

Modular Simulation of Knowledge Development in Industry: A Multi-Level Framework

Rögnvaldur J. Saemundsson,¹ Kristinn R. Thórisson, Gudny R. Jónsdóttir,
María Arinbjarnar, Hilmar Finnsson, Hafthór Gudnason, Vignir Hafsteinsson,
Grétar Hannesson, Jónheidur Ísleifsdóttir, Ársaell Th. Jóhannsson,
Gunnar Kristjánsson, Sverrir Sigmundarson

Department of Computer Science
Reykjavik University
Ofanleiti 2, 103 Reykjavik, Iceland
{rjs, thorisson, gudny04}@ru.is

Abstract. Innovation is a central element of economic development. Understanding knowledge – its organization and especially its dynamics in a market – becomes therefore the main challenge when explaining economic development in general, and the competitiveness and growth of firms and industries in particular. Past research has generally treated knowledge as a monolithic object rather than a composite dynamic phenomenon. In this paper we present work on a new fine-grain, dynamic, morphogenic model of knowledge that is easy to manage, interpret and extend. This knowledge model is embedded a larger market simulation where selected elements of an economy, including employees, companies, banks and consumers, are modeled at multiple levels of abstraction, from agents to monolithic entities. We present data from early runs of the system, showing predictable results in baseline conditions and product innovation effects using the knowledge representation. The results show the model's excellent potential to address questions about emergent phenomena related to knowledge evolution, knowledge transfer and knowledge use in market innovation.

Introduction

Innovation is a central element of economic development. Paraphrasing Metcalfe (2001), innovation is the manifestation of the interaction between the growth of knowledge and the market process, driving the qualitative change of the economy in an evolutionary manner. Firms play a key role in this process as they are creative entities, the outcome of an ongoing process of organizational design that provide a framework for creating and combining different kinds of productive knowledge. Understanding knowledge, the organization of knowledge, and especially its dynamics becomes therefore the main challenge when explaining economic development in general, and the competitiveness and growth of firms and industries in particular. It is thus not sufficient to understand economic development as a macroeconomic, aggregate phenomenon, because the source of variety exists at the individual micro level and there are numerous feedbacks from higher to lower levels which shape and constrain the development. As other researchers have pointed out (c.f. Nissen and Levitt 2004), knowledge has too often been treated as an object, rather than a dynamic and continuously morphing multidimensional phenomenon.

We are developing a family of agent-based simulation models meant to study the generation, organization, development and evolution of the knowledge embedded in an industry. To provide a context for these evolutionary processes our models include agents at multiple levels of abstraction, i.e. individual, firms and other organization, and industry. In doing so we hope to extend insights obtained from other efforts that have not been focused on knowledge development in itself, e.g. the Java Enterprise Simulator described in Terna (2003) and other models of knowledge development that focus on the industry level of analysis, such as for example that developed by Cowan et al. (2003).

¹ School of Business, Reykjavik University, Ofanleiti 2, 103 Reykjavik, Iceland

Our ultimate objective is to build a platform that enables us to investigate the dynamics of knowledge development, evolution and change. One of the first topics we are looking at is that of *knowledge substitution*, i.e. when a new generic technology with superior performance potential is introduced in an industry, and its effect on company competitiveness under various market and policy conditions. We want to understand how knowledge development policies at the firm level and human resource constraints at the industry level influence the knowledge buildup in an industry. At the heart of such exploration must lie a high-granularity representation of knowledge that allows the knowledge to react naturally to the myriad of factors that may matter in its evolution, at the same time that it morphs its surroundings, as it does in the innovation process. This would allow the shape and form of the knowledge and the surrounding structures to emerge rather than being prescribed – as if developing from first principles.

We have created a morphogenic² model of knowledge that has several of these properties. In this regard we are able to model how knowledge development policies at the firm level influence knowledge development at the industry level, as well as the competitiveness of each firm. Similarly, our model can represent how resource constraints, such as the lack of talented individuals or the lack of funding, influence the knowledge buildup in an industry and eventual knowledge substitution. Our flexible design methodology (Thórisson et al. 2004) helps in making the system flexible so that many families of models can be relatively quickly constructed and compared.

In this paper we describe the basis of our approach, give a brief overview of our design methodology and detail the main design choices in the model. The model built so far is fairly large, composed of several types of modules each containing a number of decision-making policies. It is not possible to detail its every single aspect; instead we focus on the morphogenic knowledge model and explain how it operates in the simulation. We also present some data from our effort to ground the model and present early results from testing knowledge evolution in the system.

The paper is organized as follows: First we explain the model we have developed. Then we give an outline for the full simulation model, including its entities, their relationships and policies. We then present basic data on for various example runs that show the behavior of the model under baseline assumptions. Finally we give an example of model behavior when a company introduces a new product into the world.

Knowledge Model

Our representation of knowledge is a multi-dimensional vector space model derived from Saviotti and Kraft (2004). In our model each individual worker's knowledge, \mathbf{I}_k , is composed of a set of knowledge endowments or fields, \mathbf{K} , specified in a $2 \times N$ matrix, where N is the number of knowledge fields or areas of expertise that s/he possesses at any point in time.

$$\mathbf{I}_k = \{\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_n\}$$

$$\text{where } \mathbf{K} = \{l, c \mid l \in \mathbb{N} \text{ and } c \in \mathbb{N}\}.$$

For each \mathbf{K} , l represents the knowledge level while c represents the capacity of the individual to use the knowledge for production. The higher this l is, for a given \mathbf{K} , the more advanced knowledge does the individual hold for that particular knowledge area; the higher a c is the faster can the individual apply that knowledge (i.e. the more the worker can produce per time unit of things requiring knowledge represented by the associated \mathbf{K}). l is updated through deliberate training; c increases monotonically with experience as the individual uses the knowledge during production.

All \mathbf{K} s have a global maximum associated with their l , l_{\max} , representing a hard limit on how much a \mathbf{K} can grow inside each individual; no individual can achieve a higher l for a \mathbf{K} than l_{\max} for that \mathbf{K} , no matter how much training s/he receives. The value thus represents “current state of the art” in that area. The market is represented by a number of customer groups, $\mathbf{G} = \{g_1, g_2, \dots, g_n\}$, each group sharing the same preferences, \mathbf{P} .

$$g = \{s, w, f, q, r, \mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_n\}.$$

where s is a consumer group's “sweet spot”, a point in P -dimensional space ($s = \{p_1, p_2, \dots, p_n\}$, where p_1 represents the desired level for \mathbf{P}_1 , etc.), w is the size of the group, measured as Euclidian distance from s , and f is the shape of the distribution, that is, f determines how much quantity of a particular product/service is bought at any point along the radius from s to w . In our current setup w is given a

² Morphogen: An agent that controls the growth and shape of something, e.g. biological tissue.

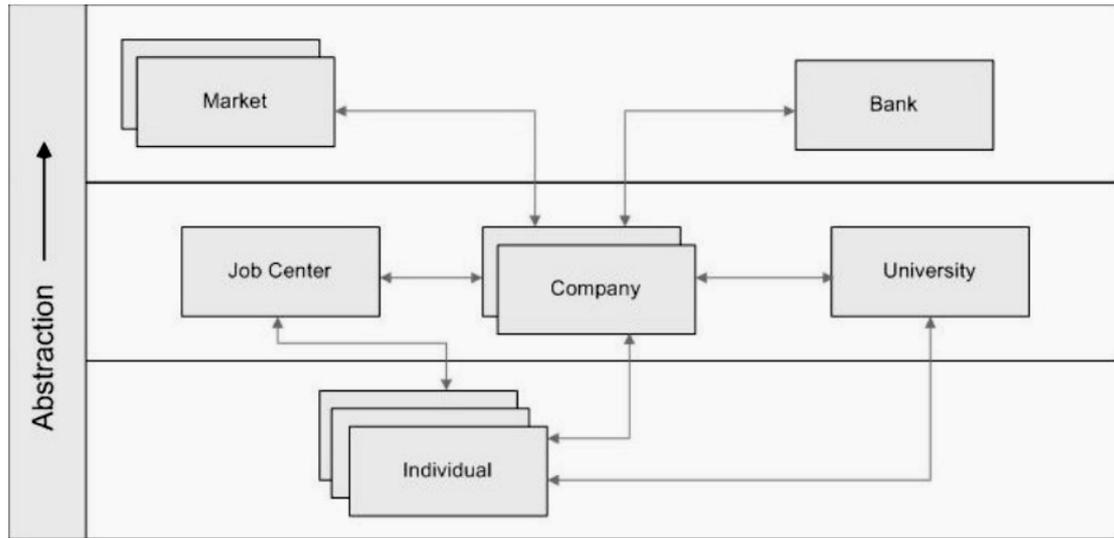


Figure 1. Agent/module types and their interaction relationships. The module types can be broadly classified into three layers, based on how much of an abstraction of the real world they represent: The Individual is the module type that is least abstract: A single Individual module corresponds to a single individual agent in the real world; the Market and Bank modules are the most abstract as they each model a large number agents in the real world. The Job Center, Company and University fall somewhere in between these two extremes.

constant distance from s (per g), and f is a linear function of distance. Maximum quantity is represented by q ; r is an upper bound on desired price range (from r down to gratis); q and r are used by f in the following way: For a product that lies perfectly in the sweet spot s and is offered at price $\leq r$, f returns q . If the price offered is greater than r , f reduces the quantity bought according to the difference between r and the price offered in a linear fashion.

Even though \mathbf{P} stands for a group's preference, it may be simplest to think of a \mathbf{P} as a feature of a product or service. Like any preference for a product in the real world, e.g. the preference for a gadget that can play MP3 music files, certain key features of a product determine the group's purchasing behavior for that product, for example, the size of the player's memory. In our model the product's features correspond to the group's \mathbf{P} s. The list of \mathbf{P} s then defines the group's feature interests for that product. The list of \mathbf{P} s can therefore be taken to stand for either a set of preferences or a set of features of a product or service – our model makes no distinction between these views.

Associated with each preference are a number of knowledge fields, \mathbf{K} s. In the real world a product or service's features require knowledge to be produced. In our model this knowledge is represented by a set of \mathbf{K} s associated with each \mathbf{P} :

$$\mathbf{P} = \{\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_n\}.$$

From the companies' perspective, any of these \mathbf{K} s can be used to produce the particular \mathbf{P} . As described above, for a given group g , the multidimensional point s defines the group's desired, "ideal" product. For a \mathbf{P}_x of g_x , p_x is a company's "target" when choosing a \mathbf{K} to produce \mathbf{P}_x : A company will look at their employees' \mathbf{K} s from the set \mathbf{P} and pick the \mathbf{K} whose l is closest to p_x . Any employee with a \mathbf{K}_x with $l \geq p_x$ will be able to produce \mathbf{P}_x according to market wishes. (If $l < p_x$ the product would be produced with an inferior \mathbf{P}_x to what the market wants, and a lower quantity – or nothing – would be bought of the product, as determined by f .)

In our current simulations, s is typically held very high (well above what can be achieved with l_{\max} for each \mathbf{K}). This is like in the real world where the market is never fully satisfied – the invisible mobile phone with the everlasting battery is still somewhere in the future waiting to be invented.

The \mathbf{K} s can thus fulfill \mathbf{P} to various extents; a \mathbf{K}_n which can fulfill a \mathbf{P}_n may only reach a level of 20 at time t_1 , because $l_{\max} = 20$ at t_1 for \mathbf{K}_n – yet \mathbf{P}_n 's sweet spot value $s = p_n$ for that \mathbf{K} may be different, e.g. 40. If a new \mathbf{K}_{new} comes along (through research conducted at the University) that can reach a higher level, and that is listed as being able to fulfill \mathbf{P}_n , this would represent an opportunity for companies to train their employees in \mathbf{K}_{new} , because with a worker with \mathbf{K}_{new} at 30 the company can create a better product than it could before and hopefully can better reach the sweet spot for \mathbf{P}_n .

The product/service innovation in this model is represented by listing the necessary **Ks** for a preference, even future **Ks** that may not be available to the worker population when the simulation is started. All **Ks** that are to enter the simulation during its runtime, however, need to be defined before the simulation starts. The firms only have knowledge about **Ks** through their employees. Information on new **Ks** can only be gained by sending employees to training at the University and those employees bring new **Ks** back; thus, l_{\max} for \mathbf{K}_{new} is not common knowledge. This means that companies with more aggressive training policies are quicker to spot opportunities arising with \mathbf{K}_{new} while less aggressive companies might not even know that \mathbf{K}_{new} exists, until they perhaps notice a product being sold on the market that exceeds their current production capabilities.

To compute fulfillment, the Euclidean distance between the product offered and the group's sweet spot s is calculated and the result compared to the group's w . If the value is greater than w the group buys 0 units of that product; if the result is less than w f determines how many units are bought. f operates by linearly trading off between the price offered (price is determined by the company offering the product) and the amount of fulfillment of the group's \mathbf{Ps} .

In our model we do not represent technology explicitly. A “consumer product” such as a mobile phone, or a particular service in a service market, is completely and exclusively defined by the set of \mathbf{Ps} listed for a particular group g . The \mathbf{Ps} , in turn, are completely and exclusively defined by the **Ks**, i.e. *the knowledge available to produce it*. By decoupling preferences for a service or product, such as a smaller mobile phone, from the knowledge that goes into producing a small mobile phone, and representing the combined technology/product/service/preferences as a set of **Ks**, we are able to significantly simplify the simulation model as a whole while maintaining a highly atomic and modular model, and maintaining control of all parameters that we are interested in exploring.

Simulation Model

World Model. Rather than build a simulation from first principles, or model the whole system as a single agent, our system employs a modular approach at many levels of abstraction, individual, firm and other organizations, and market. The model is composed of many types of modules each containing a number of decision-making policies. Here we give an outline of the main elements of each module type. In most cases the policies mentioned are of a fairly simple kind; in every case we tried to mirror their natural counterpart to a first approximation.

Figure 1 shows a block diagram of the main agents/modules in the system, and their interaction relationships. The individuals represent the module type that is least abstract, in that a single Individual module corresponds to a single individual agent in the real world; the Market module is the most abstract, in that a single Market module represents thousands of consuming agents in the real world. The Companies lie there in between, being partly represented by Individuals (the employees) and partly by monolithic rule sets that determine their policies.

Individuals. Individuals are atomic agents in the system, containing mechanisms for knowledge acquisition, retention and application, salary negotiation and employment search. They can be initialized to have different **Ks** in the beginning of a simulation. They have a learning rate that describes how quickly/slowly they learn, that is, acquire new **Ks** and increase the l of their already-acquired **Ks**. Individuals evaluate their salaries in connection with the rest of the world (average salaries, employment ads) and decide if they think they are more valuable than their salaries state, using a heuristic comparing their own salary with that of the market for each **K**. This mechanism is realized by the Job Center module.

The Individuals differ in their ability to learn; in our simulation runs we currently have the learning rate of individuals distributed normally over the whole pool of employees.

Companies / Firms. Individuals are employed by firms that offer services on the market based on the knowledge endowments of their employees. Firms compete with other firms on the service market, as well as in factor markets (employees and funding). Decision-making in the firm is based on predetermined policies local to each firm, including a training policy, an alertness policy and development policy. The alertness policy determines how often a company analyzes what products are profitable on the market; companies with high alertness will regularly analyze what products are profitable on the market and try to copy those products.

Firms are thus partly based on emergent features stemming from interactions between employees (individuals) and interactions with other organizations, as well as pre-determined rules related to hiring, layoffs and other decision-making such as salary negotiations with employees.

Each firm is able to innovate based on its knowledge endowments, which are a function of the endowments of the individuals they hire as well as the firms' knowledge development policies. There are two ways for a firm to add products to its suite of products: Development and copying. If a company decides to develop a product, it picks levels of \mathbf{Ps} for the new product based on their

employee's **Ks**, as well as on previously sold products. A new product requires that the company produce a specific amount before it can start selling it, simulating that it both takes time and resources to develop a new product.

Market. Each customer group is initialized with a particular set of values for $\{s, w, q, r\}$. The market includes various shifting policies and can be set to gradually change its size or shift for higher/lower values of l . In our current simulations these values stay unchanged throughout. At runtime, customer groups receive product advertisements from companies and make a decision of how much they will buy, according to how well the product meets their needs and the price of the product. A dialogue between the Market and the companies tells the company how much they should produce of the advertised product.

University. A single agent, the University, controls the general knowledge development within all knowledge fields, i.e. the maximum l that individuals may attain at each point in time for each of their **Ks**. It also controls the availability of new knowledge fields, which may overtake or substitute other knowledge fields in due time. The University thus determines the rate of innovation, which in our model equates to the frequency with which new **Ks** are added to the pool of **Ks** that Individuals can be trained in. In the University each **K** has an associated "difficulty" rate which determines how intrinsically complex that knowledge is. This interacts with the learning rate of Individuals, which varies. The University holds l_{\max} for all **Ks**.

Job Center. A single agent called Job Center manages all hiring of employees. Companies advertise for individuals with specific levels of **Ks**. Unemployed individuals can answer these advertisements and the Job Center finds the most suitable employee through a method that takes into account the salary demands of the individuals. The Job Center also has a role in calculating average salaries which individuals use when evaluating their own salaries.

Bank. The Bank is a simple loaning institution that the companies can apply for a loan from. The Bank is especially important when starting a simulation, especially if companies are bootstrapped as startups instead of having funding from the outset. In our current simulation runs we initialize all companies with some amount of starting cash – they will only apply for loans if they run out of cash. This part of the model is among the simplest mechanisms and we intend to increase its functionality so as to better be able to study the role of startup capital in knowledge formation.

Methodology. The model was built using an adaptation of the Constructionist Design Methodology (CDM) described by Thórisson et al. (2004). The methodology is based on the concept of multiple interacting modules; modularity in the approach is made explicit by using one or more blackboards (cf. Adler 1992) for message-passing. As has been pointed out by Thórisson et al. (2005), among others, the benefits of message-based, publish-subscribe blackboard architectures are numerous. Among the most obvious ones is that for modules whose input and output is fully defined by messages in the system, the messages embody an explicit representation of the modules' contextual behavior. In some cases the messages become a key driving force around which the modules are defined and organized. Messages in turn mirror directly the abstraction level of the full system, making the blackboard architecture a powerful tool for incremental building and de-bugging of complex, interactive systems.

The software we have chosen for constructing our models is Psyclone (CMLabs 2004). Like related middleware such as Swarm,³ Psyclone directly supports re-configuration and re-design of both non-existing and existing architectures (the feature is called "mutability" in Psyclone; "composability" in Swarm). This ability makes it easier to build systems where the exact final model is not known before construction, as is very often the case with complex simulations. Psyclone's architectural mutability is the result of flexible module APIs and ease of moving processes between computers. In Psyclone an implemented architecture can be radically changed relatively quickly; changes involving re-routing messages, temporal dependencies, re-organizing distribution across machines, and splitting up or merging modules or functionality, can be achieved with relative ease. Semantic interfaces used for specifying data flow provide great flexibility in changing layout after the initial system is built.

We use XML to specify all aspects of the architecture that undergo frequent changes or require significant tuning. At present this includes subscriptions to message types and parameter settings. The use of Psyclone's whiteboards (a type of blackboard, cf. Adler 1992) support flexible and easy-to-use real-time monitoring and statistics generation.

Although our largest model currently has less than 50 modules (40 Individuals, 5 Firms, one Market, one University, a Job Center and a Bank), our tests indicate that the number of modules should be extensible to several hundred. Parameters in our models are centralized and we are therefore able to modify them quickly, allowing for many variations and comparative runs of the models.

3 <http://www.swarm.org/wiki/>

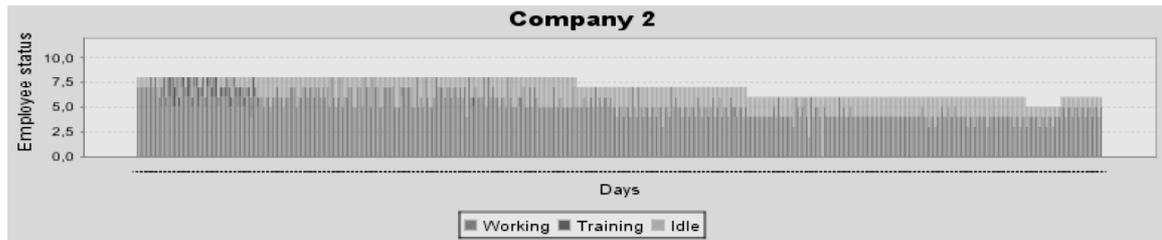


Figure 2. Example of a company's employee status over a few simulated months. As the employees' capacity increases over time, the company's staff decreases in a natural proportion to that. Light gray = idle; mid gray = working; dark gray = training.

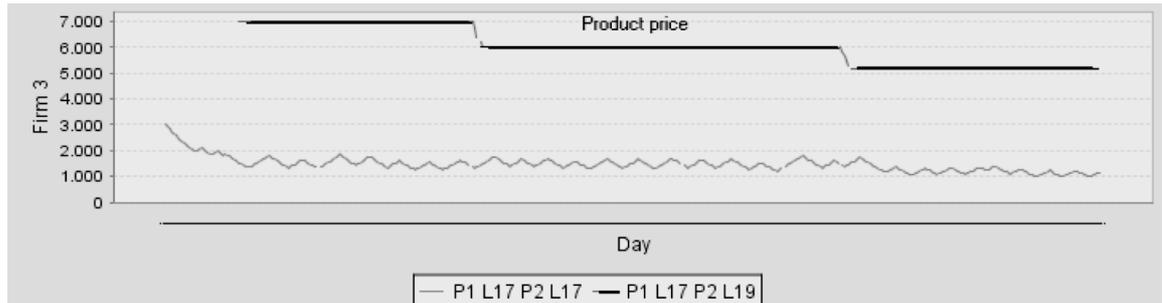


Figure 3. Companies tend to lower prices very early when there is competition in the market and are reluctant to lower prices of products that few or no other companies are selling. This graph shows two products, the top line is a product for which there is little competition (Firm 3 lower its price periodically to explore whether it can sell more of the product – in this case that strategy is working). The bottom line shows the price over time of a product from Firm 3 that is similar to other products in the market.



Figure 4. When companies are competing, the price of their products, and thus the companies' income, will reach equilibrium. The graph shows how closely the income of two competing companies, Company 1 and Company 2, will follow each other under competing conditions (the lines are inseparable).

Results

We have done dozens of runs of the described model, lasting several simulation-months each. Although most of the questions we intend to answer remain to be investigated, our results so far are highly encouraging. Here we present some results from our attempts in the initial validation of the model, as well as preliminary results from exploration of knowledge development.

Baseline. For baseline testing of the model we have tried several different scenarios. In these runs all shifting policies in the market are held stable, that is, the market contains a static number of Ps and the parameters of all Ps are static. The model shows that competing companies that start with the same product and homogenous employees always reach an equilibrium where none of the companies are making much profit. Companies tend to lower price of competitive products more quickly than of products where there is less competition. Further baseline testing is undoubtedly necessary; we intend to focus on verification of behavior at boundary conditions, testing a broader spectrum of parameter changes and a wider range of initial starting states.

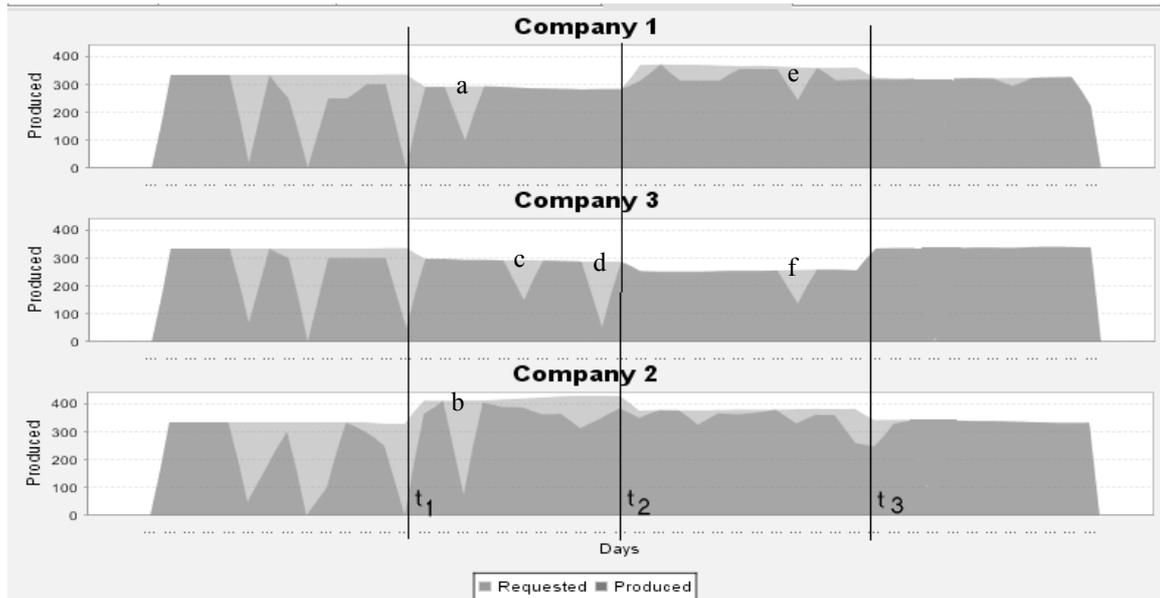


Figure 5. Example run over 49 simulated days, showing supply versus demand for three companies. Light gray: demand, dark gray: supply. t_1 marks a point where the development of a new product was initiated by Company 2; t_2 , t_3 mark when C1 and C3 copy that product, respectively. See text for further explanation.

Knowledge Development Example. We have done several dozen runs of the model where we test its behavior with regards to the effects of the addition of new knowledge. Here we give a simple example of such a run. In this example we have set up an industry with 3 firms (C1, C2, C3) and one target market group whose preferences are along \mathbf{P}_1 and \mathbf{P}_2 . Three knowledge fields (\mathbf{K}_1 , \mathbf{K}_2 , and \mathbf{K}_3) can be used to provide services fulfilling the preferences.

$$\mathbf{P}_1 = \{\mathbf{K}_1, \mathbf{K}_3\}, \mathbf{P}_2 = \{\mathbf{K}_2\}$$

$$\text{and } \{\mathbf{K}_1[l_{\max}=20], \mathbf{K}_3[l_{\max}=30], \mathbf{K}_2[l_{\max}=20]\}.$$

One customer group is set up with s at $\{\mathbf{P}_1[l=100], \mathbf{P}_2[l=100]\}$. This means that higher levels of \mathbf{P} are always better, simulating a situation where current state-of-the-art in the relevant fields is nowhere near being sufficiently advanced to produce a product that meets market demand.

All firms start with the same product $\{\mathbf{P}_1[l=17], \mathbf{P}_2[l=17]\}$ and the industry starts at an equilibrium where all firms have an equal share of the market. At t_1 company 2 introduces a new product $\{\mathbf{P}_1[l=19], \mathbf{P}_2[l=19]\}$ and subsequently gains a larger part of the market (Figure 5). The companies adapt to market needs at different speeds: C1 at t_2 and C3 at t_3 . After t_3 all companies are selling both products (initial and new) and have again gained close to even market share. The synchronized drops and rises in the production of all three companies, in the early part of the run (before t_1), show the companies' response to an artificial elimination of market demand for their products, deliberately introduced for baseline testing purposes. The drops in production after t_1 (a – f, Figure 5) is due to a time-sensitivity of the companies: If they receive a purchase order from the Market after a certain delay they will discard it. In this run the delay was likely caused by CPU overload; in later runs with longer simulated days (each day lengthened by a few seconds) this pattern has been observed to go away. We could also run the model on a faster computer. In the current model the Market is represented as a single agent. A third – and perhaps the most desirable – way of changing this behavior would be to increase the granularity of the Market such that instead of a single Market order a large number of orders would arrive in each company from different purchasers.

Conclusions

We have presented a new model for representing knowledge and innovation in markets. We find that its high degree of extensibility and the resolution afforded gives it promising capabilities in answering questions about emergent phenomena related to knowledge evolution, knowledge transfer and knowledge use in product innovation in a competitive market.

While some work remains on validating our knowledge model, we believe that its relative simplicity will enhance our ability to interpret results from worlds with complex underlying dynamics, including

a large number of independent variables. By decoupling preferences for a service or product, such as a smaller mobile phone, from the knowledge that goes into producing a small mobile phone, and representing the combined technology/product/ service/preferences as a set of **Ks**, we are able to significantly simplify the simulation model as a whole while maintaining a highly atomic and modular model, maintaining control of all parameters that we are interested in exploring.

The mutability of our framework makes it possible to build several models with different assumptions and test the outcomes against each other. Our plans to take a closer look at knowledge substitution, radical innovation and non-contiguous jumps in product and service evolution will provide us with further data on this aspect. In our runs so far we have shown the framework to produce predictable baselines, given sensible starting points. The model has a clear potential to address several classes of questions regarding the generation, organization and evolution of knowledge in economic settings.

Inherent flexibility and scalability in the framework, stemming from the use of the Constructionist Design Methodology and Psyclone, enables us to extend the model, thus adding to the complexity of a simulation model, without disturbing the overall architecture. Module types can easily be expanded without affecting the design and implementation of other modules. The current distribution of modules in the model along an axis of abstractness can be fairly easily changed through the mutability functionality in Psyclone. Coupled with the model's scalability we envision being able to substitute the Market module with several more fine-grain modules that approximates a large group of consumers more closely.

Future work includes extending several of the mechanisms related to these issues and exploring the results from running the model with various parameter settings, policies and initial states. We plan on splitting the Company module into several smaller units that become responsible for the companies' execution and decision making in a more agent-based manner than at present. We also plan on adding multiple levels in the value chain, something that is important to better understand how knowledge evolution influences, and is influenced by, vertical integration and disintegration in the market.

Acknowledgments

This research was supported by a Marie Curie European Reintegration Grants within the 6th European Community Framework Programme. The work was in part also made possible by research grants from The Icelandic Centre for Research (Rannis). We'd like to thank Anna Ingólfssdóttir for proofreading. Psyclone is a trademark of Communicative Machines Ltd., Scotland.

References

- Adler, R. (1992). Blackboard Systems. In S. C. Shapiro (ed.), *The Encyclopedia of Artificial Intelligence*, 2nd ed., 110-116. New York, NY: Wiley Interscience.
- CMLabs (2006). *The Psyclone Manual*. Communicative Machines Inc.
- Cowan, R., Jonard, N., & Özman, M. (2004). Knowledge Dynamics in a Network Industry. *Technological Forecasting and Social Change*, *71*(5):469-484.
- Metcalfe, J. S. (2001). *Technology and Economic Development: A Comparative Perspective* (CRIC Working Paper No. 10). Manchester: CRIC, University of Manchester.
- Nissen, M.E., & Levitt, R.E. (2004). Agent-based modeling of knowledge flows: illustration from the domain of information systems design. *Proceedings of the 37th Hawaii International Conference on System Sciences*.
- Saviotti, P. P., & Krafft, J. (2004). *Towards a generalised definition of competition*. Presented at the *DRUID Summer Conference 2004 on Industrial Dynamics, Innovation and Development*, Elsinore, Denmark, June 14-16.
- Terna, P. (2003). *The quest for the enterprise: jES, a Java Enterprise Simulator*. Manuscript. Torino University.
- Thórisson, K. R., T. List, C. Pennock, J. DiPirro (2005). Whiteboards: Scheduling Blackboards for Semantic Routing of Messages & Streams. In K. R. Thórisson, H. Vilhjalmsón, S. Marsella (eds.), *AAAI-05 Workshop on Modular Construction of Human-Like Intelligence*, Pittsburgh, Pennsylvania, July 10, 8-15. Menlo Park, CA: American Association for Artificial Intelligence.
- Thórisson, K. R., Benko, H., Arnold, A., Abramov, D., Maskey, S., Vaseekaran, A. (2004). Constructionist Design Methodology for Interactive Intelligences. *A.I. Magazine*, *25*(4):77-90. Menlo Park, CA: American Association for Artificial Intelligence.