13th ICCRTS: C2 for Complex Endeavors

"Group Decision Making by Insects and Humans"

Modeling and Simulation
Cognitive and Social Issues

Michael I. Bell

Alidade Incorporated

31 Bridge Street
Newport, Rhode Island 02840

(301) 871-6247

mike.bell@alidade.net

Group Decision Making by Insects and Humans

Abstract

A model developed by biologists to describe nest site selection by social insect colonies has been generalized and used to explore the group decision-making behavior of humans in organizations such as command staffs, integrated process teams, and other task groups. This decision process consists of an exploration phase, in which team members discover potential solutions to the problem at hand, a recruitment phase, in which advocates for a given solution attempt to persuade their colleagues to support that option, and a "voting" phase, in which a decision may be reached. A main focus of this work is the trade-off between speed and accuracy. For decisions made by a plurality vote with a specified quorum, speed may be increased by lowering the decision threshold (quorum) or by raising the effectiveness of recruiting. In principle, either approach reduces the average decision quality by favoring options discovered early in the process. Perhaps surprisingly, accurate decisions can often be made by much less than a majority vote; a quorum as small as 20% of the group may be sufficient. The process does not require individuals to make direct comparisons of the alternatives or to make highly accurate assessments of their quality.

Keywords: social choice, decision making

Introduction

Social choice, the process of converting individual preferences into group decisions, has been widely studied but is far from completely understood. Recent work on group decision making in animals (see Franks *et al*. 2002) has yielded process models (Passino and Seeley 2006) that, with careful interpretation, can yield insights into human activities.

Biological Background

The model on which the current work is based was developed to describe the process of nest site selection by honey bees (*Apis mellifera*). Very similar models have been used to describe food source location by honey bees and "house hunting" by colonies of ants (*Leptothorax albipennis*, Franks *et al*. 2003).

Nest site selection, or house hunting, is a problem that social insects must solve when a successful colony outgrows its nest or hive. In honey bees, this typically occurs in the spring, when the queen bee and roughly half the workers (as many as 10,000 bees) leave the hive in a swarm. This swarm settles on a temporary resting place, often a tree branch, while several hundred scout bees explore for a suitable nest site, typically a cavity in a tree. Over the course of several days, perhaps a dozen potential sites may be found, from which the swarm must make a choice. After each exploratory flight, the scouts return to the swarm, and those that have discovered potential nest sites communicate the location

of their discoveries to other scouts with a variant of the "waggle dance" used to transmit the location of food sources (Visscher 2003). This dance is performed on the surface of the swarm, and its duration (and hence the likelihood of contact with other scouts) increases according to the assessed quality of the candidate site. (Major quality factors are cavity volume, height above the ground, and size and orientation of the entrance hole. It is not known how, or how consistently, the bees weight these factors.) Other scouts are "recruited" by these dances and fly off (with the recruiter) to examine the proposed sites and may return to recruit yet other scouts. When the number of scouts visiting a particular site exceeds a certain threshold (typically 10 to 20 bees), somewhat loosely called a "quorum," the bees at that site return to the swarm, and instead of continuing to recruit they produce a vibrational signal ("piping"), which stimulates the rest of the swarm to warm up their flight muscles in preparation for a flight to the chosen site (Seeley and Visscher 2003).

A key feature of the recruitment process in honey bees is that when bees return to a site for which they have recruited and find fewer than a quorum of scouts at that site, their commitment to that site decreases, and they recruit less vigorously for it when they return to the swarm. Eventually, if no quorum is reached for that site, the scout will cease visiting or recruiting for it and potentially become available for recruitment to another option. This so-called "expiration of dissent" (Seeley 2003) serves both to avoid deadlocks in which no choice can achieve a quorum and to prevent poor solutions discovered early in the process from being selected before there has been an opportunity to discover better options.

Passino-Seeley Model

A quantitative description of nest site selection in honey bees has been developed (Passino and Seeley 2006), based on a stochastic, discrete time model in which scout bees may belong to various groups, exploring, observing the dances of others, or resting, and be committed or uncommitted to a nest site. The probabilities for transitions between these groups were derived, as much as possible, from field observations. We refer to this as the Passino-Seeley (PS) model and use it as the basis for further investigation.

Relevance to Human Groups

A honey bee swarm is very vulnerable in its exposed interim location, and the ultimate selection of a nest site is so vital to its survival that the process has been described as "a form of crisis management" (Franks *et al*. 2002). The trade-off between speed and decision quality is crucial, and it should not be surprising that natural selection has produced a process that achieves an excellent balance between the two. The physical and cognitive limitations of the individual bees, relative to the complexity of the task at hand, make it advantageous to employ a collective, distributed process involving a high degree of self-organization (Seeley 2002). When this decision-making process is reduced to its basic elements, removing features specific to social insects, the result is an abstract process model with obvious relevance to group decision making in other species, including humans.

The Model

The process model we have examined is shown schematically in Figure 1. We will neglect the issue of problem definition, despite its potential impact on human decision making, and assume that, as in social insects, the problem is well known and understood identically by all the participants. Our attention will focus on the other three principal steps of the process: exploration, recruitment, and voting (indicated by rectangular boxes in Figure 1) as well as the overall process flow.

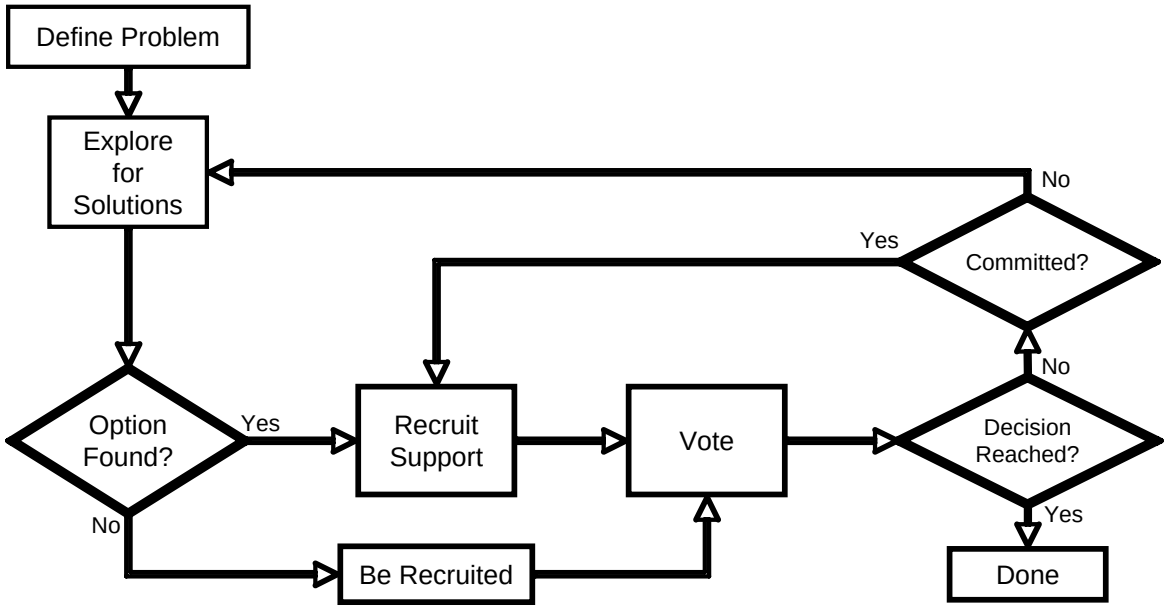

Figure 1. Group decision making process

Exploration

This is the only phase of the process that confronts the external "real world." The number of potential solutions to the problem at hand and the ease with which they can be discovered are generally beyond the control of the decision-making group, and constitute what we term the "option space." We assume implicitly that there are a finite number of distinguishable options facing the group. This is almost certainly the case for social insects, and among humans it is often one of the goals of problem definition to make it highly likely. For the purpose of modeling, we make three other key assumptions. 1. All group members are equally likely to find a possible solution. 2. Solutions are discovered by random search of the option space, so that each solution is equally likely to be discovered during each exploration phase. 3. There exists a "ground truth" measure of the quality of each proposed solution. These represent reasonable assumptions for house-hunting honey bees and a starting point for the generalization to humans.

Recruitment

With no bias in favor of high-quality solutions during the discovery phase, the only mechanism for making good (or better than random) selections is to have recruiting effectiveness increase with solution quality. Honey bees accomplish this by making the duration of the waggle dance proportional to the perceived nest site quality. The intensity with which humans advocate for a particular course of action is influenced by (among other things) the perceived value of that option, and they seem fairly adept at assessing both the intensity and the sincerity of the opinions of others. For simplicity in modeling, we assume that all group members are equally effective recruiters for a solution of given quality, and we assume that all group members are equally accurate in assessing the quality of proposed solutions.

The recruiting process may depend on one-to-one or one-to-many interactions within the group. When honey bees dance to advertise a nest site, the range over which they can be observed is very limited, and the interactions are essentially one-to-one. Humans employ both mechanisms. From the point of view of modeling, it is sufficient to implement only one-to-one interactions if all interactions are between recruiters and potential new recruits. Interactions among uncommitted individuals that alter the probability of their recruitment are possible, but these are beyond the scope of the present work.

Recruiting interactions may be identical between all possible pairs of individuals, or they may be governed by an organizational or social structure that gives special characteristics to certain individuals or interactions. In social insects, there is no evidence that any of the scouts are significantly different from the others or that any particular structure governs their interactions. This is clearly not the case in human groups, but dealing with individual variations in recruiting effectiveness and the influence of organizational structure or social networks has been left to future work.

Voting

We assume that the decision criterion does not change during the process and that the decision depends only on the votes of the group members. Other decision rules, such as "convincing the boss" (or a hierarchy of bosses) or restricting the vote to an "inner circle" are possible but will not be discussed here. It is worth noting that Figure 1 makes no assumption about where each phase of the process takes place. In honey bees, voting and recruitment occur at different locations, the former at the candidate nest sites and the latter at the swarm. In humans, both may take place at meetings of the group or be mediated by one or more forms of communication if the group is spatially dispersed.

Continuing the Process

As long as no decision has been reached, both exploration and recruiting continue. Voters may remain loyal to their previous choices (but with diminished intensity) or begin to explore and/or become available for recruitment to other options. Two aspects of the PS model have been omitted from Figure 1 for clarity but are retained in our work. First, empirical studies indicate that up to 10% of the scouts (but not other members of the

swarm) may die during the course of a migration. Passino and Seeley assumed a corresponding per-sortie attrition rate, but found that it had little effect on the results. (In our model, individuals of this type are described as "retired" rather than "dead.") Second, a scout that ceases recruiting for a site once its recruiting intensity has decayed to zero may rest for some period of time before becoming available to explore or be recruited. The rate at which such scouts stop resting is a parameter of the model, but it has little effect, since we find that very few of these bees will ever discover or be recruited to another site.

Timing and Synchronization

The process of Figure 1 is clearly a sequence of activities for any individual, but an important question is whether the entire group is synchronized as the members pass through each of the phases. Observation of house hunting in insects indicates that all of these activities are taking place simultaneously. In honey bees, however, the swarm forms in a coordinated manner and the first expedition of scouts probably departs during a reasonably well-defined time interval. The PS model treats the scouts as essentially synchronized and measures time in units of the expedition (one sortie for each scout that is not resting or dead), lasting roughly 30 minutes. Over time, there is likely to be a loss of coherence among the scouts, and this will produce some error in quorum counting by the model, since in reality not all of the bees committed to a given site will be present at the site simultaneously. Also, the PS model imposes a time limit of 64 cycles to reach a decision, on the assumption that after 32 hours (plus delays for darkness and weather), the accumulated risk to the swarm is too great and the search should be considered a failure. These questions may be less relevant to human decision-making groups, which are notoriously effective at devising timetables and maintaining synchronization. The question of deadlines, however, and changes in behavior caused by their approach, is important although beyond the scope of this work.

Agent-Based Simulation

At this point, one might reasonably ask whether the process shown in Figure 1 is too generic to be interesting. We have developed an agent-based computer implementation of this model that not only reproduces the results of Passino and Seeley (2006) but also exhibits a variety of behaviors relevant to human group decision making, some of which are surprising and counter-intuitive.

We have used the NetLogo agent-based modeling environment (Wilensky 1999) to construct a numerical simulation of the model of Figure 1. As a starting point we have taken the parameters used in the baseline case of the PS model. These are listed in Table 1. Some of the entries in Table 1 require further explanation, which will be provided in the context of the phases of the model to which they apply.

Quantity	Definition	Value
N	Number of scouts	100
A	Number of search areas	441
S_1	Option space	{0.1, 0.3, 0.35, 0.5, 0.55, 1}
p_m	Maximum recruiting probability	0.037
Δ_r	Recruiting decay rate	0.1
Q_0	Quorum (fraction)	0.2

Table 1. Baseline model parameters

Strictly speaking, the absolute values of these parameters, which have been derived from observation of honey bee house hunting, have no special significance. The behavior of the model depends only on their relative values, and wherever possible we will frame the discussion in terms of ratios and probabilities.

Exploration

The absolute number of scouts is significant only to the extent that it affects the run-to-run statistical fluctuations in the model results. We found it convenient to use the value of 100 employed in the PS model. Based on field observations, Passino and Seeley divided the nest site search area into a 21 by 21 grid of sub-areas, each of which can be searched by a single scout during one sortie (time step). They assumed the existence of six potential nest sites, so that the probability of a scout discovering a site on any given sortie is $p_d = 6/441 = 0.0136$. Despite this small per-sortie discovery probability, the initial expedition of 100 scouts has probability $1 - (1 - p_d)^{100} = 0.75$ of discovering at least one of the sites and probability 0.2 of discovering the best one.

The six nest sites have relative quality given by the option space S_1, which is a list of the quality values of the sites. (The best site is always assigned a quality value of one.) Note that this option space provides a "clear winner," in the sense that the best site is nearly twice as good as the second best. Passino and Seeley also explored an excursion using a "Lake Woebegon" option space $S_2 = \{0.5, 0.6, 0.7, 0.8, 0.9, 1\}$, where nearly all choices are "above average." We have examined two other families of option spaces, and the full list of spaces available in the model is shown in Table 2.

Space	No. of Options	Options	Description
S_1	6	{0.1, 0.3, 0.35, 0.5, 0.55, 1}	"Clear Winner" (PS Baseline)
S_2	6	{0.5, 0.6, 0.7, 0.8, 0.9, 1}	"Lake Woebegon" (PS Excursion)
S_3	N_s	$\{1/N_s, 2/N_s, \ldots 1\}$	Uniform
S_4	N_s	$\{Q_2, \ldots Q_2, 1\}$	Binary

Table 2. Option spaces

The uniform option space is useful because it eliminates the possibility of bias in the results produced by having a high or low average quality of the options. The average solution quality for S_3 is $(N_s + 1) / 2N_s$. This approaches 0.5 for sufficiently large N_s. The binary option space S_4 is useful in studies of the discrimination capability of the process. The quality Q_2 of the second-best solution can be varied, along with other model parameters, to determine the point at which the process can no longer distinguish between the two solutions. Also, by varying N_s we can introduce replicates of the second-best option. These make discrimination more difficult by attracting scouts away from the best site. A study in progress is designed to map out the combinations of N_s and Q_2 that permit a given rate of success in discrimination.

Recruitment

In order to determine the recruitment probability for uncommitted scouts, the Passino and Seeley describe recruitment as two-step process. Scouts will wait to be approached by a recruiter but will eventually give up and go exploring themselves if forced to wait too long. This behavior is described by an empirical equation for the rate at which observers transition to exploration as a function of the total amount of recruiting, with parameters chosen to agree with field data. The remaining observers are recruited to the various candidate sites in proportion to the intensity of recruiting for each site.

In our agent-based model, it is more natural to base recruiting on interactions between the agents. For simplicity, we assume that each returning scout that has discovered a new candidate site attempts to recruit all the other uncommitted scouts and that he succeeds with a probability $p_r = p_m Q_s$, where Q_s is the assessed quality of the site. On subsequent returns from the site, there is a decrease in recruiting intensity, with the result that the recruiting probability becomes the greater of $p_r = p_m (Q_s - n \Delta_r)$ or zero for the nth return. Thus, p_m is the maximum probability of recruiting success, and its value has been chosen to give a scout returning from the first expedition after discovering the best nest site the same expected number of new recruits as in the PS model.

Voting

Passino and Seeley chose the baseline quorum value of $Q_0 = 0.2$ (corresponding to 20 scouts) based on field observations. Since both their study and the present one explore the effects of varying Q_0, this value is significant in only one respect. in that such a relatively small quorum is required in honey bee nest site selection, while most human decision-making groups seem to assume that a majority or super-majority is needed for "consensus" decisions.

Simulation Results

Figure 2 is a view of the user interface of the NetLogo simulation at the end of a typical model run, providing an indication of the input variables that can be adjusted by the user and the outputs than can be tracked. A history of all the internal variables can be maintained, if necessary. The animation of house-hunting bees has entertainment value only, and is disabled in the interest of execution speed during production runs. NetLogo also provides a means of automating multiple model runs ("BehaviorSpace"), allowing the user to program many replications of a given run, vary the input parameters, and log the results. Unless otherwise noted, all the results presented represent the average of 1000 replications of the model for a given set of parameters.

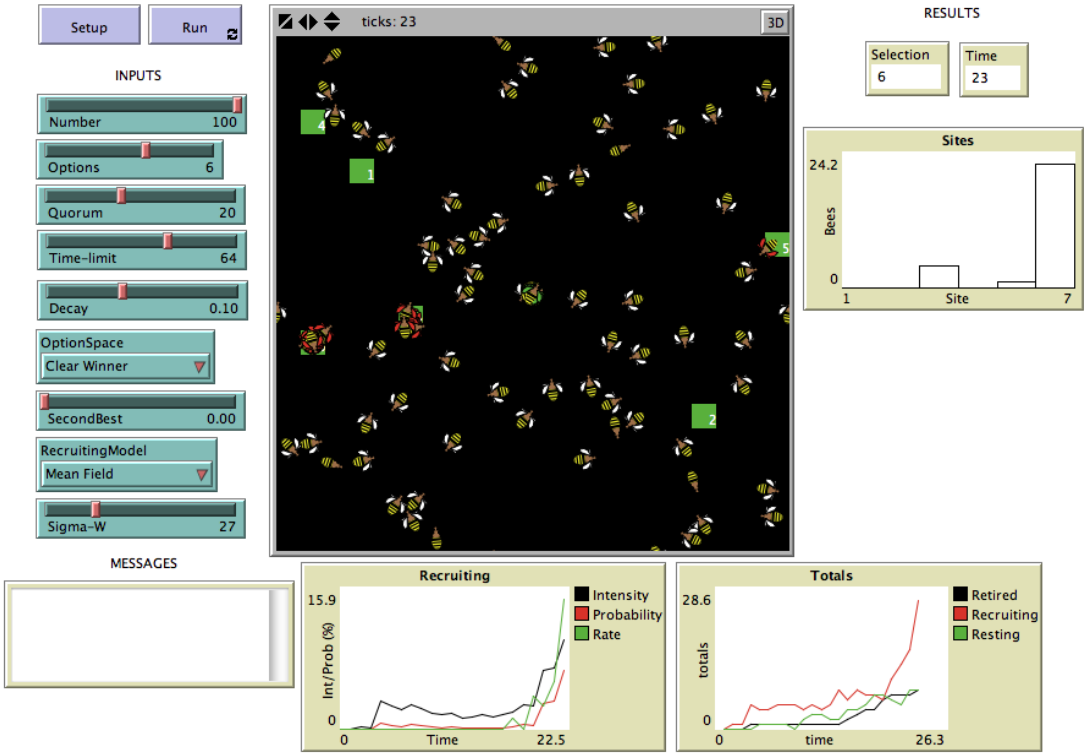

Figure 2. NetLogo simulation user interface

Space limitations prevent us from presenting all of the results obtained from this simulation, but the following sections describe some of the more interesting and thought-provoking findings. In presenting the results, we will make indiscriminate use of terms from the domains of modeling, insect behavior, and human decision making. Hence, the terms "agent," "scout," and "individual" will be used interchangeably, as will "time-step," "expedition," and "ballot."

Speed-Accuracy Trade-off

Once the option space, discovery probability p_d and recruiting mechanism have been selected, the main influence on speed and accuracy is the quorum fraction Q_0. Figure 3 shows how often each of the sites in option space S_1 is chosen as function of Q_0, over the range between 0.5 and 0.5.

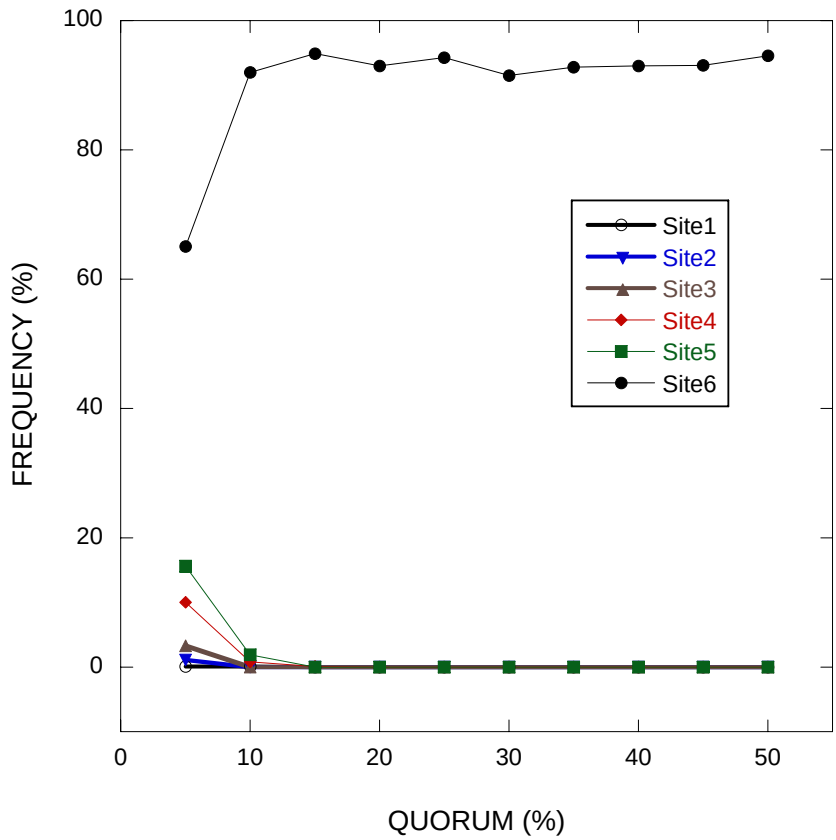

Figure 3. Site selection frequency as a function of quorum size for the "Clear Winner" option space S_1

It is striking that the best option is chosen about 95% of the time once the quorum is greater than 0.15. The quorum fraction required to achieve this result is somewhat higher for the uniform option space S_3 and higher still for the "Lake Woebegon" space S_2, but in no case does it exceed 0.25. This supports the claim by Passino and Seeley that the model explains the small quorum size observed in honey bees.

The second aspect of the speed-accuracy relationship, the time required to reach a decision, is shown in Figure 4. Since the decision time continues to rise with quorum size well beyond the point where the decision quality has reached its limit, it is reasonable to conclude that the best choice of quorum size is obtained when the decision quality has reached it maximum, somewhere in the range of 0.25 to 0.25.

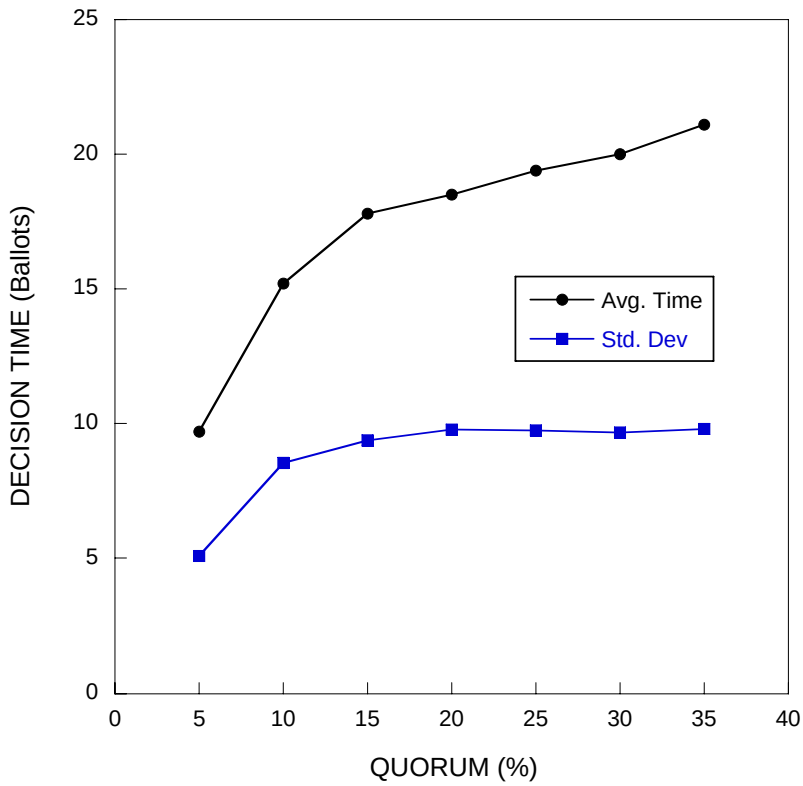

Figure 4. Decision time (in ballots or expeditions) as a function of quorum size for the "Clear Winner" option space S_1.

Similar results for the "Lake Woebegon" option space S_2 are shown in Figure 5. Here we have plotted average decision quality (defined as the quality of the option chosen averaged over all 1000 iterations of the model) against the average decision time. It is not always possible to do this effectively, for reasons that will be discussed in the next section. In this case we can clearly see that the average decision quality begins to saturate in the vicinity of $Q_0 = 0.25$. The result is similar to that for the S_1 option space, but the changes in decision quality are much smaller. This is due to the fact that the six options in S2 have an average quality of 0.75, so that even random selection would perform at that level.

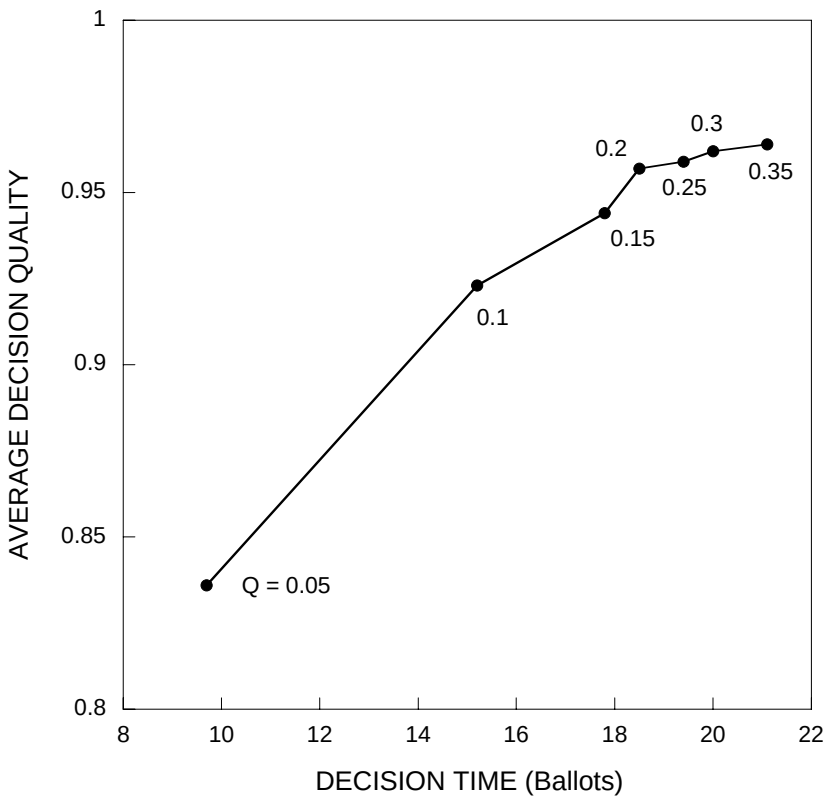

Figure 5. Average decision quality as a function of decision time for the "Lake Wobegon" option space S_2

Distribution of Decision Times

The average decision time used in Figures 4 and 5 was obtained as the mean value for 1000 replications of the model. In examining all four of the option spaces (Table 2) over a range of parameter values, we often found that the decision speed-accuracy relation was not as well behaved as in Figure 5, particularly when large changes in accuracy were observed. This was traced to the fact that the distribution of decision times is not well represented by its mean. Figure 6 shows the distribution of decision times obtained for the uniform option space S_3 with $Q_0 = 0.2$. The distribution is clearly skewed toward longer times. The mean is 21.4 ballots; the median is 19 ballots, and the mode is only 16.7 ballots. Work is in progress to determine where there exists a statistical description of this distribution that can lead to a more reliable measure of decision speed. Until such a measure is available, it is preferable to base studies of the speed-accuracy trade-off on results such as those in Figures 3 and 4.

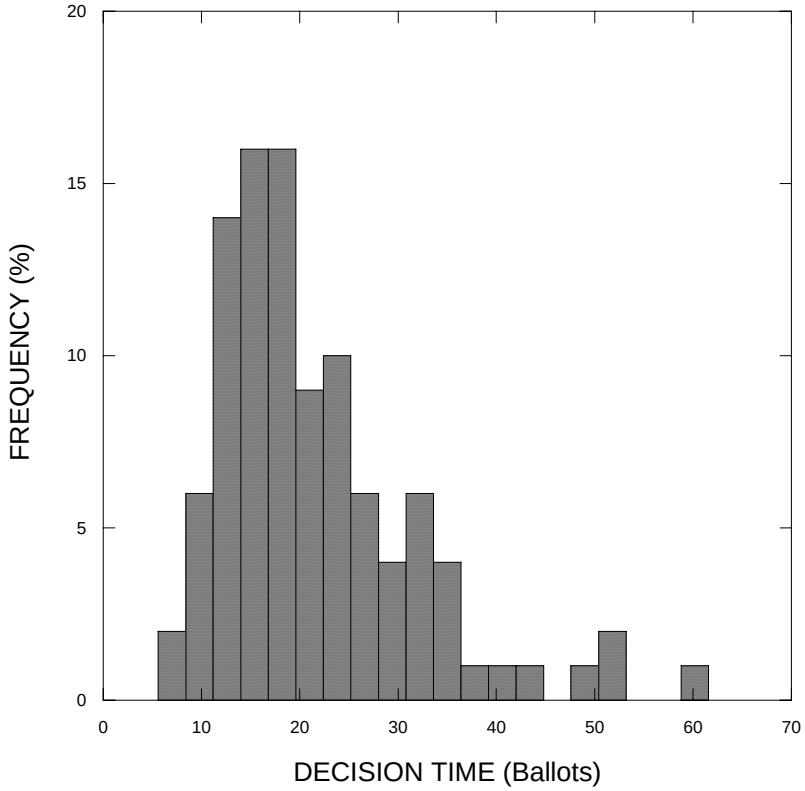

Figure 6. Distribution of decision times for the uniform option space S_3 with quorum fraction $Q_0 = 0.2$.

Assessment Accuracy

Work with the PS model has shown that results of the kind exhibited in Figures 3 through 5 are obtained even when the scouts exhibit an uncertainty in their assessments of site quality of ± 0.1. (Note that this represents 10% of the quality of the best site, not 10% of the quality of the site being assessed.) This is particularly striking in the case of the "Lake Woebegon" option space S_2, where the difference in quality between the best and second best site is only 0.1. As noted above, a study is under way using the binary option space S_4 to determine the limits beyond which assessment errors will begin to degrade the accuracy of decision making.

Discussion

The model described here makes no assumptions about cognition or the underlying mechanisms of social interaction. As a result, it can tell us nothing new about these subjects. On the other hand, the simple process shown in Figure 1, together with rules for some of the process steps based on the behavior of social insects, yields a description of collective decision making that accurately describes insect behavior and has implications for the study of human organizations and processes. Perhaps the most striking of these

implications is related to how small a quorum is sufficient to achieve accurate and timely decisions. Human task groups and command staffs often expend considerable effort to secure high levels of participation in each decision and to reach a "consensus" representing a majority or even a super-majority of the members. The present results suggest that the effort (and time) required to do this, while it may yield benefits in terms of individual satisfaction and group solidarity, may not lead to higher decision quality.

An aspect of the PS model that we have mentioned only indirectly is that few if any scouts visit more than one candidate nest site. Yet, accurate decisions are made without direct comparison between alternatives. Humans are often suspicious of opinions that are not based on careful examination of more than one option. Our work shows that if individuals can make reasonably accurate, or at least consistent, assessments of the quality of options, based on shared criteria, and can communicate those assessments accurately, then direct comparisons are not necessary.

The brief discussion of assessment accuracy in the previous section points to the fact that the process in Figure 1 appears to be remarkably tolerant of assessment errors. This is somewhat different from the "wisdom of crowds" (Surowiecki 2004) derived from the central-limit theorem, because the averaging of assessments that takes place is of a very limited and special nature. A scout that is recruited as the result of a mistakenly high assessment of the quality of a nest site may compensate for that error by making its own assessment and recruiting accordingly, but under-assessment leads to fewer recruits than are merited by the site quality, and the process provides no mechanism to compensate for this error.

Finally, the concept of "expiration of dissent" may have special relevance to collective problem solving in humans. When advocates for a particular course of action do not moderate their recruiting effort over time, they may risk either creating deadlock or driving the group to the premature selection of a low-quality solution. Strong, consistent support for one's opinions is often regarded as an aspect of good leadership, but willingness to rely on the assessments and recruiting of others may lead to improvement in the quality of collective decisions.

Questions for Further Study

A number of important issues not resolved in the present study will be subject of future work. Several of the simplifying assumptions made here must be relaxed in order to investigate the questions described below.

Exceptional Individuals

Both the PS model and the present one allow some variation in the behavior of individuals, but these are random deviations from average behaviors that are uniform across the population. Among humans (and perhaps honey bees) there are exceptional individuals who are likely to perform consistently better or worse than the general population. Some variations are worth exploring at each stage of the decision-making process.

During exploration, certain individuals may be more perceptive, that is, more likely to discover potential solutions, and others may be more accurate in assessing solution quality. When either type ceases to explore by becoming committed to a single option, the effectiveness of further exploration may be reduced.

During recruitment, unusually persuasive individuals are more likely to succeed at recruiting, while stubborn ones may persist longer than usual in recruiting for an option that fails to achieve a quorum. Both these types may have negative effects on decision quality by promoting the selection of poor-quality solutions discovered early in the process. Further work is needed to determine whether or not this is the case and, if it is, how tolerant the process is to the presence of such individuals.

During voting, powerful individuals may be able to command blocs of voters who were not formally recruited, or their votes may be given increased weight by the rules of the organization. The opportunity to recruit such individuals is clearly important and may in turn be regulated by the organizational or social structure.

Individuals may combine exceptionally good performance in some of these areas with exceptionally poor performance in others. The consequences of such a mixed population should be studied, since the process may be more tolerant of some "personalities" than others. Virtually nothing is known about how honey bee scouts are selected, other than that they are mature workers, so this study would have to rely on numerical experimentation, perhaps combined with data mining or factorial experiment design, to identify combinations of characteristics for detailed examination.

Attrition

As noted above, exploration is a relatively high-risk occupation for honey bees. The physical risk to humans in decision-making groups is probably less, but groups that deliberate for extended periods may experience significant attrition (with or without replacement), and the effects of this remain to be examined.

Multiple Decision Criteria

It is unusual for either insects or humans to have a single quality metric against which options can be evaluated. As noted above, honey bees seem to employ at least four factors in their assessments of nest sites. The well-established techniques of multiple criteria decision making (Zeleny 1982) are directed as much toward identifying and quantifying the criteria used by human decision makers as they are toward facilitating good decisions. A question to be addressed is how the process described in Figure 1 will respond to variations in the criteria or criteria weights employed by different individuals in their assessments.

Deception and Exploitation

It has been suggested that the evolution of dancing as a form of communication in honey bees was driven in part by the possibility of exploitation by predators or competitors of signals such as sound or odor which can be sensed over long distances or leave extended,

persistent traces (Nieh *et al*. 2004). This raises the possibility that the decision-making process might be compromised by the deliberate introduction of misinformation or by the presence of individuals exhibiting various forms of deviant behavior.

Adaptation, Failure, and Deadlock

It has been found experimentally that ants are able to vary the quorum required for nest site selection according to prevailing conditions (Franks *et al*. 2003). Specifically, under harsh conditions, the decision threshold is lowered to achieve greater speed at the possible expense of accuracy. No similar behavior has been found in honey bees, and the mechanism in ants is not understood quantitatively. Some such behavior may well occur in a group of humans faced with an approaching deadline. An added risk of reducing the decision threshold is that several options may achieve a quorum simultaneously, producing deadlock. Insect colonies have been known to divide under these circumstances, but this behavior may result in a selective disadvantage.

Conclusion

We have defined and modeled quantitatively a collective decision-making process based on field observations of nest site selection in social insects. This process achieves a very favorable balance of speed and accuracy and has characteristics of considerable interest in the study of human group decision making. Our principal findings stand in contrast to much of the "conventional wisdom": accurate decisions can be made without majority support, without direct comparison of alternatives, and without lasting commitment by advocates for the chosen alternative.

References

Franks, Nigel R., Stephen C. Pratt, Eamonn B. Mallon, Nicholas F. Britton, and David J. T. Sumpter. 2002. Information flow, opinion polling and collective intelligence in house-hunting social insects. Philosophical transactions of the Royal Society of London Series B, Biological Sciences 357: 1567-1583.

Franks, Nigel R., Anna Dornhaus, Jon P. Fitzsimmons, and Martin Stevens. 2003. Speed versus accuracy in collective decision making. Proceedings Royal Society of London: Biological Sciences 270: 2457-2463.

Nieh, James C., Lillian S. Barreto, Felipe A. L. Contrera, and Vera L. Imperatriz-Fonseca. 2004. Olfactory eavesdropping by a competitively foraging stingless bee, Trigona spinipes. Proceedings Royal Society of London: Biological Sciences 271: 1633-1640.

Passino, Kevin M and Thomas D. Seeley. 2006. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. Behavioral Ecology and Sociobiology 59: 427-442.

Seeley, Thomas D. 2002. When Is Self-Organization Used in Biological Systems? Biological Bulletin 202: 314-8.

Seeley, Thomas D. 2003. Consensus building during nest-site selection in honey bee swarms: the expiration of dissent. Behavioral Ecology and Sociobiology 53: 417-424

Seeley, Thomas D. and P. Kirk Visscher. 2003. Choosing a home: how the scouts in a honey bee swarm perceive the completion of their group decision making. Behavioral Ecology and Sociobiology 54: 511-520

Surowiecki, James. 2004. The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations. New York: Doubleday.

Visscher, P. Kirk. 2003. How self-organization evolves. Nature 421: 799-800.

Wilensky, Uri. 1999. NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Zeleny, Milan. 1982. *Multiple Criteria Decision Making*. New York: McGraw-Hill.