

Polya's bees: A model of decentralized decision-making

Russell Golman,¹ David Hagmann,^{1*} John H. Miller^{1,2}

2015 © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC). 10.1126/sciadv.1500253

How do social systems make decisions with no single individual in control? We observe that a variety of natural systems, including colonies of ants and bees and perhaps even neurons in the human brain, make decentralized decisions using common processes involving information search with positive feedback and consensus choice through quorum sensing. We model this process with an urn scheme that runs until hitting a threshold, and we characterize an inherent tradeoff between the speed and the accuracy of a decision. The proposed common mechanism provides a robust and effective means by which a decentralized system can navigate the speed-accuracy tradeoff and make reasonably good, quick decisions in a variety of environments. Additionally, consensus choice exhibits systemic risk aversion even while individuals are idiosyncratically risk-neutral. This too is adaptive. The model illustrates how natural systems make decentralized decisions, illuminating a mechanism that engineers of social and artificial systems could imitate.

INTRODUCTION

Ants and bees appear to rely on decentralized decision-making for critical choices. For example, in choosing a new nest site—a decision that has huge implications for the survival of the group—decisions must be made without central control and with no single individual evaluating the total available information (1) or any one individual making direct comparisons of the available options (2–5). Although individual agents follow simple rules that allow them to uncover very limited and local information, the colony as a whole must efficiently integrate the resulting flow of information into a high-quality, final decision (6–10).

Consider how a swarm of honey bees, *Apis mellifera*, chooses a new hive location (8–14). When a swarm abandons the old hive, it temporarily gathers at a tenuous location. About 5% of the bees are scouts, and after exploring the surrounding area for possible hive sites, a scout may perform a waggle dance that indicates the location of the site it discovered (15, 16). The likelihood of performing a waggle dance, and its duration, depends on the quality of the site that was investigated (9, 13, 14). The waggle dances serve to recruit additional scouts to further investigate the “advertised” site. The longer the dance, the more likely that the new scouts will investigate the site and bring back independent evaluations. Over time, positive feedback loops are generated (8, 13), and once the number of scout bees at a particular site reaches a quorum threshold (of about 30 to 40 bees), those scouts return to the swarm and lead it to the new site (12).

Leptothorax (Temnothorax) albipennis ants choosing a new nest site behave similarly (1, 6, 17). When a scout finds a higher-quality site, it quickly returns to the old nest site and recruits a nestmate by tandem running, a tedious process that entails the scout teaching the recruit the route to the new site (18–21). The speed of recruitment is tied to the quality of the site (1, 3), with better sites inducing quicker responses. As before, positive feedback arises when recruits become recruiters. Finally, when the number of ants at the new site reaches a quorum threshold, the recruiting ants switch from tandem running to the much faster process of carrying their remaining nestmates from the old to the new site (22).

Other social organisms make collective decisions with mechanisms reminiscent of those of ants and bees. Social spiders coordinate their emigrations to a new nest (23) with silk draglines, allowing positive reinforcement of existing routes much like the pheromone trails of ants (23–25). Cockroaches are more likely to remain in shelters when other cockroaches are nearby, leading to a collective choice of a single home (26, 27). Even bacteria share information and detect quorums, allowing for collective decisions regarding sporulation, virulence, and gene exchange (28).

There is some speculation that primate brains use a similar decentralized decision mechanism. No single neuron is solely responsible for the brain's decision. In a visual discrimination task, for example, a subset of specialized neurons integrates sensory signals from other neurons and allows the brain to make a decision to trigger other neurons to initiate a motor response (29). Complex, decentralized information processing can be achieved with a cell assembly, a recurrent circuit of neurons that becomes active when stimulation spreads with positive feedback (30, 31). Neurons are generally understood to accumulate information and fire when the stimulus hits a threshold to implement a decision (32–34). In this regard, the primate brain may function analogously to a colony of social insects (35–37).

Similar mechanisms may even be at work in large-scale social processes (38). For example, consider the choice of a personal MP3 player. Consumers who purchase such players “advertise” them when they use them (particularly if the players have some distinctive feature, for example, white ear buds). Moreover, consumers who enjoy their players are more likely to use them. Someone new to the MP3 market may observe the players that others use and purchase on the basis of these observations. This kind of direct marketing is often a major driver of consumer demand for new products, especially when competing brands have not yet established distinct reputations (39). At some point in the evolution of the market, however, a critical mass of consumers may choose the same product and fundamentally change the market dynamics (say, by adopting a particular technological innovation, by fueling economies of scale in the production process, or by enticing suppliers or producers of complements to enter into exclusive agreements) so that only the leading product can survive in the market (40, 41). Similarly, people collaborating to make a group decision also tend to share information that favors options that already have popular support

¹Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA. ²Santa Fe Institute, Santa Fe, NM 87501, USA.

*Corresponding author. E-mail: hagmann@cmu.edu

while hesitating to share information that favors unpopular options (42), and they commonly reach a consensus choice through plurality voting (43).

Any decision mechanism must trade off exploring new options versus exploiting the best option known to date. Although further exploration helps identify superior options, it comes at the cost of not acting on a known option. Thus, too much exploration may lead to indecisiveness and thus harm fitness, whereas too little may imply the acceptance of suboptimal choices. Given the inherent tradeoff between speed and accuracy, the diffusion model of decision-making (34)—which does not incorporate positive feedback in the accumulation of evidence—is efficient for binary choice, but it requires the agents to compare the options or at least to inhibit activity for competing options (35, 44). When this is not possible, positive feedback in the process of exploration may prove useful for making reasonably good, quick decentralized decisions.

The model

We present a model of a two-part process for decentralized decision-making, involving search and recruitment with positive feedback in the first phase and quorum detection to trigger a consensus choice (without centralized processing) in the second phase. First, individual agents randomly search over a set of feasible options, biased by the quality of the options revealed during previous searches. Then, a final consensus choice is triggered when a quorum of agents investigating any one particular option forms. [A consensus choice is a choice that the entire system must abide (45), but of course, it is not necessarily a unanimous choice via “consensus sensing” (10).] We capture this process with an urn scheme that runs until hitting a threshold.

Search and recruitment

Agents (scouts) explore one of many possible options and then return home to recruit additional agents to further explore that option. A Poisson process for each agent governs its trips home, and the recruitment of additional scouts on each trip depends on the quality of the explored option.

We assume that there are C possible options and that some number of agents w_c^t investigate each $c \in C$ at any time t . We refer to w_c^t as the weight on c at time t . All weights are initially set to the same, positive value w^0 and then accumulate over time. Each agent is equally likely to return home at any time, and when it does, it recruits additional agents to join it and further explore the same option. The chance of an agent recruiting for option c at time t is simply proportional to the weight w_c^t . Each option has a set of immutable attributes that defines its quality, and the extent of recruitment for c depends on its quality. The agents investigating c will recruit ν_c additional agents to continue exploring it when they return. We think of the number of recruits per return trip home, ν_c , as an ordinal measure of the quality of c .

Quorum detection

The search process above generates a distribution of agents investigating each possible option at any given time. Given the decentralized nature of these systems, there must be some feasible trigger that ends the search process and finalizes the consensus choice. One possible solution to this problem would be to have the search probabilities converge to zero or one—that is, have all of the probability concentrated on a single option. Forcing such a unanimous decision on the system is problematic, because it may form extremely slowly, perhaps leading

to a serious loss of fitness. Moreover, we have empirical evidence, at least in the case of honey bees (46, 47), ants (6), and stickleback fish (48, 49), that unanimity is not what triggers a consensus choice. Instead, a final choice is made once the number of agents in favor of a particular option reaches a quorum (45, 50).

On the basis of the above arguments, we incorporate into our model a *quorum threshold*, τ , that triggers—as the final decision—any option that is being investigated by at least that number of agents. This threshold is effectively the finish line in the race for each option to accumulate weight. The decision is determined by the first passage of $w_c^t \geq \tau$. The level of the quorum threshold has important implications for the decisions that arise in the system. If the threshold is set too high, then a quorum may not be reached for a long time, resulting in prolonged inaction. If the threshold is set too low, then a quorum might be achieved for a relatively low-quality option. Thus, the optimal quorum threshold depends on a tradeoff between speed and accuracy in the decision-making process. From a normative standpoint, a good threshold allows the system to withstand various transients in the probability distribution while still remaining responsive to the acquired information in a timely manner.

The urn scheme

We use a simple (Polya) urn process to model this decision mechanism. This process is easy to visualize. Assign to each of the C options a unique color, and place w^0 balls of each color into an urn. The number of balls of a particular color in the urn corresponds to the number of agents investigating the associated option. Each ball has the same rate at which it may be randomly drawn from the urn. When a ball with color c is drawn, it is immediately placed back into the urn along with ν_c identically colored balls. This process continues until a threshold number of balls τ is reached. We assume that the supply of balls available to enter the urn is large relative to the quorum threshold, because the number of scouts in a swarm of bees or a colony of ants is typically much larger than the number required to achieve a quorum.

Analytical results

Our analysis aims to characterize the behavior of this decision process, that is, to determine the choice probabilities and the (distribution of) decision times. First, we must understand the accumulation process. We precisely characterize how the composition of balls in the urn evolves over time. Lemma 3.1 of (51) gives us the distribution of balls of each color at any time t (in the absence of a threshold for stopping the process).

Lemma 1. The moment-generating function $\phi_c(t, s) = \mathbf{E}[e^{s w_c^t}]$ is given by $\phi_c(t, s) = \left(\frac{e^{\nu_c(s-t)}}{e^{\nu_c(s-t)} - e^{\nu_c s} + 1} \right)^{\frac{w^0}{\nu_c}}$ (where s is the argument of the moment-generating function).

Proof. We can describe this Polya process with a diagonal $C \times C$ matrix with the ν_c values along the diagonal and 0's elsewhere. The evolution of the number of balls of a given color is independent of the evolution of other colors (until the threshold is hit). Thus, Lemma 3.1 of (51) directly applies.

In principle, this moment-generating function fully characterizes the distribution of weights, w_c^t , where s is the argument of the moment-generating function. In practice, however, calculating the likelihood of hitting a threshold τ at a given time t is complicated.

An asymptotic result is simple to obtain. Suppose the threshold τ is infinite so that the Polya process can run forever. Eventually, almost all of the weight converges on the choice with the highest quality.

Theorem 1. If there is a unique optimal choice $c^* = \operatorname{argmax}_c v_c$, then:

$$\lim_{t \rightarrow \infty} \frac{w_c^t}{\sum_j w_j^t} = \begin{cases} 1 & \text{if } c = c^* \\ 0 & \text{otherwise.} \end{cases}$$

Proof. As $t \rightarrow \infty$, $\frac{w_c^t}{e^{v_c t}} \xrightarrow{D} \text{Gamma} \left(\frac{w^0}{v_c}, v_c \right)$ [see Theorem 3.1 of (51)].

In the infinite time limit, we recover the unanimity decision rule, and the probability of a mistake (that is, selecting an option other than c^*) vanishes.

Although the asymptotic properties of the urn process are informative, feasible decentralized systems must make decisions in finite time and require a finite threshold. For any given threshold τ , we would like to describe the probability $p_c(\tau)$ of selecting each possible choice c as well as the waiting time $T(\tau)$ until the decision is made.

We can characterize the waiting time $T_c(\tau)$ until the number of agents exploring option c would hit the threshold τ (independent of recruitment for other options). Let λ denote the intensity of the Poisson process for each agent's return home.

Lemma 2. The waiting time $T_c(\tau)$ has the Hypoexponential $(\lambda_0, \lambda_1, \dots, \lambda_n)$ distribution with:

$$\lambda_i = (w^0 + iv_c)\lambda \text{ for all } i, \tag{1}$$

and

$$w^0 + nv_c < \tau \leq w^0 + (n + 1)v_c, \tag{2}$$

which implies

$$n = \text{ceiling} \left(\frac{\tau - w^0}{v_c} \right) - 1. \tag{3}$$

Proof. We have an Exponential(λ) distribution for the time until a given agent returns home, and thus, at any time t , we have an Exponential($w_c^t \lambda$) distribution for the time until additional agents are recruited to explore option c . Thus, the waiting time $T_c(\tau)$ until the number of agents exploring option c hits the threshold τ is the sum of independent, exponentially distributed variables with arithmetically increasing parameters.

The hypoexponential density function is $f(t) = \sum_{i=0}^n C_{i,n} \lambda_i e^{-\lambda_i t}$ with $C_{i,n} = \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i}$. Taking n and $\lambda_1, \dots, \lambda_n$ to be functions of c and τ (given by Eqs. 1 and 3), this gives us the probability density function $f_{c,\tau}(t)$ for each $T_c(\tau)$.

Modeling the decision process in terms of accumulation to a fixed threshold, we can think of $-T_c(\tau)$ as a stochastic utility for each alternative c . The system selects the alternative that maximizes this utility. [Of course, because this process is stochastic, maximizing this utility does not necessarily align with maximizing quality. The design of the system reflects dual objectives of aligning utility with quality (to more often choose higher-quality options) and maximizing this utility (to make quicker decisions).] The density functions for the $T_c(\tau)$ variables thus determine the quantities of interest in the system: the probability $p_c(\tau)$ of selecting each possible choice c as well as the waiting time $T(\tau)$ until the decision is made.

Theorem 2. The waiting time $T(\tau)$ distribution and the choice probabilities $p_c(\tau)$ are determined by the $T_c(\tau)$ distributions given by Lemma 2. The time until a decision is made by the decentralized system

is $T(\tau) = \min_c T_c(\tau)$. The probability that the eventual decision is for choice c is $p_c(\tau) = \Pr[T_c(\tau) < \min_{c \neq c} T_c(\tau)]$.

Proof. The urn process runs until the first time that balls of any one color accumulate to the threshold, and the probability of selecting any given choice is simply the probability that balls of the corresponding color reach that threshold first.

Theorem 2 characterizes the (distribution of) time(s) it takes to make a decision and the choice probabilities as functions of the quorum threshold, given any menu of possible choices. We can, of course, calculate the minimum of a set of random variables, as the theorem requires us to do, but there is no simple, closed-form expression for this. Given the lack of a closed-form solution, we use computation to gain additional insight about this process.

Computational results

To explore the effects of parameter variation and the introduction of noise into the process, we run computational experiments of the proposed mechanism. Although the process described above runs in continuous time, we identify discrete time steps every time an agent returns home to recruit (that is, every time a ball is drawn from the urn). Let the index μ count the number of agents that have returned home for a visit, and denote the time when the μ th agent returns home as t_μ . When there are $w = \sum_c w_c^\mu$ agents exploring the set of possible options, the expected time until the next agent returns home is $\frac{1}{w\lambda}$. Setting $\lambda = 1$, which normalizes the units of time, we have:

$$\mathbf{E}[t_{\mu+1} - t_\mu | w] = \frac{1}{w}.$$

We generally set $w^0 = 1$ for simplicity. Each computation reports the average time until decision $T(\tau)$ and the probability of a mistake $p_{\sim c^*}(\tau) = 1 - p_{c^*}(\tau)$ as a function of the quorum threshold τ .

We see the tradeoff between speed and accuracy across varying quorum thresholds by viewing the expected time until decision and the mistake probability as parametric functions of the threshold. Allowing the threshold to vary, we have a Pareto-efficient frontier along which the speed of the decision mechanism cannot be improved without sacrificing accuracy (and vice versa).

Parameter variation

Increasing the number of possible options C makes for a less accurate decision, but a slightly quicker one as well. (This is shown in the appendix in fig. S1.) More options provide more opportunities for sub-optimal options to accumulate a quorum, leading to more mistakes and less decision time. But then, to reach the same level of accuracy, the system needs a higher threshold, and this increases the time required to make the decision (as shown in Fig. 1). Intuitively, more possible options make for a more difficult decision.

Increasing the quality of the optimal choice, v_{c^*} , makes the decision easier (as shown in the left plot in Fig. 2). (These Pareto frontiers are derived from simulations shown in figs. S2 and S3.) As the quality of the optimal choice increases, the decision can be made faster and with less chance of error. Recruitment becomes more effective, so the agents accumulate at this option more quickly and the system achieves the quorum sooner.

Increasing the quality of a suboptimal choice, however, does not have such straightforward consequences. It has three effects: (i) it makes the decision process quicker, (ii) it increases the probability of selecting the suboptimal choice, and (iii) it lowers the cost of making

a suboptimal decision. Thus, the net impact on the ultimate quality of the decision could go in either direction. The right graph in Fig. 2 shows that increasing the quality of a suboptimal option hurts the accuracy of the decision that can be made within a given time.

Noise and risk aversion

Assuming that each option has some absolute quality helped keep our model analytically tractable, but we can extend the model by allowing the quality, v_c , to be a random variable. At a conceptual level, we might think of the process of search as inherently noisy, because there could be natural variation in the agents' perceptions of quality. Alternatively, we might think of the process of recruitment as inherently noisy, with variation in the ability of agents to recruit other agents. In either case, the formal treatment is the same.

One potential source of noise is in the perception of the quality of options when sampled across multiple attributes. If the individual agents cannot investigate all attributes, they may produce a noisy estimate of overall quality by sampling a single (or a few) attribute(s). We can associate random sampling of attributes with noise in the overall quality of the option. [If agents were to specialize in sampling particular attributes and were also more likely to recruit their own type of specialist, then the dynamics of the decision mechanism would be more complicated. Natural systems (for example, a swarm of bees or a colony of ants) do not (to our knowledge) exhibit such behavior, but it could perhaps arise in human-engineered systems.] The decision mechanism is generally robust to noise (see fig. S4), but capable of distinguishing when one option has noisier quality than another. We can think of an option with noisier quality as riskier.

We compare the attractiveness of risky and safe options (that is, options with the same expected quality but more or less noise, respectively) in Fig. 3, which shows that (for a fixed quorum threshold of 100) the safe option (with $v_{\text{safe}} = 2$) is more likely to be selected than a risky option (with $v_{\text{risky}} = \left\{1, \frac{R-2}{R-1}; R, \frac{1}{R-1}\right\}$), and it is increasingly preferred to even riskier options (that is, as R increases). There is nothing special about the threshold of 100, and the result holds for almost all thresholds (possible exceptions being low thresholds that can be reached by a single draw of the risky option, due to recruitment having discrete increments), as shown in fig. S5. The effect persists with high thresholds because noise in the process of search and recruitment does not inevitably balance out; rather, positive feedback in the process makes it more difficult for the risky option to overcome early indications of low quality. [We prove in the appendix (Proposition 1) that the probability of selecting a safe option with quality $v_{\text{safe}} = 1$ over a risky option with quality $v_{\text{risky}} = \left\{0, \frac{R-1}{R}; R, \frac{1}{R}\right\}$, for a quorum threshold of $\tau = R + 1$ is $RB\left(R, 1 + \frac{1}{R}\right)$ (where B is the Euler beta function), which is an increasing function graphed in fig. S6.] Thus, the decision mechanism exhibits a systematic degree of risk aversion.

Risk aversion is defined for deterministic choice models as a preference against mean-preserving spreads and is conventionally represented with concave utility functions. However, risk aversion can get more complicated for stochastic choice models. A strong condition of risk aversion with stochastic choice would require the entire distribution of the stochastic utility to shift downward for mean-preserving spreads. We do not obtain such universal risk aversion, noting occasional exceptions to this general pattern of preferences at low thresholds (see fig. S5). Instead, we define risk aversion in our context as occurring if a mean-preserving spread of an alternative's quality decreases the probability of choosing that alternative, for sufficiently high thresholds. We observe this in our computational results.

The intuition behind the emergence of risk aversion is that the positive feedback in the search and recruitment process allows small advantages to be self-reinforcing, so an option that consistently appears relatively good fares better than one that occasionally appears either great or lackluster. In a world with natural selection, where an entire population can be decimated if a risky choice turns out badly, it may well be adaptive to use a decision mechanism that inherently favors safer choices (52). Moreover, our mechanism permits the system's consensus choice to be risk-averse even when individual agents




Fig. 1. Pareto frontiers of mistake probability and expected waiting time with 2 and 4 options. The optimal choice has quality $v_{c^*} = 2$, whereas the suboptimal choices have quality $v_c = 1$ for all $c \neq c^*$. As an artifact of specifying recruitment (that is, choice quality) so precisely, there are thresholds for which the decision is both slower and less accurate than for a threshold one unit smaller. The corresponding points on the graph are clearly not on the Pareto frontier, but they are shown for completeness.





Fig. 2. Pareto frontiers of mistake probability and expected waiting time with varying choice quality. Left: Varying optimal choice quality. There are $C = 2$ possible choices, and the quality of the suboptimal choice is $v_{c^*} = 1$. Right: Varying suboptimal choice quality. There are $C = 2$ possible choices, and the quality of the optimal choice is $v_{c^*} = 4$.




Fig. 3. Probability of selecting the safe option over a risky option with the same expected quality for a fixed quorum threshold of 100. There are $C = 2$ options. The safe option has quality $v_{\text{safe}} = 2$. The riskiness of the risky option is indexed by the potential reward R such that the quality of the risky option is $v_{\text{Risky}} = \left\{ 1, \frac{R-2}{R-1}; R, \frac{1}{R-1} \right\}$, that is, it has expected quality 2 and variance $R - 2$.

are risk-neutral. Thus, the system could simultaneously be risk-averse for systemic risk and risk-neutral for idiosyncratic risk (when consensus choice is not required), which would be evolutionarily adaptive (53).

Discovery and disruption

We can enrich our model by allowing agents to discover the possible options on their own and to be disrupted from search and recruitment by outside forces. We assume that a Poisson process, with intensity β_c , governs the discovery of each possible option c and that each agent has an exponentially distributed lifetime for search and recruitment with hazard rate δ of disruption. With discovery, a natural initial condition is $w^0 = 0$, that is, the urn is initially empty and agents need to discover an option for recruitment to begin. The index μ for discrete time steps now must count all events, that is, every time an agent discovers a possible choice or falls prey to disruption, as well as when an agent returns home to recruit. The expected time step becomes:

$$E[t_{\mu+1} - t_{\mu}|w] = \frac{1}{w(\lambda + \delta) + \sum_c \beta_c}.$$

Figure 4 shows that decision speed and accuracy are fairly robust to the introduction of option discovery and disruption of search and recruitment. Increasing the rate of discovery (for all options) speeds up the decision and reduces mistakes by limiting the sensitivity to initial advantage, making it easier for the optimal choice to catch up when it gets discovered later on. Increasing the rate of recruitment also speeds up the decision, but can lead to more mistakes by reinforcing initial advantages (that is, when a suboptimal option is discovered first). Increasing the rate of disruption slows down the decision, yet it counteracts initial advantages, allowing the optimal option more time to get ahead through stronger recruitment. (Figure S7 shows the distinct effects on mistake probability and expected time until decision.) In all cases, the quorum threshold could be adjusted to efficiently navigate the new speed versus accuracy tradeoff. The net effects, which are shown in Fig. 4, are that increasing the rate of disruption harms the decision, whereas increasing the rate of discovery or the rate of recruitment improves the decision. Thus, although recruitment introduces positive feedback that can reinforce suboptimal options, it speeds up the decision process enough that at higher thresholds the system can make better, quicker decisions. On the other hand, disruption slows down the decision process so much that despite the opportunities for




Fig. 4. Pareto frontiers of mistake probability and expected waiting time with varying rates of discovery, recruitment, and disruption. There are $C = 2$ options. The optimal choice has quality $v_{c^*} = 2$, and the suboptimal choice has quality $v_{c^*} = 1$. The rate of discovery β_c is the same for both options, and it varies across the columns. The agents' hazard rate of disruption δ varies across the rows. The agents' rate of recruitment λ varies within each panel.

error correction (by effectively allowing agents to occasionally change their minds), at lower thresholds the system makes slower, worse decisions. The error-correcting feature of disruption could nonetheless be helpful if, contrary to our assumption, a system were constrained to a small number of agents and thus forced to have a low threshold.

DISCUSSION

From ants (54) to bees (46) to neurons in the brain (55), a variety of systems productively use decentralized decision mechanisms. Our general notion of decentralized decision-making assumes that no single agent has direct access to information across all of the choices or the ability to make, and communicate, a final decision. Although each agent does have the ability to make limited judgments and decisions, it is the system as a whole that must integrate these activities into a final choice and action. Our model of decentralized decision-making abstracts beyond any one of these systems and aims to provide a deeper understanding of how such mechanisms behave.

A Polya urn scheme running until it hits a finite threshold parsimoniously captures a decentralized decision mechanism in which agents gather local information about possible options through search and recruitment with positive feedback, and the system then makes a consensus choice when it detects a quorum in support of a particular option. In this approach, we add to the literature that uses the Polya urn process to model positive feedback in firm growth (56, 57), technology lock-in (40, 58), the common law legal system (59), the evolution of social and political institutions (60, 61), and the design of medical trials (62).

Analytically, we characterized the waiting time to make a decision and the choice probabilities for any quorum threshold, and we identified an inevitable tradeoff between the speed and the accuracy of the decision. Numerical experiments showed that the system's ability to make reasonably good, quick decisions is robust to parameter variation, noise, and disruption. Moreover, the computation reveals that the decision mechanism naturally exhibits systemic risk aversion.

Additional assumptions about the cost of waiting and the relative values of possible options would be necessary to evaluate the exact tradeoff between a decision's speed and accuracy. At the extremes, an infinite quorum threshold requires infinite waiting time, and a minimal quorum threshold corresponds to uniformly random choice, so the optimal threshold lies in between. The optimal threshold depends on the particular decision context, and there is some evidence that natural systems tune their thresholds in response to the decision context to make better tradeoffs between speed and accuracy (17, 63).

The fact that many natural systems independently evolved similar decentralized decision mechanisms suggests that such mechanisms may represent a robust solution to the general problem of making good, group-level decisions in the absence of centralized control. Indeed, we suspect that evolutionary forces are sufficient to form such natural systems and, over evolutionary time, tune their performance. The decentralized decision mechanism we described here may also prove useful in the design of new social and artificial systems. Novel applications range from improving human organizations to applying such techniques to artificial systems like algorithmic search and the control of swarms of robots or networked computers.

MATERIALS AND METHODS

We used a Monte Carlo method based on 100,000 trials of a Polya urn process with a range of parameters using code written in Python 3.3. The results are confirmed with numeric calculations of the exact distributions in those cases where such computations are possible (typically, thresholds up to about 20). The standard errors are negligibly small for both the mistake probability (no greater than 1% of the estimate for any threshold) and the expected time (not exceeding 0.1% of the estimate). The simulation source code, along with the data used to produce our figures, is available at <http://files.dhagmann.com/papers/polyaBees.zip>.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/1/8/e1500253/DC1>

Fig. S1. Mistake probability and expected waiting time as a function of the quorum threshold with 2 and 4 possible options.

Fig. S2. Mistake probability and expected waiting time as a function of the quorum threshold, varying the quality of the optimal choice.

Fig. S3. Mistake probability and expected waiting time as a function of the quorum threshold, varying the quality of the suboptimal choice.

Fig. S4. Mistake probability and expected waiting time as a function of the quorum threshold in noisy and noiseless environments.

Fig. S5. Probability of selecting a safe option with quality $v_{\text{safe}} = 2$ over a risky one with quality $v_{\text{risky}} = \{1, 50\%; 3, 50\%\}$, as a function of the quorum threshold.

Fig. S6. Probability of selecting a safe option with quality $v_{\text{safe}} = 1$ over a risky one with quality $v_{\text{risky}} = \left\{0, \frac{R-1}{R}; R, \frac{1}{R}\right\}$, for a quorum threshold of $\tau = R + 1$.

Fig. S7. Mistake probability and expected waiting time as a function of the quorum threshold, with varying rates of discovery, recruitment, and disruption.

An Ancillary Result

REFERENCES AND NOTES

1. E. Mallon, S. Pratt, N. Franks, Individual and collective decision-making during nest site selection by the ant *Leptothorax albipennis*. *Behav. Ecol. Sociobiol.* **50**, 352–359 (2001).
2. P. K. Visscher, S. Camazine, Collective decisions and cognition in bees. *Nature* **397**, 400 (1999).
3. N. R. Franks, E. B. Mallon, H. E. Bray, M. J. Hamilton, T. C. Mischler, Strategies for choosing between alternatives with different attributes: Exemplified by house-hunting ants. *Anim. Behav.* **65**, 215–223 (2003).
4. R. Jeanson, J.-L. Deneubourg, A. Grimal, G. Theraulaz, Modulation of individual behavior and collective decision-making during aggregation site selection by the ant *Messor barbarus*. *Behav. Ecol. Sociobiol.* **55**, 388–394 (2004).
5. E. J. H. Robinson, O. Feinerman, N. R. Franks, How collective comparisons emerge without individual comparisons of the options. *Proc. Biol. Sci.* **281**, 20140737 (2014).
6. S. C. Pratt, E. B. Mallon, D. J. Sumpter, N. R. Franks, Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albipennis*. *Behav. Ecol. Sociobiol.* **52**, 117–127 (2002).
7. C. List, Democracy in animal groups: A political science perspective. *Trends Ecol. Evol.* **19**, 168–169 (2004).
8. S. Camazine, P. K. Visscher, J. Finley, R. S. Vetter, House-hunting by honey bee swarms: Collective decisions and individual behaviors. *Insectes Soc.* **46**, 348–360 (1999).
9. T. D. Seeley, S. C. Buhrman, Group decision making in swarms of honey bees. *Behav. Ecol. Sociobiol.* **45**, 19–31 (1999).
10. T. D. Seeley, P. K. Visscher, Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making. *Behav. Ecol. Sociobiol.* **54**, 511–520 (2003).
11. T. D. Seeley, S. C. Buhrman, Nest-site selection in honey bees: How well do swarms implement the “best-of-N” decision rule? *Behav. Ecol. Sociobiol.* **49**, 416–427 (2001).
12. N. F. Britton, N. R. Franks, S. C. Pratt, T. D. Seeley, Deciding on a new home: How do honeybees agree? *Proc. Biol. Sci.* **269**, 1383–1388 (2002).
13. T. D. Seeley, Consensus building during nest-site selection in honey bee swarms: The expiration of dissent. *Behav. Ecol. Sociobiol.* **53**, 417–424 (2003).
14. T. D. Seeley, P. K. Visscher, Group decision making in nest-site selection by honey bees. *Apidologie* **35**, 101–116 (2004).
15. K. von Frisch, *The Dance Language and Orientation of Bees* (Harvard Univ. Press, Cambridge, MA, 1967).
16. T. D. Seeley, How honeybees find a home. *Sci. Am.* **247**, 158–169 (1982).
17. A. Dornhaus, N. R. Franks, R. M. Hawkins, H. N. S. Shere, Ants move to improve: Colonies of *Leptothorax albipennis* emigrate whenever they find a superior nest site. *Anim. Behav.* **67**, 959–963 (2004).
18. M. Möglich, Social organization of nest emigration in *Leptothorax* (Hym., Form.). *Insectes Soc.* **25**, 205–225 (1978).
19. M. Möglich, B. Hölldobler, Social carrying behavior and division of labor during nest moving in ants. *Psyche J. Entom.* **81**, 219–236 (1974).
20. N. R. Franks, T. Richardson, Teaching in tandem-running ants. *Nature* **439**, 153 (2006).
21. T. O. Richardson, P. A. Sleeman, J. M. McNamara, A. I. Houston, N. R. Franks, Teaching with evaluation in ants. *Curr. Biol.* **17**, 1520–1526 (2007).
22. S. C. Pratt, Quorum sensing by encounter rates in the ant *Temnothorax albipennis*. *Behav. Ecol.* **16**, 488–496 (2005).
23. L. Avilés, Nomadic behaviour and colony fission in a cooperative spider: Life history evolution at the level of the colony? *Biol. J. Linn. Soc.* **70**, 325–339 (2000).
24. F. Saffre, R. Furey, B. Krafft, J. L. Deneubourg, Collective decision-making in social spiders: Dragline-mediated amplification process acts as a recruitment mechanism. *J. Theor. Biol.* **198**, 507–517 (1999).
25. R. Jeanson, J.-L. Deneubourg, G. Theraulaz, Discrete dragline attachment induces aggregation in spiderlings of a solitary species. *Anim. Behav.* **67**, 531–537 (2004).
26. J.-M. Ame, C. Rivault, J.-L. Deneubourg, Cockroach aggregation based on strain odour recognition. *Anim. Behav.* **68**, 793–801 (2004).
27. R. Jeanson, C. Rivault, J.-L. Deneubourg, S. Blanco, R. Fournier, C. Jost, G. Theraulaz, Self-organized aggregation in cockroaches. *Anim. Behav.* **69**, 169–180 (2005).
28. E. Ben Jacob, I. Becker, Y. Shapira, H. Levine, Bacterial linguistic communication and social intelligence. *Trends Microbiol.* **12**, 366–372 (2004).
29. M. N. Shadlen, W. T. Newsome, Motion perception: Seeing and deciding. *Proc. Natl. Acad. Sci. U.S.A.* **93**, 628–633 (1996).
30. D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (John Wiley & Sons Inc., New York, 1949).
31. S. Kaplan, M. Weaver, R. French, Active symbols and internal models: Towards a cognitive connectionism. *AI & Soc.* **4**, 51–71 (1990).
32. J. R. Busemeyer, J. T. Townsend, Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432–459 (1993).
33. M. Usher, J. L. McClelland, The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).

34. P. L. Smith, R. Ratcliff, Psychology and neurobiology of simple decisions. *Trends Neurosci.* **27**, 161–168 (2004).
35. J. A. R. Marshall, R. Bogacz, A. Dornhaus, R. Planqué, T. Kovacs, N. R. Franks, On optimal decision-making in brains and social insect colonies. *J. R. Soc. Interface* **6**, 1065–1074 (2009).
36. D. R. Hofstadter, *Gödel, Escher, Bach: An Eternal Golden Braid* (Basic Books, New York, 1999).
37. K. M. Passino, T. D. Seeley, P. K. Visscher, Swarm cognition in honey bees. *Behav. Ecol. Sociobiol.* **62**, 401–414 (2008).
38. A. Kirman, Ants, rationality, and recruitment. *Q. J. Econ.* **108**, 137–156 (1993).
39. E. M. Rogers, *Diffusion of Innovations* (Simon and Schuster, New York, 2010).
40. W. B. Arthur, Competing technologies, increasing returns, and lock-in by historical events. *Econ. J.* **99**, 116–131 (1989).
41. J. M. Utterback, W. J. Abernathy, A dynamic model of process and product innovation. *Omega* **3**, 639–656 (1975).
42. G. Stasser, W. Titus, Pooling of unshared information in group decision making: Biased information sampling during discussion. *J. Pers. Soc. Psychol.* **48**, 1467–1478 (1985).
43. R. Hastie, T. Kameda, The robust beauty of majority rules in group decisions. *Psychol. Rev.* **112**, 494–508 (2005).
44. R. Bogacz, E. Brown, J. Moehlis, P. Holmes, J. D. Cohen, The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
45. L. Conradt, T. J. Roper, Consensus decision making in animals. *Trends Ecol. Evol.* **20**, 449–456 (2005).
46. T. D. Seeley, *Honeybee Democracy* (Princeton Univ. Press, Princeton, NJ, 2010).
47. T. D. Seeley, When is self-organization used in biological systems? *Biol. Bull.* **202**, 314–318 (2002).
48. D. J. Sumpter, J. Krause, R. James, I. D. Couzin, A. J. Ward, Consensus decision making by fish. *Curr. Biol.* **18**, 1773–1777 (2008).
49. A. J. W. Ward, D. J. T. Sumpter, I. D. Couzin, P. J. B. Hart, J. Krause, Quorum decisionmaking facilitates information transfer in fish shoals. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 6948–6953 (2008).
50. D. J. Sumpter, S. C. Pratt, Quorum responses and consensus decision making. *Philos. Trans. R. Soc. B Biol. Sci.* **364**, 743–753 (2009).
51. S. Balaji, H. M. Mahmoud, Exact and limiting distributions in diagonal Pólya processes. *Ann. Inst. Stat. Math.* **58**, 171–185 (2006).
52. J. Seger, H. Brockmann, *Oxford Surveys in Evolutionary Biology*, P. Harvey, L. Partridge, Eds. (Oxford Univ. Press, Oxford, UK, 1987), vol. 4, pp. 182–211.
53. R. Zhang, T. J. Brennan, A. W. Lo, The origin of risk aversion. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 17777–17782 (2014).
54. D. M. Gordon, *Ant Encounters: Interaction Networks and Colony Behavior* (Princeton Univ. Press, Princeton, NJ, 2010).
55. J. I. Gold, M. N. Shadlen, The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
56. G. Bottazzi, A. Secchi, A stochastic model of firm growth. *Physica A* **324**, 213–219 (2003).
57. R. Golman, S. Klepper, (Under review) (RAND J Econ), Spinoffs and clustering (2015).
58. P. A. David, Clio and the economics of QWERTY. *Am. Econ. Rev.* **75**, 332–337 (1985).
59. O. A. Hathaway, Path dependence in the law: The course and pattern of legal change in a common law system. *Iowa Law Rev.* **86** (2001).
60. J. Mahoney, Path dependence in historical sociology. *Theory Soc.* **29**, 507–548 (2000).
61. S. E. Page, Path dependence. *Q. J. Polit. Sci.* **1**, 87–115 (2006).
62. L. J. Wei, The generalized Pólya's urn design for sequential medical trials. *Ann. Stat.* **7**, 291–296 (1979).
63. S. C. Pratt, D. J. T. Sumpter, A tunable algorithm for collective decision-making. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 15906–15910 (2006).

Funding: This work was not supported by outside funding. **Author contributions:** All authors contributed equally to this manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** The Python source code for the simulation and all data used to create the figures in the paper are available at <http://files.dhagmann.com/papers/polyaBees.zip>. Supplementary materials consist of additional figures and an ancillary proof, referenced in the manuscript.

Submitted 26 February 2015

Accepted 4 August 2015

Published 18 September 2015

10.1126/sciadv.1500253

Citation: R. Golman, D. Hagmann, J. H. Miller, Pólya's bees: A model of decentralized decision-making. *Sci. Adv.* **1**, e1500253 (2015).

Polya's bees: A model of decentralized decision-making

Russell Golman, David Hagmann and John H. Miller

Sci Adv 1 (8), e1500253.

DOI: 10.1126/sciadv.1500253

ARTICLE TOOLS

<http://advances.sciencemag.org/content/1/8/e1500253>

SUPPLEMENTARY MATERIALS

<http://advances.sciencemag.org/content/suppl/2015/09/15/1.8.e1500253.DC1>

REFERENCES

This article cites 54 articles, 4 of which you can access for free
<http://advances.sciencemag.org/content/1/8/e1500253#BIBL>

PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2015, The Authors