Artificial Swarms find Social Optima

(Late Breaking Report)

Louis Rosenberg
Unanimous AI
San Francisco, CA, USA

Gregg Willcox
Unanimous AI
San Francisco, CA, USA

***Abstract*— in the natural world, many social species amplify their collective intelligence by forming real-time closed-loop systems. Referred to as Swarm Intelligence (SI), this phenomenon has been rigorously studied in schools of fish, flocks of birds, and swarms of bees. In recent years, technology has enabled human groups to form real-time closed-loop systems modeled after natural swarms and moderated by AI algorithms. Referred to as Artificial Swarm Intelligence (ASI), these methods have been shown to enable human groups to reach optimized decisions. The present research explores this further, testing if ASI enables groups with conflicting views to converge on socially optimal solutions. Results showed that "swarming" was significantly more effective at enabling groups to converge on the Social Optima than three common voting methods: (i) Plurality voting (i) Borda Count and (iii) Condorcet pairwise voting. While traditional voting methods converged on socially optimal solutions with 60% success across a test set of 100 questions, the ASI system converged on socially optimal solutions with 82% success (p<0.001).**

Keywords— Swarm Intelligence, Artificial Swarm Intelligence, Collective Intelligence, Human Swarms, Artificial Intelligence.

I. INTRODUCTION

Artificial Swarm Intelligence (ASI) connects groups of networked individuals into real-time closed-loop systems modeled after biological swarms. Using swarming intelligence algorithms and dynamic feedback loops, ASI enables distributed human populations to answer questions, make predictions, express opinions, and reach decisions by thinking together as a unified emergent intelligence (i.e. as a "hive mind"). Prior research studies have shown that "human swarms" can generate significantly more accurate solutions than traditional methods of harnessing group intelligence [1-4]. In one recent study, researchers at Unanimous AI and Oxford University tasked swarms of financial traders with predicting four common market indices (SPX, GLD, GDX, and CRUDE). Across three months of weekly trials, results showed a 26% increase in forecast accuracy for swarm-based predictions (p=0.001) [5].

While prior studies show that Artificial Swarm Intelligence can amplify the predictive accuracy of human populations, an important question still remains: *Why do swarms outperform traditional aggregation methods?* Because swarms are real-time closed-loop systems, it has been hypothesized that they enable populations to converge upon solutions that combine conflicting views in optimal ways. To explore this hypothesis, the current study compares real-time human swarms with traditional voting methods for aggregating conflicting views.

Specifically, the present study compares human swarming to three common voting methods: (i) Plurality voting, (ii) ranked Borda Count voting, and (iii) Condorcet pairwise voting. The study uses conflicting cash payouts as the motivator to ensure that human groups have highly conflicting views when trying to reach a decision that best represents their collective interests. The goal is to assess which method finds the socially optimal solution at the highest success rate.

"SWARMS vs CROWDS"

In traditional crowd-based instruments like votes and polls, *respondents* are simply that – a source of *responses* which are captured as isolated data points and combined statistically with data from other isolated respondents. While such methods are often referred to as crowd-sourcing, the "crowd" is essentially a statistical metaphor for data aggregation. Even prediction markets, which are more interactive than traditional polls and surveys, do not enable full populations to converge as real-time systems. That's because each market transaction is conducted between a single "buyer" and a single "seller," which are then executed in sequence to engage a full population. The serial nature of markets is very different from the parallel structure of swarms. In swarms, all members of a population interact at the same time, converging in synchrony on unified decisions.

In swarm-based systems, the full population of participants are connected by real-time feedback loops and governed by intelligence algorithms. Often referred to as "hive minds," a prime objective of these systems is to enable diverse populations to converge on solutions that best represent the conflicting opinions, interests, and/or intelligence of the group as a whole. The swarming process is generally modeled on biological systems such as schools of fish, flocks of birds, and swarms of bees. The present research effort employs *Swarm AI* technology from Unanimous AI Inc, which is inspired largely by honeybee swarms. This particular biological model was chosen because honeybee swarms have been shown to greatly amplify the collective intelligence of bee colonies, enabling their population to converge on optimal solutions to complex problems. [6]

The decision-making processes that govern the behavior of honeybee swarms have been studied since the 1950s and have been shown to be remarkably similar to the decision-making processes in neurological brains [6,7]. Both employ large networks of simple excitable units (i.e., bees and neurons) that work in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in synchrony. In both, outcomes are arrived at through a real-time competition among sub-populations of excitable units. When a sub-population exceeds a threshold level of support, the corresponding solution is chosen. In honeybees, this enables complex multi-variable problems to be solved as a "hive mind," converging on optimal solutions. For example, honeybee colonies have been shown to use swarm intelligence to select the best possible homesite from a set of available homesites with high precision [8,9,10].

The similarity between “brains” and “swarms” becomes even more apparent when comparing decision-making models that represent each. For example, the decision process in primate brains is often modeled as mutually inhibitory leaky integrators that aggregate incoming evidence from competing neural populations [11]. A common framework for primate decision is the Usher-McClelland model in Figure 1 below.

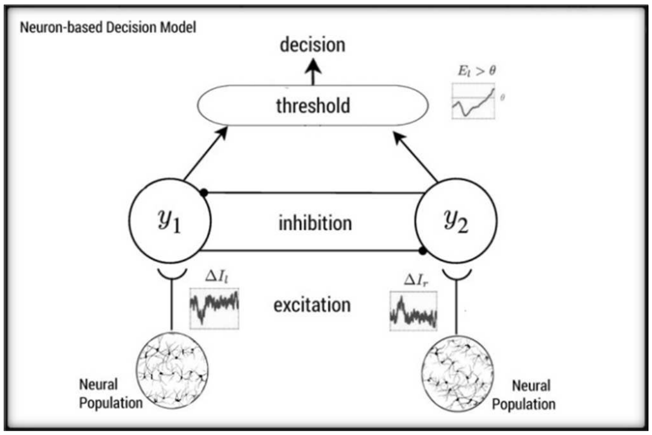

Fig. 1. Usher-McClelland model of neurological decision-making

This neurological decision model can be directly compared to swarm-based decision models, for example the honeybee model represented in Figure 2 below. As shown, swarm-based decisions follow a similar process, aggregating input from sub-populations of swarm members through mutual excitation and inhibition, until a threshold is exceeded.

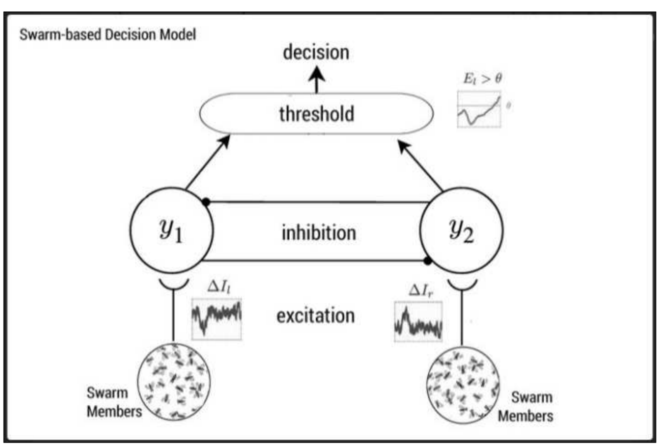

Fig. 2. Mutually inhibitory decision-making model in bee swarms

In this context, brains can be viewed as *systems of neurons* structured so intelligence emerges, while swarms can be seen as *systems of brains* structured so amplified intelligence emerges. In other words, a swarm can be thought of as a “brain of brains.” The question thus remains, how can we most effectively build swarms of networked human participants connected over the internet? One approach for building tightly integrated systems of networked human groups is described as follows:

II. ENABLING “HUMAN SWARMS”

Unlike many other social species, humans have not evolved the natural ability to form real-time swarms. That’s because we lack the subtle connections that other organisms use to establish feedback loops across members of a population. Schooling fish detect vibrations in the water around them. Flocking birds detect subtle motions propagating through the group. Swarming bees use complex body vibrations called a “waggle dance” [12]. To enable real-time swarming among populations of networked humans, we must replicate these natural processes.

To address this need, a cloud-based software system called *swarm.ai* was developed that enables distributed human groups to login from anywhere in the world and form closed-loop systems [1,2]. Modeled after the decision-making process of honeybee swarms, this online platform allows groups of distributed users to work in parallel to (a) integrate noisy evidence, (b) weigh competing alternatives, and (c) converge on decisions in synchrony, while also allowing participants to perceive and react to the changing system in real-time, thereby closing a feedback loop around the full population.

A screenshot from the *swarm.ai* platform is shown below in Figure 3. The image depicts a “human swarm” comprised of approximately 100 networked users, simultaneously connected from diverse locations around the United States. As shown, the swarm is in the middle of answering a political question. A decision like this generally takes between 10 and 60 seconds, during which time the population works to move a graphical puck to select among a set of available answer options.

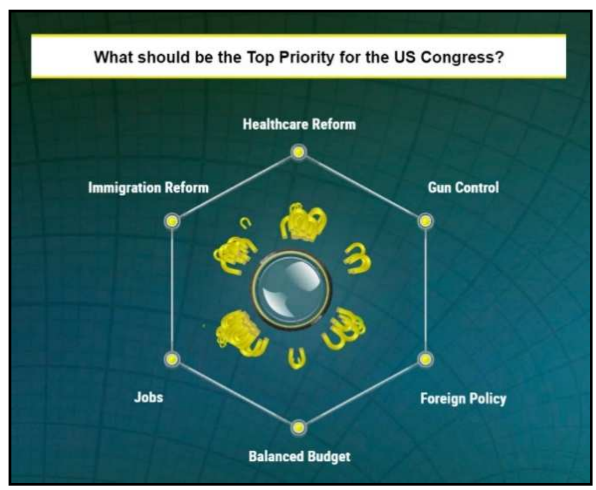

Fig. 3. A human swarm answering a question in real-time

Each participant provides input by manipulating a graphical magnet with a mouse or touchscreen. By positioning their magnet with respect to the movable puck, participants impart their personal intent. The input from each user is not a discrete vote, but a stream of directional vectors that varies freely over time. Because every member can adjust their intent continuously in real-time, the overall swarm moves, not based on the input of any individual, but based on the dynamics of the full system. This enables the group to collectively explore the decision-space and converge on the most agreeable solution in synchrony.

It is important to note that participants do not only vary the direction of their intent, but also modulate the magnitude by adjusting the distance between their magnet and the puck. Because the puck is in continuous motion across the decision-space, users need to continually move their magnet so that it stays close to the puck's outer rim. This is significant, for it requires participants to be engaged continuously throughout the decision process, evaluating and re-evaluating their intent as they convey their contribution. If they stop adjusting their magnet with respect to the changing position of the puck, the distance grows and their applied sentiment wanes.

Thus, like bees vibrating their bodies to express sentiment in a biological swarm, or neurons firing activation signals to express conviction levels within a biological neural-network, the participants in an artificial swarm must continuously update and express their changing preferences during the decision process, or lose their influence over the collective outcome. This is similar to the leaky-integrator process found in honeybee swarms. In addition, real-time intelligence algorithms monitor the behaviors of all members during the process, predicting their conviction levels based upon their actions over time. This process reveals a range of behavioral characteristics within the population and weights contributions accordingly.

III. Social Optima Study

Social Choice Theory has its roots in the 18th century, when Nicolas de Condorcet observed that voting methods based on majority rule can lead to irrational outcomes, even when the preferences of the individual voters are completely rational [13]. Over the centuries since, researchers have explored a wide range of methods for aggregating the preferences of populations, often with the goal of reaching optimized decisions [14,15]. Popular methods include ranked voting (e.g. Borda count) and pairwise voting (e.g. the Condorcet method). While used in countless applications from electing political candidates to determining business priorities, these methods fail to find socially optimal solutions under many conditions [16,17].

The present study compares human swarms to traditional voting methods, testing the ability of each to reach socially optimal solutions. To conduct this comparison, a simple decision task was designed using financial incentives (i.e. cash payouts) to dictate preference levels within human groups. Specifically, participants were asked to choose among three fruit types, each of which was assigned a different cash payout that would be earned if that fruit was selected by the group. Of course, if everyone in the group had the same payout structure, the decision would be easy to arrive at – the group would simply choose the fruit with the highest payout every time. The goal of this study, however, was to explore decision tasks where the population has conflicting interests.

To address this requirement, the population of participants in each decision task were assigned to two distinct sub-groups, each of which was given a different payout structure. The payouts were designed to ensure that any decision reached would involve conflicting interests. A set of 100 decision tasks were defined in this way, crafted to enable rigorous testing of how often each aggregation method enabled the group to converge on the solution that maximized their collective return (i.e. how often each method found the Social Optima).

A population of 170 randomly selected human subjects were employed for this study. The participants were divided into seven test groups. Each test group was then split into two sub-populations. Each sub-population was assigned a unique set of financial incentives (i.e. payouts). The incentives were crafted to ensure that the two sub-populations would have conflicting interests when asked to select one fruit type from a set of three fruit types. A set of 100 test questions were defined in this way and evaluated using four distinct methods – three common voting methods (Plurality, Borda, and Condorcet) as well as by real-time ASI (i.e. human swarm). The voting methods were assessed using the commonly accepted social-choice function associated with each. These are described as follows:

Plurality Voting:

- Each voter casts a single vote for their preferred option from the set of three available options (A, B, C).
- The option with the most votes is selected.

Borda Count Voting:

- Each voter submits a ranked ordering of the three available options (A, B, C).
- Each option is assigned a score based on the ranked ordering provided by the population of participants.
- The option with the highest score is selected.

Condorcet Pairwise Voting:

- Each voter considers pairs of options in head-to-head matchups, picking one for each (A-v-B, B-v-C, A-v-C).
- Options are scored by the number of pairwise victories minus the number of pairwise defeats.
- The option with the highest score is selected.

In each test question of this study, the participants were tasked with choosing a preferred fruit from a set of three fruit options, each option associated with a real cash payout. The payouts were structured so that participants had conflicting interests (i.e. fruits associated with high payouts for some participants were associated with low payouts for others). Each participant was provided a payout table that indicated how much they would personally earn if the overall group selected each of the fruits in the available set. Individuals had no knowledge of the payout tables given to other participants and thus could only be motivated by their own personal interests. Payouts ranged from \$0.00 to \$0.25, in 5 cent increments. For swarm-based decisions, if the group failed to reach any answer in 60 seconds, the question was repeated once. If the swarm failed a second time in 60 seconds, no payout was given.

An example decision task is shown graphically in Figure 4 below. In this task, all participants in the group were asked to choose a preferred fruit from a set of three options (Orange, Apple, Grape). The group was split into two sub-populations, each with different payout tables to ensure conflicting interests. The members of Population A had payouts that preferred Apple most (\$0.20), Orange next (\$0.05), and Grape least (\$0.00). At the same time, the members of Population B had interests which preferred Orange most (\$0.15), Grape next (\$0.10), and Apple least (\$0.05). Each participant only knew their own payouts and had no knowledge about the payouts of others.

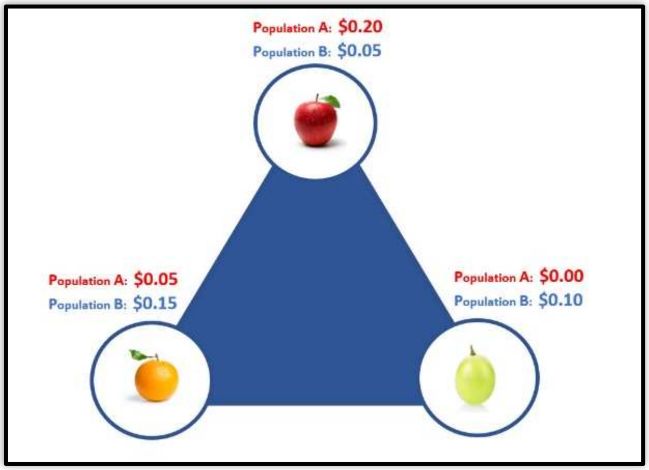

Fig 4. Example Payout Structure in a Fruit Selection Task

Using the three-fruit structure above, but varying the payout tables, a set of 100 test questions were defined. The payout sets were varied to ensure a wide range of conflicting interests and selected so that every question had an optimal solution – i.e. a selection that maximized the payout across the full population. The 100-question set also varied the relative size of the payouts. For example, in some payout structures, the social optima for the group was far from the maximum payout for some participants, making it a more difficult solution.

IV. RESULTS

As a benchmark, the set of 100 questions were assessed using three voting methods described above: (i) Plurality voting, (ii) ranked Borda Count, and (iii) Condorcet pairwise voting. As participants would select their preferred options based entirely on the payout structure provided to them, the success rate in finding the Social Optima is highly predictable. For these three benchmark methods, the success rate is shown below in Table 1.

Method	Success	Failure	Tie
Plurality Vote	35%	10%	55%
Condorcet Method	49%	27%	24%
Borda Count	52%	33%	15%

Table 1: Benchmark Performance via Voting Methods on Question Set

As shown, each voting method has three potential outcomes, either, correctly finding the Social Optima, selecting an option that was not the Social Optima, or producing results that were a tie among two or more options. In the instance when a voting method produced a tie among two options, one of which was the Social Optima, it was assumed that the voting method would, by random statistical chance, achieve the correct solution in 50% of those cases. That produces benchmark success rates as follows:

Method	Success
Plurality Vote	62.5%
Condorcet Method	60.3%
Borda Count	58.3%

Table 2: Benchmark Performance via Voting Methods on Question Set

The above benchmarks were then compared against the experimental results generated by participants working together as an ASI systems. As described above, each of the seven test groups were tasked with selecting a preferred fruit from sets of three fruit options, with participants assigned conflicting payout structures. Across the seven groups, and 100 test questions, the swarm-based system correctly identified the socially optimal solution 82% of the time. Figure 5 below shows this result as compared to the three benchmark voting methods. This is a significant improvement, equivalent to making 52% fewer sub-optimal decisions than the three common voting methods. The probability that the voting methods scored lower than the human swarms by chance was extremely low ($p=4.4x10^{-6}$), indicating a highly significant result.

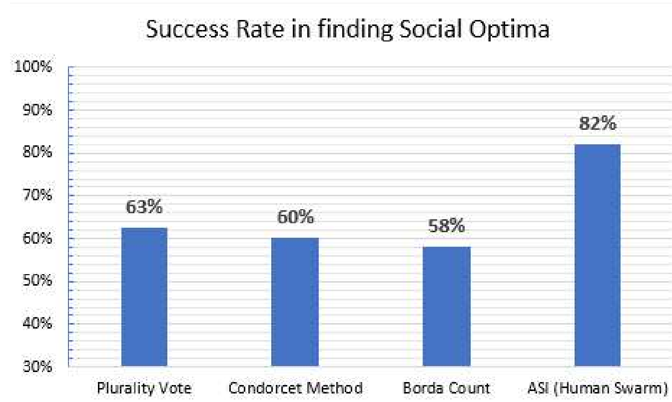

Fig 5 Optimality Comparison across Decision Methods

Framing these results in terms of payout, the differences between optimal payout and chosen payout were computed for each method. As shown in Figure 6, the swarm-based system selected payouts that were significantly closer to optimal than the three voting methods. Using Plurality voting, the group could expect to receive, on average, 11.54% below optimal payout across the question set. Using Condorcet and Borda Count voting, the group could expect, 10.31% and 10.63% below optimal payout respectively. When working as a swarm, the group averaged only 4.72% below optimal payout across the question set. This translates to 54% less money lost due to poor decision-making as compared to the best voting method tested. This confirms that when working as a swarm-based system, the population made reliably better decisions. The probability that voting gave a lower average payout by chance, as compared to ASI, was very low ($p=5.5x10^{-4}$), indicating a significant result.

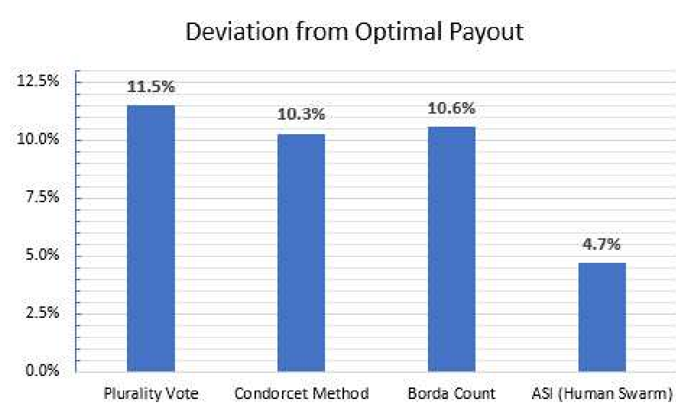

Fig 6. Deviation from Optimal Payout across Decision Methods

V. CONCLUSIONS

Does the process of Artificial Swarm Intelligence enable human groups to reach optimal decisions at higher rates than traditional methods for aggregating group input? The results of this study suggest that forming real-time ASI systems among human groups can be significantly more effective at finding socially optimal solutions than using common voting methods, including Plurality voting, ranked Borda Count voting, and Condorcet pairwise voting. While traditional voting showed success rates of approximately 60% across the 100-question test set, the ASI system converged on optimal solutions 82% of the time. In addition, human swarms achieved average payouts that were 54% closer to the optimal than traditional methods.

This is a significant result and suggests that human swarming may be an effective path, not only for amplifying the intelligence of human populations, but also for enabling human groups with conflicting views to find solutions that maximize their collective interests and achieve higher levels of overall satisfaction. These results may also explain why human swarms have been found to be significantly more successful than traditional polls and surveys at combining the knowledge, wisdom, and opinions of diverse populations. Swarming appears to aggregate the input of human groups in a more efficient manner, finding the optimal solutions at significantly higher rates.

ACKNOWLEDGMENT

Thanks to Jules Sulpico and Jeffrey Do for their help in this research effort. Also, thanks to Unanimous AI for the use of the Swarm.ai platform for this ongoing work.

REFERENCES

[1] Rosenberg, L.B., “Human Swarms, a real-time method for collective intelligence.” Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[2] Rosenberg, Louis. “Artificial Swarm Intelligence vs Human Experts,” Neural Networks (IJCNN), 2016 International Joint Conference on. IEEE. J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.

[3] Rosenberg, Louis. Baltaxe, David and Pescetelli, Nicollo. "Crowds vs Swarms, a Comparison of Intelligence," IEEE 2016 Swarm/Human Blended Intelligence (SHBI), Cleveland, OH, 2016, pp. 1-4.

[4] Baltaxe, David, Rosenberg, Louis and N. Pescetelli, “Amplifying Prediction Accuracy using Human Swarms”, Collective Intelligence 2017. New York, NY ; 2017.

[5] Rosenberg, Louis, Pescetelli, Niccolo, and Willcox, Gregg. “Human Swarms Amplify Accuracy in Financial Predictions,” Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON), IEEE Annual, New York, NY, 2017.

[6] Seeley T.D, Buhrman S.C 2001 “Nest-site selection in honeybees: how well do swarms implement the ‘best-of-N’ decision rule?” Behav. Ecol. Sociobiol. 49, 416–427

[7] Marshall, James. Bogacz, Rafal. Dornhaus, Anna. Planqué, Robert. Kovacs, Tim. Franks, Nigel. “On optimal decision-making in brains and social insect colonies.” Soc. Interface 2009.

[8] Seeley, Thomas D., et al. "Stop signals provide cross inhibition in collective decision-making by honeybee swarms." Science 335.6064 (2012): 108-111.

[9] Seeley, Thomas D. Honeybee Democracy. Princeton Univ. Press, 2010.

[10] Seeley, Thomas D., Visscher, P. Kirk. “Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making.” Behavioural Ecology and Sociobiology 54 (5) 511-520.

[11] Usher, M. McClelland J.L 2001 “The time course of perceptual choice: the leaky, competing accumulator model.” Psychol. Rev. 108, 550–592

[12] Couzin, Iain. “Collective Minds.” Nature 445 (15) 2007. 715-716.

[13] Condorcet, Nicolas de, 1785, Essay sur l'Application de l'Analyse à la Probabilité des Décisions Rendue à la Pluralité des Voix. Paris.

[14] Boutilier, Craig. “Optimal Social Choice Functions: A Utilitarian View.” Artificial Intelligence, Elsevier, 25 June 2015.

[15] Anshelevich, Elliot and Postl, John. 2017. Randomized social choice functions under metric preferences. J. Artif. Int. Res. 58, 1 (January 2017), 797-827.

[16] Arrow, Kenneth (1951), Social Choice and Individual Values, Yale University Press.

[17] Edwards, W. (1954). The theory of decision making. Psychological Bulletin, 51(4), 380-417.