students for tuition and fees, and how much comes from other sources such as state, local, and federal governments, private gifts, endowment earnings, auxiliary enterprises (athletics, dormitories, meal plans, etc.), and other sources. The numerator for c_t is constructed as follows:

- 1997-2000: total current revenue must be computed as the sum of current-fund revenue for public and private institutions, from the DES.
- 1976-1996: total current revenue equals “current-fund revenue of institutions of higher education” from the DES.
- 1932-1975: total current revenue equals “current-fund revenue of institutions of higher education” in “120 Years of American Education: A Statistical Portrait”.
- 1908-1930: total current revenue equals “total receipts exclusive of additions to endowment” for colleges, universities, and professional schools, from the BSE.
- 1900-1908: total current revenue equals “total receipts exclusive of additions to endowment” for colleges, universities, and professional schools, and is computed as (income per student)*(total students, excluding duplicates) from the BSE. Continuity with later years can be verified using the “income per student” series, which was published from 1890-1920.

The denominator for c_t is constructed as follows:

- 1946-2000: total fall enrollment for institutions of higher education, from the DES.
- 1938-1946: resident college enrollments, from the BSE. Continuity with the later series can be verified in that year 1946 data matches in both.

- 1900-1938: total students, excluding duplicates, in colleges, universities, and professional schools, from the BSE. Continuity with the later series can be verified in that year 1938 data matches in both.

Second, I construct two time series which estimate the share of annual college costs paid out-of-pocket by students. One measure, λ_t , includes only tuition and fees paid by students, and the other measure, ϕ_t includes tuition, fees, room, and board. In each year λ_t equals total tuition and fees paid by all students divided by total current revenue received by institutions of higher education. Similarly, ϕ_t equals total tuition, fees, room, and board aid by all students divided by total current revenue received by institutions of higher education. In each case, the measure of total current revenue is the same time series as was used above in constructing c_t . The time series for λ_t is constructed as follows:

- 1997-2000: current fund revenues from tuition and fees for all institutions of higher education is computed as the sum of the series for public and private institutions, from the DES.
- 1976-1996: current fund revenues from student tuition and fees, from the DES.
- 1930-1975: current fund revenues from student tuition and fees, from “120 Years of American Education: A Statistical Portrait”.
- 1918-1930: receipts of universities, colleges, and professional schools for student tuition and fees, from BSE.
- 1900-1918: I am unable to obtain proper data for these years.

The time series for ϕ_t is constructed as follows:

- 1976-2000: Average tuition, fees, room, and board paid by full-time equivalent (FTE) students is obtained from the DES. I multiply this by enrollment of FTE students, also from the DES, and divide by the current fund revenues to compute ϕ_t .
- 1960-1976: I am unable to obtain proper data for these years.
- 1932-1958: Data available biennially on total revenues from student tuition and fees, as well as revenue from auxiliary enterprises and activities (room and board), in the BSE. ϕ_t computed as the sum of these, divided by total current revenue.
- 1900-1930: ϕ_t computed as total revenue from student fees (included tuition, fees, room, and board) divided by total current revenue.