


Modeling multi-robot task allocation with limited information as global game

Anshul Kanakia^{1,2} · Behrouz Touri^{1,3} ·
Nikolaus Correll^{1,2} 

Received: 3 August 2015 / Accepted: 21 April 2016 / Published online: 28 April 2016
© Springer Science+Business Media New York 2016

Abstract Continuous response threshold functions to coordinate collaborative tasks in multi-agent systems are commonly employed models in a number of fields including ethology, economics, and swarm robotics. Although empirical evidence exists for the response threshold model in predicting and matching swarm behavior for social insects, there has been no formal argument as to why natural swarms use this approach and why it should be used for engineering artificial ones. In this paper, we show, by formulating task allocation as a global game, that continuous response threshold functions used for communication-free task assignment result in system level Bayesian Nash equilibria. Building up on these results, we show that individual agents not only do not need to communicate with each other, but also do not need to model each other's behavior, which makes this coordination mechanism accessible to very simple agents, suggesting a reason for their prevalence in nature and motivating their use in an engineering context.

Keywords Threshold-based task allocation · Swarm robotics · Social insects · Game theory · Global games

1 Introduction

Task allocation subject to communication constraints is ubiquitous in nature in many different organisms, ranging from cellular systems (Yoshida et al. 2010; Suzuki et al. 2015) to social insects (Robinson 1987; Gordon 1996; Bonabeau et al. 1998; Theraulaz et al. 1998) to large animal herds (Conradt and Roper 2003, 2005) and human society (Raafat et al. 2009).

✉ Nikolaus Correll
ncorrell@colorado.edu

¹ College of Engineering and Applied Sciences, University of Colorado Boulder, Boulder, CO, USA

² Department of Computer Science, University of Colorado Boulder, Boulder, CO, USA

³ Department of Electrical, Computer and Energy Engineering, University of Colorado Boulder, Boulder, CO, USA

Inter-agent communication in large systems is not always possible or desired, either due to physical limitations at the agent level (cellular/insect systems) or properties of the task itself (adversary behavior in humans). Here, we show that formalizing task allocation problems as a global game, a concept from the field of game theory, reveals that a simple threshold strategy leads to Bayesian Nash equilibria (BNE) despite the absence of communication between agents. This result not only provides a hypothesis about the inner workings of a wide range of systems with limited communication between agents but also provides a formal analysis tool for threshold-based task allocation in social insects. In particular, we show how noise in the perception apparatus of individual agents leads to commonly observed sigmoid response threshold functions that control the trade-off between exploration and exploitation (Bonabeau et al. 1997) in natural systems and can be used to design engineered systems ranging from swarm robotics (Martinoli et al. 1999; Krieger et al. 2000; Kube and Bonabeau 2000; Mataric et al. 2003; Gerkey 2004) to smart composites (McEvoy and Correll 2015), made of computational elements with very low complexity.

1.1 Related work

Task allocation is a canonical problem in multi-robot systems (Gerkey 2004; Brambilla et al. 2013) and is the problem of allocating individuals among different tasks, usually to satisfy some metric of optimality. Task allocation is distinct from task partitioning (Pini et al. 2013), which is concerned with breaking a larger problem into consecutive and parallel subtasks. The simplest possible task allocation approach is to use fixed probabilities for individuals to switch between tasks. The resulting Markov chains lead to equilibria, which correspond to the ratios of probabilities (Correll 2008). Whereas robots equipped with sufficient means for computation and communication might employ more sophisticated means for task allocation, e.g., using market-based approaches (Amstutz et al. 2008; Vig and Adams 2007; Choi et al. 2009) or using leader–follower coalition algorithms (Chen and Sun 2011), probabilistic algorithms are of particular interest for swarm robotics with individually simple controllers (Dantu et al. 2012) and little to no communication ability. Recruitment of an exact number of robots to a particular task has been extensively studied using the “stick pulling” experiment (Lerman et al. 2001; Martinoli et al. 2004). The problem of distributing a swarm of robots across a discrete number of sites/tasks with a specific desired distribution has been studied in Berman et al. (2009) and Correll (2008). We showed in Kanakia and Correll (2016) that these problems are related, as they can be cast as instances of sigmoidal threshold task allocation problems with varying slopes. Similarly, Mather et al. (2010) present a stochastic approach that is a hybrid between the work in Berman et al. (2009) and Martinoli et al. (2004), allowing allocation to tasks requiring a varying number of robots. Yet, formally understanding probabilistic task allocation strategies is an open challenge. Although there exist differential equation models for specific classes of foraging problems (Lerman et al. 2006; Liu and Winfield 2010) that allow one to study the relationship between robot control parameters and resulting dynamics, results are purely phenomenological and leave it open whether there exist better or worse approaches for task allocation given the individual agent’s capabilities. Other works address this question empirically, by performing comparisons of probabilistic and deliberate strategies for certain classes of task allocation problems (Kalra and Martinoli 2006; Correll 2007). Here, the general insight is that stochastic task allocation is beneficial under the influence of noise, which leads to failure of “optimal” assignment techniques, thereby making them comparable with a stochastic strategy under such conditions. Most importantly, however, there currently exists no formal underpinning for threshold-based task allocation mechanisms.

There exists also a rich body of work that uses techniques from game theory to multi-agent coordination (Parsons and Wooldridge 2002; Nisan et al. 2007) and task allocation in particular (Shehory and Kraus 1998). For games with perfect information and specific structures, such as potential games, weakly acyclic games, generalized potential games, or anti-coordination games, one can utilize learning algorithms, such as log-linear-learning and fictitious play with inertia to reach efficient strategies (Tumer and Wolpert 2004; Arslan et al. 2007; Marden et al. 2009; Grenager et al. 2002). These approaches are not applicable here, where we consider systems with imperfect information and no communication between agents. For this class of systems, there exists a branch of game theory in which information about characteristics of the other players is incomplete. These class of games are known as *Bayesian games* (Harsanyi 2004) and are of particular interest to multi-agent systems problems that involve uncertainty. Here, *global games* (Carlsson and Van Damme 1993) are a subset of Bayesian games, in which players receive possibly correlated signals of the underlying state of the world, making them an interesting, but yet unexplored, avenue for modeling and understanding stimulus-response task allocation systems.

1.2 Outline of this paper

This paper provides a formal analysis framework for threshold-based task allocation by formulating it as a “global game” (Sect. 2) and formally showing that a simple, communication-free threshold policy (Sect. 3) leads to a Bayesian Nash equilibrium (Sect. 4). We then show that such a policy is equivalent to a sigmoid response threshold policy, which is commonly observed in social insects and used in swarm engineering, in Sect. 5.

2 Global games: a brief overview

Game theory is the study of strategic interactions among multiple agents or players, such as robots, people, and firms where the decision of each party affects the payoff of the rest. A fundamentally important class of games is one with incomplete or imperfect information where each agent’s utility depends not only on the actions of the other agents, but also on an underlying fundamental signal that cannot be accurately ordained by the agents. The class of global games with incomplete information was originally introduced in Carlsson and Van Damme (1993) where two players are playing a game and the utility of the two players depends on an underlying fundamental signal $\tau \in \mathbb{R}$, but each agent observes a noisy variation of this signal, x_i . For example in a firefighting task, this fundamental signal τ is the *magnitude* of the task of putting out the fire, i.e., the number of robots needed to do so. The size and intensity of the fire along with environmental and other site-specific factors all play a major role in determining whether an agent should begin the task or wait for more help to arrive.

While we use the term *magnitude* to describe τ in the above example, it is a stand-in for a simplified representation of a more abstract quality of any task. All tasks demand completion and the act of completion requires resources, be it time and/or energy of some form. In swarms of minimalist agents with limited capabilities, the resource required to collaboratively complete a task is invariably quantized into the number of agents attempting to complete that task.

3 Task allocation as global games

Consider a group of agents performing a task contributing to a common goal, which we refer to as a concurrent benefit. This benefit is related to a stimulus τ that can be observed by all agents, albeit subject to sensing noise. Agents do not share any information. All agents decide, for themselves, whether or not to engage in the task. A task is successfully attempted if a critical mass of agents is willing to participate in it. Otherwise, the attempt fails.

Situations like this arise in a number of different fields including neurology (Yoshida et al. 2010; Suzuki et al. 2015), ethology (Robinson 1987; Gordon 1996; Bonabeau et al. 1998; Theraulaz et al. 1998), sociology (Raafat et al. 2009), economics (Morris and Shin 2000), and robotics (Martinoli et al. 1999; Krieger et al. 2000; Kube and Bonabeau 2000; Pynadath and Tambe 2002; Gerkey and Mataric 2003; Matarić et al. 2003; Gerkey 2004; Kanakia and Correll 2016). All of these multi-agent scenarios share the common notion of a joint action or response to a commonly observed stimulus. The task can take on many forms ranging from neurons simply firing in concert, collective decision problems like flocking, herd grazing, and colony defense to individual actions based on the environment and other agents' beliefs like foraging, bank runs, and political revolutions.

In the case of a bank run (Morris and Shin 2000), τ is an aggregate stimulus parameter that represents the strength of the economy of a nation. Here, agents decide when to withdraw their assets from banks based on their own noisy estimate of the economy together with a simple threshold. In the case of social insects foraging for food (Bonabeau et al. 1996; Theraulaz et al. 1998; Krieger et al. 2000), τ represents a number of environmental cues such as the (imperfect) measurement of food stores in a colony, pheromone levels (Robinson 1987), or the waiting time for food transfer from one agent to another (Seeley 1989). A complex combination of these internal and external cues (Gordon 1996) tempers an agent's perception of the magnitude of a task. In an engineering context, τ can be seen as the magnitude of a fire (heat intensity and area covered) as sensed by a robot using onboard instruments in an automated firefighting scenario (Kanakia and Correll 2016). Figure 1 illustrates each of these three examples with their corresponding stimulus parameters.

The group dynamics in the above examples may seem orthogonal at first; while adversarial behavior between agents drives bank runs, collaborative behavior between robots is essential for the automated firefighting scenario. Both scenarios, however, share the notion that to be successful an agent not only needs to assess the magnitude of the task itself but also the likelihood of the other agents to act. This is because only acting in concert leads to the desired group action, be it because using up water resources to put out a fire is futile before critical mass is reached, or disengaging from the banking system is non-desirable unless there is a major crisis. In a system with multiple tasks, such as an ant colony, coordination is required to achieve a desirable proportion between tasks.

We build on results from global games (Carlsson and Van Damme 1993) to show that the observed behavior in all these scenarios can be effectively emulated by assuming that each agent makes their individual decision on whether or not to perform a task based on some internal threshold value which is compared to their noisy estimates of the collective task's stimulus τ . This was shown for the canonical bank run example (Morris and Shin 2000). While the classical global game assumes each agent must predict the other agents' behavior, it turns out that agents can reach an equilibrium without this capacity. This, and the fact that agents do not need to communicate, makes this approach widely applicable to a wide range of multi-agent systems.

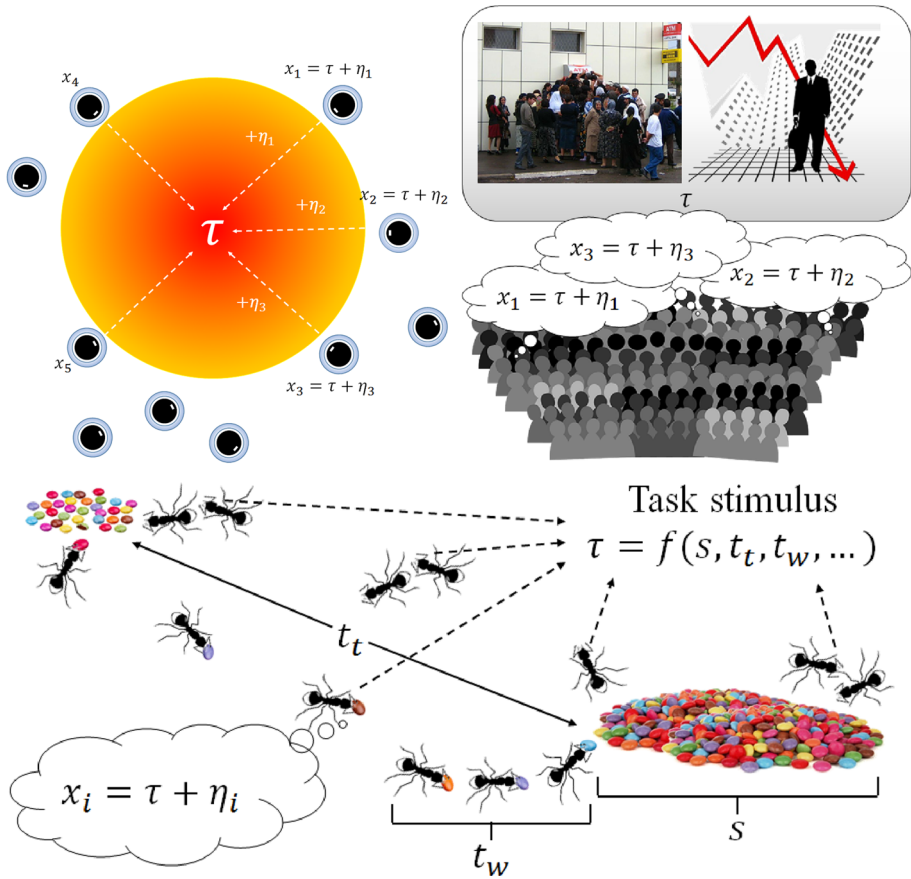


Fig. 1 Robotic firefighting, ant foraging, and bank run scenarios presented as global games. Each player’s imperfect estimate of the task is represented by x_i , comprising of the global stimulus parameter τ and noisy sensor measurements η_i . In the robot firefighting scenario, τ is representative of the magnitude of the fire, while in the case of a bank run τ is indicative of an agent’s current level of trust in the nation’s economy. For the ant foraging scenario, τ represents an ant’s willingness to take part in the foraging task based on a number of internally measured parameters such as the distance to the food source (t_t), the wait time to deliver food (t_w), and the food stores currently at the nest (s), among others

Consider a set of n agents and suppose that each agent has an action set $A_i = \{0, 1\}$ where 0 represents not participating in the task and 1 represents participating in the task. Every agent is also aware of the total number of other agents, n in the system. For the purpose of analysis, we assume the decision to act or not to act is made by all agents at the same time, i.e., this is a one-shot game with no notion of time. We let the stimulus τ be a real number that belongs within the interval $E = [c, d]$ in \mathbb{R} . Finally, we let $u_i : A_i \times \mathbb{Z}^+ \times \mathbb{R} \rightarrow \mathbb{R}$ be the utility of the i th agent, where $u_i(a_i, g, \tau)$ is the utility of the i th agent when g other agents have decided to participate in the task. In general, the utility of each agent depends on the joint actions of the rest of the agents. For simplicity, we assume the utility to be proportional to the number of agents participating in the activity.

The utility function discussed throughout this paper has the following properties:

- (a) $u_i(1, g, \tau) - u_i(0, g, \tau)$ is an increasing and continuous function of τ for any g . We further assume that $|u_i(1, g, \tau) - u_i(0, g, \tau)| \leq \tau^p$ for some $p \geq 1$.
- (b) For extreme stimulus ranges, taking part in the activity is either appealing or repelling, i.e., there exists $\underline{\tau}, \bar{\tau} \in (c, d)$ with $\underline{\tau} \leq \bar{\tau}$ such that for any $\tau \geq \bar{\tau}$ we have $u_i(1, g, \tau) > u_i(0, g, \tau)$ where all agents participate in very easy tasks, and for $\tau \leq \underline{\tau}$ we have $u_i(1, g, \tau) < u_i(0, g, \tau)$ so the only equilibrium of the game is for all agents to not participate as the task is too difficult.

Note that in order to have a task with such an utility, we need the above conditions to hold for all the agents, i.e., for all $i \in \{1, \dots, n\}$. An example of a utility function that would satisfy such conditions is a function $u_i(a_i, g, \tau) = a_i(1 - e^{-(g+1)} + \tau)$.

The main challenge in devising task allocation strategies is that the true value of τ is not easily accessible to the agents, for example, due to limited perception capabilities and sensor noise. We model this imperfect knowledge by assuming that agent i observes $x_i = \tau + \eta_i$ where η_i is a Gaussian $\mathcal{N}(0, \sigma_i^2)$ random variable. Note that this makes the game a Bayesian game and in this case, the type of each player is represented by the random variable x_i . Throughout our discussion, we assume that the task stimulus τ is a Gaussian random variable and is independent of η_1, \dots, η_n . This analysis is extendable to a larger class of random variables, but for the simplicity of the discussion, we consider Gaussian random variables here. Given these constraints, the question is what strategy the agents should follow to reach a BNE, that is an outcome in which no agent has the incentive to deviate from its current strategy.

A strategy s_i for the i th agent is a measurable function $s_i : \mathbb{R} \rightarrow A_i$, mapping measurements (observations) to actions. Strategy s_i prescribes what action the i th agent should take given its own measurement (type) x_i . Given this, consider a set of agents with strategies s_1, \dots, s_n . Let us denote the strategies of the $n - 1$ agents other than the i th agent by the vector $S_{-i} = \{s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n\}$. We say that a strategy s_i is a *threshold strategy* if $s_i(x) = \text{step}(x, \Upsilon_i)$, i.e., the step function with a jump from 0 to 1 at Υ_i , where Υ_i is the internal threshold value of the i th agent. For the i th agent, we define the best-response $\text{BR}(S_{-i})$ (to the strategies of the other agents) to be a strategy \tilde{s} that for any $x \in \mathbb{R}$:

$$\begin{aligned} \text{BR}(S_{-i})(x) &= \tilde{s}(x) \in \underset{a_i \in A_i}{\text{argmax}} E(u_i(a_i, g, \tau) \mid x_i = x) \\ &= \underset{a_i \in A_i}{\text{argmax}} E \left(u_i \left(a_i, \sum_{j \neq i} s_j(x_j), \tau \right) \mid x_i = x \right), \end{aligned}$$

where $E(\cdot|\cdot)$ is the conditional expectation of u_i given the i th agent’s observation. The best response of player i simply is the best course of action for agent i given that the strategies of the other players is given. The expression $\underset{a_i \in A_i}{\text{argmax}} E(u_i(a_i, g, \tau) \mid x_i = x)$ is the set of best actions that player i can take given its information, the aggregate action g of the other players, and the intensity τ . Note that given the i th agent’s observation x_i , the observations of the other agents, and hence their actions, would be random from the i th agent perspective, i.e., given x_i and τ , all $s \in S_{-i}$ are effectively random variables with respect to the i th agent. A strategy profile $S = \{s_1, \dots, s_n\}$ is a *sensible strategy*, if it leads to a BNE (Fudenberg 1998), given $s_i = \text{BR}(S_{-i})$ for all $i \in \{1, \dots, n\}$.

4 Communication-free threshold-based task allocation strategy

Any task with concurrent benefit admits a threshold strategy BNE—meaning it is sufficient for the agents to follow a simple algorithm:

1. Compare your noisy measurement x_i to a threshold value \mathcal{Y}_i ,
2. If the measurement is above \mathcal{Y}_i take part in the collaborative task, otherwise hold off.

This algorithm is extremely simple and can be implemented on systems with a wide range of capabilities, yet leads to a BNE as we will show below.

To show that there exists a sensible threshold strategy for the class of tasks with concurrent benefit leading to Theorem 1, we will first show that the best response to threshold strategies is a threshold strategy (Lemma 1) and then show that there exists an equilibrium of threshold strategies (Carlsson and Van Damme 1993; Morris and Shin 2000) (Lemma 2).

Lemma 1 *Let $S = \{s_1, \dots, s_n\}$ be a strategy profile consisting of threshold strategies for a task with concurrent benefit. Let $\tilde{s}_i = BR(S_{-i})$. Then \tilde{s}_i is a threshold strategy.*

Proof We first show that if for some observation $x_i = x$, we have $BR(S_{-i})(x) = \tilde{s}_i(x) = 1$, then $\tilde{s}_i(y) = 1$ for $y \geq x$. To show this, we note that $P(x_j \geq \tau_j \mid x_i = x)$ is an increasing function of x as $x_j - x_i$ is a normally distributed random variable. Therefore, using the monotone property of concurrent tasks and the fact that $x_i = \tau + \eta_i$, we conclude that:

$$\begin{aligned}
 & E \left(u_i \left(1, \sum_{j \neq i} s_j(x_j), \tau \right) \mid x_i = y \right) \\
 & \quad - E \left(u_i \left(0, \sum_{j \neq i} s_j(x_j), \tau \right) \mid x_i = y \right) \\
 & > E \left(u_i \left(1, \sum_{j \neq i} s_j(x_j), \tau \right) \mid x_i = x \right) \\
 & \quad - E \left(u_i \left(0, \sum_{j \neq i} s_j(x_j), \tau \right) \mid x_i = x \right) \geq 0.
 \end{aligned}$$

Therefore $\tilde{s}_i(y) = 1$. Similarly, if for some value of x , we have $\tilde{s}_i(x) = 0$, then it follows that $\tilde{s}_i(y) = 0$ for $y \leq x$. Therefore, \tilde{s}_i would be a threshold strategy. □

We can view the best response of threshold strategies as a mapping from \mathbb{R}^n to \mathbb{R}^n that maps n thresholds of strategies to n thresholds of the best-response strategies. Denote this mapping by $L : \mathbb{R}^n \rightarrow \mathbb{R}^n$. To show that there exists a BNE threshold policy, we show that the mapping L has a fixed point. Note that not all the functions from \mathbb{R}^n to \mathbb{R}^n admit a fixed point. For example, the function

$$f(x) = \begin{cases} -1 & x > 0 \\ 1 & x \leq 0 \end{cases}$$

does not admit a fixed point value, i.e., a point x^* such that $f(x^*) = x^*$. In order to show that L has a fixed point, we use the fact that L is continuous.

Lemma 2 *The mapping L that maps the threshold values of threshold strategies to the threshold values of the best-response strategies is a continuous mapping.*

Proof Let $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ be the vector of observations of $n - 1$ agents except the i^{th} agent. Note that the vector (x_{-i}, τ) given $x_i = x$ is a normally distributed random vector with some continuous density function $f_x(x_{-i}, \tau)$. Now, let $\{\alpha(k)\}$ be a sequence in \mathbb{R}^n that is converging to $\alpha \in \mathbb{R}^n$. Let $\{\beta(k)\}$ be the sequence of thresholds corresponding to the best-response strategy of the strategy with threshold vector $\alpha(k)$. Let s be the threshold strategy corresponding to the threshold vector α and let α^* be the threshold strategy corresponding to the $BR(\alpha)$. By the definition of the best-response strategy, $\beta_i(k)$ is a point where

$$\int_{\mathbb{R}^n} f_{\beta(k)}(z, t) \left(u_i \left(1, \sum_{j \neq i} u^{\alpha_j(k)}(x_j), \tau \right) - u_i \left(0, \sum_{j \neq i} u^{\alpha_j(k)}(x_j), \tau \right) \right) d(z \times t) = 0.$$

Using the fact that f has a Gaussian distribution and is continuous on all its arguments and the fact that $|u_i(\cdot, \cdot, \tau)| \leq \tau^p$, by taking the limit $k \rightarrow \infty$ and the dominated convergence theorem:

$$\int_{\mathbb{R}^n} f_{\beta}(z, t) \left(u_i \left(1, \sum_{j \neq i} u^{\alpha_j}(x_j), \tau \right) - u_i \left(0, \sum_{j \neq i} u^{\alpha_j}(x_j), \tau \right) \right) d(z \times t) = 0,$$

where u^r is a threshold strategy with threshold r . Therefore, the $\lim_{k \rightarrow \infty} L(\alpha(k)) = L(\alpha)$ for a sequence $\{\alpha(k)\}$ that is converging to α . □

Using these lemmas, we can show the existence of a threshold strategy for global games with concurrent benefit.

Theorem 1 *For a concurrent benefit task T , suppose that the stimulus parameter τ is a Gaussian random variable. Also, suppose that $x_i = \tau + \eta_i$ where η_1, \dots, η_n are independent Gaussian random variables. Then, there exists a strategy profile $S = (s_1, \dots, s_n)$ of threshold strategies that is a BNE.*

Proof By Lemma 1, the best response of a threshold strategy is a threshold strategy, and hence, it induces the mapping L from the space of thresholds \mathbb{R}^n to itself. Also, by Lemma 2, this mapping is a continuous mapping. Now, if \mathcal{Y}_i is a sufficiently large threshold, then the second property of concurrent benefit tasks implies that the $\tilde{\mathcal{Y}}_i \leq \mathcal{Y}_i$ because a large enough measurement x_i implies that agent i itself should take part in the task. Similarly, for sufficiently low threshold \mathcal{Y}_i , we will have $\tilde{\mathcal{Y}}_i \geq \mathcal{Y}_i$. Therefore, the mapping L maps a box $[a, b]^n$ to itself, where a is a sufficiently small scalar and $b > a$ is a sufficiently large scalar. Since the box $[a, b]^n$ is a convex closed set, by the Brouwer’s fix point theorem (Border 1990) we have that there exists a vector of threshold values α^* such that $\alpha^* = L(\alpha^*)$ and hence there exists a BNE for the concurrent benefit task T . □

5 From discrete thresholds to sigmoidal response functions

Observations in ethology suggest sigmoid threshold functions (Bonabeau et al. 1996), rather than fixed thresholds as suggested by our analysis. Also, roboticists have started using sigmoid-shaped threshold functions to engineer swarm systems (Bonabeau et al. 1996; Theraulaz et al. 1998; Krieger et al. 2000), as tuning the shape of a sigmoidal response threshold function allows balancing between exploration, i.e., performing a random action, and exploitation, i.e., using all available information in decision making such as a fixed threshold. We argue that this behavior can be a direct result of using a simple discrete threshold under the influence of perception noise. Indeed, one can show that a sigmoid threshold function is the outcome of deterministic threshold functions on noisy observations. Suppose that all agents share the same utility function $u(a_i, g, \tau)$ and also assume that the observation noise of the n agents (η_1, \dots, η_n) are independent and identically distributed (IID) $\mathcal{N}(0, \sigma^2)$ Gaussian random variables. Then, it is not hard to see that there exists a BNE with threshold strategies that have the same threshold value γ (Morris and Shin 2000).

Consider a realization of $\tau = \hat{\tau}$ and suppose that we have a large number of agents n observing a noisy variation of $\hat{\tau}$. Take, for example, the case of firefighting agents, and let $\hat{\tau}$ be the magnitude (including type, intensity, and area) of the fire. Then, since the observations of the n agents are IID given the value of τ , they will be distributed according to $\mathcal{N}(\hat{\tau}, \sigma^2)$. Now consider the relative number of agents taking part in the activity given $\hat{\tau}$ as defined by

$$N_{\text{rel}}(\hat{\tau}) := \frac{\text{\#agents with } x_i \geq \gamma}{n}.$$

We can now show that for the relative number of agents $N_{\text{rel}}(\hat{\tau})$, we have

$$\lim_{n \rightarrow \infty} N_{\text{rel}}(\hat{\tau}) = \Phi\left(\frac{\hat{\tau} - \gamma}{\sigma^2}\right) \tag{1}$$

where Φ is the cumulative distribution function (cdf) of a standard Gaussian, which is illustrated numerically in Fig. 2 and shown in Theorem 2:

Theorem 2 *For the relative number of agents $N_{\text{rel}}(\hat{\tau})$, we have*

$$\lim_{n \rightarrow \infty} N_{\text{rel}}(\hat{\tau}) = \Phi\left(\frac{\hat{\tau} - \gamma}{\sigma^2}\right) \tag{2}$$

where Φ is the cumulative distribution function (cdf) of a standard Gaussian.

Proof Note that $N_{\text{rel}}(\hat{\tau}) = \frac{\sum_{i=1}^n \mathbf{I}_{x_i \geq \gamma}}{n}$ where $\mathbf{I}_{i \geq j}$ is the indicator function for $i \geq j$. For a given $\hat{\tau}$, x_i are IID $\mathcal{N}(\hat{\tau}, \sigma^2)$ random variables and hence, $\mathbf{I}_{x_i \geq \gamma}$ are IID random variables. This follows from the fact that if X_1, \dots, X_n are independent random variables, then $g_1(X_1), \dots, g_n(X_n)$ are independent random variables for a deterministic (measurable) functions $g_i : \mathbb{R} \rightarrow \mathbb{R}$ [cf. Theorem 2.1.6 in Durrett (2010)]. Also, we have

$$E(\mathbf{I}_{x_i \geq \gamma}) = P(x_i \geq \gamma) = \Phi\left(\frac{\hat{\tau} - \gamma}{\sigma^2}\right).$$

Therefore, by the law of large numbers, it follows that:

$$\lim_{n \rightarrow \infty} N_{\text{rel}}(\hat{\tau}) = \Phi\left(\frac{\hat{\tau} - \gamma}{\sigma^2}\right).$$

□




Fig. 2 Visualization of Theorem 2 as N_{rel} estimates $\Phi(\cdot)$. The plot was generated by running Eq. 1 10,000 times for each point in $\hat{t} = 1-10$ in increments of 0.1. $n = 10$, $\Upsilon = 5$ and $x_i = \hat{t} + \eta_i$ ($\eta_i \sim \mathcal{N}(0, \sigma^2)$). Each solid line in the plot is generated by sweeping $\sigma^2 = \{0.1, 1, 2, 10\}$, with $\sigma^2 = 0.1$ being close to a step function and $\sigma^2 = 10$ having the flattest slope. The shaded region provides a difference comparison between the N_{rel} estimate of $\Phi(\cdot)$ and $\Phi(\cdot)$ itself, which is plotted using dotted lines

The final step to explain the prevalence of sigmoid functions in multi-agent settings is to note that:

$$\left| \Phi\left(\frac{\hat{t} - \Upsilon}{\sigma^2}\right) - \frac{1}{1 + e^{-d\left(\frac{\hat{t} - \Upsilon}{\sigma^2}\right)}} \right| \leq 0.01,$$

for all $\hat{t} \in \mathbb{R}$ and some optimal value $d \approx 1.704$ as described in Camilli (1994). This means that the aggregate behavior of the agents following deterministic threshold strategies would closely follow (to within a constant error term) the shape of the commonly observed logistic sigmoid function whose drift is directly proportional to Υ and the slope is inversely proportional to σ^2 .

Therefore, despite agents using deterministic threshold strategies, their aggregate behavior would appear to an outside observer as a continuous sigmoid threshold function instead of numerous discrete thresholds.

6 Discussion

We have formally shown that a simple response threshold method as it is customary in swarm robotics and social insects will lead to a Bayesian Nash equilibrium, without the need for individuals to communicate, but solely requiring perception of a common stimulus signal. This makes this approach simple to implement and is another argument for its ubiquity in nature and engineering.

Although we have shown how a simple threshold-based policy for task allocation can resemble behavior observed in social insects, we caution that all proofs in this paper assume Gaussian distributions for the stimulus parameter and its noisy observations. While these assumptions can easily be accommodated in an engineering system, they do not necessarily hold in a biological context. Studying the true distributions in systems of interest and extending the results presented here for other distributions is therefore subject to further work.

We note that the proposed task allocation mechanism is not optimal in terms of the allocations it can generate, but simply the best strategy for an individual agent given that other agents are using the same policy, which is the definition of a Nash equilibrium. This strategy is therefore only interesting if communication between agents is impossible or otherwise undesirable. In this case, the presented results provide an analysis framework, which makes such an approach viable in an engineered system in which a Bayesian solution is acceptable. Yet, further work is needed to formally show what the lower performance bounds of such a policy are.

The existence of a BNE in a global game is contingent on individuals not communicating with each other. This is immediately obvious as the availability of information to some agents, but not to others, might allow them to improve some global metric by changing only their own policy. As both engineered and natural systems could easily achieve better results by communicating (see related work on optimal task allocation), we interpret probabilistic response threshold task allocation mechanisms to be a suitable baseline strategy. When observed in natural systems, their existence could serve as a clue that individuals do not communicate directly for one or the other reason.

Results from this paper assume that the response threshold is constant, which is not necessarily the case in natural systems or swarm robotics (Castello et al. 2016), where adaptive thresholds can improve performance. Formally extending the methods presented here to include adaptive thresholds is subject of future work.

7 Conclusion

We are studying a class of task allocation problems that require joint action of a certain average number of agents in order to be successful and in which agents do not communicate, but only have access to a common (noisy) stimulus signal. These assumptions are very basic, making this approach accessible to a wide range of platforms.

Due to the limited information available to individuals in such systems, we are employing a probabilistic response threshold algorithm that is prominent in both engineered and natural systems such as social insects that have limited computational and communication abilities. We show that the information requirements of the problem studied in this paper are akin to those encountered by human beings during bank runs and revolutions. This insight allows us to leverage methods and tools from the field of game theory, which describes such a situation as a global game. Building up on existing proofs, we show that such a policy leads to a Bayesian Nash equilibrium, that is, on average, no player has anything to gain by changing only their own strategy. Albeit this result does not imply that such a policy results in an optimal allocation, it provides—for the first time—a formal understanding for why a probabilistic response threshold strategy is actually desirable.

Furthermore, we show how a noisy perception of a global, deterministic threshold signal can be interpreted as a deterministic observation of a sigmoidal threshold function, which are commonly observed or applied in natural or engineered swarms, respectively. Seemingly sigmoidal response functions can be explained by a noisy observation of a deterministic threshold and therefore do not need to be hardcoded in the system. Rather, the resulting balance between exploration and exploitation that is common to swarm systems can be tuned by increasing or decreasing noise in the perception apparatus.

In future work, we are interested in better understanding the relationship between the existence of BNE and various optimality criteria for task allocation. Here, the growing field

of games with imperfect information is a promising direction to possibly show that response threshold functions lead not only to Nash equilibria, but also to globally optimal task allocations in case of communication between individuals either not being permissive or possible.

Acknowledgments A. Kanakia and N. Correll have been supported by NSF CAREER Grant #1150223. We are grateful for this support.

References

- Amstutz, P., Correll, N., & Martinoli, A. (2008). Distributed boundary coverage with a team of networked miniature robots using a robust market-based algorithm. *Annals of Mathematics and Artificial Intelligence*, 52(2–4), 307–333.
- Arslan, G., Marden, J. R., & Shamma, J. S. (2007). Autonomous vehicle–target assignment: A game-theoretical formulation. *Journal of Dynamic Systems, Measurement, and Control*, 129(5), 584–596.
- Berman, S., Halász, Á., Hsieh, M. A., & Kumar, V. (2009). Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics*, 25(4), 927–937.
- Bonabeau, E., Sobkowski, A., Theraulaz, G., & Deneubourg, J.-L. (1997). Adaptive task allocation inspired by a model of division of labor in social insects. In *Biocomputing and emergent computation: Proceedings of BCEC97* (pp. 36–45). World Scientific Press.
- Bonabeau, E., Theraulaz, G., & Deneubourg, J.-L. (1996). Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 263(1376), 1565–1569.
- Bonabeau, E., Theraulaz, G., & Deneubourg, J.-L. (1998). Fixed response thresholds and the regulation of division of labor in insect societies. *Bulletin of Mathematical Biology*, 60(4), 753–807.
- Border, K. C. (1990). *Fixed point theorems with applications to economics and game theory*. Cambridge: Cambridge Books.
- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.
- Camilli, G. (1994). Teachers corner: Origin of the scaling constant $d = 1.7$ in item response theory. *Journal of Educational and Behavioral Statistics*, 19(3), 293–295.
- Carlsson, H., & Van Damme, E. (1993). Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, 61(5), 989–1018.
- Castello, E., Yamamoto, T., Dalla Libera, F., Liu, W., Winfield, A. F., Nakamura, Y., et al. (2016). Adaptive foraging for simulated and real robotic swarms: The dynamical response threshold approach. *Swarm Intelligence*, 10(1), 1–31.
- Chen, J., & Sun, D. (2011). Resource constrained multirobot task allocation based on leader–follower coalition methodology. *The International Journal of Robotics Research*, 30(12), 1423–1434.
- Choi, H.-L., Brunet, L., & How, J. P. (2009). Consensus-based decentralized auctions for robust task allocation. *IEEE Transactions on Robotics*, 25(4), 912–926.
- Conradt, L., & Roper, T. J. (2003). Group decision-making in animals. *Nature*, 421(6919), 155–158.
- Conradt, L., & Roper, T. J. (2005). Consensus decision making in animals. *Trends in Ecology and Evolution*, 20(8), 449–456.
- Correll, N. (2007). Coordination schemes for distributed boundary coverage with a swarm of miniature robots: Synthesis, analysis and experimental validation. PhD thesis, Ecole Polytechnique Fédérale, Lausanne, CH.
- Correll, N. (2008). Parameter estimation and optimal control of swarm-robotic systems: A case study in distributed task allocation. In *IEEE international conference on robotics and automation (ICRA)* (pp. 3302–3307). IEEE.
- Dantu, K., Berman, S., Kate, B., & Nagpal, R. (2012). A comparison of deterministic and stochastic approaches for allocating spatially dependent tasks in micro-aerial vehicle collectives. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 793–800). IEEE.
- Durrett, R. (2010). *Probability: Theory and examples* (4th ed.). Cambridge: Cambridge University Press.
- Fudenberg, D. (1998). *The theory of learning in games* (Vol. 2). Cambridge, MA: MIT Press.
- Gerkey, B. P. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *The International Journal of Robotics Research*, 23(9), 939–954.

- Gerkey, B. P., & Mataric, M. J. (2003). Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. In *IEEE international conference on robotics and automation* (Vol. 3, pp. 3862–3868). IEEE.
- Gordon, D. M. (1996). The organization of work in social insect colonies. *Nature*, *380*(6570), 121–124.
- Grenager, T., Powers, R., & Shoham, Y. (2002). Dispersion games: General definitions and some specific learning results. In *AAAI innovative applications of artificial intelligence conference (IAAI)* (pp. 398–403). AAAI.
- Harsanyi, J. C. (2004). Games with incomplete information played by Bayesian players, I-III Part I. The basic model. *Management Science*, *50*(12–supplement), 1804–1817.
- Kalra, N., & Martinoli, A. (2006). Comparative study of market-based and threshold-based task allocation. In *Distributed autonomous robotic systems 7* (pp. 91–101). Springer.
- Kanakia, A., & Correll, N. (2016). A response threshold sigmoid function model for swarm robot collaboration. *Distributed and autonomous robotic systems (DARS), volume 112 of the series springer tracts in advanced robotics* (pp. 193–206). Heidelberg: Springer.
- Krieger, M. J., Billeter, J.-B., & Keller, L. (2000). Ant-like task allocation and recruitment in cooperative robots. *Nature*, *406*(6799), 992–995.
- Kube, C. R., & Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, *30*(1), 85–101.
- Lerman, K., Galstyan, A., Martinoli, A., & Ijspeert, A. (2001). A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life*, *7*, 375–393.
- Lerman, K., Jones, C., Galstyan, A., & Mataric, M. J. (2006). Analysis of dynamic task allocation in multi-robot systems. *The International Journal of Robotics Research*, *25*(3), 225–241.
- Liu, W., & Winfield, A. (2010). Modelling and optimisation of adaptive foraging in swarm robotic systems. *The International Journal of Robotics Research*, *29*(14), 1743–1760.
- Marden, J. R., Arslan, G., & Shamma, J. S. (2009). Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control*, *54*(2), 208–220.
- Martinoli, A., Easton, K., & Agassounon, W. (2004). Modeling swarm robotic systems: A case study in collaborative distributed manipulation. *The International Journal of Robotics Research*, *23*(4–5), 415–436.
- Martinoli, A., Ijspeert, A. J., & Mondada, F. (1999). Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems*, *29*(1), 51–63.
- Mataric, M. J., Sukhatme, G. S., & Astergaard, E. H. (2003). Multi-robot task allocation in uncertain environments. *Autonomous Robots*, *14*(2–3), 255–263.
- Mather, T. W., Hsieh, M. A., & Frazzoli, E. (2010). Towards dynamic team formation for robot ensembles. In *IEEE international conference on robotics and automation (ICRA)* (pp. 4970–4975). IEEE.
- McEvoy, M., & Correll, N. (2015). Materials that couple sensing, actuation, computation, and communication. *Science*, *347*(6228), 1261689.
- Morris, S. E., & Shin, H. S. (2000). *Global games: Theory and applications*. New Haven, CT: Cowles Foundation for Research in Economics.
- Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic game theory* (Vol. 1). Cambridge: Cambridge University Press.
- Parsons, S., & Wooldridge, M. (2002). Game theory and decision theory in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, *5*(3), 243–254.
- Pini, G., Gagliolo, M., Brutschy, A., Dorigo, M., & Birattari, M. (2013). Task partitioning in a robot swarm: A study on the effect of communication. *Swarm Intelligence*, *7*(2–3), 173–199.
- Pynadath, D. V., & Tambe, M. (2002). Multiagent teamwork: Analyzing the optimality and complexity of key theories and models. In *Proceedings of the first international joint conference on autonomous agents and multiagent systems (AAMAS): Part 2* (pp. 873–880). ACM.
- Raafat, R. M., Chater, N., & Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, *13*(10), 420–428.
- Robinson, G. E. (1987). Modulation of alarm pheromone perception in the honey bee: Evidence for division of labor based on hormonally regulated response thresholds. *Journal of Comparative Physiology A*, *160*(5), 613–619.
- Seeley, T. D. (1989). Social foraging in honey bees: How nectar foragers assess their colony's nutritional status. *Behavioral Ecology and Sociobiology*, *24*(3), 181–199.
- Shehory, O., & Kraus, S. (1998). Methods for task allocation via agent coalition formation. *Artificial Intelligence*, *101*(1), 165–200.
- Suzuki, S., Adachi, R., Dunne, S., Bossaerts, P., & O'Doherty, J. P. (2015). Neural mechanisms underlying human consensus decision-making. *Neuron*, *86*(2), 591–602.

- Theraulaz, G., Bonabeau, E., & Deneubourg, J.-L. (1998). Response threshold reinforcements and division of labour in insect societies. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 265(1393), 327–332.
- Tumer, K., & Wolpert, D. (2004). A survey of collectives. In *Collectives and the design of complex systems* (pp. 1–42). Springer.
- Vig, L., & Adams, J. A. (2007). Coalition formation: From software agents to robots. *Journal of Intelligent and Robotic Systems*, 50(1), 85–118.
- Yoshida, W., Seymour, B., Friston, K. J., & Dolan, R. J. (2010). Neural mechanisms of belief inference during cooperative games. *Journal of Neuroscience*, 30(32), 10744–10751.