

An agent-based model that relates investment in education to economic prosperity

Yuqing Tang¹, Simon Parsons², and Elizabeth Sklar²

¹ Department of Computer Science
Graduate Center, City University of New York
365, 5th Avenue, New York, NY 10016, USA
ytang@gc.cuny.edu

² Department of Computer & Information Science
Brooklyn College, City University of New York
2900 Bedford Avenue, Brooklyn, NY 11210, USA
{parsons, sklar}@sci.brooklyn.cuny.edu

Abstract. We examine the relationship between the investment that a society makes in education in one generation, and the outcome in terms of the health of the society’s economy in ensuing generations. We have employed multiagent simulation techniques in which equation-based models from the economics literature are used to derive parameterized sets of agent behaviors and environmental characteristics. Agents are divided into three chronological categories: students, adults and pensioners; and each responds to and affects the environment in different ways, in terms of both human and physical capital. We explore the effects of different parameter settings on the education investment of a society and the resulting economic growth.

1 Introduction

We are working towards creating tools that can be used in determining the effects of particular choices in education policy. Our aim is to be able to use such tools to inform the debate about initiatives like the US “No Child Left Behind” Act [5], and illuminate the controversies that such initiatives have created. To this end we have been extracting equations from sets of data related to human education and implementing agent-based models that we can use for experimentation, as shown here, and prediction, as shown in some of our earlier work [7, 8].

Typically, education data is collected in one of two ways. Either very large, aggregate data sets over entire populations (like whole cities, school districts, states or provinces) are gathered or very small, localized experimental samples are collected. In both cases, the data is usually analyzed using standard statistical methods. Often, the most highly publicized statistics are the simplest, for example the mean and standard deviation of standardized test scores in mathematics and language arts. These values are frequently the ones used to make policy decisions. Occasionally, analysis is performed that examines how multiple factors influence each other, such as the relationship between student-teacher ratios and test scores, dollars per student and test scores, or class size and test scores.

In the literature where these types of data is used to construct models, they are typically formulated in traditional terms, as sets of interrelated differential equations. In contrast to such models, commonly called *equation-based models* (EBMs), we are building agent-based models (ABMs) which are designed in terms of a set of autonomous interacting entities. Such models have been successfully used to generate useful predictions about the behavior of populations made up of individuals [6], especially where such individuals make their own decisions about how to act [9]. These types of models can also be used to run experiments that allow us to explore variations on the policies, not only to make predictions but also to gain a deeper understanding of the model and the interplay between sets of variables in the model.

In this paper, we describe some results of our work developing and experimenting with one specific agent-based model, a model derived from an equation-based model published in the economics literature [4]. The model relates the effectiveness of education to economic productivity and money spent on education to the effectiveness of education. It therefore provides a means to tie models like that described in [7, 8] into the wider economic picture. The remainder of this paper is organized as follows. Section 2 describes the model that we have implemented, both the equation-based original and the agent-based model that we derived from it. Section 3 then describes a set of experiments that we performed using the model. Section 4 gives the results and section 5 analyzes them. Section 6 then concludes.

2 The model

The model that we have implemented is taken from the work of Laitner [4]. We were looking for a way to relate the effects of education to the economy, specifically seeking a model that connects an individual’s economic productivity to their level of education. In the long term, we would like to be able to show, for example, that if the government of a state puts emphasis on education and invests in raising a generation of children to be well-educated, then when these children grow up, the economy of that state should prosper. Laitner’s model is a particularly good choice because it is expressed in such a way as to make it possible to assign “credit” to individuals. Rather than having to make broad interpretations of statistical equations and infer behaviors, we can relate aspects of Laitner’s equations directly to agent representations of individual humans within a population. This section describes the main features of Laitner’s model and the main aspects that we adapted to create our agent-based model.

2.1 Laitner’s equation-based model

The setting for the model is a simple economy that has two sectors. Each of these sectors produces one good. The goods are:

- units of education that are used to train individuals in the population; and
- units of a numeraire good.

“Numeraire” is defined as “a basic standard by which values are measured, as gold in the monetary system” [2]. In [4] it is a good that is produced (see below) and then




Fig. 1. Time periods in the model.

traded for things that individuals consume. Presumably these things are produced by a different economy that trades with the one we are studying.

The individuals that inhabit this economy live for three time periods, periods in which they are first students, then adult workers and finally pensioners (see Figure 1). Consider an individual who is a student during period $t - 1$. She spends this period living with her parent³ and studying. Parents provide the numeraire good that supports the child during this period, but the child selects her own units of schooling, borrowing the money to finance this. In the period t , the now adult individual forms her own household, rears a child (paying for the child’s consumption but not the child’s schooling), and chooses how much of the numeraire good, c_t^1 , that she earns during this period will be consumed by the household during the same period. In the period $t + 1$, the individual is a pensioner, and chooses her consumption for that last period, c_{t+1}^2 , from the numeraire good than she has saved. An individual’s utility, u , is:

$$u = (1 - \alpha) \ln(c_t^1) + \alpha \ln(c_{t+1}^2) \quad (1)$$

where $\alpha \in (0, 1)$, and all individuals have the same α .

While working (period t in our example), our individual earns w_t for each unit of human capital that she possesses. Her human capital depends on her *ability* and the amount of schooling she chose as a child. An individual with ability a who purchased e_{t-1} units of education will have human capital:

$$h = a \left(\frac{(e_{t-1})^\gamma}{\gamma} \right) \quad (2)$$

where $\gamma \in (0, 1)$ and all individuals have the same γ . The relationship between e and a allows education to raise human capital, but in a way that is subject to the law of diminishing returns. Ability is randomly assigned from a stationary distribution given in [4].

The model does not include inheritance and bequests, so every individual has to pay for her consumption and education out of what she earns during her working period t . If r_t is the interest rate on savings made during period $t - 1$ and held until period t , every individual is constrained by:

$$c_t^1 + \left(\frac{c_{t+1}^2}{r_{t+1}} \right) + p_{t-1} r_t e_{t-1} \leq w_t a \left(\frac{(e_{t-1})^\gamma}{\gamma} \right) \quad (3)$$

³ In this model, every individual has one child and raises that child alone—in terms of real world accuracy, this is equivalent to the model in [3], implemented as an agent-based model in [8], which assumes every family is a “perfect” nuclear family of a mother and a father and a son and a daughter.

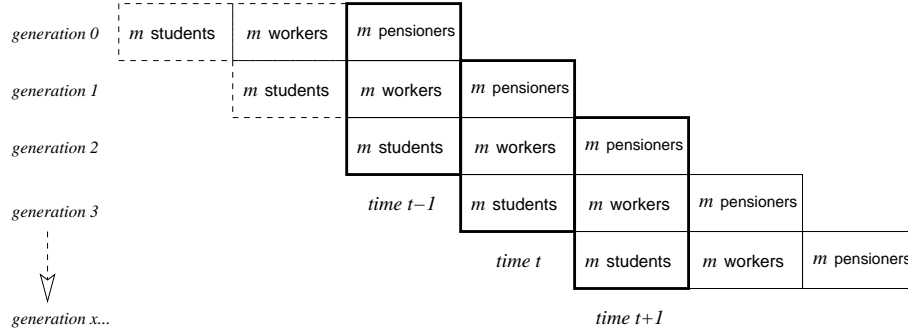


Fig. 2. Flow of generations across time periods in the model.

where p_{t-1} is the cost of a unit of education in period $t - 1$. In other words, the total amount that an individual consumes, including their education, suitably discounted over time, must be less than their earnings.

In every period, m individuals are born, and so there are $3m$ individuals in total in every period in time— m of these are being educated, m are working, and m are retired. This scheme is illustrated in Figure 2. At any given time in the simulation, there are three concurrent “generations” acting in parallel. We initialize the system in generation 2, when there are already 3 populations active.

Considering the sector of the economy that produces the numeraire good, the model follows the law of diminishing returns, so that the output per individual in a given generation is:

$$n = (K_t^n)^{\beta^n} (\lambda_t^n H_t^n)^{1-\beta^n} \quad (4)$$

where $\beta^n \in (0, 1)$, K_t^n is the physical capital the sector has per working individual at time t , and H_t^n is the average human capital per individual in the generation that is currently working. $\lambda^n > 1$ models the tendency of technological change to increase the effect of human capital in the sector of the economy that generates the numeraire good. The other sector of the economy produces education. Here we have:

$$e = (K_t^e)^{\beta^e} (\lambda_t^e H_t^e)^{1-\beta^e} \quad (5)$$

and the model allows for λ^e and β^e to be different from, or the same as, λ^n and β^n , their counterparts in the numeraire sector of the economy. For both the numeraire and education sectors, the assumption is that all physical capital is consumed in a single period, so the numeraire good produced in period t has to equal all consumption plus the physical capital used at time $t + 1$.

Taken together, these equations and the values of the constants provided in [4] provide a rather standard economic model.

2.2 The agent-based model

As described in [1, 6, 8], it is possible to generate agent-based models from equation-based models, by equipping individuals with processes that make decisions in line with the equations. For this model, the decisions faced by an individual are:

1. How much education to purchase⁴.
2. What proportion of wages to save.

The first of these is, in essence, an investment decision. Given (2), for a given level of ability, the more education that an individual purchases, the greater their productivity. All other things being equal—and in particular, the performance of the other individuals in the economy—this greater productivity will turn into greater production of numeraire goods, and, once the cost of education is paid off, greater utility for the individual. Because (2) captures diminishing returns, an individual who spends too much on education will not recoup their investment. The second decision is the same decision faced by anyone who has considered their own retirement—how much of one’s lifetime earnings, minus cost of living and any debts accumulated, should be saved for retirement rather than spent while one is working.

In addition to these decisions, there are decisions faced by the economy as a whole. In the current version of the model, these decisions are taken by a single agent, representing the government. These are:

1. What proportion of numeraire production should be turned into physical investment.
2. What proportion of the physical investment should be put into the numeraire sector rather than the education sector.
3. How to allocate workforce between the education and numeraire sectors.

The first of these decisions can be considered as the effect of taxation—some money is taken out of the income pool and is distributed by the government.

The second decision determines how much of this taxation is invested in education rather than into the production of numeraire goods—this provides the K_t^n and K_t^e in (4) and (5). Since this investment amortises over a single time period—which is reasonable given that each time period represents a third of a lifetime, or approximately 25 years given average life expectancy in the United States—it needs to be renewed in every time period. With this second decision under the control of some central authority, the model looks like a command economy. A topic for future work is the exploration of a more capitalistic model, in which firms compete for individual investors and use that capital to provide physical investment for the numeraire sector while leaving the government to deal with education investment.

Allowing the government to determine directly what proportion of workers to place in each sector also looks like something one would expect to find in a command economy. However, all governments exercise some control over aspects like this through their actions—in most economies the government has a large say in the organization of

⁴ Since every unit of education that an individual undergoes must be paid for from its later wages, it seems appropriate to think of choices about education as a purchase.

the education sector and can encourage people to work in the education sector by, for example, spending money raised through taxation to increase the wages of education workers. We have implemented a number of ways in which each of these decisions can be made. These are explored in the next section, which also gives a description of the experiments we have run.

3 Experiments

The experiments that we will describe here were intended to explore the properties of the model described in the previous section, examining whether our agent-based version could run successfully. That is, whether the decision-making functions with which we had equipped the model were sufficient to create a healthy economy and to approximate the behaviors of real economies. Before we explain how the experiments were run, we need to say what the decision functions are.

In our current implementation, individuals only have one decision to make because the proportion of wages that are saved is kept fixed. The decision they have to take, then, is how much education e_d to purchase, and the implementation provides two ways for individuals to do this:

1. Randomly: e_d is drawn from a normal distribution with mean 13 and standard deviation 2.1. This is the distribution used in [3, 8], and was originally taken from recent US census data.
2. Maximum utility: e_d is chosen by:

$$e_d = \operatorname{argmax}_e \left(w_t \cdot a \cdot \frac{e^\gamma}{\gamma} - p_{t-1} \cdot r_t \cdot e \right) \quad (6)$$

As described above, the government has to decide:

1. $pr_{n,k}$: the proportion of numeraire production to be used as physical investment in the next period;
2. $pr_{k,e}$: the proportion of physical investment to be put into the education sector; and
3. $pr_{h,e}$: the proportion of the working population to move into the education sector.

Our implementation provides several ways that each of these decisions can be made, as a basis for comparison between Laitner's model and the experimental versions we explored. In each case, we used Laitner's settings and also selected one or two other settings for exploration. Our values were chosen to fit within the scale of the selected parameters and so that the overall model would not collapse (through experimentation we found that certain combinations of parameter settings would cause the simulation to produce invalid results).

There are three ways to decide on $pr_{n,k}$:

1. Constant proportion: $pr_{n,k}$ is set to 0.4.
2. Self adjustment: if numeraire production exceeds demand then $pr_{n,k}$ is decreased by 5%, otherwise $pr_{n,k}$ is increased by 1%.
3. Z policy: The policy that [4] uses for this decision.

Laitner's Z policy first computes an intermediate variable Z_t which describes a relationship between physical investment K_t and capital value H_t , then the policy computes an estimate of H_{t+1} from the education students have received at time t and their ability, and then computes the physical investment K_{t+1} from the estimate of Z'_{t+1} . These computations are the following:

$$Z_t = \frac{K_t}{\lambda^t \cdot H_t} \quad (7)$$

$$Z'_{t+1} = \theta \cdot (Z_t)^\xi \quad (8)$$

$$\theta = \left(\frac{\alpha \cdot (1 - \gamma)}{\lambda^{1/(1-\gamma)}} \cdot \frac{1 - \beta}{1 + \gamma \cdot [(1 - \beta)/\beta]} \right)^{1-\gamma} \quad (9)$$

$$\xi = (\beta - 1) \cdot (1 - \gamma) + 1 \quad (10)$$

$$K'_{t+1} = Z'_{t+1} \cdot \lambda^{t+1} \cdot H_{t+1} \quad (11)$$

$$pr_{n,k} = \frac{K_{t+1}}{n} \quad (12)$$

The implementation includes two methods for choosing $pr_{k,e}$:

1. Constant proportion: $pr_{k,e}$ is set to 0.1.
2. Self adjustment: if education production exceeds demand, $pr_{k,e}$ is decreased by 5%, otherwise $pr_{k,e}$ is increased by 1%.

and there are two methods implemented for choosing $pr_{h,e}$:

1. Constant proportion: $pr_{h,e}$ is set to 0.1.
2. Self adjustment: if education production exceeds demand, then $pr_{h,e}$ is decreased by 5%, otherwise $pr_{h,e}$ is increased by 1%.

We ran experiments for each combination of these 24 decision mechanisms, as illustrated in table 1.

4 Results

The results of these experiments, which were run over 100 time periods, or just over 30 generations, are given in Figures 3 and 4, which show, for each economy:

- The average utility of individuals.
- The total earnings of all individuals in the economy, along with their savings for retirement, and the unpaid debt for their education.
- The education that is produced, per individual in the economy, along with the average demand for education.
- The number of numeraire goods that are produced, per individual in the economy, along with the average demand. Demand is measured by the amount of goods and individual chooses to consume, an amount that may not be satisfied if the economy does not produce enough.
- The wage rates, broken down across the numeraire and education sectors.

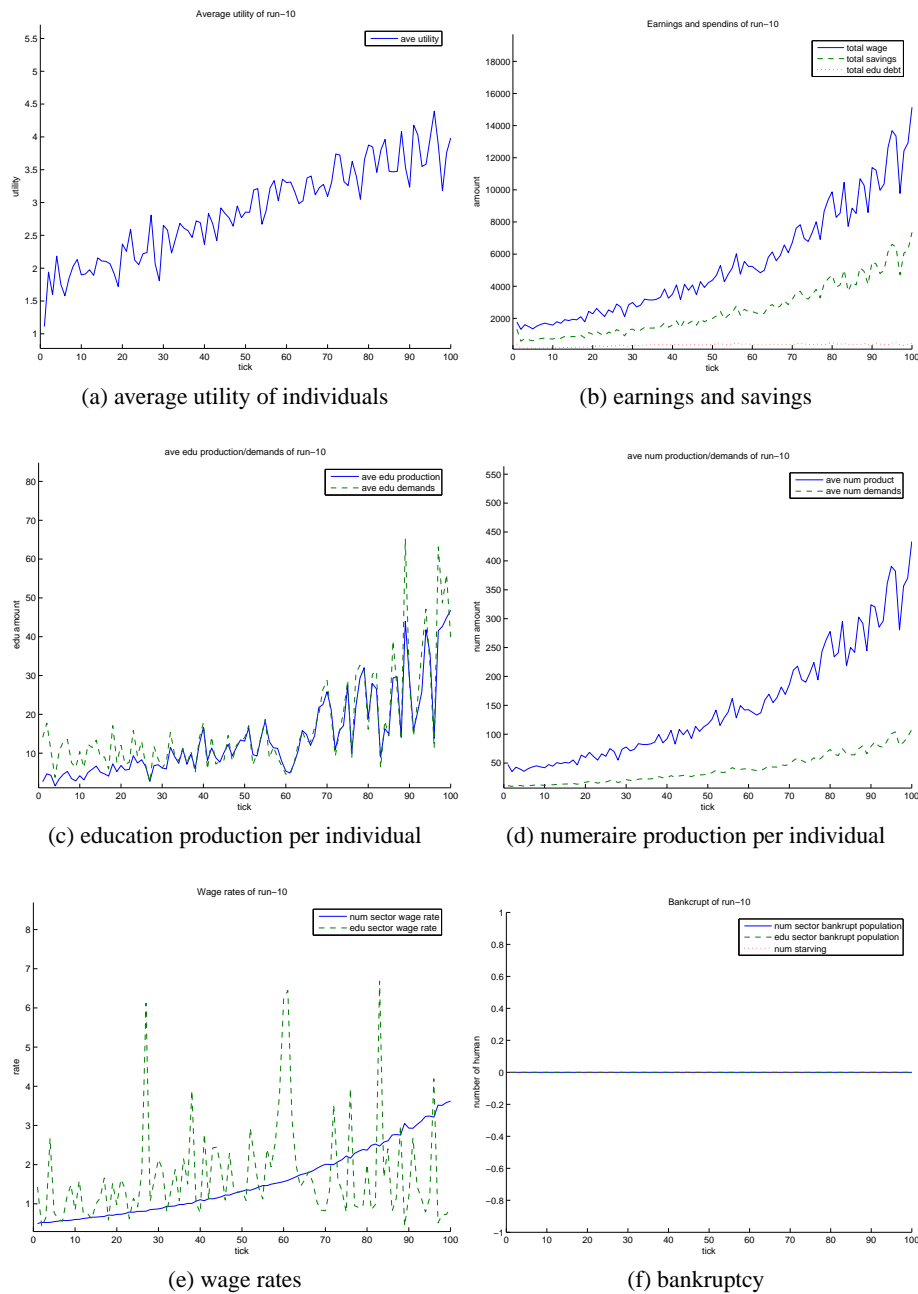


Fig. 3. Experimental run number 10, an example of a healthy economy.

Legend: (b) The solid line shows total earnings. The dashed line shows total savings. The dotted line shows debt due to education. (c) and (d) The solid line shows actual production. The dashed line shows demand. (e) The dashed line shows wages in the education sector. The solid line shows wages in the numeraire sector. (f) The solid line shows the number of workers in the numeraire sector who are bankrupt. The dashed line shows the corresponding number for the education sector. The dotted line shows the number of individuals who cannot afford to consume.

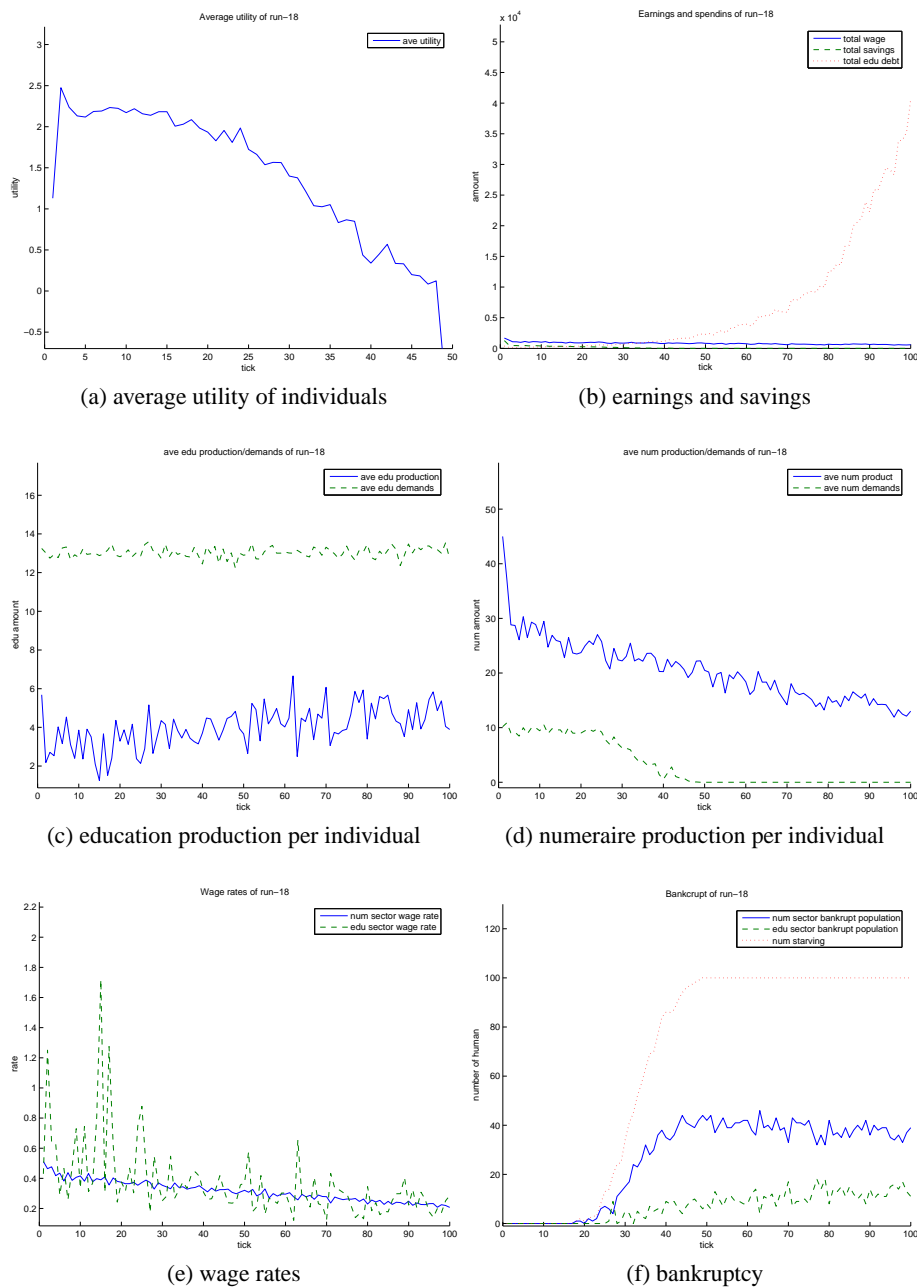


Fig. 4. Experiment run number 18, an example of an unhealthy economy.

Legend: (b) The solid line shows total earnings. The dashed line shows total savings. The dotted line shows debt due to education. (c) and (d) The solid line shows actual production. The dashed line shows demand. (e) The dashed line shows wages in the education sector. The solid line shows wages in the numeraire sector. (f) The solid line shows the number of workers in the numeraire sector who are bankrupt. The dashed line shows the corresponding number for the education sector. The dotted line shows the number of individuals who cannot afford to consume.

Table 1. Policy combinations used in experiments.

Two runs are highlighted and discussed in Section 4. Run 10 is an example of a healthy economy. Run 18 is an example of an unhealthy economy.

run	$pr_{n,k}$	$pr_{k,e}$	$pr_{h,e}$	e_d
1	Z	Constant	Constant	Random
2	Constant	Constant	Constant	Random
3	Self adjust	Constant	Constant	Random
4	Z	Self Adjust	Constant	Random
5	Constant	Self Adjust	Constant	Random
6	Self adjust	Self Adjust	Constant	Random
7	Z	Constant	Constant	Max utility
8	Constant	Constant	Constant	Max utility
9	Self adjust	Constant	Constant	Max utility
10	Z	Self Adjust	Constant	Max utility
11	Constant	Self Adjust	Constant	Max utility
12	Self adjust	Self Adjust	Constant	Max utility
13	Z	Constant	Self Adjust	Random
14	Constant	Constant	Self Adjust	Random
15	Self adjust	Constant	Self Adjust	Random
16	Z	Self Adjust	Self Adjust	Random
17	Constant	Self Adjust	Self Adjust	Random
18	Self adjust	Self Adjust	Self Adjust	Random
19	Z	Constant	Self Adjust	Max utility
20	Constant	Constant	Self Adjust	Max utility
21	Self adjust	Constant	Self Adjust	Max utility
22	Z	Self Adjust	Self Adjust	Max utility
23	Constant	Self Adjust	Self Adjust	Max utility
24	Self adjust	Self Adjust	Self Adjust	Max utility

- The number of individuals who cannot generate enough wages during their lifetime to pay for their education and their consumption as a worker or as a retiree, broken down across the numeraire and education sectors.

By all these measures, the economy in Figure 3 (experimental run number 10) is healthy. The overall utility of individuals grows over time, as do wages (which reflect production). Education production fluctuates over time, but fits well with demand—note that when demand exceeds supply, then individuals only receive a proportion of the education they want, and the surplus demand is spread across the population. Numeraire production grows over time. Wages in the numeraire sector grow steadily over time, as do those in the education sector, but these latter are also affected by spikes in demand. Finally, no individuals go bankrupt.

In contrast, the economy in Figure 4 (experimental run number 18) is dramatically unhealthy. Once we get past the start-up effects, which are responsible, for example, for the same modest jump in average utility in both Figure 3 and 4 (note that Figure 3 (a) and Figure 4 (a) are on rather different scales), utility enters a long slump, total earnings

are static while debt mounts, demand for education consistently outstrips supply by a factor of around 3, average wages have a downward trend, and after about six generations (20 time periods) become insufficient to support the whole population—indeed after around 15 generations (40–50 time periods) the entire population cannot meet its needs. The only apparent bright spot is that numeraire production exceeds demand, but this is because individuals do not have enough money to consume any of the goods—at the end, production is 40 times less than that in the healthy economy.

5 Discussion

The results in the previous section are taken from only two examples of the 24 outlined in Section 3, but they are typical. To back up this claim, Figure 5 gives the average production of numeraire good and the demand for that good (which is a useful measure of economic health) against the demand and supply of education for 10 of the models. The results are presented in pairs, so Figure 5 (a) and (c) give numeraire production and education production for one model, Figure 5 (b) and (d) for the next model, and so on. The broad trends shown in Figures 3 and 4 are repeated in these other models—the results in Figures 3 and 4 are those in Figure 5 (f) and (h) and Figure 5 (n) and (p) respectively. All of the other runs have results that fall into the same two broad classes—not only are all healthy economies healthy in exactly the same way, but all unhealthy economies are unhealthy in the same way.

The question, of course, is “why do the failing economies fail?”, and it seems to us that the reason for the failure is clear from Figure 5. All the economies that fail have a consistently unmet demand for education. Over time, if economies lack the ability to educate the workforce, productivity falls, there is no basis for capital investment, and so demand for education remains unmet.

Of course, this feedback effect is written into the equation-based model, so it is no great surprise that it surfaces in the agent-based model. Indeed, we would be worried if it did not. However, note that in all the economies, even the successful ones, the demand for education initially outstrips supply. It is those economies responding to this mismatch by pumping resources into education and thus growing education production, that manage to bootstrap themselves out of the initial surplus demand for education (which will, of course, limit the productivity of the education sector since future educators themselves will be less productive if their education demands are not met). Interestingly, all the economies in Figure 5 that fail are economies that use the self-adjustment mechanism to set investment. This mechanism is much more short-term than the others, cutting investment at the first suggestion that production exceeds demand. It is tempting to interpret the failure of this approach in the models depicted as a failure for short-termism in economic policy, but we need to run more experiments before we can be confident in making any judgement on this.

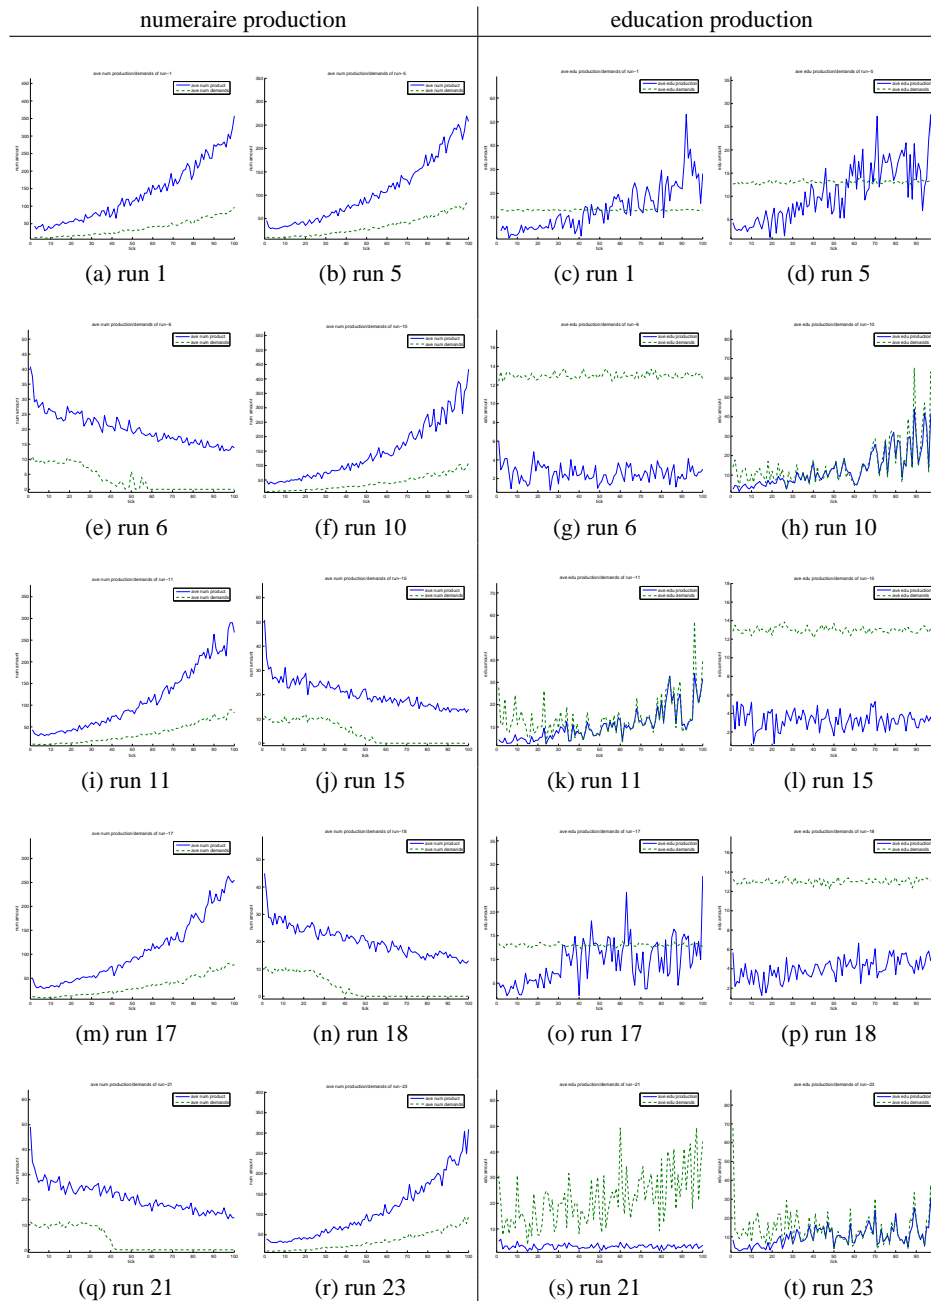


Fig. 5. Relationship between the education produced by the economy and the production of the numeraire good for representative experiments.

Legend: In all graphs, demand is given by a dotted line, supply by a solid line. The left two columns give the numeraire production, the right two columns the education production. (a) and (c) are taken from the same economy, as are (b) and (d), and so on.

Successful economies are those in which the demand for education is met and the numeraire production increases over time (e.g., runs 1, 5, 10, 11, 17 and 23); whereas in failed economies, numeraire production decreases and the demand for education falls to zero because no one has any money with which to pay for schooling (e.g., runs 6, 15, 18 and 21).

6 Summary

This paper has described the creation of an agent-based model of an economy from an equation-based model, and the results of some experiments intended to establish the behavior of the model under a range of conditions. These experiments have shown that the model tightly couples investment in education to production, and, through production, to the overall health of the economy.

Our next step with this model is to extend it towards the policy evaluation tool that we described in the introduction. To do this, we first envisage combining it with the model described in [3], converted to an agent-based model along the lines already achieved by Tang *et al.* [8], to give us a mechanism that individuals use to determine their level of education that is a good fit for real data. With that done, we want to couple in models like that in [7] which relate policy changes in education, like class size, to the quality of education that is provided.

Acknowledgments

This work was partially supported by NSF #ITR-02-19347 and by PSC-CUNY #68525-00-37.

References

1. E. Bonabeau. Agent-based modelling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Science*, 99(3):7280–7287, May 2002.
2. S. B. Flexner, editor. *The Random House Dictionary of the English Language*. Random House, New York, NY, 1987.
3. M. Kremer. How much does sorting increase inequality? *The Quarterly Journal of Economics*, 112(1):115–139, Feb. 1997.
4. J. Laitner. Earnings within education groups and overall productivity growth. *The Journal of Political Economy*, 108(4):807–832, Aug. 2000.
5. No Child Left Behind. <http://www.whitehouse.gov/news/reports/no-child-left-behind.html>, July 2001.
6. H. Van Dyke Parunak, Robert Savit, and Rick L. Riolo. Agent-based modeling vs. equation-based modeling: A case study and users' guide. In *Proceedings of the Workshop on Multiagent-based Simulation*, pages 10–25, 1998.
7. Elizabeth Sklar, Mathew Davies, and Min San Tan Co. SimEd: Simulating Education as a MultiAgent System. In *Proceedings of the Third International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2004)*, pages 998–1005, 2004.
8. Y. Tang, S. Parsons, and E. Sklar. Agent-based modeling of human education data. In *Proceedings of the Fifth International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2006)*, 2006.
9. M. Wooldridge, N. R. Jennings, and D. Kinny. The Gaia methodology for agent-oriented analysis and design. *Autonomous Agents and Multi-Agent Systems*, 3(3):285–312, 2000.