Knowledge Bases and Information Public Goods: A Multi-Agent Model

Maksim Tsvetovat, Craig Schreiber, Kathleen Carley
maksim@cs.cmu.edu, craigs@andrew.cmu.edu, kcarley@ece.cmu.edu

Carnegie Mellon University**

Abstract. Recent advancements in information technology have provided organizations with new opportunities for communication and knowledge sharing. Among them are discretionary knowledge bases that store experiential knowledge contributed by employees of the organization. This paper looks in-depth upon the social mechanisms that affect the functioning of such knowledge bases and attempts to explain them through emergent behaviours of self-interested autonomous agents.

1 Introduction

Recent advancements in information technology have provided organizations with new opportunities for communication and knowledge sharing. Much of the current work in organizational theory is focusing on effects of these new technologies have on organizations as well as the underlying mechanisms at play in their use.

One specific example of a new information technology is a discretionary database. A discretionary database is a collection of shared data to which agents can voluntarily contribute and from which agents can voluntarily retrieve.

It has been shown that discretionary databases and, more largely, information technologies act as a public good [10][8][7][2]. In addition, a review of knowledge management practices within large corporations has shown two types of data that are mainly used in discretionary databases. These two types of data are task-related data and referential data. Referential data is useful in establishing the knowledge about expertise carried by different members of the organization.

In this paper, we present the theoretical views upon the creation and growth of public goods, and propose that the macro-level behaviours described are *rooted in* and *emerge from* individual-level decision schemas. We proceed to propose a multi-agent simulation methodology that can realistically model the low-level social processes involved in creation of public goods. As a result of the simulation experiments, we present the findings of emergent organizational level behaviour.

** This work is supported in part by the IGERT Grant, National Science Foundation and DARPA

2 Public Goods Theory - an Overview

The concept of a public good has two defining characteristics. The first characteristic is *jointness of supply* [1][9] and the second is *impossibility of exclusion* [4].

Jointness of supply is simply this - if one agent consumes the good then that consumption does not reduce the amount of consumption available to any other agents. Now it is true that most public goods rarely have a pure jointness of supply. Most public goods have the possibility of experiencing "crowding" [1]. Crowding occurs when there is an over-consumption of a public good and the subsequent consumption available to others is adversely affected. In the case of knowledge management and the use of discretionary databases, it does not matter how many times information in the database is accessed and used; the availability of information is not dependent upon consumption.

Impossibility of exclusion is a concept by which all agents of the considered public have the opportunity to consume or benefit from the good. As long as all of the agents of the considered public have access to the good then this requirement is met. In looking at discretionary databases as a public good, the unit of analysis needs to be examined for determining if this requirement is fulfilled.

2.1 Characteristics of Public Goods

There are two characteristics of the gbood - *production function*(divisibility) and *heterogeneity.*

The production function is denoted by $P(R)$ whereas the level of provision of the public good, P, is a function of the total resources contributed, R. There are four types of production functions that can describe the interaction between provision and resources for a public good. Which production function best describes a particular public good depends on how divisible the good is. The following is a list of the four production functions matched up with their public good divisibility:

- Step production function, non-divisible goods
- Linear production function, continuously divisible goods
- Accelerating production function, additional resources produce proportionately more value
- Decelerating production function, additional resources produce proportionately less value

The second characteristic of a good is *heterogeneity.* Heterogeneity is the degree to which the public good offers different provisions that may vary in interest for different agents of the public. A higher degree of differentiation translates into a higher degree of heterogeneity. Traditionally heterogeneity is viewed as increasing the likelihood of collectively producing the good.

2.2 Discretionary Databases

Discretionary databases are defined[10] as "a shared pool of data to which several participants (individuals, departments) may, if they choose, separately contribute information." This definition is consistent with that of communal public goods and a discretionary database is a communal public good. Since discretionary databases are communal it is assumed that they have an accelerating production function. Examples of a discretionary database are shared knowledge databases, distribution lists and bulletin boards.

Connelly and Thorn predict discretionary databases to be undersupplied with information because of information sharing. Information sharing is defined as taking proprietary information and sharing it with the collective thus making it communal, thus incurring some cost For instance, if the shared information gave the agent a competitive advantage then the advantage would disappear when this information is shared. It does not make rational sense for an agent to contribute when they get no immediate benefit or when they can free-ride off of the system. This is especially true in the early stages of development when the overall benefits of the system are low for all agents.

An interesting phenomenon occurs with respect to discretionary databases because of this problem. This phenomenon is contributions under the *norm of reciprocity.* The norm of reciprocity is making a contribution and incurring costs in the belief that other agents will gain benefit and return the favor with a contribution of their own. This norm is assuming future benefit over immediate benefit and is a factor that influences agent contributions.

Large collectives have more opportunity for contributions based upon their relative size. But in discretionary databases the driving force is the norm of reciprocity. For large collectives the norm of reciprocity is much less visible. Because of the larger size it is possible that an agent may not be able to tell if their contribution has been reciprocated therefore deterring any future contributions. Alternatively, an agent receiving benefit may not be compelled to make a contribution of their own knowing that others will not be able to notice the lack of reciprocation. This poses a very strong opportunity for free-riding. In contrast, a small collective has a much higher visibility for reciprocity. It is easier to notice a reciprocal contribution, which will motivate more contributions, and it is easier to notice free-riding, which will deter such loafing.

2.3 Characteristics of the Agents

Interest and *resources* are the two main characteristics of the agents. *Interest* is defined as the value that an agent perceives he or she will receive as provided by the good. Perception is used in this definition because the value is subjective for each individual agent. This perception of value is a function of the production function, the overall level of provision provided by the total resources contributed.

Resources are anything that can contribute to the production of a public good. Examples are equipment, money, time and expertise. Associated with resources are the costs of contributing. Contribution costs are a function of the resources provided by the agent and represented. The relationship between the cost of contributions and the resources provided is assumed to be linear but that need not be the case as incentives can reduce the cost of contributing and may change the linear relationship.

Incentives act as inducements to the agent to contribute resources to the good by reducing costs. Incentives are not characteristics but they do have an effect on the agent's decision so the subject is discussed under this section. Incentives can be of three types - material, solidary and purposive. Material incentives are in the form of money (bonuses or fines), gifts, certifications, job titles, penalties, demotions and the like. Solidary incentives come from the social ties that an agent has. Purposive incentives are based on feelings. An agent may contribute based on feelings of morality or self-worth. All three types of incentives, material, solidary and purposive, are based upon the subjective viewpoints of the individual agents and will not affect each agent equally, making the relationship between cost and resources contributed nonlinear.

Given these low-level assumptions, we set out to implement a multi-agent model that will show the emergence of macro-level organizational behaviour, such as public good production functions.

3 Multi-Agent Network Model Methodology

The Multi-Agent Network Model paradigm is based upon the following assumptions:

- The simulation consists of agents
- Agents are independent, autonomous entities endowed with some intelligence
- Agents are cognitively limited
- Agents can learn knowledge about the world and referential knowledge about other agents, with a limited learning capacity
- Agents can forget
- Agents communicate asynchronously and deal with asynchronicity (i.e. deadlocks, delays, etc) in an autonomous manner.
- Agents do not have accurate information about the world
- Agents do not have accurate information about other agents
- Unless required by the simulation domain, there is no central mediating entity to resolve the conflicts.
- Unless required by the simulation domain, the agents do not use predefined geometrical locations or neighborhoods.

The simulation paradigm is task-independent. The task is merely defined as a function that maps a problem vector and agent's knowledge vector unto a result vector. Thus, the simulation can be easily adapted to different simulation domains, from military simulation to electronic marketplace simulation.

The agents within the system are implemented as non-deterministic finite automata, with states of the automaton representing low-level behaviours and transitions governing the way the agent switches between them. Some transitions are deterministic, others rely on probabilistic equations.

3.1 Properties of Multi-Agent Networks

Structural Realism The artificial-life based simulations are built upon the concept of agents (or cellular automata) located on a grid of a specified shape (square, infinite, toroidal, 3-dimensional, etc). Interactions are based upon the concept of proximity, defined by the agent neighborhood on the grid. Unfortunately, the choice of grid shape and type of neighborhood is arbitrary and does not carry any face validity. Yet it can significantly alter the behaviour of the system and thus, simulation results.

While physical proximity in the real world carries some importance, the majority of human interactions are based upon social networks - arbitrary graph structures not constrained by the concepts of grids or grid neighborhoods. To increase the face validity of social simulations, the simulation architectures must use similar structures to describe the social structure of the simulation domain.

Moreover, while the social network of an organization can be studied objectively by an outsider, none of the participants of the network actually have an accurate view of the interaction structure of the organization. They do, however, have beliefs about that structure, and use them to guide them through the interactions. This point is important because these beliefs are often inaccurate, and change rapidly as information is processed.

In a multi-agent network simulation paradigm, the agents' interactions are governed by the formal structure of the organization, and agents' beliefs about the informal structure.

The formal structure of the organization is specified as a directed weighted graph that specifies the communication channels that are open as well as their throughput or cost of communication. The directed nature of the graph allows one to specify one-way relationships and chain-of-command relationships.

The beliefs about the informal structure are individual to every agent, and also consist of a weighted directed graph. However, when an agent joins a network, its informal relationship graph is empty, and it must learn about the informal network before it can be used for communication.

Information Flow Realism In a multi-agent network, the agents do not have perfect knowledge about the world. The only way to obtain information about either the world or other agents is to ask about it and then learn the results of the query - or obtain the information as a result of information exchange interaction. Agents may or may not communicate their beliefs truthfully and can be strategic about their communication. This is required for simulation of domains where the agents are competitive (i.e. electronic markets) or hostile (i.e. military simulations)

4 Mutli-Agent Model of Public Goods

The model of information diffusion that we used to simulate discretionary databases is based upon the CONSTRUCT information diffusion model [6][3].

The agents in the model perform a classification task that is information-intensive (i.e. requires a large amount of knowledge to complete without guessing). In the beginning of the simulation, agents are endowed with relatively little knowledge and must engage in learning behaviours in order to increase their task performance. Agents learn by interaction: trading facts with other agents or asking direct questions in hope of getting an accurate answer. Agents also forget little-used facts.

4.1 Operationalization of Information Diffusion

In a brute-force scenario where every agent tries to communicate with every other agent, the amount of interaction necessary to find relevant information makes the task intractable. Thus, a question of "Who do I interact with" becomes very relevant.

The answer to this question comes from psychological literature. It has been shown that there are two distinct behaviours that humans engage in while looking for information: interaction with a peer group (i.e. people who are somewhat similar to the person in question) or interaction with an expert who is more likely to know what one needs to learn.

We define peer groups to be based on a measure of relative similarity between agent i and agent j: the amount of knowledge that i and j have in common divided by the amount i shares with all other agents, or

$$RS_{i,j} = \frac{\sum_{k=0}^{K}(S_{ik} * S_{jk})}{\sum_{j=0}^{I}\sum_{k=0}^{K}(S_{ik} * S_{jk}) : wq}$$

where S_{ik} is 1 if agent i knows fact k and 0 otherwise.

Relative expertise is defined as RE_{ij} = how much agent i thinks j knows that i does not know divided by how much i thinks all others know that i does not know, or

$$RE_{ij} = \frac{\sum_{k=0}^{K}((1 - S_{ik}) * S_{jk})}{\sum_{j=0}^{I}\sum_{k=0}^{K}((1 - S_{ik}) * S_{jk})}$$

In both cases, agents operate on their beliefs about what the other agents know. Thus, their predictions of relative expertise or similarity can be inaccurate. However, as interaction progresses and agents learn more and more about each other, they learn an increasingly complete picture of their world.

4.2 Operationalization of Knowledge Bases

The literature distinguishes two types of data that were of consistent use for discretionary databases. The two types that are widely used are task-related data and referential data.

Task-related data are information related to the performance of a task or a solution to a problem. The Eureka database for Xerox service representatives is a good example of this type of data [5].

Variable	Description	Values
Organization Size	Number of Agents	20,50,100
Knowledge Amount	Number of facts relevant to tasks	2,4,6 times the number of agents
Database	Type of database used	no database, task database, referential database

Table 1. Experimental Design

Referential data 'refers' an agent to an expert in the topic of interest. In application, an agent would search for a topic on which they need some knowledge. The search would give the agent a list of names of experts in that area of knowledge. The agent can then contact the expert(s) and find the knowledge that is needed. Referential data differs from task-related data not only in that it gives a name reference rather than raw information but in the fact that the expert should know the raw information and how to apply it.

5 Hypotheses

The operational hypotheses of this work are aimed at establishing the relationship between the size of organization, its performance in the task and the size of its knowledge base (which is measured in absolute terms as well as the amount of contributions to the public good by members of the organization.

Hypothesis 1: Small collectives will have a higher contribution rate to the database than will large collectives [10].

Hypothesis 2: Information will diffuse faster and more completely when an organization uses a database.

Hypothesis 3: Organizations that use databases will outperform those that do not use databases.

Hypothesis 4: Referential databases will have more of a positive impact than will task-related databases.

6 Experimental Design

The simulation is run on a 3x3x3 design, shown in table 1 Each collective was simulated 100 times for 150 time periods. There is one communal database in each collective. The database is set at an initial knowledge base of 5

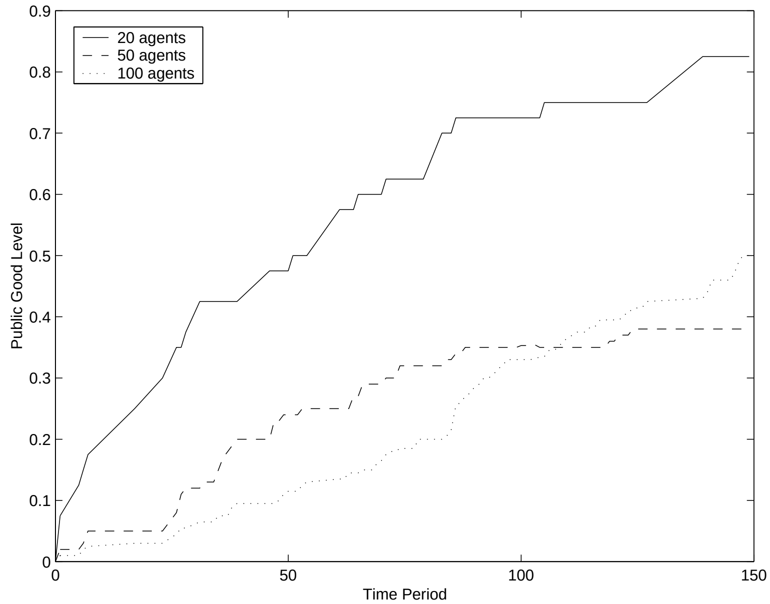

Fig. 1. Growth of Knowledge Bases and Organization Size

7 Results

As figure 1 shows, in a smaller organization, the public knowledge base grows quicker and achieves a greater degree of completeness then in larger organizations, which confirms our *Hypothesis 1*.

Use of databases dramatically increases knowledge diffusion in an organization (see figure 2). Availability of knowledge in an easily accessible place allows agents to find the information they need for task performance without incurring the time penalties for peer-to-peer communication.

The simulation also shows that task databases speed up information diffusion more then referential databases. This is due to the fact that accessing the database takes time and effort. When an agent accesses a task database with a query, it receives an immediately applicable answer within one interaction. In the same time, if the agent chooses to interact with a referential database, the agent has to perform two interactions (with the database and with the peer it has been referred to) in order to learn a fact. The time spent in these interactions is still less then time spent randomly querying peer agents in search of information, but it is more efficient to access the data directly.

However, there is no clear improvement in organizational task performance from introduction of databases (see figure 3. Actually, for small organizations, existence of a databases decreases task performance slightly. Thus, *Hypotheses 3 and 4* are not supported by the results of the simulation.

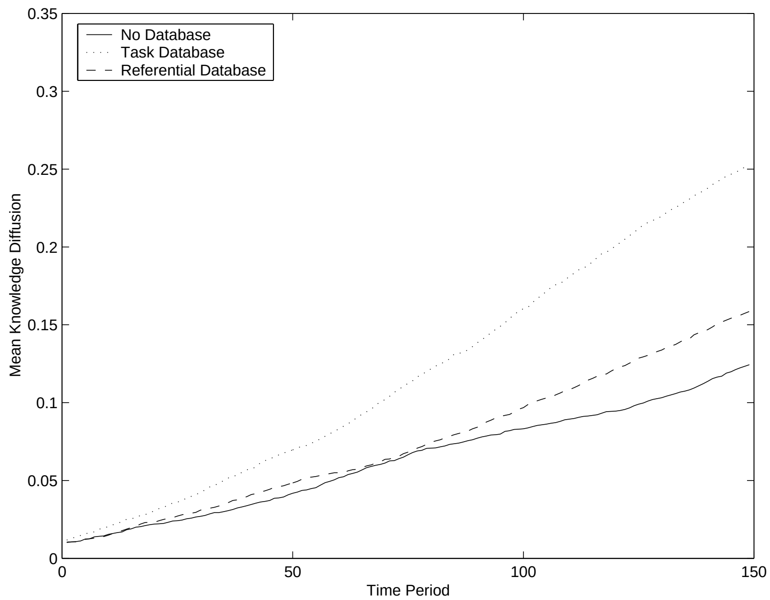

Fig. 2. Effect of Databases on Task Performance

This can be explained by the fact that in a densely connected small organization a knowledge base poses a distraction and takes time away from task performance and local interaction. However, the differences in task performance are fairly slight and don't show a consistent trend.

8 Conclusions and Future Work

In the past, simulations have been used in proving and strengthening theoretical findings. In this paper, we have presented an initial approach to simulating emergent organizational behaviours starting from low-level cognitive models specifying the way agents interact. Study of these emergent behaviours can become a powerful theory-building tool.

References

1. Barry B. and Hardin R., *Rational man and irrational society*, Sage Beverly Hills CA, 1982.
2. Monge P. Fulk J. Kalman M. Flanagin A. Parnassa C. and Rumsey S., *Production of collective action in alliance-based interorganizational communication and information systems.*, Organization Science (1998), no. 9, 411–433.
3. Kathleen Carley, *A theory of group stability*, American Sociological Review **56** (1991), no. 3, 331–354.

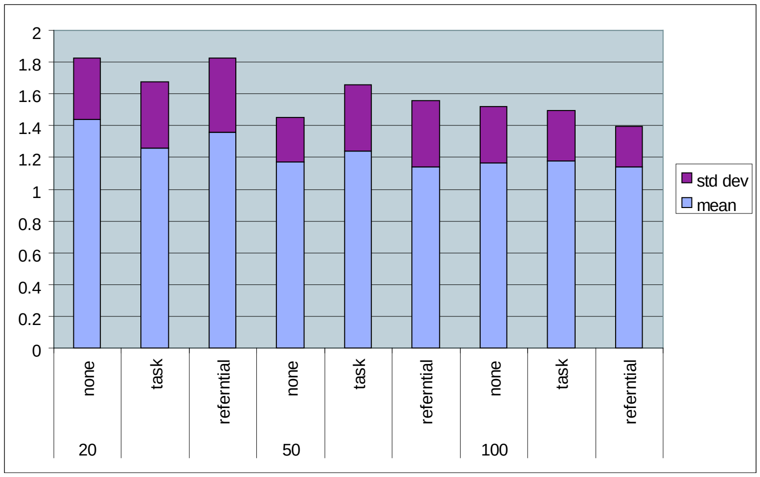

Fig. 3. Effect of Databases on Task Performance

4. Head J.G., *Public goods: The polar case.*, R. M. Bird and J. G. Head (Eds.) Modern fiscal issues: Essays in honour of Carl S. Shoup, pp. 7–16, University of Toronto Press, Toronto Ontario, 1972.
5. Brown J.S. and Duguid P., *Balancing act: How to capture knowledge without killing it.*, Harvard Business Review **78** (2000), no. 3, 73–80.
6. Carley Kathleen, *Group stability: A socio-cognitive approach.*, lawler e. markovsky b. ridgeway c. and walker h. (eds.) ed., Advances in Group Processes: Theory and Research., vol. VII, pp. 1–44, CN:JAI Press, 1990.
7. Markus M. L., *Toward a 'critical mass' theory of interactive media.*, Fulk J. and Steinfeld C. (Eds.) Organizations and Communication Technology, pp. 194–218, Sage, Toronto Ontario, 1990.
8. Fulk J. Flanagin A. Kalman M. Monge P. and Ryan T., *Connective and communal public goods in interactive communication systems.*, Communication Theory (1996), no. 6, 60–87.
9. Harding R., *Collective action*, John Hopkins University Press, 1982.
10. Connolly T. and Thorn B.K., *Discretionary databases: Theory data and implications.*, Fulk J. and Steinfeld C. (Eds.) Organizations and Communication Technology (1990), 219–234.