(The Evolution of) Post-Secondary Education: A Computational Model and Experiments

Andreas Ortmann
Sergey Slobodyan
Samuel S. Nordberg

CERGE-EI

Charles University
Center for Economic Research and Graduate Education
Academy of Sciences of the Czech Republic
Economics Institute

WORKING PAPER SERIES (ISSN 1211-3298)
Electronic Version

208

(The Evolution of) Post-Secondary Education: A Computational Model and Experiments[1]

Andreas Ortmann
CERGE-EI, Charles University
and
Academy of Sciences of the Czech Republic,
Prague, Czech Republic
andreas.ortmann@cerge-ei.cz

Sergey Slobodyan
CERGE-EI, Charles University
and
Academy of Sciences of the Czech Republic,
Prague, Czech Republic
sergey.slobodyan@cerge-ei.cz

Samuel S. Nordberg
Bowdoin College, Brunswick, USA

February 14, 2003

[1]We thank participants of the WHEIA workshop in Trieste, Italy, and seminar participants at CERGE-EI for valuable comments. Special thanks to Dmitry Kolyuzhnov. Special thanks also to Martin Dieringer who helped with early versions of the computer code. The first author acknowledges financial support by the Alfred P. Sloan Foundation and the Center for Adaptive Behavior and Cognition of the Max-Planck-Institut fuer Bildungsforschung in Berlin, Germany, which he was visiting while working on early versions of the computer code.

Abstract

We propose a computational model to study (the evolution of) post-secondary education. “Consumers” who differ in quality shop around for desirable colleges or universities. “Firms” that differ in quality signal the availability of their services to desirable students. Colleges and universities, as long as they have capacity, make offers to students who apply and qualify.

We study the dynamics and asymptotics for three nested variants of this matching model: the first variant replicates the Vriend (1995) model, the second stratifies both firms and consumers by quality, while the third variant of our model additionally equips some firms with economies of scale. The third variant of our model is motivated by the entry of for-profit providers into low-quality segments of post-secondary education in the USA and empirical evidence that, while traditional nonprofit or state-supported providers of higher education do not have significant economies of scale, the new breed of for-profit providers seems to capture economies in core functions such as advertising, informational infrastructure, and regulatory compliance. Our computational results suggest that this new breed of providers is likely to continue to move up the quality ladder.

Our model also lends itself to the study of such issues as the consequences of opportunistic behavior of firms (admittance of unqualified students for fiscal reasons) and the emergence of behaviorally different consumers (traditional “patronizers” vs “hoppers”), among others. Our computational results suggest that opportunism is a poor long–run strategy and that consumers are rather heterogenous in their shopping behavior but that the mix of behaviorally different consumers is unaffected by the presence of for–profits or opportunistically behaving firms.

1 Introduction

Post–secondary education in the USA, formerly known there as higher education, has undergone dramatic changes over the past decade (Ortmann 1997). The new label reflects new realities such as the increasing orientation of traditional higher education providers toward vocationalism (Breneman 1994), and the emergence of a new breed of higher education providers — publicly traded, degree–granting providers of post–secondary education (Ortmann 2001), (Ortmann 1998), and (?) that we shall call for–profits from here on. These for–profit "mutants" now represent about 10 percent of the post–secondary education institutions in the USA[1].

That for–profits have managed to invade the higher education sector as we knew it is little short of sensational. Higher education in the USA was, and for the most part still is, a heavily subsidized industry whose private not–for–profit and public segments were, and still are, subsidized through significant tax and regulatory breaks, (Facchina, Showell, and Stone 1993), as well as significant donations. In addition, not–for–profit and public institutions of higher education in the USA do not have to pay investors a reasonable return. For–profits were, and are, thus clearly handicapped. How then could they succeed? This is the first question we address below.

For–profits invaded higher education initially by providing services to market niches such as information technology training and continuing education/workplace training for adults, see (Ortmann 1998). In terms of the classification proposed by (Zemsky, Shaman, and Iannozzi 1997), for–profits entered post–secondary education through segments in which one typically also finds community colleges. Over the past few years, for–profits have successfully moved up to segments in which one typically also finds state universities. It is thus an interesting question whether this invasion of ever higher segments of higher education through for–profit "mutants" can be stopped, or whether for–profits will ultimately invade the "brand–name segments". In other words, *could* a liberal arts college – arguably the paragon of the brand-name segment – be organized as a for–profit institution?[2] This is the second question we address below.

[1] The major publicly traded, degree-granting providers of post-secondary education in the USA (by way of their stock market symbols, APOL, UOPX, CECO, COCO, DV, EDMC, ESI, STRA, WIX) will generate in excess of $4 billion in revenue in 2002 which represents about 2% of the higher education market as traditionally understood. The divergence between market share in terms of number of institutions and revenue reflects the particularities of the ways for-profits operate: Typically they have centralized administrative and curricular development facilities. Indeed, teachers' curricular liberties are severely restricted and teachers' role may be best described as learning facilitators. "Campuses"/learning centers are no frills and located for easy access. For more details, see Ortmann (1998, 2001) and Ruch (2001).

[2] We are agnostic on the issue of whether a liberal arts college *should* be organized as a for-profit institution. The issue, however, is of some relevance as rudimentary forms of liberal arts education offered by for-profit providers such as Edison Schools, Mosaica Education, and Chancellor Beacon Academies have made, for better or worse, inroads into primary and secondary education (?).

In addition to understanding how for–profits managed to invade higher education as we knew it, and what the future of these "mutants" is, we are interested in studying the consequences of opportunistic behavior of colleges and universities (e.g., admittance of unqualified students for fiscal reasons), viable quality improvement strategies for such firms, the emergence of behaviorally different consumers (traditional "patronizers" versus "hoppers"), and various other issues explained below.[3]

Toward those ends, we propose a computational model that we ultimately intend to calibrate with data from post–secondary education in the USA (e.g., the data on which the VIRTUAL U simulation is based)[4]. Our model is a progression of three increasingly refined ("nested") variants. Following exhortations in the literature to concatenate new computational models with predecessors, e.g., (Axelrod 1997), the first variant of our model "reverse–engineers" and somewhat generalizes (especially the classifier system) an influential model of decentralized markets consisting of locally interacting boundedly rational and heterogeneous agents (Vriend 1995). This variant of our model is meant to establish a baseline and reference point that ties the other two variants of our model to the literature by using his basic set–up and parameterization. Indeed, we have been able to replicate reasonably well Vriend's results (e.g., the service ratio approaching 1, approximately one third of consumers patronizing previously attended firms, etc.) Since his model presented a decentralized market with buyers and sellers not stratified by quality (as buyers and sellers of post–secondary education in the USA surely are), we introduce in our second variant (from here on "Q–model") stratification by quality both of buyers and sellers. Our third variant (from here on "QES–model") adds to the Q–model a new kind of firm that is distinguished from other firms by its cost configuration, namely economies of scale. The QES– and Q–models are the computational laboratories in which we study the invasion into traditional higher education by publicly traded, degree–granting providers, their likely future trajectories, and the various other issues already mentioned.

The paper is structured as follows: Section 2 presents the matching model, a discussion of our computational agents, and the details of experimental design and implementation. Section 3 presents, among other things, findings on the equilibrium distributions of firms across the quality spectrum under various treatments. Section 4 provides a brief discussion of related literature. In Section 5 we proffer some concluding remarks addressing objections to agent-based modeling and summarizing what we have accomplished and what remains to be done. Appendices contain an analysis of an analytically tractable simplified version of our model and the pseudo–code of one of our programs.

[3] While our study is motivated by recent developments in post-secondary education in the USA, similar developments can also be observed in countries such as Germany that historically were much less open to curricular and other educational innovations.

[4] ... but haven't yet. VIRTUAL U is an ambitious attempt to build a Sim City like simulation of higher education in the US. It draws on real-world data in parameterizing the underlying simulation machines. See http://www.virtual-u.org/ for more details. Here we simply propose a computational laboratory, explore its properties, and relate it to the simple theoretical model shown in Appendix A.

2 Structure of the matching model

2.1 Summary of the matching model

Buyers (prospective students and/or their parents) and sellers (colleges and universities) of post-secondary education try to match optimally in a decentralized market for a number of periods.[5] In the first period, buyers are randomly and uniformly distributed along a quality spectrum that is normalized to the interval [0, 100]. Likewise, in the first period sellers are randomly and uniformly distributed along a quality spectrum that is normalized to the interval [0, 100].

Buyers and sellers are modelled as boundedly rational decision makers that sometimes "tremble", i.e. they select actions that, in their experience, worked best but they select these actions probabilistically. Buyers and sellers are characterized by preferences and internal states, behavioral rules (= rules for selecting actions out of the current choice set), the number of behavioral rules, internal rules (= rules for selecting and modifying rules), and specifications of the decision makers' interactions with the world. Table 1 summarizes these characteristics which are discussed in more detail below.

[5] For the remainder of the text we use as synonyms the words buyers, consumers, and students, on the one hand, and sellers, firms, and colleges and universities, on the other hand.

Table 1

CONSUMERS	FIRMS
PREFERENCES	
Δ^F	Δ^C
INTERNAL STATES	
Weights; own Q, firm attended last period, list of schools that are desirable	Weights; own Q, demand, avg. Q of consumers, target number of consumers, profit, average profit
BEHAVIORAL RULES	
IF (SAT, no SAT, indifferent to SAT) **AND** (INFO, no INFO, indifferent to INFO) **THEN** (PATR, KNOWN, RAND)	(production, signal)
NUMBER OF RULES	
18 (27)	20
INTERNAL RULES	
Rules for Selecting Rules stoch. auction; reinforcement	*Rules for Selecting Rules* stoch. auction; reinforcement *Rules for Changing Rules* production and signaling adjustment, GA
MATCHING PROTOCOL	
Specification of firm selection	Specification of consumer selection

The preferences of buyers and sellers are defined by the minimal quality of a counterpart they are willing to consider: Buyers will go only to firms that meet a given quality threshold (defined as own quality Q minus Δ^F); sellers are interested only in those consumers who meet a given quality threshold (defined as own quality minus Δ^C).

The internal states of buyers and sellers are defined as follows: Buyers keep track of the "strengths" (to be explained presently) of their behavioral rules (also to be explained presently), their own quality Q[6], the index of the firm which

[6] Currently, our consumers do not change their quality, i.e., exactly what school they attend has no consequence for their educational outcomes. Firms thus face a fixed distribution of consumers in quality space. Given our current focus nothing seems lost through this restriction which could be relaxed easily. We note that there is quite some discussion about the value that colleges and universities add to human capital formation, see (Altonji and Dunn 1996),

they attended last period, and a list of schools that are desirable (i.e., have a minimum quality $Q - \Delta^F$); sellers analogously keep track of the "strengths" of their behavioral rules, their own quality Q[7], the number of consumers in their market niche, the realized demand for their services, and current as well as (trailing) average profit.

Turning to behavioral rules, buyers maintain lists of rules, each with a conditional and an action part (a Classifier System). The conditional part determines if a rule will be activated (to be explained below) given the current state of the world while the action part encodes possible actions. Specifically, rules have the following form: '**IF** (SAT, not SAT, indifferent to SAT) **AND** (INFO, not INFO, indifferent to INFO) **THEN** (PATR, KNOWN, RAND)'. Here SAT means satisfaction (served) last period, and INFO a presence of signals from firms in a current period. Buyers have three actions available to them: they can try to patronize the firm they attended (PATR) or go to a firm that signalled them (KNOWN) or randomly choose some firm (RAND) with probability $\frac{1}{N_{firms}}$.[8] If a buyer has to take into consideration the firm's quality, as in our Q– or QES–models, the RAND action does not apply and is not available. Buyers who are not able to take the PATR, KNOWN, or RAND actions (because firms do not accept them), do not get matched. Buyers start with a complete set of $3 \times 3 \times 2(3) = 18(27)$ rules which remain unchanged over the course of our computational experiments.

Sellers also maintain lists of rules. However, unlike the classifier system representing the behavior of buyers, the behavioral rules for sellers encode pairs of integer numbers, one representing the number of units produced and the other the number of signals to be sent. Every integer is coded by a bitstring of length 10; therefore, numbers from 0 to 1023 can be coded. Each period sellers produce slots which they then signal to prospective (and desirable) buyers. Note that rules in this sense translate directly into actions. There are twenty such rules that are initialized randomly (for every bit in a string, a fair coin is tossed to determine whether it is 0 or 1) so as to represent various production–signaling combinations.[9]

(Behrman, Rosenzweig, and Taubman 1996), and (Tamura 2001).

[7] A firm's quality is updated according to the following rule, $Q = w_1 \cdot Q_{avg} + w_2 \cdot \pi$, where Q is the firm's quality, Q_{avg} the average quality of its consumers, π the firm's profits, and w_1 and w_2 are weights. Since the weight on profits is rather low, we essentially model the quality of a college or university as the average of the quality of its students. This follows well–established precedence in the literature, e.g., (**?**), (**?**). We conjecture that factoring in the quality of faculty would not affect our qualitative results for all reasonable parameterizations.

[8] (Adelman 2000) is an eminently readable sketch of the emerging "parallel universe of postsecondary credentials ... an education and training enterprise that is transnational and competency–based, confers certifications not degrees, and exists beyond governments' notice or control." We note that competency–based certifiction is also propagated by institutions such as Western Governors University which has made considerable headlines by offering its prospective students that skills and knowledge acquired at other universities, on the job, or just through life may be counted toward one's WGU degree. What all these developments point to is a new kind of student – "hoppers" we call them – who takes classes here and there and then consolidates her or his portfolio at a school of her or his choice. In our model, hoppers are modelled as consumers who never use the PATR action.

[9] Thus, we initialize with widely off–equilibrium quantity–signalling pairs, because both

The internal rules of buyers and sellers – the rules for selecting and modifying rules – are defined and explained in the next subsection.

Table 2:

FIRMS	CONSUMERS
– make production and signaling decisions	
– signal	
	– choose firms, apply to one
– accept or reject consumers	
	– if rejected, choose another firm
...	...
– calculate profits, adjust quality	

Table 2 details top–down the timeline of interactions of sellers and buyers, or matching protocol, in each period. Every period, firms first make production and signalling decisions. Next, firms signal potential buyers by picking a random buyer (with replacement) and checking its quality. Only those buyers get signalled that are within a pre–specified range of quality ($[Q - \Delta^C, Q + \Delta^F]$), up to the pre–determined number of signals that the firm has chosen to send in that period. This reflects the practice of colleges and universities to admit only those students that fulfill certain minimum quality standards and to diligently track the yield of various advertising and recruiting channels (i.e., not to waste recruiting efforts on candidates that can be expected to be out of reach or undesirable.)

Consumers then choose their firm from among the offers. Only those firms become candidates that are above a pre–specified quality that equals buyers' own quality minus Δ^F. This reflects the practice of the overwhelming number of students not to go to colleges and universities that are significantly worse than they are. It can also be interpreted as the result of decision making under constraints such as time or knowledge.

Since typically a student will be signalled by several colleges or universities, the question arises how he or she prioritizes among multiple offers. We assume that consumers collect all their offers and put those firms that satisfy a minimum quality on the list of desired firms. For consumers who PATRonize, this list consists of only one firm whose quality they do not check because they must have done so at some point in the past[10] and because quality changes typically do not happen suddenly.

number of slots and number of signals are drawn from an initial distribution whose support is the integers between 0 and 1023, but the number of buyers per firm is set to 100, and we have not observed a firm serving more than 200 consumers in equilibrium.

[10] Consumers are forced to take a KNOWN action in the initial period; hence they are able to PATRonize from the second period on.

All consumers (including those who PATRonize) then "apply" to their desired firm(s). This matching process involves two random processes, as follows: First, a consumer is randomly drawn. Second, that consumer then randomly draws a firm from her list of desired firms. As soon as such a firm can and wants to provide, a match is accomplished and it is another randomly drawn consumer's turn [11]. Firms do not discriminate between consumers who patronize or those responding to offers.[12]

2.2 Experimental implementation: Stochastic auction, reinforcement, and evolution of rules

Recall that buyers and sellers are modeled as boundedly rational decision makers that sometimes "tremble", i.e. they select their best actions only probabilistically. How do they do it? In a nutshell, buyers and sellers select probabilistically a rule from the sets of rules available to them. This probabilistic selection is implemented in all variants of our model through a stochastic auction into which all relevant rules are entered. However, these rules are not entered equally weighted. Rather, their weights reflect their past "success": the better they performed in the past, the more weight they get.

When participating in such an auction, every rule submits a total "bid" equal to b_1 times its weight or "strength", w, plus a random number, ε, drawn from the normal distribution. The basic bid $b_1 \times w$ may be thought of as a stake that a rule is willing to sacrifice for the right to be chosen in the auction. This stake will be higher, the higher a rule's weight is (b_1 is just a scaling parameter). Following (Holland 1992), the winning rule pays an activation fee equal to its basic bid. With a small probability we call discard probability, every rule's total bid can be discarded. This procedure makes sure that the "best" rule typically wins the auction but that inferior rules have a small chance of winning, too. In the following the winning rule is called the active rule. Our procedure operationalizes the fact that real–life buyers and sellers are boundedly rational and make their decisions under incomplete information, time pressure, or other cognitive constraints (Gigerenzer, Todd, & the ABC Research Group 1999); (Todd and Gigerenzer 2000); (Payne, Bettman, and Johnson 1988). Obviously that implies that they do not always make the optimal decision.

[11] Think of a student who collects all the information she gets in a large folder and on D–day takes the first one that fulfills her or his aspiration level. If this attempt fails, the student randomly selects another firm of sufficiently high quality out of the folder. Other procedures are, of course, thinkable. For example, rather than selecting firms randomly, consumers might call on schools according to their quality. We doubt, however, that the majority of students make their selections with that kind of high–level rationality or that they have the information that would allow them to optimize in such a sophisticated manner, see (Boylan 1998). In any case, these specifics of the matching process should not matter in equilibrium. Note that this latter statement is easily verifiable in our set–up.

[12] This is, in a sense, in contrast to (Kirman and Vriend 2001) where loyal customers could receive more or less preferential treatment. The implications of loyalty on the part of sellers remains an issue for future research. Again, we believe that this issue is not of material relevance for the issues we are interested in here.

Selection of the initial strength of a rule, its possible range (from zero to one in our case), the standard deviation of the auction's error term, ε, the speed with which that standard deviation decreases over time[13], and the discard probability, all influence two characteristics of the stochastic process generated by the stochastic auction: expected number of active rules (not more than three or four in our case) and the variance of the number of rules that will be called to duty on a regular basis.

As can also be seen from the pseudo–code in Appendix B, strengths of rules are restricted to [0,1]. This, together with the discard probability and the decreasing standard deviation of the auction's error term, ε, is done to prevent early in the simulation the emergence of "runaway" rules that might lead to premature convergence.

After the stochastic auctions have determined the actions to be performed by buyers and sellers in the current period, matching is implemented as described in the previous subsection and payoffs to buyers and sellers are realized. For a buyer, the payoff equals one if she is served this period and zero otherwise. For a seller, the payoff equals the ratio of the current profits to average profits over the last 200 periods[14], times δ, where $1 > \delta > 0$.[15] Next, each buyer's and seller's payoff is multiplied by $(1 - b_2)$ and this product is added to the active rule's weight.[16]

Note that the stochastic auction and reinforcement mechanism described above closely resembles various forms of probabilistic enforcement learning recently proposed in the literature, e.g., (Goeree and Holt 2001), (Goeree and Holt 1999); (Camerer, Ho, and Chong 2001); see also (Bush and Mosteller 1951), (Bush and Mosteller 1955); but see also (**?**), as an approach that formalizes experimental results on human decision making.

However, our evolutionary programming technique is more than simple individual reinforcement learning. To model the behavior of firms, we used a combination of the Steepest–Ascent Hill–Climbing algorithm and GENITOR algorithm[17]. This evolutionary technique is arguably the simplest program-

[13] This is being done by Vriend (1995) to reduce stochastic disturbances of the system as time progresses.

[14] This is motivated, first, by the parameterization in (Vriend 1995) and, second, by our desire to stabilize our computational model within a reasonable runlength.

[15] Multiplication by δ means that rules which consistently produce profits equal to the average cannot achieve the maximum strength of "1"; instead, they converge to δ. This construction is meant to reflect the never ending emergence of strategies that aim to beat the average performance. In our set–up, no seller rule will be used forever and eventually new combinations of (production, signalling) pairs will be experimented with. This "new broom effect" facilitates adjustment to a rapidly changing environment if other firms are behaving in an out–of–equilibrium manner.

[16] The remainder of the payoff, payoff times b_2, goes to the previous active rule ("bucket brigading"; see Holland 1992)

[17] In the GENITOR algorithm, rules are ranked according to their fitness, and the probability of selecting a particular rule is proportional to its rank. Every $n - th$ period, two evolutionary operators (crossover *and/or* mutation) are applied to produce a new rule, which is inserted into the existing ranking and replaces an old rule. One of the advantages of the GENITOR algorithm, according to (Chattoe 1998) and (Whitley 1989), is the relative stability of the ranking, which results in stable actions. (Chattoe 1998) argues furthermore that the

ming technique and as such is a desirable baseline, e.g., see (Chen, Duffy, and Yen 2002) and the critique of (Valente 2002).

The Steepest–Ascent Hill–Climbing part of our algorithm is implemented as follows: firms update both parts of their rules in every period, taking into account such perceived characteristics of the relevant market niche as actual demand and maximally possible demand.[18] If a firm's demand (number of consumers that applied to a firm in the current period) differs from its production, production is adjusted by 10% of the difference or 1 unit if 10% of the gap is less than 1. To adjust the number of signals, firms ignore PATRonizers and assume that consumers who showed interest were signalled in this period. Firms calculate the expected marginal revenues of additional signals and their marginal costs. To avoid making occasional losses, firms cap the number of signals by a value that allows them to break even, assuming that every unit that is produced is sold. Firms adjust towards their optimal expected signal 10% of the gap if their demand in this period is insufficient (less than production). If the current demand is higher than current production, firms cut 5% of the current signal level[19].

The GENITOR part of the algorithm is implemented by generating one completely new rule every 50 periods. Every firm's rules are ordered by their weight and two "parent" rules are selected from the top quarter (top five rules). A standard uniform crossover operator is applied to the binary strings—parent rules—and one of the two "children", randomly selected, is retained. Then, we mutate every bit of the child string and replace a randomly selected rule from the bottom half (bottom ten rules) with the child which is assigned a weight equal to the average of its parents' weights.

Since consumers' classifier system is complete, there is no need to evolve it further. Unlike firms, consumers have a set of rules that does not change. (Of course, the strengths of the rules might well change.)

After buyers and sellers have been matched, firms compute their revenues, costs, and profits. They also update their quality as the weighted average of the quality of students who have chosen to enroll and current profits, with weight on profits being relatively small.[20] While this approach to determining the quality of colleges and universities — essentially defining the quality of a school as the average of the quality of the students that it attracts — is admittedly simplistic, it captures the most important aspect of what determines the quality of an institution. Specifically, it allows us to study the trade–off any typical college

GENITOR algorithm closely resembles the real–world decision–making process in companies and humans.

[18] Adjustment of production was implemented in Vriend (1995), while signal adjustment is introduced by us.

[19] As can be verified by looking at the FOCs of the profit function, the derivative of profit with respect to signals is negative if demand is greater than production (although is does not give us quantitative guidance); if demand is less than production then the FOC allows us to compute the optimal adjustment.

[20] The formula for updating the quality is $Q = w_1 \cdot Q_{avg} + w_2 \cdot \pi$. $w_1 = 0.95$ in all simulations, and w_2 is calibrated by requiring the average firm quality to be equal to 50 which produces $w_2 \approx 0.1$. Other ways of calibration are, of course, possible but seem less natural.

faces on the margin of admitting a rich but not so smart instead of a poor but brilliant student. Below we call such admittance of unqualified students for fiscal reasons opportunistic behavior.

This process repeats round after round. The matching process, in other words, is a dynamic process that evolves over a number of periods. The dynamic process is defined algorithmically in terms of the behavioral rules of our agents, their internal states and preferences, their repeated interactions, and — through internal behavioral rules that govern how rules are selected and changed — the evolution of rules toward some stable outcome.

The program code consists of 10 modules: MAIN.CPP, PARAMETER.H; RULE.H, RULE.CPP; AGENT.H, AGENT.CPP; CONSUMER.H, CONSUMER.CPP; FIRM.H, FIRM.CPP.

Consumers and firms are defined (= declared) in the respective .H modules and implemented in the respective .CPP modules. Both, consumers and firms are instantiations of the Agent class declared in AGENT.H and implemented in AGENT.CPP. This super–class declares data components of agents such as the given number of rules and their initial weights (parameterized in PARAMETER.H) and implements them.

The RULES modules define and implement the data components (such as the number of segments = actions, and the number of bits per segment) as well as the functions components (such as the crossover and mutation operators).

2.3 Experimental design: Parameters and treatments

Following exhortations in the literature to concatenate new computational models with predecesors, e.g., (Axelrod 1997), we parameterize our baseline Q–model almost completely with the parameters Vriend (1995) chose, with one notable but inconsequential modification (signaling cost). Table 3 below details all relevant parameters common to our three treatments (the Q–model, the Q–model with moral hazard, and the QES–model) and also relates these parameters to those employed by (Vriend 1995).

Table 3

Run-length	3000	Vriend (1995)
Production cost C_Y (maximal quality)	.25	— · — · —
Signal cost C_S (maximal quality)	.025	Vriend = .08
Price P (maximal quality)	1	Vriend (1995)
Average number of consumers per firm	100	— · — · —
Maximum acceptable quality gap, consumers	10	NA
Initial rule weight, firms	0.3	Vriend (1995)
Initial rule weight, consumers	0.5	— · — · —
Steady state weight δ of an average rule, firms	0.65	NA
Stand dev of the auction error term, firms, $\mathrm{N}(0, R)$	0.075→0.03	Vriend (1995)
Stand dev of the auction error term, consumers, $\mathrm{N}(0, R)$	0.00875	— · — · —
Parameter b_1, firms	0.25	— · — · —
Parameter b_1, consumers	0.1	— · — · —
Parameter b_2, firms	0.4	— · — · —
Parameter b_2, consumers	0.1	— · — · —
Mutation probability	0.01	
Uniform crossover probability	0.50	
Discard probability	0.025	Vriend (1995)

The first of the three treatments that we used (from here on T1) employs the Q–model in order to generate baseline equilibrium distributions of firms across the quality spectrum. It was, importantly, used to calibrate our other treatments. The second of the three treatments (from here on T2) still employed the Q–model but inserted a small fraction of opportunistic firms in the set–up. Such firms accepted consumers whose minimum quality was 12 rather than 10 points below their own quality. In other words, Δ^C was 12 rather than 10 for these firms. The last treatment (T3) employs the QES–model to study the emergence of for-profits in post–secondary education. Specifically, to recall, we equip a subset of firms with economies of scale once they have reached a minimum efficient scale.[21]

Since scaling effects are notorious, we controlled for them by implementing treatments T1 through T3 with combinations of 10 firms/1000 consumers (Scale0 from here on), 12 firms/1200 consumers (Scale1) and 24 firms/2400

[21] The exact profit function for a normal firm is

$$\pi = (P \cdot \min(Y, D) - C_Y \cdot Y - C_S \cdot S) \cdot \frac{Q}{100},$$

where Y is the firm's production, D the realized demand, S the number of signals, and Q the firm's quality. For a for–profit mutant, the cost term $C_Y \cdot Y + C_S \cdot S$ is multiplied by $\frac{MES}{Y} \cdot 0.8$ where $MES = 50$ for all runs reported in this paper. Following (?), we do not handicap our for–profits with taxes although doing so would very likely not change our qualitative results; it would, however, slow down their moving up the quality ladder.

consumers (Scale2). Table 4 below summarizes our 3x3 design, detailling the number of runs in each cell and the number of mutants for treatments T2 and T3 across all scales.

Table 4

treatments\scales	10 firms/1000 consumers	12 firms/1200 cons	24 firms/2400 cons
T1: Q—model	20 runs	20 runs	20 runs
T2: T1 + MH	20 runs (1 mutant)	20 runs (1 mutant)	20 runs (3 mutants)
T3: QES—model	20 runs (1 mutant)	20 runs (2 mutants)	20 runs (3 mutants)

3 Results

For the most part we refrain from too detailled a summary and restrict ourselves to what we consider the essential characteristics of all runs in a treatment cell. A set of figures presenting all 180 runs may be obtained from the authors upon request.

3.1 Equilibrium distributions of firms across the quality spectrum

We first look at the distribution of firms after 3000 iterations.[22] As we will see, firms occupy well–defined "slots" in the quality spectrum, or market niches characterized by quality ranges, which we shall call, following (Zemsky, Shaman, and Iannozzi 1997), "segments". As we will also see, these segments are typically occupied by clusters of firms. We shall use the terms "segments" and "clusters" interchangeably throughout the remainder of this manuscript.

Baseline treatment T1. In Appendix A we show that, theoretically, we should have 6 clusters for all firm numbers modulo 6.[23] We also show that other numbers of firms (such as 10 in Scale0) lead to less stable configurations of clusters for small numbers of firms.

Looking at 10 firms and 1000 consumers (Scale0), we find indeed that firms "flock" into 6 to 8 clusters, with the clear mode being 7, as in Figure 1, and 8 being a not too distant second, as in Figure 2. Switching to Scale1 (12 firms/1200

[22] As we document in the following subsection, convergence to relatively stable configurations occurs in T1 within the first few hundred iterations. Even in T2 and T3, the distribution of firms stabilizes between 500 and 700 iterations (in a sense that we shall make more precise below). Recall that we initialize production randomly on [0,1023] and therefore far off the equilibrium of 100 units per firm. We could, alternatively, initialize production randomly on [0,255] speeding up convergence significantly to not more than 200 periods (which could be thought of as semesters or trimesters or quarters or some such time unit). Doing runs of 3000 allows us to estimate the likelihood of disturbances and switching behavior. More details below.

[23] We note that this number is a function of the width of the quality range and the width of the segment (to be made precise later). Ceteris paribus, increasing the quality range leads monotonically to higher number of clusters. We speculate that this relation can be described by the formula $\left[\frac{[Q]}{\Delta^F+\Delta^C}\right]+1$, where [] denotes taking an integer part and $[Q]$ denotes the quality range.

consumers) and Scale2 (24 firms/2400 consumers) we observe 6 clusters of 2 and 4, as theoretically predicted, see Figures 3, 4. The number of firms in each cluster is eseentially constant, with occasional eruptions and displacements reflecting the probabilistic nature of our modeling technique, Figure 5. Interestingly, but in light of our calculations in Appendix A not surprisingly, such displacements regularly result in an exchange of members of adjacent clusters. We note that similar results pertain for exploratory runs with scales of 20/2000 firms/consumers, as well as 40/4000, 48/4800, and 120/12000. This suggests that the design laid out in Table 4 is sensible. We note, finally, that clusters are distributed approximately equidistantly, again as predicted by the calculations in Appendix A. This result is also independent of scale.

From the above it follows that scale is important in two respects. First, only scales modulo 6 can be accurately described by our symmetric steady state calculations. In other words, there is a large degree of freedom for scales that are not of modulo 6, especially if the number of firms is rather small. As we increase the number of firms, it becomes less important whether the number of firms is modulo 6 or not. This is good news because it means that the computational model that we propose here is rather insensitive to integer constraints. Second, as we increase scales, we find — somewhat contradicting our initial intuition — a rather stable configuration of six clusters or segments which attract whatever number of firms populate our computational laboratory.

Moral hazard treatment T2. Looking at 10 firms and 1000 consumers (Scale0), we find that a sole opportunistic firm almost never increases its position in the quality spectrum after the initial adjustment process. In fact, firms roughly maintain their position, as in Figure 6, or drift down in the quality spectrum with about equal probabilities, as in Figure 7. In Scales 1 and 2 the opportunistic firms never manage to markedly increase their position in the quality spectrum after the initial adjustment process. In fact, these firms roughly maintain their position of drift down in the quality spectrum with about equal probabilities (but they sometimes do so quite dramatically, see Figure 8). We note that this detrimental effect is particularly pronounced for firms that start with very high quality. We observed instances of firms losing 60 quality points. We note also that drift downward is truncated for firms that start with low quality, hence for all scales firms actually are somewhat more likely to drift down than to maintain their position. The results reported here emerge from the very mild parameterization of moral hazard that we chose; increasing the moral hazard parameter systematically shifts more weight to downward drift. In fact, if the moral hazard parameter is doubled (i.e., decreasing the quality of the worst student from $Q-12$ to $Q-14$) in Scale2, the offending firm nearly always (more 90%) goes to the bottom of the quality spectrum.

For–profit invasion treatment T3. Looking at 10 firms and 1000 consumers (Scale0), we find that a sole for–proft firm never lowers its position in the quality spectrum after the initial adjustment process. In fact, for–profits increase their quality 80% of the time, often dramatically so, as in Figure 9. In three of the four cases where they did not, they started out being the firms with the lowest quality; in the remaining case the one with the second lowest. Looking at Scale2

(24 firms, 2400 consumers, three mutants), we find that the for–profits always move to the top cluster, displacing in the process incumbent high quality firms. The key difference lies in the timing; some for–profits take longer than others — in a few cases nearly the complete run, as in Figure 10. More often than not, though, they move up the quality spectrum with amazing speed (within a couple of hundred iterations). Looking at Scale1 (12 firms, 1200 consumers, two mutants), we find that about three out of four for–profits move all the way to the top, with the remainder almost always moving up but often getting trapped in segments below the top.

There are various metrics that could quantify the trends above. Table 5 summarizes one such metric.

Table 5

		$\overline{Q}_{end} - \overline{Q}_{start}$		
		min	avg±std	max
Scale0	T1, average for all firms	-3.0	3.1±4.2	9.8
	T2, opportunistic firm	-20.3	-4.5±8.8	8.9
	T3, for–profit mutant	-0.3	24.0±20.1	68.3
Scale1	T1, average for all firms	-3.6	3.5±5.3	12.6
	T2, opportunistic firm	-10.6	-0.4±3.4	5.3
	T3, for–profit mutant	-1.03	29.6±25.1	84.8
Scale2	T1, average for all firms	-9.0	1.8±5.0	14.2
	T2, opportunistic firms	-56.7	-8.5±11.6	3.9
	T3, for–profit mutants	0.5	40.9±25.7	85.6

$\overline{Q}_{end}$ and $\overline{Q}_{start}$ denote average quality of a firm during the last 500 periods and periods 100 — 500, respectively[24]; therefore $\overline{Q}_{end}$ - $\overline{Q}_{start}$ is a measure of the change of a firm's position in the quality spectrum over time. This measure quantifies in particular the default outcome of opportunistic firms moving down and for-profits moving up in quality for all investigated scales discussed above. Compare, for example, row 1 of Scale0 with rows 2 and 3 respectively. Clearly the average quality change for all firms in T1 (3.1) is larger than that of opportunistic firms in T2 (-4.5) and smaller than that of for–profits in T3 (24.0). Along similar lines, note that the quality change range has increased dramatically in for–profits, going from 9.8 to 68.3 at the upper limit. Similar effects can be observed for Scale1 and Scale2.

Turning to the difference across Scale1 and Scale2 (in order to avoid the confounding influence of integer constraints), the key result is that opportunistic firms tend to fare much worse in Scale2 (lower end of range being -56.7) than in Scale1 (lower end of range being -10.6). Relatedly, we see a stronger average

[24] We do not incorporate the first 100 periods because several hundred periods are needed for the initial noise to get worked out of the system. Including the frist 100 periods makes the data noisier but does not change any of the qualitative results. Excluding more initial periods would not leave enough periods for averaging before the increasing returns to scale regime of for–profits takes effect in period 501.

downward movement of opportunistic firms for Scale2 (-8.5; in contrast to -0.4 for Scale1). And somewhat analogously, we see a stronger upward movement of for-profits for Scale2 (40.9; in contrast to 29.6 for Scale1).

The correlation between firms' mobility and scale has a straightforward rationale: when a firm, for some reason, manages to get more than its equilibrium share of its segment, its increment in quality will be proportional to profits, which are in turn proportional to the number of consumers. Grabbing an additional 5% of a segment with 300 consumers adds 0.18 quality units to an average firm; an additional 5% of a segment with 100 consumers adds just 0.06 quality units to an average firm. Analogously, competitive advantage (of for–profits) or disadvantage (of opportunistic firms) translates readily into more pronounced quality changes and hence into more turbulent environments as the number of firms and consumers per segment increases. Less stability creates, of course, more opportunities, both positive and negative, for mutants.

3.2 Convergence toward equilibrium distributions

Especially in baseline treatments T1, the segment in which a firm will find itself typically depends on its initial rank: If a firm is, to take the Scale2 example, one of the top four firms initially, it is likely to end up in the top segment even if its initial quality lies significantly below the predicted quality of the cluster. The implicit adjustment process takes roughly between 200 (Scale1) and 400 (Scale2) iterations.

Thus, while convergence to the equilibrium location is relatively fast, we do observe occasional eruptions and displacements in quality even in the absence of opportunistic firms or entrants with economies of scale. A firm that moves up or down the quality spectrum typically dislodges another firm from the segment it invades. While the number and location of segments is relatively stable, there is quite some jockeying going on for those segments.

Opportunistic firms (T2) or entrants with economies of scale (T3) complicate the picture, generating more eruptions and displacements and slower convergence toward the equilibrium distribution. In fact, we often see cascade–like sequential convergence toward the equilibrium distribution, see Figures 10, 11.[25]

3.3 Signaling, production, and demand trends

Even though we initialize with widely off–equilibrium quantity–signalling pairs (recall that both number of slots and number of signals are drawn from an initial distribution whose support is the integers between 0 and 1023), production and demand tend to converge to their equilibrium values (independent of scale 100 for slots, dependent on scale $450-950$ for signals) within the first 500 periods,

[25] We have computed a measure of deviations of firms from the theoretical equilibrium — essentially the sum of squares of deviations — for all runs mentioned in Table 4. These computations give a measure of convergence beyond our informal discussion above; they are available from the authors upon request.

both in the aggregate and for individual firms (see a typical aggregate picture in Figure 12). Signaling, however, converges much more slowly and is much more volatile: a typical stochastic fluctuation in the firms' demand (10 to 15%) can lead to a much larger change in the perceived optimal signal level. As a result, adjustments to the signaling level are much larger than adjustments in production. An additional source of uncertainty arises because of the local nature of the information that firms collect. To calculate their optimal signal level, firms have to estimate their "target audience", or the number of consumers who can potentially accept a firm's offer. Even if the firm correctly knows the consumers' preferences (the maximum acceptable quality gap), as we currently assume, it still has to know the number of consumers in its segment to correctly calculate the target. We assume that the firms know only the total number of consumers and estimate the share of their "target audience" during the process of signal allocation. This estimation introduces an additional error term into the calculation of optimal signal level, which can be easily corrected if desired.

3.4 Consumers: "patronizers" versus "hoppers"

We now turn our attention to the emergence of behaviorally different consumers (traditional "patronizers" versus "hoppers".)

Rule weights versus rule usage. There is no real difference regarding rule weights and rule usage across treatments (T1 – T3). Specifically, the same four KNOWN and the same four PATRonizing rules have some significant, albeit differing, weights and usage across all treatments as well as all scales.[26]

Rule weights and usage of KNOWN and PATRonizing rules. The usage of PATRonizing rules decreases while the usage of KNOWN rules increases with the number of firms. For large numbers of firms (such as Scale2 in our computations), usage of PATRonizing and KNOWN rules bifurcates quickly but is ultimately stable (at roughly .3 and .7).[27] For smaller numbers of firms (such as Scale0 and Scale1 in our computations), the usage of these two kinds of rules converges more quickly than for larger numbers. Additionally, PATRonizing rules are used more often.

Rule weights and usage of specialized and generic rules[28]. While the usage of specialized and generic rules is apparently independent of scale, specialized

[26] Out of 18 rules, 10 include 'not SAT', 'no INFO', or both. These contingencies are very rare. Average satisfaction rate is 0.96–0.97 across scales and treatments, leaving 3–4% to cases of 'not SAT'. The average number of signals per consumer varies between 4.5 for Scale1 and 9.5 for Scale2, so the likelihood that a consumer will not be signaled is slim. In a symmetric steady state, since all consumers can be potentially served by a firm, it is just as improbable that a consumer would find herself outside of all firms' segments. The remaining 8 rules are those actually used by the consumers.

[27] The same result was obtained for exploratory runs with 120 firms and 12000 consumers. We note that this result coincides with that of (Vriend 1995) where consumers were patronizing approximately 1/3 of a time. This is interesting because the stochastic auction in (Vriend 1995) was skewed towards KNOWN rules (they were given two tries in a stochastic auction).

[28] We call a rule "specialized" if it is not indifferent to both SAT and INFO in its condition part. If the condition part of the rule includes at least one "indifferent to..." statement, we call such a rule "generic".

rules are used far less than generic ones (roughly 1 out of 5 times). This result seems due to the fact that only eight rules are typically used, of which two are specialized. The usage pattern, then, appears for the most part determined by the relative number of relevant specialized and generic rules.

Heterogeneity of consumers as measured by rule usage. Some consumers never use KNOWN and some never use PATRonizing rules. Some consumers use one of the eight "good rules" about half of the time. This is true for all treatments and all scales. The latter result is due to the rather small standard deviation of the error term in the auction for consumers. A larger standard deviation would lead to somewhat more varied rule usage.

4 Related literature

The computational matching model presented above has at least three reference points in the literature.

First, there is the classic work by (Gale and Shapley 1962) on college admissions and the stability of marriage and later related work on two-sided matching, e.g., (Roth and Sotomayor 1990), (Roth and Xing 1994), (Roth and Xing 1997), (**?**); (Simao and Todd 2002); (Pingle and Tesfatsion 2001); (Vriend 1995); (Kirman and Vriend 2001); (Weisbuch, Kirman, and Herreiner 2000). This literature has theoretically illustrated the heavy mathematical machinery necessary to model matching processes; it has also provided compelling evidence both theoretically and empirically on the importance of institutional arrangements that prevent, for instance, lower–ranked market participants from "jumping the gun" on other (higher–ranked) market participants and, exploiting well–known psychological phenomena such as loss aversion, pushing them into decisions that they might come to regret. In the context of post–secondary education in the USA, this issue is currently on the front burner as a number of colleges and universities are considering throwing out their early admissions policies, see (Schemo April 26, 2002).

Second, there is work that documents important changes in higher education. What little is out there in academic journals has already been mentioned in the introduction above; for the time being much of the relevant information on those developments is currently available only in official SEC forms or in research reports of investment houses; for details, see (Ortmann 1998). The situation is slowly changing though, e.g., (**?**). Three other academic papers deserve mention here. (**?**) study peer effects and show theoretically why it is imperative for colleges and universities to give out need–based financial aid. Only by doing so, will they be able to attract those bright (but poor) students that are an indispensable input in the production process of those students that can pay (but are not so smart). There is thus nothing altruistic about giving out need–based financial aid. Not surprisingly, colleges and universities routinely monitor the comparative attractiveness of their own financial aid packages and those of their close competitors. This rather insightful paper motivated our moral hazard treatment. (**?**) study theoretically and computationally the competition between tax–financed, tuition–free public schools and competitive, tuition–financed private schools in primary and secondary education, the impact of vouchers, and peer–group effects when students differ by ability and income. The equilibrium of their model shows that schools stratify along the quality spectrum and that students in private schools, dependent on their marginal productivity, either receive tuition discounts or have to pay tuition premia, as suggested in (?). (?) also study the competition between public and private schools in an agent–based model of school choice.

Third, there is a literature on modeling social processes through GAs and related evolutionary programming techniques. (Arthur 1994), (Arthur 1991) persuasively argues the case for agent–based models of interactions of boundedly rational and heterogeneous agents. (Arthur 1994) points out that such models are grounded in plenty of evidence. Indeed, much of the recent evidence in experimental economics, e.g., (Nagel 1995); (Stahl 1996); (Stahl 2000); (Stahl and Wilson 1995); (Costa-Gomes, Crawford, and Broseta 2001) and experimental psychology, e.g., (Cosmides and Tooby 1996); (Gigerenzer, Todd, and ABC Research Group 1999) has reinforced the impression that Arthur gets it right — people (whether real or fictitious, such as organizations) are "intuitive statisticians" (Cosmides and Tooby) who inductively keep track of the performance of a set of plausible, simple models of the world that they can cope with. When it comes time to make choices, people act upon the most credible and possibly most profitable one. The others they keep in the back of their mind, so to speak (Arthur 1994, p. 407, sligthly edited). (Arthur 1991) makes a similar argument but also stresses the importance of calibrating computational agents so as to accurately reflect how human agents learn. Not much attention has been paid during the last decade to this exhortation, although recent developments comparing the performance of human and computational agents in more or less identical settings, e.g., (Chen, Duffy, and Yen 2002), (Pingle and Tesfatsion 2001), are encouraging signs.

What (Arthur 1991) does not stress is the importance of embedding both computational and human agents in environments that mimic the essential features of the real–world situation that they attempt to study. (Plott 1987) has proposed, in a different context, a "parallelism postulate" — the challenge to experimental economists, especially if they give advice to policy makers, to create as test–beds small–scale versions of the situation that they study. (Simon 1956) has captured the need to understand both — agents and environment — in his metaphor of the two being like the blades of a pair of scissors; one without the other is of little use.

5 Concluding remarks

We have proposed a computational model to study (the evolution of) post–secondary education. Although our model is motivated by developments in the USA, the insights it generates should be easily transferable to related developments in other countries. While we intend to calibrate our model with data from the USA (or other countries, for that matter), and while we believe that our model captures key aspects of post–secondary education, we currently use our model primarily as a computational laboratory. It is useful to conceptualize such a laboratory as culture–dish, as (Tesfatsion 2002) does, that allows us to explore how macro regularities might emerge through the repeated local interactions of boundedly rational, heterogeneous agents from the bottom up.

5.1 Objections to agent-based modeling

A standard objection to agent-based modeling is, why not model the matching process the good old-fashioned way, i.e. using an equilibrium search model with perfectly rational agents? Our answer is two-fold. Firstly, drawing on Arthur's arguments and the experimental literature already mentioned, we believe it self-evident that agents (including aggregate agents such as firms) are not perfectly rational. Rather, they are boundedly rational in their reasoning and heterogeneous in their behavior. Secondly, we simply do not see a way to model the issues we have addressed above in the good-old fashioned way. (That said, see our analytic results in the appendix for a simplified version of the problem we analyze here.)

We do, however, acknowledge that the sensitivity of agent-based modeling is an important question. This sensitivity has two aspects, that of external validity and that of internal validity.

External validity, apart from the question of parallellism that we discussed above, addresses the question of the basic behavioral assumption that goes into our evolutionary programming, such as the way students make choices. One of the nice features of our technique is that we can easily incorporate competing behavioral assumptions, and that we can do so even by modelling various types. Think of this feature as a list of assumptions that is initially given to a reader who then might (dis)agree with their (un)reasonableness. To illustrate, recall how we have presently conceptualized the consumer's decision: Consumers randomly draw a firm from their lists of desired firms that have signaled them. One might argue in favor of a higher degree of rationality and have consumers select the best firm instead. It is obvious that checking the sensitivity of the model to this change in assumption would require only a couple of key strokes and another set of runs.

Internal validity addresses the question of the specifics of the evolutionary modelling technique, such as the specific parameter values of the standard deviation of the auction's error terms, the speed with which that standard deviation decreases over time, and the discard probability. Specifically, a number of authors such as (Michalewicz 1999), (Mitchell 1996), (Chattoe 1998), have voiced concern that the degrees of freedom inherent in evolutionary modelling techniques — similar to the degrees of freedom of the design and implementation of human experiments — subject any computational model to the real danger of being a mere example, an example for that matter that may be rather unrepresentative as regards the complete set of sensible parameterizations. This issue (which is basically that of replication) is admittedly an important and tricky one. For now, we have solved it by relying almost exclusively on the implementation details in Vriend (1995).

(Axelrod 1997) has enumerated some of the problems that complicate replication of computational simulations (and re–engineering of extant model). Our own experience supports Axelrod's exhortation to make as unambiguous and complete a model description and presentation of results as possible, and to facilitate other researchers' attempts to re–engineer one's model; see also the related discussion in (Valente and Andersen 2002) although we have our own reservations about the approach they propose. Replicability, in our view, is the hallmark of good science among experimental economists and psychologists alike ((**?**) and the commentaries on that article) and it seems worthwhile to establish it as a fundamental methodological tenet in agent-based modeling too[29]; for other tenets see (Hollenbeck 2000).

5.2 What we have accomplished

In the introduction we enumerated the key questions that motivated our study. These questions were prompted by recent developments in post–secondary education in the USA. If indeed our computational model speaks to that issue, then preliminary answers to our initial questions are as follows:

For–profits were, and are, clearly handicapped. How then could they succeed? Can the invasion of ever higher segments of higher education through for–profit "mutants" be stopped, or will for–profits ultimately invade segments 1 — 3, i.e. the "brand–name segment"?

The answer is clear: If firms manage to produce beyond their minimum efficient scale, they are bound to move up in the quality spectrum. The speed of this process is moderated by the initial quality of the mutant, by the degree of economies of scale (and taxes), and by the degree of competitiveness of the environment (what we called "scale" above).

What are the consequences of opportunistic behavior of colleges and universities (e.g., admittance of unqualified students for fiscal reasons)?

The answer, again, is clear: Recall that opportunistic firms are those that admit more than their fair share of unqualified students for fiscal reasons. So far we have observed that opportunistic firms tend to drift down the quality spectrum. In fact, they do so remarkably quickly for what seems rather small degrees of opportunism (e.g., already for a move from a maximum acceptable gap of 10 to one of 12 and 14, i.e., by an increase of the admittance interval of 10% and 20%, respectively.) In the long run, the equilibrium level of quality for opportunistic firms is at the bottom of the quality distribution.

[29] In this spirit, we will make available our code to interested researchers.

The computational model proposed here also improves on existing matching models in the Evolutionary Programming literature e.g., (Vriend 1995); (Kirman and Vriend 2001); (Tesfatsion 2001); (Pingle and Tesfatsion 2001) by introducing differences in quality on both the demand and supply side.Our computational laboratory enables us to do comparative statics exercises of changes in key parameters like the degree of opportunism that colleges and universities allow in their admission policies. Likewise, it allows us to understand better the likelihood of for–profits making inroads into post–secondary education. Clearly, our model lends itself to a wide range of such explorations.

We thus have proposed a computational model that allows us to investigate a wide range of issues pertaining to the (evolution) of post-secondary education, and we have explored its properties and related it to a simple theoretical model.

5.3 What remains to be done: A roadmap

We prefer to think of our computational model as a mere representation of existing or actual processes rather than an accurate model of post–secondary education. We do believe that our computational agents' decision making is a reasonable approximation of real agents' decision making. Rendering our computational model a more reliable laboratory of post–secondary education requires, in our view, not so much a more refined calibration of our computational agents as a more refined mapping of post–secondary education to our computational model.

Toward such a mapping, we shall implement in future work a distribution of consumer quality that mirrors more closely the distribution of students in the USA (which surely is not uniform, as we have assumed so far; we intend to proxy it through measures of quality distribution such as GRE or SAT scores that are readily available in the literature.) One interesting observation will be how different distributions of consumer quality affect rankings of firms. And, how in turn these rankings will affect the distribution of consumers if we were to endogenize quality.

Another refinement we plan to introduce is a turnover in buyers and sellers, especially continued entry and exit of buyers. While sellers in higher education tend to exist for remarkably long periods of time, it would be useful to allow firms to switch their status. In the current set-up, however, the outcomes of such refinements are fairly predictable.

Varying population size, in contrast, does not strike us as a problem worth pursuing since enrollments – while steadily growing – are rather stable, even through business cycles.

Yet another refinement we intend to study is the effect of (differential) levels of knowledge. We are particularly interested in understanding what happens if firms do not know consumers' quality or preferences (as parameterized by Δ^F) well, or if consumers systematically overestimate their own quality, a psychological fact for individuals well–documented in the literature. We could also have more complicated situations of asymmetric information such as firms or buyers with heterogeneous perceptions. All of these studies are easily implementable in our current computational laboratory.

Extensions of our model would, however, be required for other intriguing issues such as early admission: Exploring such jumping-the-gun strategies might require an increase of the dimensionality of the action space for both buyers and sellers, and thus definitely would change, and complicate the matching protocol.

References

Adelman, C. (2000): "A Parallel Universe: Certification in the Information Technology Guild," *Change*, May/June, 20 – 29.

Altonji, J. G., and T. Dunn (1996): "The Effects of Family Characteristics on the Return to Education," *Rev. Econ. Statis.*, 78, 692–704.

Arthur, W. B. (1991): "Designing Economic Agents That Act Like Human Agents: A Behavioral Approach to Bounded Rationality," *American Economic Review*, 81(2), 353–59.

——— (1994): "Inductive Reasoning and Bounded Rationality," *Complexity in Economic Theory*, 84(2), 406–411.

Axelrod, R. (1997): *The Complexity of Cooperation: Agent-Based Models of Conflict and Cooperation.* Princeton, N.J.: The Princeton University Press.

Behrman, J. R., M. R. Rosenzweig, and P. Taubman (1996): "College Choice and Wages: Estimates Using Data on Female Twins," *Rev. Econ. Statis.*, 78, 672–85.

Boylan, J. F. (1998): *Getting In.* Warner Books, New York.

Breneman, D. W. (1994): *Liberal Art Colleges: Thriving, Surviving, or Endangered?* Washington, DC: The Brookings Institution.

Bush, R. R., and F. Mosteller (1951): "A Mathematical Model for Simple Learning," *Psychological Review*, 58, 313–23.

——— (1955): *Stochastic Models for Learning.* Wiley, New York.

Camerer, C. F., T.-H. Ho, and J. K. Chong (2001): "Behavioral Game Theory: Thinking, Learning, and Teaching," .

Chattoe, E. (1998): "Just How (Un)Realistic are Evolutionary Algorithms as Representation of Social Process?," *Journal of Artificial Society and Social Simulation*, 1(3).

Chen, S.-H., J. Duffy, and C.-H. Yen (2002): "Equilibrium Selection Via Adaptation: Using Genetic Programming to Model Learning in a Coordination Game," *The Electronic Journal of Evolutionary Modeling and Economic Dynamics*, 1, 1–44.

Cosmides, L., and J. Tooby (1996): "Are Humans Good Intuitive Statisticians After All?: Rethinking some Conclusions of the Literature on Judgment under Uncertainty," *Cognition*, 58, 1–73.

Costa-Gomes, M., V. Crawford, and B. Broseta (2001): "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69, 1193–1235.

Facchina, B., E. Showell, and J. Stone (1993): "Privileges and Exemptions Enjoyed by Nonproft Organizations," *University of San Francisco Law Review*, 28, 85–121.

Gale, D., and L. Shapley (1962): "College Admissions and Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.

Gigerenzer, G., P. M. Todd, and ABC Research Group (1999): *Simple Heuristics That Make Us Smart*. Oxford University Press, Oxford, UK.

Goeree, J. K., and C. A. Holt (1999): "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96, 10564–7.

——— (2001): "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *The American Economic Review*, 91(5), 1402–22.

Holland, J. H. (1992): *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA.

Hollenbeck, J. R. (2000): "Consequences of Organizational Reward Systems," in *Computational Modeling of Behavior in Organizations : The Third Scientific Discipline*, ed. by D. R. Ilgen, and C. L. Hulin. APA Books, Washington, D.C.

Kirman, A. P., and N. J. Vriend (2001): "Evolving Market Structure: An ACE Model of Price Dispertion and Loyality," *Journal of Economic Dynamics and Control*, 25, 459–502.

Michalewicz, Z. (1999): *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag Berlin Heidelberg New York.

Mitchell, M. (1996): *An Introduction to Genetic Algorithms*. MIT Press/Bradford, Cambridge, MA.

Nagel, R. (1995): "Experimental Results on Interactive Competitive Guessing," *American Economic Review*, 85, 1313–26.

Ortmann, A. (1997): "How to Survive in Post-Industrial Environments: Adam Smith's Advice for Today Colleges and Universities," *Journal of Higher Education*, 68, 483–501.

——— (1998): "The Emergence of a For-Profit Higher Educational Sector: Recent Developments.," Presentation at the ARNOVA Conference, Seattle, WA.

——— (2001): "Capital Romance: Why Wall Street Fell in Love with Higher Education," *Education Economics*, 9(3), 293–311.

Payne, J. W., J. R. Bettman, and E. J. Johnson (1988): "Adaptive Strategy Selection in Decision Making," *J. Exp. Psych. Learning, Memory, and Cognition*, 14, 534–52.

Pingle, M., and L. Tesfatsion (2001): "Non-Employment and the Evolution of Worker-Employer Cooperation: Experiments with Real and Computational Agents," *ISU Economic Report No.*, 55.

Plott, C. R. (1987): "Dimensions of Parallelism: Some Policy Applications of Experimental Methods," in *Laboratory Experimentation in Economics. Six Points of View*, ed. by A. E. Roth. Cambridge University Press, Cambridge, MA.

Roth, A., and X. Xing (1994): "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions," *American Economic Review*, 84, 992–1044.

Roth, A., and X. Xing (1997): "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," *Journal of Political Economy*, 105, 284–329.

Roth, A. E., and M. A. O. Sotomayor (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press.

Schemo, D. J. (April 26, 2002): "A Key University Abolishes Early Decision Admissions," *The New York Times on the Web*.

Simao, J., and P. M. Todd (2002): "Mate Choice in Monogamous Mating Systems with Courtship," *Submitted to Adaptive Behavior*.

Simon, H. A. (1956): "Rational Choice and the Structure of Environments," *Psychological Review*, 63, 129–39.

Stahl, D. O. (1996): "Boundedly Rational Rule Learning in a Guessing Game," *Games and Economic Behavior*, 16, 303–30.

——— (2000): "Local Rule Learning in Symmetric Normal-Form Games: Theory and Evidance," *Games and Economic Behavior*, 32, 105–38.

Stahl, D. O., and P. Wilson (1995): "On Players Models of Other Players: Theory and Experimental Evidance," *Games and Economic Behavior*, 10, 213–54.

Tamura, R. (2001): "Teachers, Growth, and Convergence," *Journal of Political Economy*, 109(5), 1021–59.

Tesfatsion, L. (2001): "Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search," *Journal of Economic Dynamics and Control*, 25, 419–57.

——— (2002): "Agent-Based Computational Economics: Growing Economies from the Bottom Up," *ISU Economics Working Paper No.*, 1.

Todd, P. M., and G. Gigerenzer (2000): "Precis of Simple Heuristics That Make Us Smart," *Behavioral and Brain Sciences*, 23, 727–80.

Valente, M. (2002): "Comments of the Paper "Equilibrium Selection Via Adaptation: Using Genetic Programming to Model Learning in a Coordination", by Chen, Duffy and Yeh," *The Electronic Journal of Evolutionary Modeling and Economic Dynamics, http://www.e-jemed.org*, 102.

Valente, M., and E. S. Andersen (2002): "A Hands-on Approach to Evolutionary Simulation: Nelson-Winter Models in the Laboratory for Simulation Development," *The Electornic Journal of Evolutionary Medelling and Economic Dynamics, http:// www,e-jemed.org*, 1003.

Vriend, N. J. (1995): "Self-Organization of Markets: An Example of a Computational Approach," *Computational Economics*, 8, 205–231.

Weisbuch, G., A. Kirman, and D. Herreiner (2000): "Market Organisation and Trading Relationships," *The Economic Journal*, 110, 411–36.

Whitley, D. (1989): "The GENITOR Algorithm and Selection Pressure: Why Rank-Based Allocation of Reproductive Trials is Best," *Schaffer, J. David (Ed) the Proceedings of the Third International Conference on Genetic Algorithms, George Mason University, 4-7 June 1989 (San Mateo, CA: Morgan Kaufman)*, pp. 116–21.

Zemsky, R., S. Shaman, and M. Iannozzi (1997): "In Search of Strategic Perspective: A Tool for Mapping the Market in Post Secondary Education," *Change*, pp. 23 – 38.

A Theoretical equilibrium configurations

In the following we describe the symmetric steady state of the Q–model. "Symmetric steady state" denotes situations where every firm serves the same number of agents, where every firm has the same profit share (which is indeed what we observe empirically), and where quality of the firm equals average quality of its consumers. We calibrate the model so that average firm quality $\overline{Q}$ equals average consumer quality of 50 (which given our assumption of uniform distribution of consumers along the quality spectrum [0,100] is what we can expect on average.) We note that symmetric steady state implies $\overline{Q} = 50$ but that the reverse implication does not necessarily hold. Since firm quality is defined as a weighted sum of both average consumer quality and profits, we begin with the profit weight calibration before proceeding with an analysis of the equilibrium number of clusters and, in fact, the exact location of the clusters (cluster configuration).

A.1 Profit weight calibration

A firm's quality is updated according to the following rule,

$$Q = w_1 \cdot Q_{avg} + w_2 \cdot \pi, \tag{1}$$

where Q is the firm's quality, Q_{avg} the average quality of its consumers, π the firm's profits, and w_1 and w_2 are weights. Symmetric steady state profits are given by

$$\pi = \frac{N}{[Q]} \alpha Q,$$

where N is the number of consumers per firm (100 in all runs), $[Q]$ the quality range (100 in all runs), and α the profit share (average α is $0.46 - 0.48$ for different configurations, with a standard deviation $0.02 - 0.03$). Note that α is determined experimentally.

The requirement that $Q = Q_{avg}$ amounts to $Q(1 - \alpha w_2) = Q w_1$, or $w_2 = (1 - w_1)/\alpha$. The empirical value of w_2 which prompts $\overline{Q} = 50$ is indeed very close to the one just derived. For example, with 24 firms and $\alpha = 0.48$, $w_1 = 0.95$, the derived value $w_2 = 0.107$, while $\overline{Q} \approx 50$ requires an empirical value of $w_2 \sim 0.104$.

A.2 Equilibrium number of clusters

In this subsection we show why the configuration that we observe in most runs with number of firms modulo 6 (6 relatively tight clusters of firms) is a stable symmetric steady state for our choice of the quality range. For the sake of argument, assume that firms' quality is adjusted according to (1) with $w_1 = 1$ and $w_2 = 0$, *i.e.,* a firm's quality equals the average quality of its consumers. (Runs with this quality adjustment rule reveal the same distribution of 6 relatively tight clusters of firms.) Additionally assume that if a given number of T

consumers can be served by n firms, then T/n of them will be served by every firm, that is, competition leads to even distribution of consumers among firms in equilibrium. A firm will accept customers who are at least of quality $Q - \Delta$, where Q is the firm's quality. A customer will accept a firm that has at least quality of $Q_{cust} - \Delta$ where Q_{cust} is the customer's quality. Therefore, a firm with quality Q can serve only customers in the quality interval $[Q - \Delta, Q + \Delta]$.

Assume that all firms are in steady state. Assume next that one of them has, by some random disturbance, its quality adjusted upwards by dq. The firm under consideration loses some consumers at the lower end of its segment at quality $Q - \Delta$ but also obtains some consumers at the upper end of its segment at quality $Q + \Delta$. If n other firms are competing at the lower end and $n + j$ other firms can serve consumers at the upper end, the number of consumers lost and obtained are respectively $\frac{1}{n+1}\frac{dq}{[Q]}N_{tot}$ and $\frac{1}{n+j+1}\frac{dq}{[Q]}N_{tot}$, where N_{tot} is the total number of consumers. Thus, the new average quality of the firm if given by

$$\widetilde{Q} = \frac{\sum Q - \frac{1}{n+1}\frac{dq}{[Q]}N_{tot} \cdot (Q - \Delta) + \frac{1}{n+j+1}\frac{dq}{[Q]}N_{tot} \cdot (Q + \Delta)}{N - \frac{1}{n+1}\frac{dq}{[Q]}N_{tot} + \frac{1}{n+j+1}\frac{dq}{[Q]}N_{tot}},$$

where $\sum Q$ is the sum of the firm's consumers' quality in steady state and equal to $Q \cdot N$ by assumption. Dividing the numerator and denominator by $\sum Q$ and N respectively and using the fact that $\frac{1+x}{1+y} \approx 1 + x - y$ for $x \ll 1$, $y \ll 1$, one obtains

$$\widetilde{Q} \approx Q + \frac{dq}{[Q]}\frac{N_{tot}}{N}\left\{\frac{Q+\Delta}{n+j+1} - \frac{Q-\Delta}{n+1} - \frac{Q}{n+j+1} + \frac{Q}{n+1}\right\},$$

or

$$dq' = \widetilde{Q} - Q \approx \frac{dq}{[Q]}\frac{N_{tot}}{N} \cdot \Delta \cdot \frac{n+1+n+j+1}{(n+1)(n+j+1)}.$$

The steady state is stable if random fluctuations in quality are dampened over time, or $|dq'| < |dq|$[30]. Therefore, stability of the steady state depends on the magnitude of the following term

$$\frac{\Delta}{[Q]}N_f \cdot \frac{n+1+n+j+1}{(n+1)(n+j+1)}, \tag{2}$$

where $N_f = \frac{N_{tot}}{N}$ is the number of firms in the economy.

Let us consider some special cases of (2). Suppose that a firm in steady state does not have any competition at the lower end of its segment, $n = 0$. Then (2) becomes $\frac{\Delta}{[Q]}N_f \cdot \frac{2+j}{1+j}$, and for parameter values ($\Delta = 10$, $[Q] = 100$) this expression is greater than one for any $j > 0$, and any $N_f \geqslant 10$. In other words, any steady state that implies no competition at the lower end is not stable, because a random upward quality movement is amplified. Similarly, suppose that there is no competition at the upper end of a firm's segment. In this case,

[30] In other words, we want the eigenvalue of the difference equation $Q_{n+1} = f(Q_n)$, linearized around the steady state, to be less than one. It is always positive, therefore an oscillating dynamics around the steady state is impossible.

$j = -n$, and (2) is $\frac{\Delta}{[Q]} N_f \cdot \frac{2+n}{1+n}$ which is again greater than one for any $n > 0$, and any $N_f \geqslant 10$. Therefore, a steady state involving zero competition at the upper end cannot be stable.

The preceding result demonstrates that steady states with fewer than five clusters are unstable, because they necessarily involve zero competition either at the lower or at the upper end of the quality segment. How about five segments then? Assume a steady state with five firm clusters, numbered in ascending quality order. Given the parameter values that we used for our treatments, $\Delta = 10$, $[Q] = 100$, the five firm clusters will be located at qualities 10, 30, 50, 70, and 90. Suppose now that clusters number 2 and 4 move down and up, respectively. In this case, a firm from cluster 3 that randomly increased its quality by dq will have $dq' > dq$, while the one that had its quality decreased by dq will have $|dq'| > |dq|$. (Recall that dq' is the deviation from steady state after one iteration.) In other words, cluster 3 will be torn apart by any non–negligible simultaneous movements of clusters 2 and 4. Therefore, a configuration with 5 clusters is stable but the associated basin of attraction is very small. In numerical simulations with $N_f = 10$ we have observed stable constellations with 5 clusters of firms only once or twice every 100 runs.

Why, then, do we observe constellations of 6 clusters for runs with large number of firms, say 24 and 48? And why do we observe constellations with between 6 and 8 clusters for runs with 10 firms? Compare two steady states, one with C clusters and another with $C + 1$, where $10 > C > 5$. A firm that moved up by dq faces the same competition at its lower end from members of its own cluster and the lower one, with the total number of other firms given by $\frac{N_f}{C} - 1 + \frac{N_f}{C}$ (disregarding integer constraints). On the other hand, at the upper end of its segment competition from members of its own cluster disappears, and only that from the upper cluster remains. Therefore, $j = 1 - \frac{N_f}{C}$. (2) is now proportional to $\frac{\Delta}{[Q]} N_f \cdot \frac{3\frac{N_f}{C}+1}{2\frac{N_f}{C}\cdot\left(2\frac{N_f}{C}+1\right)}$ or

$$\frac{\Delta}{[Q]} N_f \cdot \frac{C \cdot (3N_f + C)}{2N_f \cdot (2N_f + C)}. \tag{3}$$

The partial derivative of the preceding expression with respect to C is proportional to

$$\frac{2N_f \cdot \left(6N_f^2 + 4CN_f + C^2\right)}{4N_f^2 \cdot (2N_f + C)^2},$$

which is always positive. Therefore, the movement to a higher number of clusters implies a larger eigenvalue, and hence a less stable steady state.[31]

Summarizing the results, we see that configurations with 4 clusters are unstable and those with 5 clusters are likely to be destroyed even by small fluctuations. Furthermore, configurations with more than 6 clusters are less stable than those with 6, and indeed they are increasingly less stable as the number of

[31] A similar result is true for any number of clusters. The math, however, becomes tedious.

clusters goes up. Therefore, in numerical simulations one is likely to observe a configuration with 6 clusters.

Finally, observe that with $C = 6$, (3) equals 0.42 for $N_f = 12$, 0.433 with $N_f = 24$, and approaches 0.45 as $N_f \to \infty$. This means that configurations of 6 clusters are always stable for any number of firms.

A.3 Cluster configurations

Having established theoretically the most likely distribution of clusters, we next calculate their exact location in the symmetric steady state with C clusters. We assume that there is an equal number of firms in each cluster. Under the symmetric steady state assumptions spelled out in the previous subsection, calculations are the same for one or n firms in a cluster; we thus restrict our discussion to one firm per cluster.

Order quality locations in a symmetric steady state in ascending order from Q_1 to Q_C. For $10 \geqslant C \geqslant 5$, the first firm (remember we restrict our discussion to one firm per cluster) has no competition at its lower end and competition from the second firm only at the upper end. Denote as D the density of customers per unit of quality. Then the first firm will serve customers located in $[0, Q_2 - \Delta]$ alone and those in $[Q_2 - \Delta, Q_1 + \Delta]$ together with the second firm. Since in symmetric steady state the average quality of consumers equals its own quality, we have

$$\begin{aligned} Q_1 &= \frac{D \int\limits_0^{Q_2-\Delta} Q dQ + \frac{1}{2} D \int\limits_{Q_2-\Delta}^{Q_1+\Delta} Q dQ}{D \int\limits_0^{Q_2-\Delta} dQ + \frac{1}{2} D \int\limits_{Q_2-\Delta}^{Q_1+\Delta} dQ} = \\ \frac{1}{2} \frac{(Q_2-\Delta)^2 + (Q_1+\Delta)^2}{Q_2 - \Delta + Q_1 + \Delta} &= \frac{(Q_2-\Delta)^2 + (Q_1+\Delta)^2}{2 \cdot (Q_2 + Q_1)}. \end{aligned} \tag{4}$$

Consider now the second firm. It is the sole provider to consumers in $[Q_1 + \Delta, Q_3 - \Delta]$ and a joint provider with first and third firm in $[Q_2 - \Delta, Q_1 + \Delta]$ and $[Q_3 - \Delta, Q_2 + \Delta]$, respectively. The symmetric steady state condition then becomes

$$\begin{aligned} Q_2 &= \frac{\frac{1}{2} D \int\limits_{Q_2-\Delta}^{Q_1+\Delta} Q dQ + D \int\limits_{Q_1+\Delta}^{Q_3-\Delta} Q dQ + \frac{1}{2} D \int\limits_{Q_3-\Delta}^{Q_2+\Delta} Q dQ}{\frac{1}{2} D \int\limits_{Q_2-\Delta}^{Q_1+\Delta} dQ + D \int\limits_{Q_1+\Delta}^{Q_3-\Delta} dQ + \frac{1}{2} D \int\limits_{Q_3-\Delta}^{Q_2+\Delta} dQ} = \\ & \frac{(Q_3-\Delta)^2 + (Q_2+\Delta)^2 - (Q_2-\Delta)^2 - (Q_1+\Delta)^2}{2 \cdot [(Q_3-\Delta) + (Q_2+\Delta) - (Q_2-\Delta) - (Q_1+\Delta)]}. \end{aligned} \tag{5}$$

After some algebra, (5) transforms into

$$Q_2 = \frac{Q_1 + Q_3}{2}, \tag{6}$$

which says that the symmetric steady state location of the second firm is exactly between the first firm and the third firm. It is trivial to show that a similar result will obtain for all other firms located in the interior of the quality spectrum,

$$\begin{aligned} Q_3 &= \frac{Q_2 + Q_4}{2}, && (7a) \\ Q_4 &= \frac{Q_3 + Q_5}{2}, && (7b) \\ & \ldots && (7c) \\ Q_{C-1} &= \frac{Q_{C-2} + Q_C}{2}. && (7d) \end{aligned}$$

Finally, for the last firm C, the symmetric steady state condition is given by

$$Q_C = \frac{2 \cdot [Q]^2 - (Q_C - \Delta)^2 - (Q_{C-1} + \Delta)^2}{2 \cdot [Q] - Q_{C-1} - Q_C}. \qquad (8)$$

Combining (6) and (7) we obtain $Q_4 = 3Q_2 - 2Q_1$, $Q_3 = 2Q_2 - Q_1$ or $Q_3 = Q_2 + (Q_2 - Q_1)$, $Q_4 = Q_3 + 2(Q_2 - Q_1)$. In other words, firms are located at equal distance $\delta = (Q_2 - Q_1)$ from each other. The problem of finding symmetric steady state locations is thus reduced to solving a system of two quadratic equations, (4) and (8), in two unknowns, Q_1 and δ (remember that $Q_2 = Q_1 + \delta$, $Q_{C-1} = Q_1 + (C-2) \cdot \delta$, $Q_C = Q_1 + (C-1) \cdot \delta$). The solution can be found numerically when C, the number of clusters in symmetric equilibrium, is given.

In the previous subsection we have argued that given our parameter values Δ and $[Q]$, the symmetric steady state with 6 clusters should be the most stable one. Steady state positions with 6 clusters are given by [8.48; 25.09; 41.70; 58.30; 74.91; 91.52].

B Pseudo code

The following pseudo code presents our matching model in a manner more palatable to programmers. Subroutines and parameters are in CAPITAL LETTERS while variables programming language commands are expressed in lower case.

```
program MAIN;
begin
CREATE_FIRMS, CREATE_CONSUMERS
for iteration 1 to RUNLENGTH do
    RESET_FIRMS_AND_CONSUMERS
    FIRMS_CALCULATE_PRODUCTION_AND_SIGNALING
    FIRMS_SIGNAL_DESIRABLE_CONSUMERS
    CONSUMERS_SELECT_DESIRABLE_FIRMS
    FIRMS_ACCEPT/REJECT_CONSUMERS
    FIRMS_COMPUTE_PROFITS
    FIRMS_COMPUTE_AVE_CONSUMER_QUALITY
    FIRMS_UPDATE_QUALITY
    FIRMS_REINFORCE_WEIGHTS
    CONSUMERS_REINFORCE_WEIGHTS
    if iteration modulo 50 then run_GA on firms' rules
end

CREATE_FIRMS
    profit[0:MEMORY_SIZE]=0.
    rules[1:NUM_FIRM_RULES].weight=INIT_FIRM_WEIGHT
    RANDOMLY_GENERATE_RULES [PRODUCTION_&_SIGNALING_PAIRS]
    RANDOM_QUALITY[0:99]

CREATE_CONSUMERS
    rules[1:NUM_CONS_RULES].weight=INIT_CONS_WEIGHT
    rules[1] = 'IF not SAT AND no INFO THEN PATR'
    rules[2] ='...
    rules[5] = 'IF not SAT AND INDIFFERENT to INFO THEN PATR'
    rules[6] = '...
    rules[18] = 'IF INDIFFERENT to SAT AND INDIFFERENT to INFO
THEN KNOWN'
    RANDOM_QUALITY[0:100]

RESET_FIRMS_AND_CONSUMERS
    ADJUST_WEIGHTS_TO_[0,1]
    RESET_DESIRED_FIRMS_LIST
    RESET_SATISFACTION
    WRITE_THE_STATE_VECTOR

FIRMS_CALCULATE_PRODUCTION_&_SIGNALING
```

```
        SELECT_RULE (stochastic auction)
        CALC (production, number_of_signals)

    FIRMS_SIGNAL_DESIRABLE_CONSUMERS
    for i = 1 to number_of_signals
        SELECT_RANDOM_CONSUMER
        if MIN_CONS_QUAL ≤ CONS_QUAL ≤ MAX_CONS_QUAL
        then SEND_SIGNAL

    CONSUMERS_SELECT_DESIRABLE_FIRMS
        SELECT_RULE (stochastic auction)
        if action = 'PATR' then firm_selected = last_firm
        if action = 'KNOWN' then
            firm_selected = 'NONE'
            for i = 1 to NUM_FIRMS
                SELECT_FIRM_FROM_AMONG_THOSE_THAT SIGNALED
                if FIRM_QUAL ≥ MIN_FIRM_QUAL then
                    firm_selected = i
                    firm_desired[i] = true
                end
            end
            last_firm = firm_selected
        end

    FIRMS_ACCEPT/REJECT_CONSUMERS
        stock = production
        demand = 0
        for i = 1 to NUM_CONSUMERS
            for j = 1 to NUM_FIRMS
                if CONS.firm_desired = firm and firm.quality => MIN_FIRM_QUAL
then
                    demand = demand + 1
                    if stock > 0 then
                        SERVE
                        stock = stock - 1
                        last_firm = firm
                    end
                end
            end
        end

    FIRMS_COMPUTE_PROFITS
        profit = {price*min (production, demand) - cost(production, number_of_signals)}*qual

    FIRMS_COMPUTE_AVERAGE_CONSUMER_QUALITY
```

```
FIRMS_UPDATE_THEIR_QUALITY
    quality = weight_quality*ave_consumer_quality + weight_profit*profit

FIRMS_REINFORCE_WEIGHTS


CONSUMERS_REINFORCE_WEIGHTS

RUN_GA
For k = 1 to NUM_FIRMS
    SELECT_TWO_PARENTS_FROM_TOP 25%_OF_RULES
    DO_CROSSOVER
    DO_MUTATION
    REPLACE_RULE_FROM _BOTTOM_HALF
```

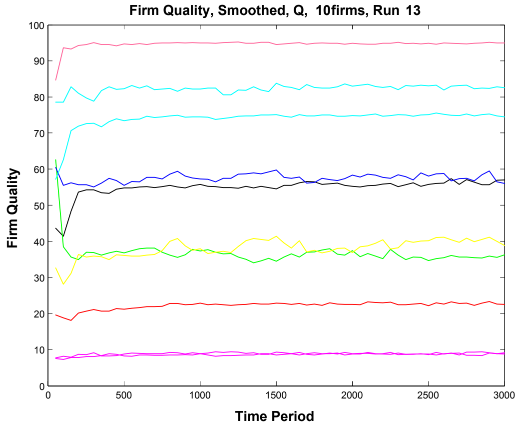

Figure 1

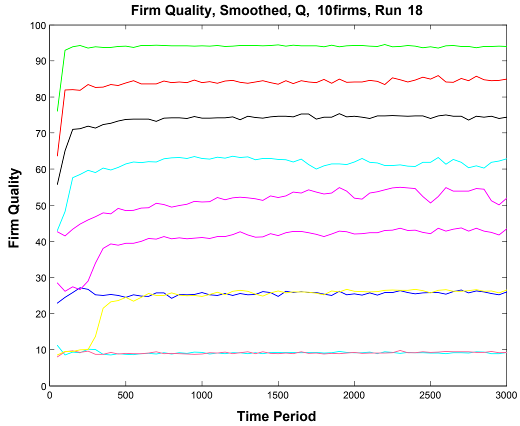

Figure 2

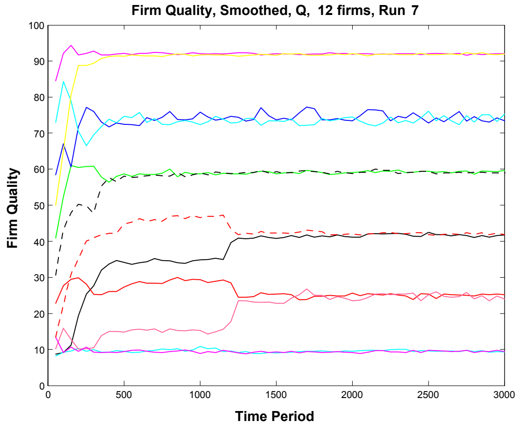

Figure 3

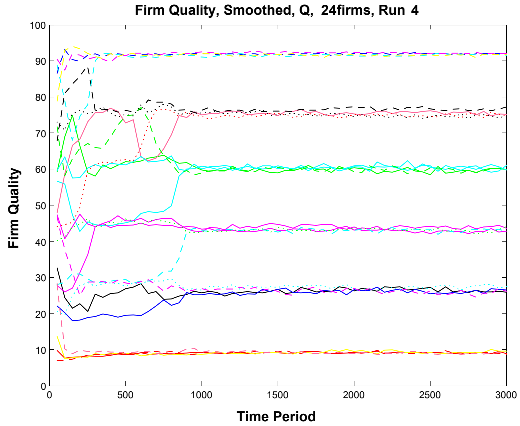

Figure 4

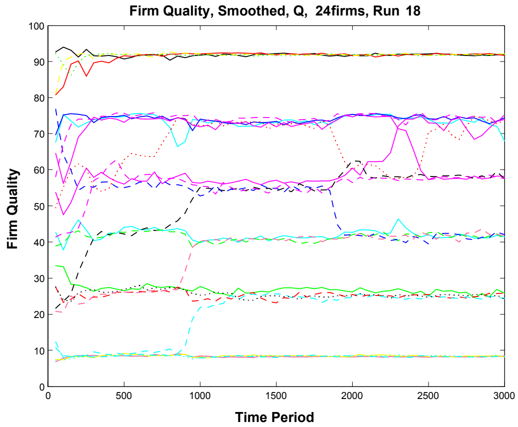

Figure 5

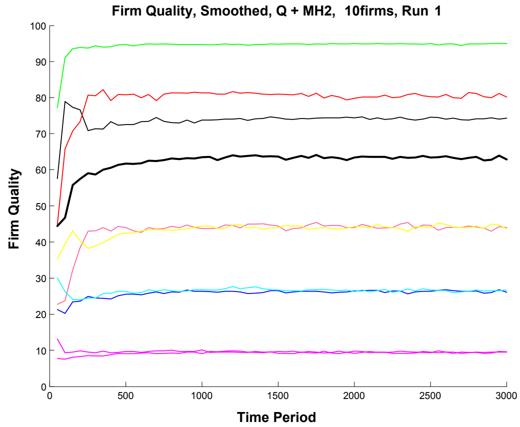

Figure 6

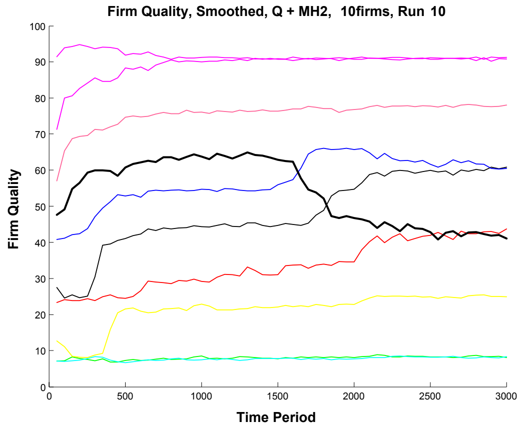

Figure 7

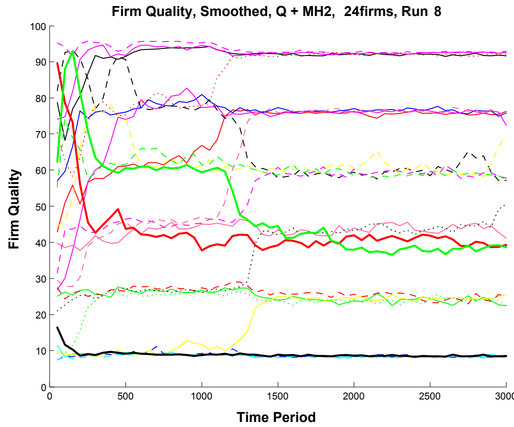

Figure 8

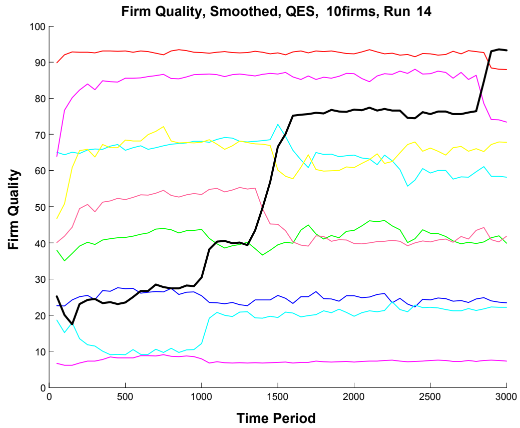

Figure 9

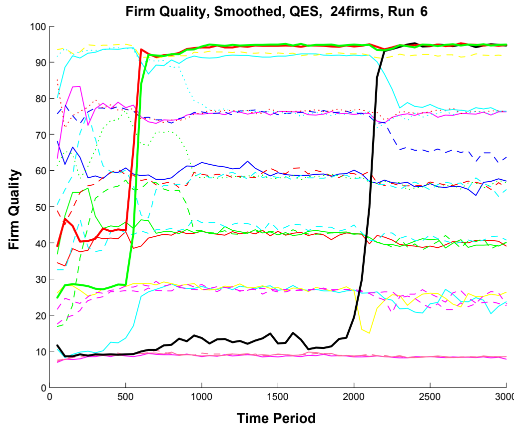

Figure 10

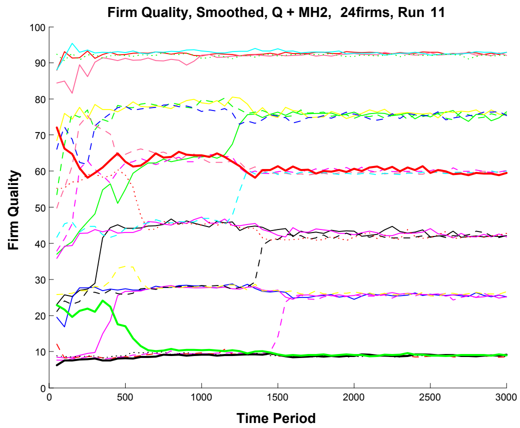

Figure 11

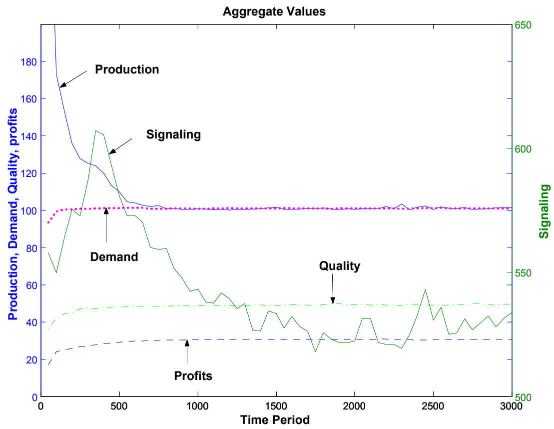

Figure 12

CERGE-EI
P.O.BOX 882
Politických vezòù 7
111 21 Prague 1
Czech Republic
http://www.cerge-ei.cz