Free Market Control for a Multi-Agent Based Peer Help Environment

Kevin Kostuik and Julita Vassileva

Department of Computer Science,
University of Saskatchewan,
57 Campus Drive, Saskatoon S7N 5A9, Canada
[kkostuik@skatter.usask.ca; jiv@cs.usask.ca]

Introduction

Electronic commerce allows not only creating new market opportunities for trading with the traditional goods, but also an opportunity for trading with new intangible types of goods, for example, knowledge, advice, teaching, or help. In this paper we discuss a distributed, multi-agent-based environment (Greer et al., 1998), which supports peer help in university or workplace educational setting and is based on a free market economic model. This environment can be viewed as a special case of electronic market because of the following features:

- There are users who possess some goods or resources (knowledge in this case) and users who need these goods / resources (asking for help or advice);
- A person interested to buy a good must find a seller who offers the good with acceptable quality and at acceptable price. In our case a user with a specific help request needs to find a competent helper;
- The buyer is willing to pay some amount of money in order to achieve the goal of gaining some knowledge, while the seller is willing to give away some knowledge in exchange for money. The goal of accumulating some resource, like money, which can be exchanged with some goods (promotion, salary increase in workplace environment, marks in University environment) creates a motivation for knowledgeable users / students to participate.
- The price of a certain good depends on the offer and the demand for this good on the market. People having exceptional and highly demanded knowledge / expertise, can put higher prices for their advice.
- There is some cost associated with supplying the buyer with the good; helping costs some time for the helper, which could be used for achieving some other goal.

The architecture of our peer help environment consists of personal agents representing the users /students. The agents maintain user models containing information about the users' goals (help requests, current goals), knowledge resources / competencies and about the relationships existing between the users. The users communicate with their agents to update their models and to assert their goals. The agents communicate with each other and with matchmaker agents to search for appropriate helpers for their users depending on their help-request and then negotiate about the price. The reasoning and negotiation mechanisms of the agents are a topic of a different paper (Mudgal & Vassileva, submitted).

In this paper we focus on the design of the agent economy. An economy is needed in order to predict and control the overall behaviour of the complex system, emerging from the actions of the individual agents and users. Without such control, a number of problems may arise which may limit or completely block some users from participation.

Designing an economic coordination mechanism is an exercise in emergent control — within an economic infrastructure, global dynamics emerge from the aggregation of local agent choices. The design task is to define an infrastructure such that local optimizations lead to desired global dynamics, subject to constraints on the welfare of individual agents. Since general design principles in the field of market-based controls are not well developed, the design must appeal to an economic metaphor: assuming that students within an intelligent tutoring system have motivations similar to consumers within a general economy, an economic infrastructure that models a real economy can coordinate the distribution of learning resources. Given an economic metaphor, established economic theory provides insights into resource allocation and control parameters for different economic models.

After analysis of various e-commerce mechanisms, like retail auctions, classified ads and stock markets (for a review and comparison among these, see (Guttman & Maes, 1998)) we decided that free-market models provide functional and extendible control mechanisms for a peer help environment. Within the context of such environment, an effective coordination mechanism must be stable, seek equilibrium quickly, and provide access to control parameters. Economic theory and practice both suggest the functionality of a free-market system. Theory proves that, in the absence of externalities[1], competitive equilibria correspond to efficient resource allocations. (Cornes, 1992: p.240) Practice demonstrates that fiscal, monetary, and trade policy can guide macroeconomic behaviour, while taxation and redistribution can support individual agent welfare. Free-market mechanisms are extendible because they provide a foundation for other economic models. For example, allowing collective ownership opens the possibility of co-operative economies, and introducing global control elements creates command-style economies.

In other words, a free market economy allows a self-organizing mechanism plus a fairly wide variety of options for controlling the economy from outside to stimulate a desirable overall behaviour of the system. This is particularly important in our case, where the goods involved are intangible and the overall goal of the environment is to support learning, i.e. the "wellness" of the economy is not measured by the money turnover, but by the accumulated knowledge of the participating users. This leads also to a number of differences between a market for peer help and a conventional market.

- In conventional market the prices emerge and develop historically, electronic markets start usually with real-world prices. In our case, due to the unusual nature of the good "knowledge", there is no price history.
- We can't expect a purely rational behavior from the sellers and buyers ("rational" means always trying to maximize their profit). Psychologists show that people may help with various motives, like for example, to increase their status / reputation in the group, to create or affirm a relationship, or just from altruism.
- One has to consider always the fact, that there are two levels of interaction in the environment: the "real world", where the players are users / students, and an "agent world" where the players are the personal agents. The economic activity: price definition, market activity, negotiation happens in the agent world (among the personal agents), while the users / students can "cash" at certain points the wealth accumulated by their agents with real world equivalent, but are not involved directly in the economic activity (in order not to become preoccupied with it and forget their main goals of learning). The real world can influence the agent world in many ways. Since users are autonomous, i.e. having final approval over their personal agent's choices (Friedman and Nissenbaum, 1997) sometimes negotiated deals between the agents can fail due to unpredictable (for the agents) circumstances. In addition, deliberate "real world" influences can be introduced in the "agent world" aimed at controlling the economy, which can be done by a system administrator, for example, introducing a taxation, free resources (on-line tutorials), banning certain users / agents from participation.

In the rest of the paper we describe the requirements to the design of an agent economy for the peer help environment and the design of a simulation which is supposed to give us an idea of how such system is going to work when applied in reality.

---

[1] In the economic sense, externalities are goods (or bads) that consumers care about, but cannot control the quantity of because there are no associated markets (Cornes 1992 p.237). For example, a free on-line tutorial introduced by an administrator in the environment would be an externality.

Requirements

Meta Level Goals

The economic coordination mechanism has meta level goals existing above the economic layer. At the meta level, the system should:

• *Involve all students*. All students should benefit from the system. Particularly strong students have much knowledge to contribute and little knowledge to gain; the system must positively motivate such students to participate. Poorer students have much knowledge to gain and few means to earn currency and buy it; the system must give such students the opportunity to participate.

• *Create general enthusiasm*. Participation should be productive, enjoyable, and easy. A game-like atmosphere would help; however, students must still focus on learning and not just devote themselves to playing the game.

• *Allow unclassified resources into the economy*. Learning resources are greatly varied, and change quickly over time. Even the most exhaustive system analysis would overlook valuable resources e.g. tutorials, FAQs, discussion forum threads, on-line materials, so the system should be flexible enough to incorporate new and unanticipated resources.

• *Require little maintenance*. In the spirit of a free-market economy, decision making should be decentralized, and the system should evolve properly without significant intervention from the system administrator.

• *Be generally applicable*. The system must be readily adaptable to intelligent data systems and adaptive organizations in the business context.

• *Obtain measurable results*. Validating the system requires demonstrating real improvements; therefore, the system must provide summary statistics that track trading behaviour and student performance. A system simulation provides an initial test platform. Rigorous testing requires sample user groups and a working prototype.

Desired Macroeconomic Behaviours

At the macroeconomic level, the coordination mechanism must:

• *Establish and maintain trade*. Whereas real economies evolve into existence, a tutoring economy begins abruptly with the system's deployment. The coordination mechanism must be responsive and robust enough to reach a stable equilibrium soon after deployment; the economy must provide students with immediate and persistent benefits to ensure the system's initial and continuing use. A currency helps to establish and maintain trade by removing the double coincidence of wants required for barter trading[2].

• *Achieve a well-behaved price level*. Efficient consumer participation in an economy requires faith in the economy's currency. A chaotic or inflationary price level erodes the currency's purchasing power and ability to store value, and forces consumers into barter trading. Furthermore, the price level signals supply and demand conditions in the market; an ill-behaved price level reveals little information, and does not help producers and consumers adjust their behaviour to market conditions.

---

[2] Unlike (Boyd, 1997) we believe that currency trades have advantages over barter trades in educational context. In barter trades, the purchaser must have expendable goods that the seller wants. With currency trades, the purchaser must only have expendable goods. The purchaser can split the trade into two more common transactions, with the currency being an intermediate unit of exchange.

• *Realize net gains from trade*. The overall knowledge level must increase as a consequence of the system. Students who sell what they are relatively good at producing and buy what they are relatively poor at producing will realize gains from trade. Both knowledge and currency may be gained from the trade.

• *Distribute benefits fairly*. Defining a fair distribution of benefits in a real economy is a difficult normative problem. The problem is even more acute in educational contexts where equality is precious and often fiercely defended. As a free-market economy, the mechanism initially distributes wealth to those who produce it. Efforts to mediate the accumulations of wealth common in free-market economies require some sort of re-distribution policy. Since the definition of fair distribution is best left to administrators of the tutoring system, the system must provide means to implement re-distribution policy.

A reasonable hypothesis is that there exists a critical complexity, or phase transition, for the economy. When many students are on-line, trade possibilities are rich, and when few students are on-line, trade possibilities are sparse. The critical complexity hypothesis states that the economy responds abruptly to the richness of trade possibilities: there exists some critical number of users above which the economy functions effectively, and below which the economy freezes.

The proposed tutoring system is complicated because not all users are on-line and actively trading at all times. While the total number of participants is expected to be on the order of hundreds, the number of students on-line at any one time is expected to be on the order of tens. Trade possibilities among all participants are rich; however, trade possibilities among on-line users are potentially sparse. A simulation scheme presented later in this paper can test the critical complexity hypothesis and suggest whether the economy will function effectively for expected user populations.

Agent Welfare Constraints

Desired macroeconomic behaviours are subject to constraints on individual agent welfare. From the perspective of individual agents, the coordination mechanism must:

• *Encourage win-win transactions*. A zero-sum game is unacceptable in an educational context; if one user's gain is another user's loss, then there is no net improvement from the system. Users ultimately make trade decisions, and are responsible for ensuring their own welfare; however, the system must also present opportunities to make value-for-value, mutually beneficial trades.

• *Not make anyone worse off*. The coordination mechanism must seek Pareto preferred states; that is, states in which at least one agent is better off and no agent is worse off. In this sense, Pareto optimal states are states from which no agent can be made better off without making some other agent worse off (Burstein, 1968: p.230). Collectivist perspectives may challenge Pareto optimality and prefer the maximization of average or minimum welfare; however, seeking collective welfare requires abandoning the de-centralized, free-market metaphor. Seeking collective welfare within a free-market structure threatens the participation of individual users / students, particularly the most productive and competent ones; only coercion ensures the participation of students made worse off by the system.

• *Avoid unreasonable wealth accumulation*. The peer help economy may freeze, if extreme wealth accumulations occur. Once a small number of magnates capture the majority of the economy's currency, the only ways to release currency back into the economy are lending, investment, and luxury expenditure. Students should focus on learning rather than on loan sharks, employers, and luxury production.

A competitive market without externalities ensures win-win transactions and Pareto optimality. Self-serving rational agents refuse all detrimental trades, so the absence of externalities ensures win-win transactions. An important theorem of welfare economics states that any competitive equilibrium is Pareto optimal (Shoven & Whalley, 1992: p.34), but only in the absence of externalities (Cornes, 1992: p.240). However, since the relevance of externalities to the learning economy is unknown, the system must have low-level checks and balances to ensure these qualities.

Protection against excessive wealth accumulation is a controversial topic in real economies, but it is less troublesome in learning economies. Real economies focus on acquisition and property, while learning economies focus on improvement and dissemination. The two focuses are different, thus the perspectives on wealth accumulation should be different. Conclusions drawn from real economies are not necessarily relevant in the educational context.

Design of a Peer Help Economy

Design Abstraction

Creating an abstraction of the peer help environment eases the design task. A useful abstraction captures the essence of economic interactions in the system, and eliminates low-level details that have little or no impact.

A suitable abstraction considers the system as a market in a single differentiated good, where knowledge is the single good. The system's actual goods are a wide range of learning materials; however, knowledge is the fundamental quality sought from all materials. Knowledge is a differentiated good because its actual and perceived quality varies — tutoring ability varies from student to student, and information correctness varies from source to source. A currency facilitates trades of knowledge, and competitive local negotiation determines prices. Furthermore, the design abstraction does not distinguish between personal software agents and their users; the abstraction treats the personal agent and its user as a single decision making agent. This decision making agent is assumed to be predominantly self-serving and independent of externalities.

Nature of Local Optimization

Considering the nature of local optimizations yield insights into the system's global dynamics. Traditional economic analysis treats producers and consumers separately, then links the two by supply and demand relationships. Since each student plays both a production and consumption role, a more unified and symmetric treatment is desirable.

Objective Function

A hybrid objective function captures students' combined production and consumption roles. The system's economic flows consist of currency and knowledge. Students seek to attain both currency and knowledge. Let M represent the amount of currency a student has, and let G be a measure of the student's knowledge. Considering money, M, and grade, G, as the two most important student goals, a student's objective function, π, is some function of that student's change in money, ΔM, and change in grade, ΔG, over the system period. The system period is the total time over which the system is used; in an educational setting the system period is an academic term. Each agent locally optimizes this objective function for both purchases and sales. The agent may choose a simple greedy optimization or some other more complex optimization strategy.

A suitable first approximation considers the objective function as a linear combination of the money and grade goals.

$$\pi = \sigma(\Delta M) + (1-\sigma)(C \bullet \Delta G) \quad (1)$$

The parameter σ ranges from 0 to 1, and represents the relative importance of the money and grade goals. C is a conversion factor between units of grades and units of money. Each student has a unique goal preference parameter and grade conversion factor.

Goal Achievement

Achievement of the money and grade goals depends upon contribution to the system and usage of the system. Consider the abstract knowledge economy presented earlier. A student's change in money balance varies with the net amount of knowledge purchased. Call the amount of knowledge sold Teaching, and the amount of knowledge bought Usage, then

$$\Delta M \propto \text{Teaching - Usage} \quad (2)$$

A student's change in grade also depends upon the amount of knowledge bought and sold, although not in a simple linear relationship. Students learn both by teaching and by being taught. Buying knowledge clearly contributes to a student's grade improvement. Selling knowledge contributes to grade improvement less obviously. Before students can sell knowledge, they must review and summarize class material; review and summary both help students to learn.

Unlike the money goal, the grade goal does not increase linearly with the amounts of knowledge bought and sold; instead, the grade goal experiences diminishing marginal returns. Teaching provides less benefit as the total amount of teaching increases; each time the student reviews material, learning something new becomes more difficult, until eventually the student masters the topic. The same argument holds for being taught. Denoting the marginally declining grade improvements of Teaching and Usage functions with stars, the change in grade is roughly

$$\Delta G \propto \text{Teaching* + Usage*} \quad (3)$$

Expected Outcome

Examining the relationships in (2) and (3) suggests that stratification may be a problem in the system. The most knowledgeable students have little knowledge to gain and much knowledge to sell. Since highly knowledgeable students cannot improve their grades by much, their goal preference parameter, σ, will be near 1. This *teacher-type user / student* will experience a significant increase in money balance with some increase in grade.

The least knowledgeable students have much knowledge to gain and little knowledge to sell. Since less knowledgeable students can improve their grade much and have little knowledge to sell, their goal preference parameter will be near 0. This *user-type user / student* will experience a significant decrease in money balance with some increase in grade.

Users / students between these two types have enough knowledge to sell, yet can still benefit from buying knowledge; their goal preference parameter will be near 0.5, and they will produce and consume in nearly equal amounts. Profits from production can buy knowledge, which can once again be incorporated and re-sold. This *dual-type user/ student* will experience a nearly constant money balance with a larger increase in grade.

In the absence of a re-distribution mechanism, it appears that benefits will fall primarily to average and better than average students: teacher-types receive money and dual-types receive better grades, while user-types receive only a limited grade improvement.

Economic Infrastructure

An economic infrastructure was designed after analyzing desired macroeconomic behaviour, agent welfare constraints, meta-level goals, and the model objective function. An effective free-market coordination mechanism requires:

• *Currency that is visible to users and has external value*. The currency should be visible to human users, rather than just being a low level unit of accounting for software agents. Making students aware of their personal agent's money balance fosters a game-like atmosphere, and encourages student involvement and enthusiasm. The currency should also be externally valuable, rather than just being meaningful within the peer help context. Assigning the currency value outside of the peer help economy motivates participation of teacher-type students, and increases confidence in the currency. System administrators can assign the currency external value by allowing students to redeem it for goods other than learning resources; for example, a store from which students can buy tangible goods at the end of the system period serves this function.

• *Tax system*. The peer help economy must dynamically re-distribute currency to mediate excessive wealth accumulation and to ensure individual agent welfare. A tax system provides a control mechanism for the economy; system administrators can implement fiscal policy by altering tax structures, and by spending tax revenues on public goods and agent welfare schemes.

• *Initial currency endowment*. The economy requires an initial seeding of currency to establish trade. Students might buy into the economy by purchasing peer help currency with real world dollars; this de-centralized currency endowment is beneficial because it implicitly assigns the currency an external value. System administrators might also inject an amount of currency with external value equal to the amount they are willing to invest in the economy; this centralized currency endowment is beneficial because it provides another economic control parameter. A combination of both endowment types is possible.

• *Well known and stable money supply*. System administrators control the amount of currency in the economy. Confidence in the currency requires a well-known and stable money supply, so administrators must implement monetary policy wisely and not change the money supply abruptly.

• *Posted trade prices and volumes*. Public knowledge of trade prices and volumes helps to establish a well-behaved price level. Past trades transmit information about the market's supply and demand conditions, and give a starting point for negotiation between agents.

• *Policing*. Users face the risks of broken agreements and poor quality information when they trade. Successful deal making requires the evaluation and minimization of such risks. A policing mechanism can minimize risk by enforcing agreements, and by notifying users of poor quality helpers. Possible policing strategies range from de-centralized reputation based approaches to centralized authority based approaches.

The economic infrastructure presented here is basic; however, it meets the system's requirements at the meta, macroeconomic, and agent levels. The most notable omission from the economic metaphor is a banking system. While banking allows more sophisticated monetary control policy, it adds an unnecessary level of complication for a simple peer help system. Note that the omission of banking does not threaten the system's generality — the addition of centralized and de-centralized banking is a natural extension of the current infrastructure.

The crux of the design problem is compromising between states that are optimum from an economic perspective and states that are optimum from a social welfare perspective. We chose a competitive market because it seeks a Pareto optimum. We chose a taxation system because it maintains economic activity and achieves a fair wealth distribution. However, these two goals have conflicting constraints. Attempts to control the economy with currency re-distributions and public goods impose externalities, and move the economy away from Pareto optimality (Shoven & Whalley, 1992: p.34) — controlling the economy imposes an overall economic cost. The implicit point is that our ultimate measure of social welfare is more complicated than simple Pareto optimality.

Simulation

Simulation is important in the early stages of developing the peer help system. We are currently developing a simulation based upon the design abstraction of a market in a single differentiated good. Assuming that the design abstraction is reasonable, and that students' objective functions are of the form shown in (1), we can verify the economic infrastructure's design. We can also test hypotheses about system behaviour such as critical complexity, and refine control policies such as taxation and policing. Verification of the design abstraction and model objective function must wait until after a working prototype and user test groups are available.

The current simulation scheme considers a population of agents within the abstract knowledge economy presented earlier. Agents exist within the designed economic infrastructure, and negotiate competitively according to the model objective function in (1). The goal preference parameter, σ, characterizes individual agents. The money balance, M, and grade, G, vary with trades according to the relationships shown in (2) and (3). Each agent's parameter may be different, so a distribution of goal preference parameters characterizes the agent society. The simulation requires a negotiation algorithm that determines prices and maps onto the model objective function; it also requires marginally declining functions that represent grade improvements. We are currently developing a suitable algorithm and realistic marginally declining functions.

Models of this type are further classified according to how the goal preference parameters are treated. An agent's goal preference parameter may be static or dynamic. Simple static models assume that the agent's parameter remains constant throughout the simulation; more complicated dynamic models allow the agent to change its parameter during the simulation. An agent's goal preference parameter may also be discrete or continuous. Simple discrete models only allow parameters from a finite set of possibilities; more complicated continuous models allow parameters that are real numbers in the range from 0 to 1.

The simulation currently being implemented is a static/discrete model. A more advanced dynamic/continuous model is also planned. The implementation uses the Swarm simulation system (Minar et al., 1996), a set of Objective C libraries developed at the Santa Fe Institute that provides a framework for constructing and observing multi-agent simulations.

Conclusions

The paper shows a design of a multi-agent economy for a distributed peer help environment. The specifics of trading with an intangible good like knowledge imposes specific requirements to the economy design. We propose using free market economy model designed for agents with a certain type of objective function. We also propose using a simulation to predict the behavior of the real system and to identify possible ways of influencing the economy to achieve a desirable global behavior, appropriate for the specifics of the application. This paper describes the requirements and the basic design of a multi-agent simulation of the economy using an available simulation tool, Swarm. Currently we are working on implementing the simulation and we hope by the time of the workshop to have some interesting empirical results to report about.

References:

Boyd G. (1997) Providing Real Learning with Virtual Currency, Proceedings of the International Conference on Distance Education, Penn State Univ., June 1997.

Brustein M. (1968) Economic Theory; Equilibrium and Change. John Wiley & Sons Ltd. London.

Cornes, R. (1992) Duality and Modern Economics. Cambridge University Press. Cambridge.

Friedman B. & Nissenbaum H. (1997) "Software Agents and User Autonomy". In Proceedings of Agents'97, also available online at: http://sigart.acm.org/proceedings/agents97/A186/A186.PDF

Greer, J., McCalla, G., Cook, J., Collins, J., Kumar, V., Bishop, A. and Vassileva, J. (1998) The Intelligent HelpDesk: Supporting Peer Help in a University Course, Proceedings ITS'98, San Antonio, Texas.

Guttman R. & Maes P. (1998) Cooperative vs Competitive Multi-Agent Negotiations in Retail Electronic Commerce, in *Proceedings of the Second International Workshop on Cooperative Information Agents (CIA'98)*. Paris, France, July 1998., also available online at: http://ecommerce.media.mit.edu/

Minar, M., Burkhart, R., Langton, C., Askenazy, M. (1996) The Swarm Simulation System: A Toolkit for Building Multi-agent Simulations. Santa Fe Institute.

Mudgal, Ch. and Vassileva J. (submitted) A Bayesian Approach to Learning about the Opponent in Negotiation. Submitted to this conference.

Shoven J. and Whalley J. (1992) Applying General Equilibrium. Cambridge University Press. Cambridge.