Exploring the Impact of a Stochastic Hiring Function in Dynamic Organizations

Geoffrey P. Morgan
Kathleen M. Carley
Carnegie Mellon University
School of Computer Science
Institute for Software Research
5000 Forbes Avenue
Pittsburgh, PA 15213
gmorgan@cs.cmu.edu, kathleen.carley@cs.cmu.edu

Keywords:
Individual Differences, Organizational Models, Behavioral Moderators

ABSTRACT: *We present an implementation and extension of James G. March's Mutual Learning Model (1991), extending the original work through implementing a stochastic selection process based on the inherent biases of individuals. We present factors that we believe critically impact selection, literature related to these factors, and our resulting equations. We discuss the simulation, and present two experiments. The first experiment was a docking study, comparing our implementation of March's work with his published results, and we found similar patterns – establishing relational equivalence. The second experiment compared organizations with and without a stochastic selection process. Organizations that stressed socialization tended to need to review more (otherwise equally qualified) applicants than organizations that did not. Finally, we discuss implications, limitations and future directions for this work.*

1. Introduction

Organizations exist within and must respond to dynamic environments. This paper examines the impact selective hiring has on an organization's ability to compete fruitfully in a competitive and turbulent market-place. March (1991) shows that turnover strongly influences an organization's capacity for adapting to a changing marketplace, as fresh blood brings new insights. Further, he states (pp 80-81):

> *"The positive effects of moderate turnover depend, of course, on the rules for selecting new recruits. In the present case, recruitment is not affected by the (organizational) code. Replacing departing individuals with recruits closer to the current organizational code would significantly reduce the efficiency of turnover as a source of exploration."*

We are interested in what impact a stochastic hiring function might have on the efficiency of turnovers within a given organization. This stochastic hiring function will be dependent not on March's construct, the organizational code, but rather on the preferences of specific individual agents that form the simulated hiring committee. The stochastic hiring function extends Morgan, Morgan, and Ritter's (2010) work modeling participation in goal oriented groups. Morgan et al. (2010) discuss seven factors that affect the probability of taking beneficial and hostile actions towards a third party. The seven factors identified in that paper are:

1. Group Size
2. Group Composition
3. Social Distance
4. Spatial Distance
5. Mutual Support and Surveillance
6. Presence or absence of legitimate authority figures
7. Task Attractiveness

The goal of Morgan and his colleague's work was to present some first steps towards a general social reflexivity mechanism – one that would work even for light-weight agent simulations. Although they demonstrated their work by modeling it in the small-group combat domain, it was a goal of that work that it could be applied broadly, and one of the goals of this paper is to demonstrate the approach in an entirely different setting.

From the seven factors presented above, we focus on the impact of Group Size, Group Composition, Social Distance, and Mutual Support and Surveillance. More precise and technical definitions will follow, but a brief summary of how these terms are used in this paper is in order. Group Size is defined as the size of the hiring committee. Group Composition is an abstraction of the diversity of the committee, relative to the total diversity available. Social Distance is the perceived difference between the job candidate and the committee member. Mutual Support and Surveillance is an abstraction of the level of urgency in the group to hire similar individuals to themselves, and this pressure is based on the current diversity of the hiring committee.

This work does not deal with the other factors identified by Morgan et al. (2010) for the following reasons. Spatial distance is at best an abstraction in this context. Furthermore, the presence of legitimate leaders and a constant value for task attractiveness seems implicit to a hiring process, allowing us to hold them constant across all test conditions.

Although a great deal of work has explored why individuals leave firms, for this paper we presume a collection of exogenous factors. We represent these factors by using a random value to determine which agents leave at each time-point. The selection of members for the hiring committee is also random, while the number of members in the hiring committee is fixed – although we see this as an interesting point for exploration.

2. Related Work

In this section, we first describe in more detail March's simulation. We follow this discussion with a brief literature review of the factors used in the hiring function.

2.1 March's Mutual Learning Model

March (1991) predicts how organizational knowledge develops from the aggregate of individual knowledge sets. His simulation model remains influential. Consequently, it is important to note early the similarities and differences between the two models.

As in March's model, we posit an external reality, which changes over time due to un-modeled exogenous factors (the accumulated effect of which can be thought as turbulence). This external reality can be thought of as both changes in the local market-place in which the firm operates, as well as new directional changes from top management to compensate for those changes. Reality, however, is considered to have some inertia in this model, and thus the conditional probabilities of changing the state favor remaining at the current value rather than flipping from the current value. The organization does not interact or learn from reality directly, but instead from high-performing individuals.

We also posit, as March does, an organizational 'code' that represents the firm's current understanding of reality. This code is developed over time based on the adoption of the values and knowledge of high performers. Although our software allows for modeling perception as an error-prone process, exploration of the critical impact of perception has been reserved for future work.

Finally, in his open system extension (showing both environmental turbulence and turnover), March uses a random function to identify both individuals to be removed and what knowledge newly hired individuals possess. This work uses a random function to identify individuals who depart, but a function informed (though still stochastic) by various variables to control who is hired. These variables are defined below.

2.2 Defining a Stochastic Hiring Function

Morgan, et al. (2010) define seven variables, introduced above, that contribute to the probability of taking a beneficial or negative action towards a third-party. They demonstrated the impact of three of these variables, group size, spatial distance, and presence or absence of leaders, in a simulation of ground combat that realistically replicated some of the social dynamics found in war. Alternatively, this work focuses on the following variables: 1) Group Size, 2) Group Composition, 3) Social Distance, and 4) Mutual Support and Surveillance. We summarize each of these factors, and how they impact the hiring process.

Group size influences individual behavior in a variety of ways. Members of larger groups tend to be able to disassociate from the results of collective action (Grossman, 1995). Large groups, despite having the capacity to do so, are less likely to help needy outsiders (Latane & Darley, 1970). Larger groups, when compared to dyads, tend to allow more confrontational language and are less concerned about actor participation (Slater, 1958).

Consequently our model associates increases in group size with decreases in the likelihood that any particular committee member will recommend a specific candidate, if all other factors are held equal.

The composition, the individuals, of a group also influences the ability of the group to take collective beneficial or negative action. Drawing on the social integration literature (Harrison et al., 1998), we distinguish between surface (superficial or cosmetic differences) and deep-level (differences in attitudes, beliefs, and skills) diversity. Because candidates are attempting to communicate their knowledge, skills, and professional outlook to the hiring committee, we choose to focus on aspects of deep-level diversity. Further, strong group performance tends to correlate more closely to similarities in beliefs than to surface-level characteristics (Terborg et al., 1976).

Our literature review suggests that there is a strong and interesting interaction between Group Composition and Mutual Support and Surveillance. Groups that tend to be diverse are likely to be more welcoming of further diversity, whereas groups where individuals tend to be very similar in attitudes and beliefs find it difficult to hire candidates who do not have similar characteristics. We consider this a group level trait, similar to group size.

Members of groups enjoy several benefits from participation: identity is provided through group norms (Cialdini, Reno, and Kallgren, 1990); rules define and structure ambiguous situations (Chekroun and Bauer, 2002), and help members predict the actions of others (Smith and Mackie, 1995). Social support may diminish the effect of stress (Caplan, 1974).

But while groups offer benefits, they also impose costs on their members. Groups encourage uniformity, and the pressure to maintain that uniformity increases both when differences between members are small, and when inclusion into the group is privileged (Dinter, 1985; Festinger, 1954).

Thus, Mutual Support and Surveillance interacts with group composition. When the group is inherently diverse, there is less pressure to maintain group norms. Candidates who are perceived as similar to the hiring committee are more likely to be hired, provided all other factors are equal. Further, hiring committees of homogenous individuals are likely to take more time and require the consideration of more candidates if the pool of candidates is itself diverse.

Social Distance can be thought of as a continuous scalar, where individuals "just like me" have very low distance scores and individuals who are "not like me" have much larger distance scores. This subscribes to the view advanced by Perloff (1993) and, loosely, to that suggested by Park (1924).

We believe that Social Distance is the feature of a dyad, in this case, the amount of perceived social distance, determined by similarity of beliefs, attitudes, and knowledge, between the recipient and the observer. Individuals with similar attributes tend to interact (McPherson and Smith-Lovin, 1987).

A small social distance contributes to a feeling of connection with the candidate, making it more likely that the committee member will suggest offering employment to that candidate, if other factors are held constant.

3.0 The Stochastic Selection Function

Based on this literature review, we define a hiring function that incorporates these factors. The overall function is a logit-transform, which has been useful in previous discrete choice models (McFadden, 1980). The complete selection function is defined here. This function is a function of functions with each sub-function defined in the following sections.

Equation 1: The probability that a particular target, *t*, will be selected by a particular committee member, *c*.

$$p_{selection} = 1 - \frac{1}{e^{1/(d_{ct})^{pressure} * \sqrt{g_{size}}}}$$

The probability of a particular actor getting hired is based on the rules of that firm.

3.1 Group Composition

This is a relative term indicating the amount of differentiation present in the group compared to the maximal amount of possible variation. A group is maximally variable, has a value $g_c = 1$, if the entire maximal span of variation is represented in the group ($gmax_i - gmin_i = max_i - min_i$) for every feature *i*. The smoothing term, *k*, is to avoid the possibility of division by 0, and should be very small.

Equation 2: Group Composition, g_c, is the amount of variability present in the group compared to the maximal amount of variability present across *n* dimensions.

$$g_c = \frac{\sum_{i=0}^{n}(gmax_i - gmin_i)^2 + k}{\sum_{i=0}^{n}(max_i - min_i)^2 + k}$$

3.2 Mutual Support and Surveillance

This term uses the group composition term, g_c, defined earlier. Because social pressure is very high when group variability is low, we use an inverse function to define social pressure. Because of the k-smoothing term in the definition of g_c, 'pressure' is always defined (although potentially very large). The constant *m* should be specific to the environment in which the equation is applied. We will use the value '.25' in this work; larger values would indicate an environment where more pressure is exerted.

Equation 3: Pressure is the inverse of calculated group composition value, g_c, mediated by the constant, m.

$$pressure = \frac{m}{g_c}$$

3.3 Social Distance

We represent social distance (d) as a Euclidean distance measure across an arbitrary number of dimensions. Given n features, the committee member, c, compares their own feature (each individual feature is c_i) and for the target, t, (the target's value for each feature is t_i). The square root of the sum of these squares produces the distance between the committee member and the target, d_{ct}.

Equation 4: The social distance between a target candidate, t, and a committee member, c, is a Euclidean distance calculated across n dimensions.

$$d_{ct} = \sqrt{\sum_{i=0}^{n}(c_i - t_i)^2}$$

4.0 The Simulation

We are replicating and then extending March's simulation. Briefly, we will present the overall process that characterizes March's model and then describe our extensions to this process.

March's model has these initial properties (pp 74-75):

> *"Within this system, initial conditions include: a reality m-tuple (m dimensions, each of which has a value of 1 or -1, with independent equal probability); an organizational code m-tuple (m dimensions, each of which is initially 0); and n individual m-tuples (m dimensions, with values equal to 1, 0, or -1, with equal probabilities).*

From these starting conditions, the model proceeds as shown in Figure 1, next column:

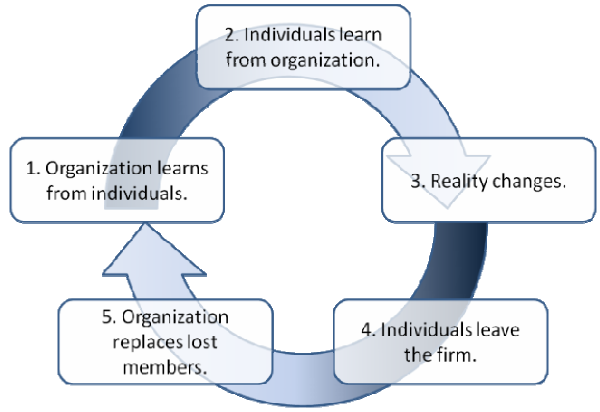

Figure 1: Overview of a single simulation turn in the Mutual Learning Model.

In step 1, *Organization learns from individuals*, the organization identifies high performers, individuals whose beliefs better reflect reality (in aggregate) than the organization's code. The dominant opinion among high performers for each portion of the m-tuple will typically be selected. This process is stochastic, and depends on the level of agreement between high performers. This process is moderated by an "organizational learning effectiveness" variable.

In step 2, *Individuals learn from the organization*, the beliefs of individuals change to reflect the organizational code. For any portion of the code m-tuple whose value is not zero, the individual may change their belief to be in accordance with the organizational code. The probability of them doing this for any portion of the m-tuple is determined by an "effectiveness of socialization" variable.

In step 3, *Reality changes*, the m-tuple of reality is probabilistically changed due to exogenous turbulence. This process is moderated by a "turbulence" variable.

In step 4, *Individuals leave the organization*, individuals are selected randomly from the organization and removed.

In step 5, *Organization replaces lost members*, new individuals join the organization. In March's model, new members of the organization are added as necessary. These new member's beliefs are initialized randomly. We extend March's model by modifying this step through incorporating a selective hiring committee, as shown in Figure 2.

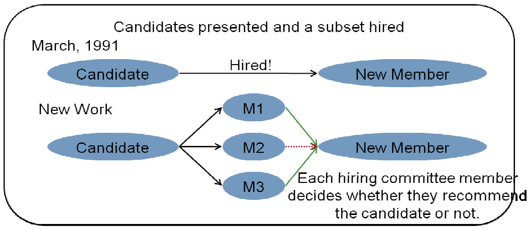

Figure 2: Extensions to the hiring mechanism.

The new work involves the assignment of members to a hiring committee. Figure 2 illustrates this new mechanism; each member of the hiring committee makes an assessment through a stochastic process of whether the applicant is a good fit for the organization. Figure 2 could imply that the committee will select candidates based on "majority rules", but that is not a firm commitment of the model – more complicated rule systems for selection could be used.

Adding this hiring committee requires some adjustments to the overall simulation turn cycle, shown in Figure 3:

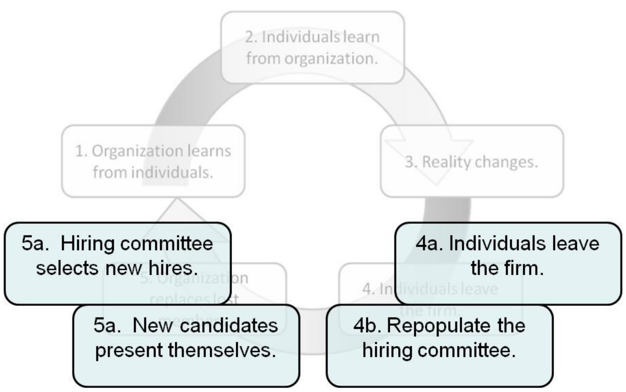

Figure 3: Modifications to March's Mutual Learning Model are in the departure and hiring phase.

The modifications to the departure process are relatively trivial; individuals leave the firm at random, as before. Individuals from the hiring committee are not exempt from this process. If a member of the hiring committee leaves, they are replaced at random with a new committee member from the larger organization.

Our additions to the hiring process are more interesting. As before, new individuals are determined randomly (following the process described for initializing the simulation), but these individuals are candidates. Each candidate is reviewed by the hiring committee, and each committee member makes their choice independently of each other, using the selection function defined previously. The aggregate of the members' individual selections is used to determine whether the candidate is allowed to join the organization as a member.

5.0 Experimental Method and Goals

This project consists of two simulation experiments. The goal of the first experiment is to replicate March's model and his results through a process known as "Docking" (Axtell, et al, 1996). The goal of docking the models is to provide a foundation for comparison, and to ensure that we have usefully implemented the original mechanisms. We compare, as March does, the impact of turnover as a counter-measure to that of environmental turbulence. The experimental variable was the amount of turnover, which was set to either .01 (each person having a 1% chance of leaving the organization each turn) or 0 (no chance of leaving). The organizational learning effectiveness variable was set to .5 and the effectiveness of socialization variable was set to .5. The "reality turbulence" variable was set to .02 (each portion of the reality M-Tuple had a 2% chance of changing). Each organization was composed of 50 actors, and there were 30 bits in the reality M-Tuple. Each simulation had 100 turns, and each condition had 200 separate simulations for a total of 400 separate simulation runs.

The second experiment considered the impact of the hiring committee – do hiring committees affect an organization's performance over time? Committee members were selected randomly from the larger population pool. We used two experimental variables: a) firm profile; and b) whether the firm used a hiring committee or not. We considered three firm profiles: 1) a firm that values exploration, allowing members of the organization to remain diverse and building knowledge slowly; 2) a firm that is exploitative in nature, where individuals rapidly conform to the organizational code, and the organization establishes opinions early; and 3) a firm with average values, neither fast nor slow to socialize employees or gain organizational knowledge. For each firm profile, a firm may or may not use a hiring committee. All other variables were held constant. There were a total of (3 x 2) six combinations; each combination ran for 200 simulation runs, for a total of 1200 simulation runs.

We expected that firms with exploitative profiles (firms focusing on conformance) would find it difficult to hire new candidates that fit their established 'type'. We anticipated this to be true because we believed that the hiring committees for these firms would be less diverse, and thus the social pressure to maintain conformity would be greater than that of the two other profiles.

Our primary performance metric is "Code Knowledge", which measures what percentage of the reality M-Tuple the organization's knowledge correctly reflects. We measure code knowledge on scale ranging from 0 to 100, with 100 being perfect performance. All organizations started with a “0” score, since they started with no opinion on any portion of the M-Tuple. Even low-performance organizations trended towards a ‘50’ or higher, as random chance perturbs reality.

6.0 Results

In this section, we discuss the results of each of the two separate virtual experiments.

6.1 Experiment 1 – The Docking Experiment

Our first experiment compared the impact of turnover on an organization; our goal for this experiment was to replicate March’s findings (i.e., that turnover is a useful explorative mechanism). Replicating March’s findings allowed us to verify that the system was coded properly, in that the organization without turnover reaches an equilibrium point relatively early. Reaching this point, the organization should no longer change, resulting in consistently decreasing performance until reaching 50%. At this point, the system’s turn-by-turn performance takes on the characteristics of a random-walk.

As indicated in Figure 4, we were able to replicate this finding (March’s chart is shown on the right). On the left chart, the compound line shows the average performance of organizations without turnover. The darker solid line shows the average performance of a firm with turnover. In both cases, the organization achieves a certain amount of knowledge, well above the random chance acquisition. Once the knowledge equilibrium is reached, however, the no-turnover organization begins to stagnate steadily, declining in the face of consistent minor turbulence.

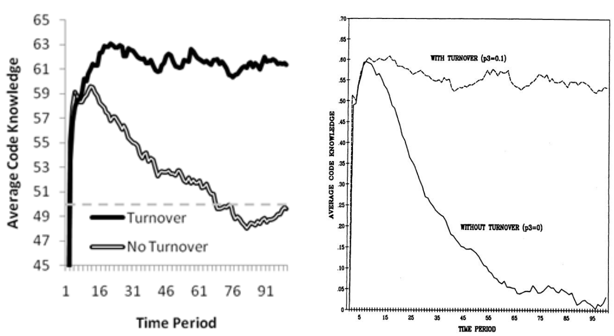

Figure 4: Both models predict that turnover is an effective mechanism for handling a reality with turbulence – Right Chart from March, 1991 (pg. 80)

From this result, we can establish that the two models share a) Component Equivalence (i.e., the models contain the same objects) and b) Relational Equivalence, where the models have similar relationships between these objects. March uses a non-obvious transform converting the code knowledge metric so that the performance dwindles toward 0, rather than .5. Because of this, we do not establish statistical nor numerical equivalence. Because we are not interested in comparing these models “head to head”, relational equivalence is sufficient for our needs.

6.2 Firm Performance

In our second experiment, we examined the impact of hiring committees by comparing firms with and without hiring committees. March predicts that hiring individuals based on their similarity to the code (and presumably to a hiring committee of people influenced by the code) would harm the efficacy of turnover as a mechanism for maintaining an organization’s performance (March 1991, pg, 81). As shown in Figure 5, although the average final performance of organizations with hiring committees is lower than those that use random draws, consistent with March’s prediction – the effect is rather subtle.

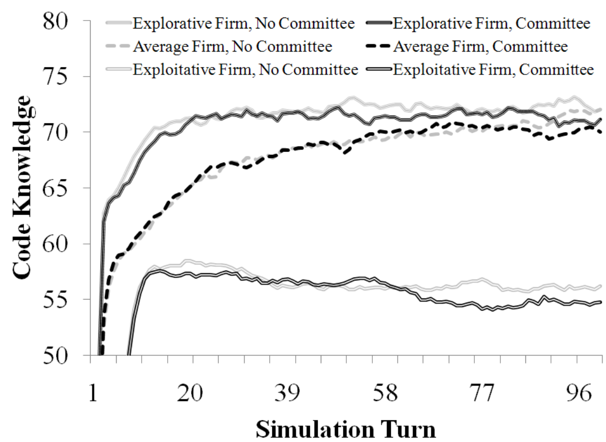

Figure 5: Committees and Random-Draw Firms tend to have similar performance characteristics although final average performance tends to be lower for firms with hiring committees.

Because the committees were extremely selective in hiring new individuals, we ended up adding an opportunity cost metric to our model that made the committee more likely to hire new individuals as it continued to reject previous applicants. Without this variable and mechanism, the committees reviewed tens and hundreds of thousands of applicants for each new position. We felt this was unrealistic, and believed an opportunity cost was an effective if not ideal solution. We, however, also believe that the implementation of this mechanism counter-balanced the expected large degradation of turn-over as a useful mechanism that March predicted. Although individuals that suited the committee, if found, were selected early, the opportunity cost metric eventually causes the majority of the hiring committee to agree to a new, diverse, candidate.

In Figure 6, we see the impact of the pressure to conform over the simulation's time course. The "effectiveness of socialization" variable influenced the diversity of the hiring committee – this was an inverse relationship. When the committee was highly diverse (unlikely in organizations that prioritize socialization), there was relatively little pressure to hire extremely similar candidates. When the committee was very similar, it became very difficult to find acceptable candidates out of the diverse candidate pool. Thus, the committee's diversity, and indirectly the stress the organization put on socialization, strongly impacted the number of candidates reviewed before finding an acceptable person for each position.

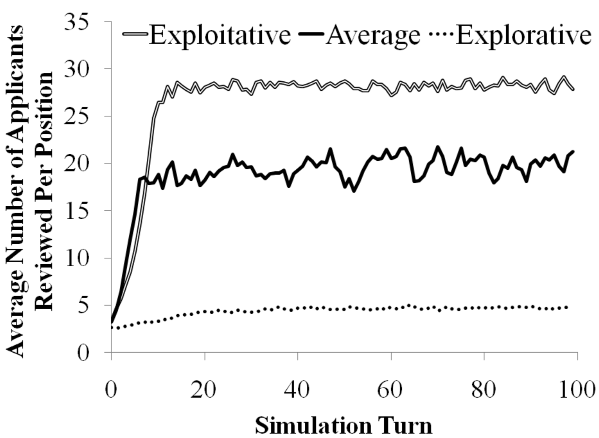

Figure 6: Firms that stressed socialization reviewed many more applicants than those that did not.

7.0 Discussion

This project is an initial attempt to extend March's powerful model by incorporating a theory of selection into the hiring process based on Morgan et al.'s work on participation (2010). Although preliminary work, this work has some interesting ramifications.

Just as Morgan, et al. (2010) showed that the decision to participate in combat was significantly affected by proximity to comrades and enemies – this simulation showed that social incompatibility among members of the hiring committee could deadlock progress. Given an n-dimensional space of reasonable size and a relatively small selection committee, the committee can rapidly find it impossible to agree to any particular candidate – each candidate receiving a single vote from the individual they most resemble. This is one reason we were forced to implement the opportunity cost mechanism

There are other limitations of this work we hope to address in the future, particularly:

One: Individuals and organizations must perceive reality. This process is error-prone and the errors are often interesting and important. The software framework is designed to support perception (as a stochastic process for apprehending reality) but further work must be done to answer some questions relating to perception. Should individuals and organizations be required to perceive rather than simply "know" themselves? Should the error-rates for various kinds of perception be different? What should inform these error rates; and what distribution should the probability model use?

Two: individuals learn, not just from the organizational code, but from each other; perpetuating knowledge both correct and incorrect over time. March (1991) abstracts this important process through his use of the organizational code construct; but in future models, we hope to include individual socialization as well as organizational socialization

Three: hiring committees are complicated. In large organizations, members of hiring committees represent various necessary roles critical to the organization. Each member is expected to weigh in on a specific portion of the applicant's credentials and fit to the organization. In future work, it would be interesting to model an existing organization and its process of hiring, to determine if various structures are more or less capable of neutralizing the challenges imposed by member bias.

Four: committee members are meta-cognitive. Members of hiring committees are aware and may attempt to control for their own biases towards similarity. Further, they are aware that their own performance will be evaluated by outside observers. Future work could involve rewards and penalties for hiring decisions using a reinforcement learning system. This may be a more effective and principled method for incorporating the opportunity cost mechanism.

Works Cited

Axtell, R., Axelrod, R., Epstein, J.M., Cohen, M.D. (1996) Aligning simulation models: a case study and results. *Computational, Mathematical, and Organizational Theory*. 1 123-141.

Caplan, G. (1974). Support systems and community mental health: lectures on concept development. Behavioral Publications, New York: NY.

Chekroun, R.B., and Brauer, M. (2002) The bystander effect and social control behavior: the effect of the presence of others on people's reactions to norm violations. *European Journal of Social Psychology* 32:6, 853-867

Cialdini, R.B., Reno, R.R., and Kallgren, C.A. (1990). Focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. *Journal of Personal Social Psychology* 58:6 1015-1026.

Dinter, E. (1985). Hero or coward: pressures facing the soldier in battle. Frank Cass and Company Limited, Totowa.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*. 7:2. 117-140.

Grossman, D.L. (1995). On killing: the psychological cost of learning to kill in war and society. Little, Brown and Company, New York: NY.

Harrison, D.A., Price, K.H., Bell, M.P. (1998). Beyond relational demography: time and the effects of surface and deep level diversity on work group cohesion. *Academy of Management Journal*. 41:1 96-107.

Latane, B. and Darley, J.M. (1970). The unresponsive bystander: why doesn't he help? Prentice-Hall, Englewood Cliffs.

March, J.G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science* 2:1, 71-87

McFadden, D. (1980). Econometric models for probabilistic choice among products. Journal of Business 53:3, S13-S29.

McPherson, J. M., Smith-Lovin, L. (1987). Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups. *American Sociological Reivew* 52 370-379.

Morgan, J.H., Morgan, G.P., and Ritter, F.E. (2010). A preliminary model of participation for small groups. *Computational, Mathematical, and Organizational Theory*. 16. 246-270.

Park, R.E. (1924). The concept of social distance as applied to the study of racial attitudes and racial relations. *Journal of Applied Sociology* 8: 339-344.

Perloff, R.M. (1993). Third-person effect research 1983-1992: a review and synthesis. *International Journal of Public Opinion Research* 5 167-184.

Slater, P.E. (1958). Contrasting correlates of group size. *Sociometry* 21 129-139.

Smith, E.R. and Mackie, D.M. (1995). Social Psychology. Worth Publishers, New York: NY.

Terborg, J.R., Castore, C. DeNinno, J.A. (1976). A longitudinal field investigation of the impact of group composition on group performance and cohesion. *Journal of Personal Social Psychology* 34 782-790.

Author Biographies

GEOFFREY P MORGAN is a PhD student in Carnegie Mellon's Computation, Organization, and Society program. Morgan has experience in developing autonomous robotic systems, self-learning systems, cognitive models, and network-enabled simulations. He is interested in organizational impacts on individual decision-making and on optimizing organizational performance.

KATHLEEN M CARLEY is the director of the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. Her research combines cognitive science, social networks and computer science to address complex social and organizational problems.