

Emerson, Patrick M.; McGough, Bruce

Working Paper

Learning about education

IZA Discussion Papers, No. 5756

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Emerson, Patrick M.; McGough, Bruce (2011) : Learning about education, IZA Discussion Papers, No. 5756, Institute for the Study of Labor (IZA), Bonn, <https://nbn-resolving.de/urn:nbn:de:101:1-201106202097>

This Version is available at:

<https://hdl.handle.net/10419/51961>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 5756

Learning about Education

Patrick Emerson
Bruce McGough

June 2011

Learning about Education

Patrick Emerson

*Oregon State University
and IZA*

Bruce McGough

Oregon State University

Discussion Paper No. 5756

June 2011

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Learning about Education

Limited human capital investment is a common characteristic of low-income countries despite the fact that estimated returns to educational investment in low-income countries are generally higher than in high-income countries. Empirical evidence suggests that income and credit constraints can only account for a small part of this underinvestment. Recent experimental evidence shows that families' misperceptions about the returns to education play a large role in their low investment levels. This paper builds a model of human capital and growth that incorporates an adaptive learning mechanism to capture the way agents form perceptions about returns to education. In an economy where human capital investments have both private and public returns, we find multiple learnable equilibria, including those which are characterized by low investment and low returns. We also find that even when the rational equilibrium corresponds to a high level of human capital investment, the learning mechanism, influenced by the agents' priors and cultural bias, may impart low human capital investment for extended periods. Policies that can speed up the learning process are examined and it is found that faster rates of growth can be achieved through interventions.

JEL Classification: O40, O15

Keywords: growth, education, learning

Corresponding author:

Patrick M. Emerson
Department of Economics
Oregon State University
303 Ballard Hall
Corvallis, Oregon 97331
USA
E-mail: patrick.emerson@oregonstate.edu

1 Introduction

One of the persistent problems developing economies face is their inability to make substantial progress in raising the average level of human capital. Low levels of human capital investment persist despite the fact that human capital investments, through education, have been shown to have high returns in low-income countries where human capital is relatively scarce. Typical explanations for the lack of optimal, or even sufficient human capital investment on the part of low-income households usually begin with income and credit constraints. However, the failure of unconditional cash transfers, microcredit and the lowering of school costs to increase educational investments suggest that income and credit constraints provide an insufficient explanation. Low school quality is also an incomplete explanation as it fails to fully reconcile the fact that returns to education are consistently estimated to be high even in areas where school quality appears to be relatively low. The failure of low income households to optimally invest in education, then, is one of the persistent puzzles in development.

New research has begun to shed light on another factor that appears to have a dramatic impact and that might solve the puzzle: low-income families may not be aware of the rate of return to investments in human capital (Jensen, 2010). Information scarcity may, in fact, be one of the key hallmarks of poor households in low-income countries. For example, other research has found that agricultural households fail to use profitable fertilizers, and that demonstrating their effectiveness can increase utilization rates (Duflo, et. al., 2004).

Uninformed actors have been found in other studies of different types of economic

decision making. Credit markets are one example where a number of studies have found that individuals underestimate the costs of borrowing (e.g. Stango and Zinman, 2007). Other studies have found that workers are not well-informed of their pension or social security benefits (e.g. Chan and Stevens, 2008; Mitchell, 1988), or do a poor job estimating the risk of smoking (Viscusi, 1990).

The fact that information can affect behavior has been supported by a number of different studies. For example, Duflo and Saez (2003) find that providing information about retirement benefits affects decisions about retirement plans, and Dupas (2006) finds that information on age and HIV infection rates can influence the risky sexual behavior of Kenyan girls.

Human capital investment decisions are based on information identifying (or at least shedding some light on) the potential returns to education; and if, for poor households in low-income countries, this information is scarce, resisted, or is perceived to be of poor quality, then understanding and modeling the process through which households update their information – i.e. learn about education – is essential to determining why certain outcomes, such as low human capital investment traps, arise. Incorporating learning into theoretical models becomes all the more critical when the economy includes complicated feedback, as is the case with human capital investment: the actions of the households themselves affect the very returns they are trying to understand.¹ Thus modeling the mechanism through which households learn about the true returns to education is critically important to understanding the economic development of a country.

¹This self-reinforcing feedback dynamic is likely an accurate description of low-income countries, which are known to have large social returns to education.

This paper builds a dynamic model of household investment in human capital, which exhibits both private and social returns. Under rationality, the model has multiple equilibria including those corresponding to high and low education outcomes. We then modify the model by assuming agents are boundedly rational, and we incorporate a set of learning mechanisms to examine the process through which households learn about the real returns to education. We find that the low education rational equilibrium may be stable under learning: if agents use our simple learning mechanism then, under certain quite general conditions, the economy will converge (in a natural, probabilistic sense) to the low human capital investment rational equilibrium. We find further, that even if agents eventually learn to coordinate on a high education outcome, the nature of the learning algorithm, as influenced by beliefs and cultural norms, may prevent rapid convergence.

Our model and results are rich enough to allow for policy prescription. We find that associated to each stable rational equilibrium is a basin of attraction; this suggests that public outreach or other similar programs designed to influence household beliefs may improve welfare by placing these beliefs in the basin of attraction of a pareto superior rational equilibrium. Analysis of this type of policy within the context of our model confirms this intuition.

Policy may also affect how rapidly an economy converges to a rational equilibrium. We identify a parameter of the learning algorithm which has two natural interpretations: “confidence in information quality;” or “strength of prior beliefs.” We show that by increasing confidence, or analogously, weakening priors, the government can increase the speed of convergence to the rational equilibrium and thereby mitigate the tendency of cautious agents to systematically underestimate the value of education.

2 Background

Average levels of education in low-income countries remain well below those of high-income countries despite large-scale educational expansion efforts over the last few decades. Barro and Lee (2001) estimated that the average years of education in developing countries for 2000 was 4.9 years whereas in advanced countries the average years of schooling was 9.8. They also find that in developing countries only 19.7% of the population over 25 have attained some secondary education and only 7.2% have attained some tertiary education whereas in advanced countries the figures are 39.4% and 29.1% respectively. They also estimate that the gains in educational attainment through time are no faster in low-income countries: from 1960 to 2000, advanced countries and developing countries had similar growth trajectories in terms of average years of schooling, in other words, low-income countries are not catching up.

These low investment levels are in stark contrast to the relatively high returns to education experienced in developing countries. In a meta-study of the received empirical evidence, Psacharopoulos and Patrinos (2004) find an average year of schooling effect on income of 10.9% for low-income countries as opposed to 7.4% for high-income countries. Returns to investment in education are similarly divergent: private returns to investment in secondary education in low-income countries are 19.9% whereas they are 12.2% in high-income countries. The returns for higher education are 26% for low-income countries and 12.4% for high-income countries.

Our model assumes both private and social returns to education - the benefits to individual educational investment that accrue to non-investors. Psacharopoulos and

Patrinos (2004) show that many studies have found substantial social returns to education in low-income countries: 21.3% for primary education, 15.7% for secondary and 11.2% for higher education, on average (though they sound a word of caution about the reliability of these estimates given the challenges involved). Hall and Jones (1999) also find significant total factor productivity (TFP) and growth effects from average education levels as do de la Fuente and Domenech (2001). And while Acemoglu and Angrist (2000) do not find evidence of social returns from high school education in the US, Moretti (2006) finds sizeable externalities associated with college education in the US. Both of these findings are reinforced by Irazano and Peri (2009) who find, using US data, social returns from high school education in the zero to one percent range but in the six to nine percent range for college education. Considering the low average level of education in developing countries, estimates of positive social returns to lower levels of educational attainment do not seem unreasonable.

Given the relatively low level of education in developing countries and the relatively high returns, researchers have been left to puzzle over explanations for the lack of investment in human capital. Explanations such as income and credit constraints, high discount rates or simply errors of bias in the measurement of returns have been explored but recent empirical research has left them wanting.²

If income and credit constraints are the explanation for the low investment levels in education in low-income countries, easing them should yield substantial returns. However, Banerjee, et. al., (2010) found evidence from a large experiment on mi-

²For example, there has been some concern over potential mis-measurement of the true returns to education. However Duflo (2001) uses a policy experiment and finds an average return to a year of schooling of 7.8% in Indonesia, and Emerson and Souza (2010) use instrumental variables techniques and find average returns to a year of education in Brazil to be 13.4%.

crocredit in India that providing families with credit did not increase educational investment or outcomes. The experience with unconditional cash transfer programs also suggests that income and credit constraints are insufficient to explain the human capital gap, as de Janvry, et. al. (2006) state: “...unconditional transfers have small effects on school choices compared to conditional transfers where the condition for the transfer is on school attendance.”

The question then arises: do families in low-income countries have complete information about the returns to human capital? Jensen (2010) finds that perceived returns to education in the Dominican Republic are very low, especially relative returns measured with earnings data. He then uses an experiment to study the effect of information about the true return to education on investment behavior. He finds that relative to students not provided with information about returns, informed students perceived dramatically increased returns. Such informed students were more likely to be enrolled in school the next academic term, and when observed four years later, those students had completed on average about 0.20 more years of schooling. He also found some evidence suggesting that students rely heavily on the earnings of workers in their own community when they formed their own expectations of earnings. There is strong evidence then that perceptions about the return to investments in human capital are both substantially low and a constraint on human capital investment in low-income countries.

The link between human capital and growth is now well-established in the empirical literature. Mankiw, Romer and Weil (1989), Barro (1991), Benhabib and Spiegel (1994), Barro (1997) and Hanushek and Kimko (2000) all find evidence of a strong causal relationship between a country’s human capital levels and its growth rates.

Human capital and growth theory dates back to the endogenous growth literature (see Aghion and Howitt (1997) for a good overview) and has been adapted to development most notably by Azariadis and Drazen (1990) which incorporates a threshold human capital externality in a model of economic growth.

Learning has a well-established place in the macroeconomics literature (see, e.g. Evans and Honkapohja, 2001) but has not been widely adopted in growth models despite the considerable evidence of incomplete information. There are now, however, a handful of papers that incorporate learning in models of economic growth. Evans, Honkapohja, and Romer (1998) use a simple learning algorithm in an endogenous growth model with research and development to explain cycles between periods of high and low growth. Arifovic, Bullard, Duffy (1997) use a genetic algorithms approach to explain the transition from low to high growth. And perhaps the work most closely related to ours is that of Steiger (2009): she adapts learning to the Azariadis and Drazen (1990) framework and finds that while the low (in her case, zero) education steady state is learnable, the high education steady state is not; interestingly, however, she finds that complex stochastic dynamics may lead the economy away from the low education steady state and induce endogenous fluctuations.

3 Model

We begin by specifying a fairly general version of the model, and then, to focus attention on issues of equilibrium coordination, we develop our analysis within the context of a highly stylized version. The common environment is a two-period, overlapping generations model with agglomeration effects. The general setting is one where young

agents have to decide on their own human capital investment, but are uncertain of the true returns. When old, true returns are realized but human capital is fixed. Returns are based on both individual human capital investment and on societal human capital investment.

We assume that, at time t , the economy is populated by many identical young agents and equally many identical old agents: population is constant. For simplicity, both young and old own a production technology and consume what they produce. This yeoman farmer assumption is a technical device which helps expose the salient features of the model; however, we could equally develop the arguments by assuming competitive goods and labor markets, and inelastic labor supply.

3.1 Equilibrium in the general model

A young agent has a unit time endowment which he may divide between goods production and human capital accumulation. Letting c_{1t} be goods consumption of the representative young agent in time t , we write

$$(1) \quad c_{1t} = f(n_t)$$

$$(2) \quad h_t = g(1 - n_t, H_{t-1}).$$

Here n is the labor supplied by the agent towards good production, h is the attained level of human capital, f is the young agent's production function, which is assumed primitive in that it takes only labor as an input, and g captures the production of human capital. Finally, H is the aggregate level of human capital, which we allow

to affect human capital production: a more educated population produces educated people more efficiently.³

An old agent also has a unit of time which he supplies inelastically to goods production. Letting c_{2t} be goods consumption of the representative agent who is old in period t , we write

$$(3) \quad c_{2t+1} = F(h_t, A_{t+1})$$

$$(4) \quad A_{t+1} = G(A_t, z_{t+1}, H_t, \varepsilon_{t+1}).$$

Here F is the old agent's more advanced production technology, which depends both on his level of human capital accumulated when young, h_t , as well as an agglomeration effect represented by A_{t+1} .⁴ Think of A_{t+1} as capturing the effectiveness of labor in time $t + 1$. We take A_{t+1} to depend on the effectiveness of labor in time t , exogenous productivity shocks z_{t+1} , the aggregate level of human capital H_t obtained by the young in the previous period, and an additional exogenous shock ε_{t+1} to be explained below.

Young agents in time t receive utility, u , from consumption in periods t and $t + 1$, as well as from the well-being of their progeny; they make time t decisions about human capital investment to maximize expected value utility subject to the constraints indicated above and conditional on available information. Let I_t be the collection of all variables dated $t - n$ for $n \geq 0$. We assume that the information

³Implicit in our production formulation is that the young agent supply labor inelastically: no value is place on leisure.

⁴In the Appendix, we allow the old agent to also access the primitive technology, and then establish conditions sufficient to guarantee that he will not choose to do so.

available to a young agent in time t is given by $\Omega_t = \{A_t, I_{t-1}\}$: notice that the young agent does not observe H_t when making time t decisions, though in an equilibrium he will be able to infer it. Let $V_t(\Omega_t)$ be the value function for the representative young agent at time t . Then $V_t(\Omega_t)$ is obtained by choosing c_{1t}, n_t, h_t , and a contingency plan $c_{2t+1}(\Omega_{t+1})$ to maximize

$$(5) \quad E(u(c_{1t}, c_{2t+1}, V_{t+1}(\Omega_{t+1})) | \Omega_t)$$

subject to the constraints (1) – (3), and taking (4) as given.

Equilibrium in the model is obtained by exploiting the assumption that all agents are identical, and identifying individual and aggregate human capital: $H_t = h_t$. Thus given exogenous processes z_t and ε_t , and initial levels of aggregate human capital H_0 and labor effectiveness A_0 , an equilibrium is any collection of stochastic processes $\{c_{1t}, c_{2t}, h_t, n_t, H_t, A_t\}$ satisfying $H_t = h_t$, constraints (1) – (4), and the representative agent's first order conditions.

3.2 Equilibrium in the stylized model

The level of generality sustained in the previous subsection is useful for defining the modeling environment and understanding the broadest set of natural assumptions; however, to make progress and obtain existence results, further assumptions are required. Because our goal is to examine issues of coordination failure in a model of human capital investment decisions, we impose restrictions explicitly designed to generate multiple equilibria, and abstract from other aspects of the model which

are tangential to these issues. We will revisit the robustness of our results in the conclusion.

Equilibrium indeterminacy is generated through increasing returns to labor effectiveness, and so we simplify other features of the model. The utility specification is modified to eliminate dependence on progeny, and to impose inter-temporal additivity and that instantaneous felicity is captured by the log form. The goods and human capital production technologies of the young are now taken to be linear and independent of lagged aggregate human capital. And, instead of assuming full rationality, agent behavior is modeled using point expectations. These simplifications lead to the following representative agent's problem:

$$\begin{aligned}
\max \quad & \log c_{1t} + \log c_{2t+1}^e \\
\text{s.t} \quad & c_{1t} = n_t \\
& h_t = 1 - n_t \\
& c_{2t+1}^e = F(h_t, A_{t+1}^e)
\end{aligned}$$

where A_{t+1}^e is the value of labor effectiveness in time $t + 1$ expected by the young agent in time t .

Because our focus is expectations formation and coordination, some discussion of our modeling technique is warranted. We think of A_{t+1} as capturing the uncertain benefit to human capital accumulation: the young know that their future consumption, and therefore well-being, somehow depends on their education level h ; however, some aspects of this dependence are unknown: we model these unknown aspects by assuming agents do not know the value of A_{t+1} when making decisions; instead, they

forecast the value of A_{t+1} and make decisions based on this forecast. In a rational equilibrium we require that their forecasts have no systematic errors.

For the remainder of the paper, we assume the behavior of the young agent is determined by the interior first order condition, which is given by

$$(6) \quad \frac{1}{1-h_t} = \frac{F_h(h_t, A_{t+1}^e)}{F(h_t, A_{t+1}^e)}.$$

In the Appendix, we examine with care the assumptions necessary to guarantee that the young agent chooses $h \in (0, 1)$. Equation (6) implicitly defines the function $h_t(A_{t+1}^e)$, which characterizes the young agent's behavior in terms of expectations.

To define equilibrium, we simplify the transition dynamics of labor effectiveness by eliminating lagged dependence and productivity shocks:

$$(7) \quad A_{t+1} = H_t + \varepsilon_{t+1},$$

where ε_t is a zero mean i.i.d. process with sufficiently small support.⁵ As suggested above, we think of ε_t as capturing the potentially stochastic (and hence uncertain) effect of aggregate human capital on labor productivity. Imposing our rationality assumption yields $A_{t+1}^e = H_t$. Thus we have the following definition:

Definition 1 *Given the process ε_t , a rational equilibrium of the stylized model is a*

⁵The need for shocks with small support is discussed in the Appendix.

collection of processes $\{A_t, H_t\}$ satisfying

$$(8) \quad \begin{aligned} \frac{1}{1 - H_t} &= \frac{F_h(H_t, H_t)}{F(H_t, H_t)} \\ A_{t+1} &= H_t + \varepsilon_{t+1}. \end{aligned}$$

Multiple equilibria obtain precisely when (8) has multiple solutions, and this possibility arises when F exhibits increasing social returns. To provide a laboratory for analysis and to gain insight into the existence of multiple equilibria, we specify a functional form for F . Our specification is guided by the following desired characteristics:

1. $F_h > 0$: individual human capital is a productive input;
2. $F_{hh} < 0$: diminishing marginal returns to individual human capital;
3. $F_A > 0$: aggregate human capital is a productive input;
4. $F_{hA} > 0$: the marginal productivity of individual human capital is benefited by increased aggregate human capital.

To capture these characteristics in a flexible function form, as well as to facilitate analysis, we take F to be given by

$$(9) \quad F(h, A) = \Gamma(A)(h - \bar{h})^{\alpha(A)},$$

where Γ and α are suitably differentiable, non-negative functions, defined on $\mathbb{R}_+$ (non-negative reals), and $\bar{h}, \alpha(A) \in (0, 1)$. The scalar $\bar{h}$ captures the minimum education necessary to access the advanced technology. A detailed analysis of the production

function, as well as a proof that it has the requisite properties, is contained in the Appendix.⁶

The function Γ captures the total factor productivity effect of aggregate human capital accumulation. Because of our log utility specification, Γ plays no decision-making role, see (6): it is used here to guarantee F has the properties listed above while allowing us to separate level effects from relative marginal effects. On the other hand, α plays a central role in our analysis. To interpret α , mentally set $\bar{h} = 0$; then the function α captures the output elasticity of individual human capital. We think of α as increasing in A ; thus a higher aggregate human capital stock increases the responsiveness of output to individual education levels.

Combining the functional form (9) with the equilibrium condition (8) yields


$$(10) \quad \alpha(H) = \frac{H - \bar{h}}{1 - H},$$

where we have removed time-subscripts because of condition's static nature. We conclude that if H satisfies (10), then $H_t = H$, $A_{t+1} = H + \varepsilon_{t+1}$ is rational equilibrium. Multiple equilibria arise when (10) has multiple solutions: see Figure 1 below.

In Figure 1, we depict α as increasing, but with varying rate. This particular shape for α may reflect a threshold effect: for low values of aggregate human capital, the elasticity is low and essentially constant, but as H crosses a threshold level, representing, say, a certain percentage of the population having a high school education,

⁶Also, it is shown in the Appendix that the young will choose $h > \bar{h}$ so that $F(h, A)$ is well-defined.

Figure 1: Multiple steady states



the elasticity jumps sharply, before again leveling off.

3.3 Learning in the stylized model

The attainment of a given equilibrium in the stylized model presented in Section 3.2 (or in the general model of Section 3.1) requires the coordination of private agents' expectations; yet the model is silent both on how this coordination is achieved and, in case of multiple equilibria, which outcome prevails. To address these concerns, we explicitly model expectations formation by taking our cue from the adaptive learning literature: see Evans and Honkapohja (2001) for a general reference. This literature dictates that we back away from the assumption that agents are omniscient, forward-looking actors; instead, we model them as using simple forecasting rules to form

expectations. The agents' expectations then predicate their actions, which, in turn generate new data; agents use these data to update their forecasting models and the process repeats. In this way, the model is coupled with a learning dynamic which may then be analyzed to assess transitions to equilibria as well as selection issues. Importantly, expectations coordination is neither assumed nor required.

Development of a learning structure begins with the identification of observables and the specification of the forecasting model. Recall that we view A_{t+1} as capturing the uncertain component of the returns to education; and the young are required to forecast this component when making decisions. We take A_t to be known to the young at time t : the young observe the returns of the old.⁷ Because rational equilibria in our model are “noisy steady states” (they take the form $A_t = H + \varepsilon_t$), the natural forecasting model is regression on a constant, that is, a simple average. To impart some flexibility, we allow agents to use a weighted average, which we may write recursively as

$$(11) \quad A_{t+1}^e = A_t^e + \gamma_t(A_t - A_t^e).$$

Here γ_t is the “gain” of the learning algorithm (11), and measures how seriously agents take new data. If $\gamma_t = \frac{1}{t}$ then (11) reduces to the recursive least squares estimator, which, in our case, is the sample mean; if $\gamma_t = \gamma > 0$ is a small constant,

⁷The value of A_t is not, in fact, the “returns of the old:” indeed, the marginal return of an additional unit of human capital investment in time $t - 1$, that is, the marginal return of the old in time t , is given by $F_h(A_t, h_{t-1})$. We focus on forecasting A_t , and we interpret this forecasting behavior as learning about returns to education, because A_t is the unique aggregate endogenous variable affecting the marginal returns to education. We could instead have assumed that agents explicitly forecast the marginal return to education and condition their behavior accordingly; however, this would complicate matters considerably, and we anticipate that because the relationship between F_h and A is exogenous, the qualitative features of our results would be unaffected.

then (11) is referred to as a “constant gain learning algorithm,” and represents geometric discounting of past data. We interpret (11) as capturing the following mental reasoning: young agents are endowed with society’s view (forecast) of the returns to education (A_t^e) and, in addition, make their own individual observations (A_t); they combine this information using (11) to obtain their forecasts (A_{t+1}^e), and then make decisions based on these forecasts.

Given A_{t+1}^e , agent behavior is dictated by $h(A_{t+1}^e)$ through condition (6).⁸ This yields

$$h(A_{t+1}^e) = \frac{\alpha(A_{t+1}^e) + \bar{h}}{1 + \alpha(A_{t+1}^e)}.$$

The assumption that all agents form expectations using the mechanism (11) implies that $H_t = h(A_{t+1}^e)$.⁹ Noting that

$$(12) \quad A_{t+1} = h(A_{t+1}^e) + \varepsilon_{t+1},$$

we may rewrite (11) as

$$(13) \quad A_{t+1}^e = A_t^e + \gamma_t(h(A_t^e) - A_t^e + \varepsilon_t).$$

Equation (13) determines the evolution of the economy and allows for the following definition:

Definition 2 *Given the process ε_t and initial expectations A_0^e , a learning equilibrium*

⁸A detailed analysis of agent behavior given arbitrary expectations A^e is provided in the Appendix.

⁹That all agents update their forecasts in the same manner is arguably a form of coordination; the literature on educative learning provides a mechanism allowing for this homogeneity assumption to be relaxed: see Evans and Honkapohja (2001) for further details.

in the stylized model is a collection of sequences $\{H_t, A_t, A_t^e\}$ satisfying (12), (13) and $H_t = h(A_{t+1}^e)$.

4 Results

The rich dynamic structure of the model's learning equilibria allows a number of different issues to be addressed. First, and perhaps most urgently, what is the relationship between rational equilibria and learning equilibria? If learning equilibria bear no resemblance to rational equilibria then we are led to question the relevance of this exercise; however if learning equilibria evolve to approximate rational equilibria in some natural sense then we may view the learning mechanism as an equilibrium selection device, providing justification for the rational requirement of expectations coordination. Second, and perhaps more ambitiously, the dynamic learning structure provides a natural environment for simple policy experiments, thus allowing us to analyze with some rigor which types of exogenous institutional actions might result in welfare gains. We address these issues in turn.

4.1 Learning to be rational

The number and nature of the stylized model's rational equilibria are determined as solutions to (8), and are depicted graphically in Figure 1, whereas learning equilibria are characterized by the dynamic system (13). Can we assess the relationship between these creatures? The problem is an a-priori difficult one: the rational equilibria are straightforward to understand and analyze; but the learning equilibria are processes

representing solutions to non-linear stochastic recursive algorithms (SRAs), and, generally speaking, SRAs are quite difficult to understand. Fortunately, the forms of the SRAs relevant for adaptive learning in economic models are amenable to asymptotic and short sequence analysis using the theory of Ljung (1977): again, see Evans and Honkapohja (2001) for applications to macroeconomics.

Our first result provides the relationship between rational and learning equilibria in case of decreasing gain.

Proposition 3 *Let H represent a rational equilibrium, that is, H solves (8). Let $\gamma_t = \frac{1}{t}$. If*

$$(14) \quad \alpha'(H) < \frac{(1 + \alpha(H))^2}{1 + h}$$

then locally $A_t^e \rightarrow H$ with probability one.

A formal statement of this result, together with its proof, are contained in the Appendix.

Proposition 3 provides a simple stability condition which, when satisfied, imparts the asymptotic equivalence of rational and learning equilibria: if the rational equilibrium satisfies (14), if initial beliefs are within the rational equilibrium's basin of attraction, and if agents form expectations by simply averaging their initial beliefs with realized data, then the learning equilibrium will converge in a natural sense to the rational equilibrium. If the stability condition is satisfied, we say that the associated rational equilibrium is stable. Somewhat remarkably, the stability condition

(14) is directly related to the graphical depiction of rational equilibria:¹⁰ stability obtains precisely when the graph of α crosses the upward-sloping curve from above. Note that, in Figure 1, both the upper and lower rational equilibria are stable: in each case, agents will learn to be rational.

For stable rational equilibria, the coordination problem is solved: agents are not required to carefully and accurately forecast the behavior of other agents (which would require forecasting the forecasts of other agents, etc); instead, they simply view past data and compute the mean. Furthermore, Proposition 3 establishes equilibrium stability as a selection mechanism in two distinct senses: first, unstable rational equilibria, such as the middle equilibrium depicted in Figure 1, should be discarded – agents cannot learn to coordinate on them; and second, given two distinct stable rational equilibria, selection by learning agents will depend on initial beliefs. As we will see below, this beliefs’ dependency may be exploited by policy-makers to achieve welfare improvement.

Our second result characterizes the finite horizon behavior in case of constant gain. The result is more technical in nature, but is useful in that it allows for precise statements about transition dynamics and allows for a good approximation to asymptotic behavior. In the Appendix it is shown that a stable rational equilibrium H corresponds to a Lyapunov stable fixed point of the following differential equation:

$$(15) \quad \frac{dA^e}{d\tau} = h(A^e) - A^e.$$

Let $\tilde{A}_t^e(A_0^e)$ be the solution to (15) corresponding to the initial condition A_0^e . The fol-

¹⁰This is remarkable because the graphical depiction of rational equilibria has no a-priori relation to the learning model’s dynamics.

lowing result provides a relationship between this solution and the associated learning equilibrium.

Proposition 4 *Let H represent a stable rational equilibrium and let A_0^e be initial beliefs which are near H . Fix T , and let $\{A_t^e(A_0^e, \gamma)\}_{t=0}^T$ be the finite-length time path of beliefs associated to the constant gain γ and initial beliefs A_0^e , as determined by the recursion (11). Then $\{A_t^e(A_0^e, \gamma)\}_{t=0}^T$ converges weakly to a random process with mean $\{\tilde{A}_t^e(A_0^e)\}_{t=0}^T$ and vanishing variance.*

Again, a precise statement of the result, together with a formal proof is given in the Appendix.¹¹

Asymptotic results in case of constant gain are not available because of indeterminacy issues; however, this result is nearly as good. Intuitively, Proposition 4 states that for any time horizon T , the expected time path of beliefs will be well-approximated by the solution to the differential equation (15), provided that the gain is small enough; and, since this differential equation locally directs paths of beliefs to a stable rational equilibrium, it follows that, locally, agents beliefs will get near the rational equilibrium and stay near it for a long time. The implication of this Proposition – and this is what we need for the paper – is that for small gain γ and finite horizon T , the solution to the differential equation (15) provides a good approximation to the expected time path of the recursive algorithm capturing the evolution of households’ beliefs: so, instead of studying the recursive algorithm (which is hard), we study the differential equation (which is easy).

¹¹As the skeptical reader might have noticed, the formal statement requires scaling time so that the continuous process $\tilde{A}_t^e(A_0^e)$ and the discrete process A_t^e are comparable: see the Appendix for details.

4.2 Implications and policy analysis

While the interaction between economic agents and policy-makers has not been formally introduced in our model, the rich structure of learning dynamics allow us to interpret exogenous changes in certain variables and parameters as directed by policy. We exploit this by exploring our model's explanations of two stylized facts of developing economies: the existence of poverty traps and the systematic errors made by youth when forecasting the returns to education.


Propositions 3 and 4, together with Figure 1 provide a simple explanation of poverty traps: agglomeration effects impart equilibrium multiplicity, and the economy is stuck in a rational equilibrium characterized by coordination failure. While this story is well-understood, the novelty of our treatment is that now we know why agents are stuck in the low human capital equilibrium: they were initially skeptical of the value of education and their simple adaptive algorithm, together with the associated forecast-based behaviors, reinforced and coordinated their skepticism.

The learning dynamics (13) suggests a mechanism through which policy-makers might dislodge the economy and place it on a welfare improving path. Referring again to Figure 1, suppose that $A_t^e = H_L$. Then $h(A_t^e) = H_L$ and so, on average, $A_{t+1}^e = H_L$: the economy is stuck near the low human capital rational equilibrium. If, through advertising or public outreach, the government is able to influence the beliefs of the youth, that is, if the policy makers can move A_t^e , then they may be able to place the economy on a learning equilibrium time path that will converge to the high human capital rational equilibrium.

As an example of this type of policy intervention, consider Figure 2. Here, we

use the same specification of α used to produce Figure 1. Figure 2 presents the solutions to the differential equation (21) for varying initial conditions. Appealing to Proposition 4, we know that these solutions approximate the expected time paths of agents' beliefs.

Figure 2: Varying initial beliefs



For low initial beliefs, the time paths converge to the low human capital rational equilibrium: skepticism is reinforced and a poverty trap obtains. However, for higher initial beliefs, the high human capital rational equilibrium is approached. As suggested by this figure, policy induced shifts in beliefs may be welfare improving.

Turning to the presence of systematic bias in forecasts of returns to education, we may examine the findings of Jensen (2010) that while observed returns to education are quite high, perceptions of returns remain very low. Our model provides precisely

the environment to explain this feature of the data. Note that for given beliefs A_t^e , the value $h(A_t^e)$ captures average realized returns. We have the following result:

Proposition 5 *If A_t^e is in the basin of attraction of a stable rational equilibrium H , and if $A_t^e < H$ then $h(A_t^e) > A_t^e$.*

This proposition follows immediately from the fact that $\alpha'(A^e) > 0 \implies h'(A^e)$. Thus when perceptions of returns are low, our model predicts that realized returns will be higher than perceptions.

A natural question remains: why would perceptions remain low for any extended period of time? Our model's explanation lies in the specification and interpretation of the gain parameter γ_t . This parameter measures the willingness of young agents to modify their beliefs; and in this sense it can be viewed as measuring the strength of the agents' priors: small values of the gain imply that the agent has strong priors and is thus unwilling to quickly alter his beliefs; large gain values imply weak conviction and the willingness to let new data dictate forecasts.¹²


Alternatively, the gain parameter can be thought of as the confidence agents have in the information they receive. If, for example, information about the true returns to education comes in the form of anecdotes about the experiences of relatives, friends and acquaintances, the agent may be hesitant to put too much weight on this information. In this sense the value of the gain parameter is a measure of the perceived quality of the information itself.¹³

¹²These interpretations hold both for decreasing gain and constant gain: Proposition 3 holds if $\gamma_t = \frac{1}{t}$ is replaced by $\gamma_t = \frac{\delta}{t}$ for any $\delta > 0$.

¹³This interpretation of the gain parameter is admittedly looser than the "strength of priors" interpretation provided just above: indeed, as modeled, the quality of the information received by

We conclude that Jensen’s results are explained by either assuming agents have strong priors, or that they don’t fully trust new information, and thus they choose small values of the gain for the updating algorithm.

Figure 3: Average dynamics for different gains



Our explanation of Jensen’s results also provides a policy prescription: policy-makers should embrace actions directed toward weakening priors and increasing the flow and quality of information - thereby increasing the rate at which agents update beliefs. Governmental strategies encouraging youth to take seriously the effect of human capital investment on productivity – i.e. place more weight on observed values of returns to education – may speed up the transition to a high human capital rational

agents – the realization of past values of A – is perfect. One can imagine modifications to the model which might better incorporate the notion of information quality: for example, agents might receive a noisy (even biased) signal of A and rely on signal extraction to form their forecasts. While we have not investigate alternatives, it is our intuition that our results would not be qualitatively altered.

equilibrium, and therefore be welfare improving. To explore this possibility via our example, we simulate constant gain learning. Figure 3 presents the average dynamics across 100 simulations of the model for three different values of the gain. Each simulation was provided the same initial condition, which was chosen to be below the high human capital rational equilibrium but within its basin of attraction. The solid horizontal line at $\approx .577$ identifies the rational equilibrium H_H . We see that, on average, weaker priors, corresponding to higher gains, yield faster convergence.

5 Conclusion

If households have inaccurate perceptions about the true returns to education, and these perceptions lead to underinvestment in human capital, then it becomes of critical importance to understand the mechanisms through which learning about the true returns takes place, and how this learning affects economic growth. This paper explored these issues in a dynamic model that explicitly incorporates learning. It is found that multiple equilibria exist and that misperceptions about the true return to education can lead to the economy becoming stuck in an equilibrium characterized by low human capital investment. Critically, if young agents place low weight on new information received about the true returns to education, they can substantially underinvest in their own human capital.

By understanding the nature of the learning process, insights about effective policy responses are gained. By actively seeking to get young agents to believe new information about the true returns to education through programs designed to undermine closely held predudices about educational investments and to increase the

flow and quality of information about the true returns to education, governments can speed up the learning process and achieve faster and more robust growth.

6 Appendix

Here we justify the various assumptions made in the body of the paper, and provide the technical proofs. To facilitate exposition, this appendix is organized into two subsections: in the first, we focus on modeling assumptions and agent behavior, and we identify with greater care the associated rational equilibria; in the second we characterize learning equilibria more formally and provide the proofs to the propositions.

6.1 The model under rationality

Set $\mathbb{R}_+ = [0, \infty)$. Let $\Gamma : \mathbb{R}_+ \rightarrow (0, \infty)$ be differentiable and increasing. Let $0 < \alpha_L < \alpha_H < 1$ and let $\alpha : \mathbb{R}_+ \rightarrow [\alpha_L, \alpha_H]$ be continuously differentiable and increasing. Let $\bar{h} \in [0, 1)$. The production function $F : [0, 1] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is given by

$$F(h, A^e) = \begin{cases} \Gamma(A^e)(h - \bar{h})^{\alpha(A^e)} & h \geq \bar{h}, A^e \geq \bar{h} \\ 0 & \text{else} \end{cases}$$

Here $\Gamma(A^e)$ is a scalar which can be thought of as capturing total factor productivity, and α is the potentially non-linear response of elasticity to the expected agglomeration measure, which is assumed bounded between α_L and α_H . Finally, $\bar{h}$ is the minimum education level needed to access the sophisticated technology F . Set

$$\hat{h} = \frac{\alpha_L + \bar{h}}{\alpha_L + 1} > \bar{h}.$$

We make the following assumption on Γ :

$$(16) \quad \frac{\Gamma'(A^e)}{\Gamma(A^e)} > -\log(\hat{h} - \bar{h})\alpha'(A^e).$$

We have the following result:

Lemma 6 *There exists a scalable family of functions increasing differentiable Γ satisfying (16). Further, if Γ satisfies (16), and if $h \in (\hat{h}, 1)$ and $A^e \geq \bar{h}$ then*

1. $F_h(h, A^e) > 0$.
2. $F_{hh}(h, A^e) < 0$.
3. $F_{A^e}(h, A^e) > 0$.
4. $F_{A^eh}(h, A^e) > 0$.

Proof. Let $\delta > 0$ and consider the differential equation

$$(17) \quad \Gamma' = \Gamma \cdot \left(\left| \log(\hat{h} - \bar{h}) \right| \alpha'(A^e) + \delta \right).$$

This ode is separable, and so may be solved using standard techniques to obtain

$$\Gamma(A^e) = \lambda e^{(|\log(\hat{h} - \bar{h})| \alpha(A^e) + \delta A^e)},$$

where λ is any positive scalar. Now simply notice that any solution to (17) also satisfies (16).

Showing that F satisfies items (1) and (2) is trivial. To see that F also satisfies (3) and (4), notice that

$$\begin{aligned} F_{A^e}(h, A^e) &= (h - \bar{h})^{\alpha(A^e)} (\Gamma'(A^e) + \Gamma(A^e)\alpha'(A^e) \log(h - \bar{h})) \\ F_{A^e h}(h, A^e) &= \alpha(A^e)(h - \bar{h})^{\alpha(A^e)-1} \left(\Gamma(A^e) + \Gamma'(A^e) \log(\hat{h} - \bar{h})\alpha'(A^e) \right) \\ &\quad + \alpha'(A^e)(h - \bar{h})^{\alpha(A^e)-1} \Gamma(A^e). \end{aligned}$$

The signs of both of these partials are positive provided that

$$\Gamma'(A^e) + \Gamma(A^e)\alpha'(A^e) \log(h - \bar{h}) > 0,$$

which is equivalent to (16). ■

We conclude that provided we restrict attention to $h \in (\hat{h}, 1)$ and $A^e \geq \bar{h}$, F satisfies the desired properties.

Now that the production function has been defined with care, we turn to agent behavior and model equilibrium. For completeness and consistency, we assume that besides having access to the technology F , old agents may also use the same primitive technology they used when young. We make the additional assumption that using F is a full-time job: if the old agent directs any labor to the primitive technology, then he relinquishes access to the advanced technology. Since we assume that the agent must base all of his decisions on expectations (i.e. he must decide whether to use the advanced technology before knowing the value of A) we know that if the agent plans to use the primitive technology in the second period, he will necessarily choose $h = 0$. Therefore, since $F(0, A^e) = 0$, we may model agent behavior as follows:

let $\delta : [0, 1] \rightarrow \{0, 1\}$ be defined by $\delta(1) = 1$ and $\delta(h) = 0$ for $0 \leq h < 1$. The representative agent solves

$$\begin{aligned} \max_{0 \leq h \leq 1} \quad & \log(c_1) + \log(c_2^e) \\ & c_1 = 1 - h \\ & c_2^e = F(h, A^e) + \delta(1 - h). \end{aligned}$$

As this formulation of the agent's problem makes clear, his choice of education, h , depends on his expectations A^e . Notice that since the marginal utility of consumption goes to infinity as consumption goes to zero, it follows that we can rule out the corner solution $h = 1$. We exogenously impose that $A^e \geq 0$. There are two cases:

Case 1: $A^e < \bar{h}$. If $A^e < \bar{h}$ then for any h , we have that $F(h, A^e) = 0$. It follows that agent the representative agent sets $h = 0$.

Case 2: $A^e \geq \bar{h}$. Now the agent has two choices. Either he can set $h = 0$, and thus receive total utility equal to zero, or the agent can choose some $h > \bar{h}$. Since we already know that $h < 1$, it follows that if $h > \bar{h}$ then h will be chosen to satisfy the interior FOC given by

$$(18) \quad h(A^e) = \frac{\alpha(A^e) + \bar{h}}{\alpha(A^e) + 1} \in (\hat{h}, 1).$$

We conclude that the representative agent chooses $h = h(A^e)$ if and only if $A^e \geq \bar{h}$ and the utility received by choosing $h = h(A^e)$ is greater than zero. But

$$(19) \quad \begin{aligned} 0 &< \log(1 - h(A^e)) + \log \Gamma(A^e) + \alpha(A^e) \log(h(A^e) - \bar{h}) \Leftarrow \\ 0 &< \log\left(\frac{1 - \bar{h}}{1 + \alpha_H}\right) + \Gamma(\bar{h}) + \alpha_L \log\left(\frac{\alpha_L \cdot (1 - \bar{h})}{1 + \alpha_L}\right), \end{aligned}$$

where the implication uses the fact that $\alpha \in [\alpha_L, \alpha_H]$ and $\Gamma' > 0$. Since Γ is scalable, we may assume that (19) holds. Thus, we assume here, and in the text, that if $A^e \geq \bar{h}$ then $h = h(A^e)$ as given by (18). We summarize these results in the following Lemma:

Lemma 7 *Assume (16) and (19) hold, and let $A^e \geq 0$. Representative agent behavior is given by*

$$h(A^e) = \begin{cases} \frac{\alpha(A^e) + \bar{h}}{\alpha(A^e) + 1} & A^e \geq \bar{h} \\ 0 & \text{else} \end{cases}$$

In the text we focused exclusively on the case $A^e \geq \bar{h}$. Note that if $A^e \geq \bar{h}$ then $h(A^e) \geq \hat{h}$; thus, whenever the representative agent is accessing the advanced technology, the associated production function has the desired properties as listed in Lemma 6.

Now recall the following definition: a *rational equilibrium* is any collection of processes $\{A_t, A_t^e, H_t\}$ so that

1. $H_t = h(A_t^e)$
2. $A_t^e = H_t$

3.

$$A_t = \begin{cases} A_t^e + \varepsilon_t & A_t^e \geq \bar{h} \\ 0 & \text{else} \end{cases}$$

where our definition of A_t has been modified to allow for $H_t = 0$. Also, to guarantee that $A_t \geq 0$, we assume that ε_t has small support: $\varepsilon_t \in (-\varepsilon, \varepsilon)$ where $0 < \varepsilon < \bar{h}$.

Note first that $H_t = 0$ identifies an “autarky” rational equilibrium: if an agent expects that no one is accumulating human capital then his best response is also to forgo education and simply adopt the primitive technology in both periods. In the text, we focused on equilibria of the form $H_t = H_{ss}$ where $H_{ss} = h(H_{ss})$. Notice that, in this case, $A^e = H_{ss} \geq \bar{h}$ so that agent behavior is captured by their interior FOC, and $h = H_{ss} \geq \hat{h}$: thus the production function has the desired properties.

6.2 Learning equilibria

In the text, we identified learning equilibria with time paths of expectations generated by a simple recursive algorithm capturing the weighted average of past observations. As is standard in the learning literature, asymptotic results yielding almost sure convergence require that the learning algorithm be modified to incorporate a projection facility. To develop the appropriate notions, recall that a learning equilibrium is characterized by the following stochastic recursive algorithm:

$$(20) \quad A_{t+1}^e = A_t^e + \gamma_t(h(A_t^e) - A_t^e + \varepsilon_t).$$

To study the asymptotic behavior of (20), we appeal to Ljung's theory, as developed in Evans and Honkapohja (2001). The idea is to approximate the mean time path of A_t^e with the solution to a differential equation. To this end, let

$$g(A^e) = E(h(A^e) - A^e + \varepsilon_t) = h(A^e) - A^e,$$

and notice that g is continuously differentiable on $(\hat{h}, 1)$. Ljung's theory provides conditions under which the differential equation

$$(21) \quad \frac{dA^e}{d\tau} = g(A^e)$$

offers locally a good approximation to the expected behavior of the stochastic process A_t^e . To make this statement precise, let H_{ss} identify a rational equilibria, and notice that H_{ss} represents a rest point of the ode (21): $g(H_{ss}) = 0$. Assume that $g'(H_{ss}) < 1$, that is, H_{ss} is a Lyapunov-stable rest point of (21). Then there exists an open set D containing H_{ss} and a twice continuously differentiable Lyapunov function $f : D \rightarrow \mathbb{R}_+$ so that, among other properties, $f(H_{ss}) = 0$ and whenever $H_{ss} \neq A^e \in D$ it follows that $f(A^e) > 0$ and $g(A^e) \neq 0$.

The function f identifies the local nature of the approximation, and tells us how to construct the appropriate projection facility: for $c > 0$, denote by $S(c)$ the lower contour set for f :

$$S(c) = \{A^e \in D : f(A^e) \leq c\}.$$

Notice that there exists c so that $c' \leq c$ implies that $S(c')$ is compact and $\text{int}(S(c')) \neq \emptyset$. Because α is continuously differentiable, it follows that h is continuously differen-

tible, and so locally Lipschitz, which is a property necessary for our stability result. Therefore, we may choose $c > 0$ so that h is Lipschitz on $S(c)$. Fix $0 < c_2 < c_1 < c$, and augment the recursive algorithm (20) as follows:

$$(22) \quad A_{t+1}^e = \begin{cases} A_t^e + \gamma_t(h(A_t^e) - A_t^e + \varepsilon_t) & A_t^e + \gamma_t(h(A_t^e) - A_t^e + \varepsilon_t) \in S(c_1) \\ \hat{A} \in S(c_2) & \text{else} \end{cases}$$

The modified learning algorithm (22) incorporates the projection facility: any time agents' expectations deviate too sharply from rationality – that is, fall outside the set $S(c_1)$ – they are placed back to an arbitrary point within the set $S(c_2)$. The employment of a projection facility is standard in the learning literature – see Evans and Honkapohja (2001) for discussion and details – and is necessary to obtain almost sure convergence: it is always possible that an unusual sequences of shocks will push expectations so far away from the steady state that return is not possible. On the other hand, the learning algorithm can be further modified so that the probability that the projection facility is used – that is, the probability that $A_t^e + \gamma_t(h(A_t^e) - A_t^e) \notin S(c_1)$ – is arbitrarily small. We are now ready to state our result:

Proposition 8 *Assume $\sum \gamma_t = \infty$ and $\sum \gamma_t^2 < \infty$. Suppose H_{ss} is a Lyapunov stable rest point of (21), that is,*

$$\alpha'(H_{ss}) < \frac{(1 + \alpha(H_{ss}))^2}{1 + h}.$$

Then A_t^e converges to H_{ss} almost surely.

Given our construction of the algorithm (22), and the assumptions imposed on

α , the proof of this proposition now follows immediately from Corollary 6.8, on page 136 of Evans and Honkapohja (2001).

Proposition 8 addresses the decreasing gain case: $\gamma_t \rightarrow 0$. Now we analyze the behavior of the learning algorithm under constant gain:

$$(23) \quad A_{t+1}^e = A_t^e + \gamma(h(A_t^e) - A_t^e + \varepsilon_t),$$

where $\gamma > 0$ is small. Since $A_t^e = G(A_{t-1}^e) + \gamma\varepsilon_t$ for appropriate G , it follows that A_t^e will never “settle down,” that is, converge to a degenerate random variable. However, results from Benveniste et al (1990) provide for short sample analysis. In particular, the results provide the precise sense in which the solution to the ode (21) approximates the expected time path of A_t^e . To compare solutions to (21) with realizations of the stochastic process A_t^e , two adjustments must be made: A_t^e must be defined for all real t , not just for integer values of t ; and the time-scale for A_t^e must be adjusted. To this end, let $A_n^e(A_0^e)$ be defined by (23) with t replaced by n and with initial condition A_0^e , and set

$$\hat{A}_t^e(\gamma, A_0^e) = A_n^e(A_0^e) \text{ whenever } \gamma n \leq t \leq \gamma(n+1).$$

So $\hat{A}_t^e(\gamma, A_0^e)$ is a step function with heights defined by A_n^e and with bins of width γ identifying the adjusted time scale.

Now let $\tilde{A}_t^e(A_0^e)$ be the solution to the ode (21) corresponding to the initial condition A_0^e . Let $S(c)$ as above, with $\tilde{A}_t^e(A_0^e) \in S(c)$ for $0 \leq t \leq T$. Finally, assume α is twice continuously differentiable. We have the following result:

Proposition 9 *There is a continuous time stochastic process y_t with $y(0) = 0$ so that*

the continuous time process

$$\gamma^{-1/2} \left(\hat{A}_t^e(\gamma, A_0^e) - \tilde{A}_t^e(A_0^e) \right)$$

converges weakly to y_t for $0 \leq t \leq T$. Furthermore, $Ey_t = 0$.¹⁴

This proposition follows immediately from our assumptions on α and from Proposition 7.8, Evans and Honkapohja (2001) page 163. Intuitively, we conclude that for small γ and finite horizons, the solution to the ode (21) well-approximates the expected behavior of the process A_t^e , provided that the time scale is appropriately adjusted.

¹⁴The process y_t is determined as the solution to the stochastic differential equation

$$dy(t) = h'(\tilde{A}_t^e(A_0^e))y(t)dt + \sigma_\varepsilon dw(t),$$

where $w(t)$ is a standard Wiener process.

References

- [1] Acemoglu, Daron, and Joshua Angrist (2000) “How Large Are the Social Returns to Education? Evidence from Compulsory Schooling Laws” (pp. 9–59), in Ben Bernanke and Kenneth Rogoff (Eds.), NBER Macroeconomic Annual.
- [2] Aghion, Philippe, and Peter Howitt (1997) *Endogenous Growth Theory*. MIT Press, Cambridge, MA.
- [3] Arifovic, J., J. Bullard, and J. Duffy (1997) “The Transition from Stagnation to Growth: An Adaptive Learning Approach,” *Journal of Economic Growth*, 2, 185–209.
- [4] Azariadis, C., and A. Drazen (1990) “Threshold Externalities in Economic Development,” *The Quarterly Journal of Economics*, 105(2), 501–526.
- [5] Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. (2009) “The Miracle of Microfinance? Evidence from a randomized evaluation,” MIT working paper.
- [6] Barro, Robert (1997). “Human capital and growth in cross-country regressions,” *Swedish Economic Policy Review*, 6(2), pp. 237-277.
- [7] Barro, Robert J., and Jong-Wha Lee (2001). “International data on educational attainment: updates and implications,” *Oxford Economic Papers*, 53(3), pp. 541-563.

- [8] Benhabib, Jess and Spiegel, Mark M. “The Role of Human Capital in Economic Development: Evidence from Aggregate Cross- Country Data,” *Journal of Monetary Economics*, October 1994, 34(2), pp. 143-74.
- [9] Bils, Mark and Peter J. Klenow (2000). “Does schooling cause growth?” *American Economic Review*, December, 90(5), pp. 1160- 1183.
- [10] Chan, Sewin and Ann Huff Stevens (2008). “What You Don’t Know Can’t Help You: Pension Knowledge and Retirement Decision-Making,” *Review of Economics and Statistics*, 90(2), pp. 253-266.
- [11] De la Fuente, Angel, and Rafael Domenech, “Schooling Data, Technical Diffusion, and the Neoclassical Model,” *American Economic Review* P&P 90:5 (May 2001), 323–327.
- [12] de Janvry, Alain, Frederico Finan, Elisabeth Sadoulet and Renos Vakis (2006) “Can conditional cash transfer programs serve as safety nets in keeping children at school and from working when exposed to shocks?” *Journal of Development Economics*, 79(2), pp. 349-373.
- [13] Duflo, Esther (2001). “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 91(4), p. 795-813.
- [14] Duflo, Esther, and Emmanuel Saez (2003). “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment,” *Quarterly Journal of Economics*, 118(3), p. 815-842.

- [15] Duflo, Esther, Michael Kremer and Jonathan Robinson (2004). “Understanding Technology Adoption: Fertilizer in Western Kenya Preliminary Results from Field Experiments.” Mimeo, Harvard University Economics Department.
- [16] Dupas, Pascaline (2006). “Relative Risks and the Market for Sex: Teenagers, Sugar Daddies and HIV in Kenya,” mimeo, Harvard University.
- [17] Emerson, Patrick M., and André Portela Souza (2010). “Is Child Labor Harmful? The Impact of Starting to Work as a Child on Adult Earnings,” Forthcoming, Economic Development and Cultural Change.
- [18] Evans, G. W., and S. Honkapohja (2001) Learning and Expectations in Macroeconomics. Princeton University Press, Princeton and Oxford.
- [19] Evans, G. W., S. Honkapohja, and P. Romer (1998) “Growth Cycles,” American Economic Review, 88(3), 495–515.
- [20] Hall, Robert E., and Charles I. Jones, “Why Do Some Countries Produce So Much More Output per Worker than Others?” Quarterly Journal of Economics 114:1 (February 1999), 83–116.
- [21] Hanushek, Eric A. and Dennis D. Kimko (2000). “Schooling, labor force quality, and the growth of nations,” American Economic Review, 90(5), pp. 1184-1208.
- [22] Iranzo, Susana and Giovanni Peri (2009) “Schooling Externalities, Technology, and Productivity: Theory and Evidence from U.S. States,” The Review of Economics and Statistics, May, 91(2): 420–431

- [23] Jensen, Robert (2010). "The (Perceived) Returns to Education and the Demand for Schooling" *Quarterly Journal of Economics*, 125 (2), p. 515-548.
- [24] Krueger, Alan B. and Mikael Lindahl (2001). "Education for growth: why and for whom?" *Journal of Economic Literature* 39.
- [25] Ljung, Leonard (1977). "Analysis of Recursive Stochastic Algorithms," *IEEE Transactions of Automatic Control*, 22, p 551 - 575.
- [26] Mankiw, Gregory N., David Romer and David N. Weil (1992). "A contribution to the empirics of economic growth." *Quarterly Journal of Economics* 107.
- [27] Mitchell, Olivia (1988). "Worker Knowledge of Pension Provisions," *Journal of Labor Economics*, 6(1), p. 21-39.
- [28] Moretti, Enrico. (2004) "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data," *Journal of Econometrics*, 121, pp. 175-212.
- [29] Psacharopoulos, George and Harry Anthony Patrinos (2004). "Returns to Investment in Education: A Further Update," *Education Economics*, 12(2), p. 111-134.
- [30] Steiger, L. Christina (2009) "Escaping the Poverty Trap: Expectations, Adaptive Learning and Human Capital Accumulation" Mimeo: University of Oregon.
- [31] Stango, Victor and Jonathan Zinman (2007). "Fuzzy Math and Red Ink: When the Opportunity Cost of Consumption is Not What it Seems," mimeo, Dartmouth College.

- [32] Viscusi, W. Kip (1990). “Do Smokers Underestimate Risks?” *Journal of Political Economy*, 98(6), p. 1253-169.