

Evolutionary Perspective on Collective Decision Making

Dene Farrell¹, Hiroki Sayama¹, Shelley D. Dionne²,
Francis J. Yammarino², David Sloan Wilson³

¹ Department of Bioengineering

² School of Management / Center for Leadership Studies

³ Departments of Biological Sciences and Anthropology
Binghamton University, State University of New York

Team decision making dynamics are investigated from a novel perspective by shifting agency from decision makers to representations of potential solutions. We provide a new way to navigate social dynamics of collective decision making by interpreting decision makers as constituents of an evolutionary environment of an ecology of evolving solutions. We demonstrate distinct patterns of evolution with respect to three forms of variation: (1) Results with random variations in utility functions of individuals indicate that groups demonstrating minimal internal variation produce higher true utility values of group solutions and display better convergence; (2) analysis of variations in behavioral patterns within a group shows that a proper balance between selective and creative evolutionary forces is crucial to producing adaptive solutions; and (3) biased variations of the utility functions diminish the range of variation for potential solution utility, leaving only the differential of convergence performance static. We generally find that group cohesion (low random variation within a group) and composition (appropriate variation of behavioral patterns within a group) are necessary for a successful navigation of the solution space, but performance in both cases is susceptible to group level biases.

1 Introduction

Collective decision making is becoming more central and indispensable in human society as modern problems increasingly involve interactivity and inseparability within large scale tasks [4]. In high-tech product and software development, for example, the amount of workers participating in a design project can be in the order of thousands as a result of a product's complexity exceeding an individual's capacity, which almost inevitably results in suboptimal outcomes [5]. More recently, online collective decision making among large populations of anonymous participants via computer mediated networks has been implemented for product rating and common knowledge base formation. Both individual behavior and the organizational structure greatly influence decision processes. The complexity of the processes is more manifested when constituents of groups are heterogeneous with regard to both their world views and behavioral propensities. Collective human decision making in such conditions is poorly understood, being one of the most significant challenges in the social sciences.

The leadership, psychology and organizational behavior/management disciplines have examined collective dynamics using both experimental and applied studies. They generally emphasize linear statistical relationships of team and individual level variables [4], without accounting for nonlinear processes, high-dimensional problem space and non-trivial social structure. Complex, nonlinear problem space has been considered in dynamical modeling studies [5], in which interdependence of aspects of

problem are considered, but not nontrivial social interactions.

Here we investigate collective decision making dynamics from a novel perspective by shifting the focus of agency from group members to potential solutions being discussed. The decision making processes are described using concepts in evolutionary theory, where evolution acts on a population of potential solutions through mechanisms of selection and variation as effected by human discussants. Group members thus serve both as an evolutionary environment and as implements of evolutionary action on a population of solutions. Within this context, several evolutionary operators can be mapped to human behaviors. Examples include replication (advocacy of an existing idea), subtractive selection (criticism against an existing idea), mutation (revision of an existing idea) and recombination (creation of a new idea by mixing existing ideas).

2 Model

2.1 Groups

We apply evolutionary framework to model simple group decision making processes within a small-sized, well-connected social network structure. We conduct multiple levels of analysis [1,7] on how homogeneities or heterogeneities of world views/goals among the participating agents, as well as group-level behavioral patterns and biases, affect the decision making dynamics and the final outcomes [2].

When group members are heterogeneous in world views, differences between individual utility functions play a crucial role in determining the group dynamics; the relevant level of analysis is within groups. Each member acts as “group parts” [1] to achieve individual objectives. Conflicts of interest make the problem space more complex than that of groups consisting of homogeneous, world perspectives. On such complex landscapes there is more possibility for populations to become stuck at local optima, detrimenting the overall adaptiveness. Contingently, the importance of variation relates to escaping from the local optima in order to reach better solutions.

If group members are homogeneous in their world view, they behave as “group wholes” [1]; the relevant level of analysis is between groups. The population of solutions evolves to adapt to a single utility function shared by all the group members, so the problem space would be simpler than with heterogeneous groups. With little conflicts of interest, selection is relatively important to adaptiveness as speeding up the convergence of discussion. Variation still holds importance, especially with complex nonlinear problems.

Our model assumes that groups are initiated with a list of randomly generated ideas, whereupon they begin to perform a set of actions on the existing population of solutions repeatedly for a fixed number of iterations. Individuals always act in the same order and groups always demonstrate a full rotation. The number of actions on the population of solutions is a product of the number of group members, N , and the number of iterations, t .

In the population, there may be multiple copies of the same type of solution, which represents the relative popularity among group members. Each action is performed on a single copy of solution, not on an equivalence class of all solution replicates.

2.2 Utility Functions

Groups are situated in an M -dimensional binary problem space, with 2^M possible solutions. For a simulation,

every solution has a utility value specified by a master utility function U that is unavailable to group members. Individuals perceive solution utility values based on their own utility functions U_j constructed by adding noise to U . We develop a semi-continuous assignment of utility values in the problem space in the following way. First, s representative solutions $S = \{\mathbf{v}_i\}$ ($i = 1 \dots s$) are generated as random bit strings, where each \mathbf{v}_i represents one solution made of M bits. One solution is assigned the maximum fitness value, 1, and another, the minimum fitness value, 0. The remaining $s - 2$ solutions are assigned a random real value between 0 and 1, ensuring that the entire range of utility values is from 0 to 1, for the sake of comparisons between simulation results.

The utility values of all possible solutions in the domain of the master utility function are defined by interpolation using the utility values of representative solutions in S . We use the Hamming distance as a measure of dissimilarity between two bit strings. With this measure, the utility value of each possible solution not present in S is calculated as a weighted average of the utility values of the representative solutions calculated as follows:

$$U(\mathbf{v}) = \frac{\sum_{i=1}^s U(\mathbf{v}_i) \cdot D(\mathbf{v}_i, \mathbf{v})^{-2}}{\sum_{i=1}^s D(\mathbf{v}_i, \mathbf{v})^{-2}} \quad (1)$$

where $\mathbf{v} \notin S$ is the solution in question, $U(\mathbf{v}_i)$ is the utility of a representative solution \mathbf{v}_i in S , and $D(\mathbf{v}_i, \mathbf{v})$ is the Hamming distance between \mathbf{v}_i and \mathbf{v} .

Each individual in a group will unconsciously have a different set of utility values for the possible solutions of the problem. Individual utility functions $U_j(\mathbf{v})$ ($j = 1 \dots N$) are generated by adding random noise to the master utility function so that:

$$U_j(\mathbf{v}) \in [\max(U(\mathbf{v}) - \nu, 0), \min(U(\mathbf{v}) + \nu, 1)] \quad (2)$$

for all \mathbf{v} , where ν is the parameter that determines the range of noise. Individuals do not access global maximum/minimum utility values, though they can retrieve a utility value from the function when a specific solution is given.

In addition to individual deviations from a common master utility function, we investigate the effect of common deviations from the “true” utility function, or group level biases. For simulating group level bias we introduce a new step in the generation of individual utility functions, in which the master utility function $U(\mathbf{v})$ differs from the original true utility function, $U_T(\mathbf{v})$. Specifically, a bias β is imposed on the true utility function both by flipping bits with probability 0.25β per bit and adding a random number ranged $[-\beta, \beta]$ to utility values. Solution sets are renormalized to the range $[0, 1]$. The master utility function is generated from the biased representative solution set. Subsequent methods follow as described above. Bias represents fidelity of information at the group level, where $\beta = 0$ denotes perfect information, and complete randomization is asymptotically approached as bias increases.

2.3 Evolutionary Operators

We identify six evolutionary operators representing individual behaviors reflecting selection or variation. Some operators use a preferential search algorithm to stochastically search the solution population, where r_p solutions are randomly selected and ranked according to their perceived utility values, and then the best or worst solution is selected depending on the nature of the operator being executed.

Replication. Replication adds an exact copy of a solution from the population of solutions back

onto the list. Solutions are chosen for replication with the preferential search algorithm. Replication therefore can neither produce a novel solution nor remove one, but it gently sways the ecology of the population by increasing the popularity of favorable existent solutions. This represents an advocacy of a particular solution under discussion.

Random point mutation. Random point mutation adds a copy of a solution with point mutations, flipping of bits at each aspect of a problem with a probability p_m . The solution on which the operator acts is chosen from the active population with a preferential search algorithm (discussed in more detail below). This represents an attempt of making random changes to the existing ideas, reflected in asking “what if” questions. Random point mutations help escape local maxima of a utility function in the problem space when a utility function is nonlinear and many-peaked.

Intelligent point mutation. A solution is selected from the population with a preferential search algorithm. It makes several (r_m) offspring of the parent solution and selects that of the highest perceived fitness for addition to the population. This represents a proposal of an improved idea derived from existing ideas under discussion. The intelligent point mutation can be useful in maximizing a utility function with one maximum by climbing monotone gradients, but it may perform poorly in a complex utility landscape.

Recombination. Recombination chooses one solution at random and one with a preferential search algorithm. It then creates two offspring from the two parent solutions. Sexual reproduction is simulated with a multiple point cross-over recombination: parent solutions are aligned by aspects, for each of which there is a probability p_s of switching their contents. Of the two offspring, that of higher perceived utility is selected and added to the population. This represents a creation of a new idea from two existing ideas.

Subtractive selection. The preferential search algorithm is used to find the solution with the worst fitness, whereupon it is singled out and deleted from the population. This represents a criticism against a bad idea. Subtractive selection is the only operator that reduces the number of existing solutions and is therefore essential to groups attempting to attain convergence in the population distribution.

Random generation. Finally, random generation of solutions adds a randomly generated solution to the population. There is no use of an individual’s utility function, nor any connection to the existing solutions “on the table” at that time. New solutions are generated utterly randomly. This represents a sudden inspiration of a totally unique idea that is unrelated to the existing ideas under discussion.

2.4 Simulation Settings

The following parameter settings were held constant for all simulations: group size $N = 6$; problem space dimensionality $M = 10$; number of sample solutions in the preferential search algorithm $r_p = 5$; number of offspring generated in the intelligent point mutation $r_m = 5$; random mutation rate per bit $p_m = 0.2$; probability of random switching in recombination $p_s = 0.4$; number of iterations $t = 60$. It was also assumed that groups were initialized with four random ideas. For each group, the noise parameter \mathcal{V} and the bias parameter β were varied from 0 to 1.2 by increments of 0.2.

2.5 Metrics of Group Performance

We use two separate performance metrics: the true utility of the mode solution at the end of group

simulation and the convergence of solutions. Convergence is based on entropy

$$H = - \sum_{i=1}^n p(x_i) \cdot \log_2 p(x_i) , \quad (3)$$

where $p(x_i)$ is a normalized frequency of the i -th type of solutions in the population. Since the maximum possible value of H is M (this is the case when there are exactly 2^M solutions in the population which are different from each other), $M - H$ is a quantitative measure that intuitively means the number of aspects of the problem on which the group has formed a cohesive opinion. For normality, we will use $(M - H) / M$ as the metric.

3 Results

We first conducted a within-group analysis examining effects of heterogeneity in world views (utility functions) within a group. Here we assumed group members were “balanced” behaviorally; in each iteration, they randomly chose one of the six operators with equal probability. Figure 1 indicates the results with several different settings of within-group variation ν and group-level bias β , plotting them in a 2-D performance space using the two metrics described above.

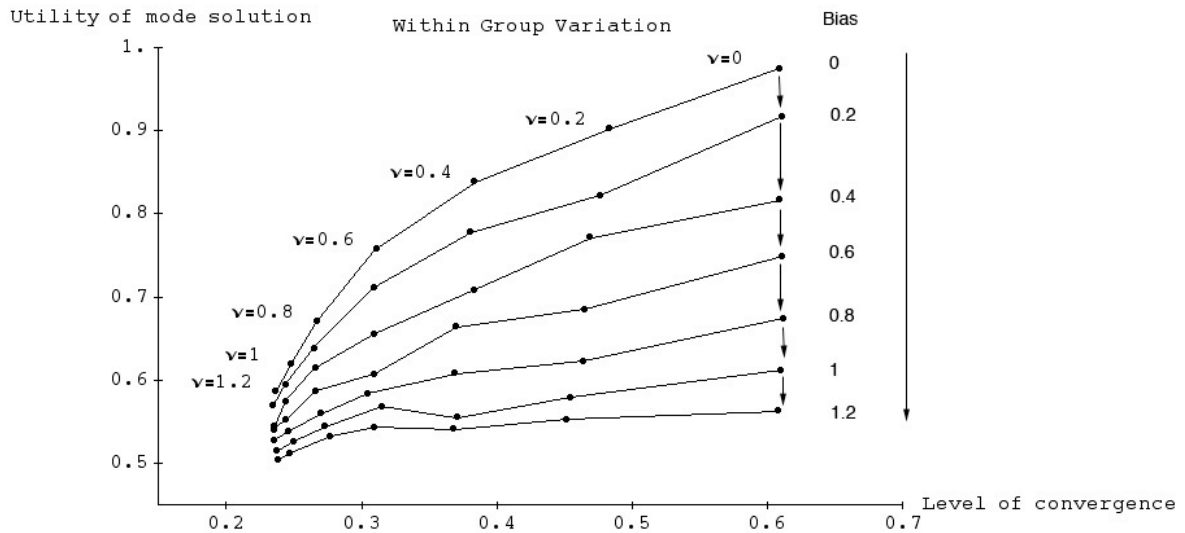


Figure 1: Simulation results showing the effects of within-group heterogeneity and group level bias. The level of convergence and the true utility value of mode solutions for several different noise levels are plotted. $\nu = 0$ represents the case of completely homogeneous groups, while larger values of ν represent more heterogeneous group cases and larger values of bias represent large discrepancies between group utility functions and true utility function.

Group-level bias affects the utility of group solutions while convergence is largely unaffected. On the other hand, within-group variation degrades both convergence and utility. Groups performed better in both performance metrics when they were homogeneous in their utility functions. As the groups’ members become more heterogeneous, the true utility value of the mode solution decreased and the final population of solutions after discussion became more diverse. The decrease of the true utility value was particularly drastic; in nonbiased conditions with no heterogeneity, the groups were able to find nearly perfect solutions

for the problem (i.e., the utility close to 1). As the groups become more heterogeneous the utility achieved dropped to just above 0.5, meaning that there was no net improvement achieved during the group discussion. This was due to the conflicts of interest among the group members.

Contrary to other findings regarding heterogeneous groups outperforming homogeneous groups on creative and intellectual problem solving tasks [3, 6], our findings indicate the opposite, which may seem to support the negative relationship reported between both surface-level (i.e., demographic) and deep-level (i.e., psychological) diversity and group functioning and performance [2]. We must note here, however, that the diversity considered in this set of experiments is about the individual utility functions only, and not about the individual behavioral patterns.

In order to explore the effects of various compositions of individual behaviors, we ran another set of experiments using the same simulation model with different behavioral patterns assumed for different groups. In forming different group properties, we modeled only a handful of potential evolutionary operators/behaviors combinations. We modeled some operators singularly (e.g., random generation was the only operator within the group), and for other groups we combined two evolutionary operators to reflect increasing complexity of group behavior (e.g., recombination and intelligent point mutation). For the former cases group members were assumed to choose the designated operator for 95% of their total actions, with 1% for each of the other five operators. For the latter combined cases they were assumed to choose each of the two operators for 48% of their total actions (96% in total), with 1% for each of the other four operators. We limited our examination to eight group types: replication and subtractive selection (Group 1); subtractive selection and random point mutation (Group 2); replication and recombination (Group 3); recombination (Group 4); recombination and intelligent point mutation (Group 5); intelligent point mutation and random generation (Group 6); random generation (Group 7); and, finally, the balanced team we used in the previous experiment as a control (Group 0).

Figure 2 shows the results of the second set of experiments comparing group performances with different group properties, plotting them in the same 2-D performance space as used for Figure 1. The effect of group-level bias is similar to that seen in Figure 1. Among the groups examined, the “balanced” Group 0 case was the best in terms of the utility value of the mode solution. Interestingly, however, we saw a variety of different group performances achieved by groups with different properties, seen as a kind of “wave front” near the upper-right corner of the performance space.

We further noticed in Figure 2 that the groups that sit along this wave front were arranged roughly in the order of the balance between selection and variation in evolutionary operators; Group 1, which was the best in terms of the convergence but poor in the mode selection utility, used the combination of replication and subtractive selection, which are both selection-oriented operators. Group 2, the second best in convergence and second worst in mode selection utility within the wave front, used the combination of subtractive selection and random point mutation, which is more variation-oriented than Group 1. Along the way toward Group 0, we saw Group 3 (replication and recombination), Group 4 (recombination only), and Group 5 (recombination and intelligent point mutation), where the qualitative shift of balance of evolutionary operators from selection-oriented to variation-oriented can be seen. It is also notable that the random generation operators (used in Groups 6 and 7) were generally not working for improving group performance.




Figure 2: Simulation results showing the effects of group-level difference of behavioral patterns. The level of convergence and the true utility value of mode solutions under varying conditions of group-level behavioral patterns and bias are plotted. Group 0 is the balanced team that uses all of the six evolutionary operators with equal probability; Group 1 uses replication and subtractive selection mostly; Group 2 subtractive selection and random point mutation; Group 3 replication and recombination; Group 4 recombination only; Group 5 recombination and intelligent point mutation; Group 6 intelligent point mutation and random generation; and Group 7 random generation only.

4 Conclusion

The application of the evolutionary paradigm is illuminating for studying collective decision making dynamics because it allows the researcher to remove themselves from the traditional teleology adopted in most simulations of human groups. We have portrayed group dynamics in a novel way by treating members of the group as constituents of an evolutionary environment in which populations of solutions evolve. In this new framework, we have characterized the properties of the population of solutions after discussion as quantitative metrics of the performance of a group. We demonstrated through simulations that heterogeneous groups with random variations in individual utility functions had a drop in both utility and convergence of solution populations compared to more homogeneous groups. We also demonstrated that variations in the compositions of individual behavioral patterns between groups resulted in a large spectrum of performance, in which groups well balanced between reductive and creative evolutionary forces yielded solutions that were highly adaptive by both performance metrics. All operators have a particular utility in appropriate circumstances, but we highlight that recombination operators are particularly important in that they demonstrate creative changes on large and small scales with a single mechanism, as are selection operators essential to promoting the best ideas by converging a solution population on the best solutions.

Bibliography

- [1] Dansereau, F., Yammarino, F. J., & Kohles, J. C. 1999. Multiple levels of analysis from a longitudinal perspective: Some implications for theory building. *The Academy of Management Review*, 24(2): 346-

- [2] Dionne, S. D., Randel, A. E., Jaussi, K. S., & Chun, J. U. 2004. Diversity and demography in organizations: A levels of analysis review. *Research in Multi-Level Issues* 3: 181-229.
- [3] Gruenfeld, D. H., Mannix, E. A., Williams, K. Y., & Neale, M. A. 1996. Group composition and decision making: How member familiarity and information distribution affect process and performance. *Organizational Behavior and Human Decision Processes* 67: 1-15.
- [4] Kerr, N. L., & Tindale, R. S. 2004. Group performance and decision making. *Annual Review of Psychology*, 55: 623-655.
- [5] Klein, M., Faratin, P., Sayama, H., & Bar-Yam, Y. 2006. An annealing protocol for negotiating complex contracts. In J.-P. Rennard, ed., *Handbook of Research on Nature Inspired Computing for Economics and Management*, vol. 2, Chapter XLVIII. Idea Group Publishing.
- [6] Nemeth, C. J. 1992. Minority dissent as a stimulant to group performance. In S. Worchel, W. Wood, & J. Simpson, eds., *Group process and productivity*, pp.95-111. Sage.
- [7] Yammarino, F. J., Dionne, S. D., Chun, J. U., & Dansereau, F. 2005. Leadership and levels of analysis: A state-of-the-science review. *The Leadership Quarterly*, 16: 879-919.