

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Collective decision making in dynamic environments

JUDHI PRASETYO, X; De Masi, Giulia; Ferrante, Eliseo

Published in:
Swarm Intelligence

DOI:
[10.1007/s11721-019-00169-8](https://doi.org/10.1007/s11721-019-00169-8)

Publication date:
2019

Document Version
Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for pulished version (HARVARD):

JUDHI PRASETYO, X, De Masi, G & Ferrante, E 2019, 'Collective decision making in dynamic environments', *Swarm Intelligence*, vol. 13, no. 3-4, pp. 217-243. <https://doi.org/10.1007/s11721-019-00169-8>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



Collective decision making in dynamic environments

Judhi Prasetyo¹ · Giulia De Masi² · Eliseo Ferrante³ 

Received: 17 December 2018 / Accepted: 15 June 2019 / Published online: 26 June 2019
© The Author(s) 2019

Abstract

Collective decision making is the ability of individuals to jointly make a decision without any centralized leadership, but only relying on local interactions. A special case is represented by the best-of- n problem, whereby the swarm has to select the best option among a set of n discrete alternatives. In this paper, we perform a thorough study of the best-of- n problem in dynamic environments, in the presence of two options ($n = 2$). Site qualities can be directly measured by agents, and we introduce abrupt changes to these qualities. We introduce two adaptation mechanisms to deal with dynamic site qualities: stubborn agents and spontaneous opinion switching. Using both computer simulations and ordinary differential equation models, we show that: (i) The mere presence of the stubborn agents is enough to achieve adaptability, but increasing its number has detrimental effects on the performance; (ii) the system adaptation increases with increasing swarm size, while it does not depend on agents' density, unless this is below a critical threshold; (iii) the spontaneous switching mechanism can also be used to achieve adaptability to dynamic environments, and its key parameter, the probability of switching, can be used to regulate the trade-off between accuracy and speed of adaptation.

Keywords Dynamic environments · Collective decision making · Best-of- n · Swarm robotics · Complex adaptive systems

1 Introduction

Collective decision making is observed in a wide variety of natural and artificial collective systems (Camazine et al. 2001; Bonabeau et al. 1999). In the context of artificial systems,

✉ Eliseo Ferrante
e.ferrante@bham.ac.uk

Judhi Prasetyo
j.prasetyo@mdx.ac.ae

Giulia De Masi
giuliademasi@gmail.com

¹ Middlesex University Dubai, Dubai, UAE

² Zayed University, Dubai, UAE

³ University of Birmingham, Dubai, UAE

collective decision making can be considered a cornerstone building block for swarm robotics collective behaviors (Brambilla et al. 2013): Many swarm robotics problems such as deciding a common moving direction to move collectively (Ferrante et al. 2012), or a common site in the environment to aggregate at (Correll and Martinoli 2011), can be seen as instances of collective decision making (Valentini et al. 2017). A special case of collective decision making is the best-of- n problem (Valentini et al. 2017), whereby individuals in a swarm need to make a collective decision and commit to an option among n discrete alternatives. Recently, we have thoroughly reviewed the best-of- n problem in a work with Valentini et al. (2017). In that review, we argue that a collective decision-making process can be influenced by two main driving forces: the agent's modulation as a response to different intrinsic qualities associated with the different options (Font Llenas et al. 2018; Valentini et al. 2014, 2015, 2016), or a biasing component due to the environment in cases where options are not symmetrically accessible, meaning options have different costs associated with them (e.g., in terms of time needed to assess them) (Montes de Oca et al. 2011; Scheidler et al. 2016; Brutschy et al. 2012).

In this paper, we consider an instance of the best-of- n problem which falls in the first of the categories presented above. A swarm of robots with minimal capabilities allowing them to interact only locally has to achieve consensus to the option associated with the best quality, among two possible alternatives ($n = 2$). Qualities are assumed to be measurable by the robots, while the environment is symmetric with respect to the distribution of the n options, meaning that all options can be evaluated on average in the same amount of time. The robots are not able to communicate the option quality. They can only advertise one option at a time, the one corresponding to their *current opinion*, and they use a *decision mechanism* to change their current opinion after observing their neighbors in local proximity. The most famous decision mechanisms used in the swarm robotics literature are the voter model (Baronchelli and Díaz-Guilera 2012; Valentini et al. 2014) and the majority rule (Montes de Oca et al. 2011). The swarm builds consensus over time via positive feedback modulation (Garnier et al. 2007), whereby fluctuations in opinion distribution will eventually produce a bias toward one of the n options, which will make that option more likely to be observed and henceforth reinforcing this bias, until consensus is reached.

The majority of the research efforts on the best-of- n for a mobile robot swarm has been put in the static environment case, whereby the environment and option qualities do not change over time, with few exceptions (Valentini et al. 2017). However, a number of real life problems may violate this condition: for example, situations where physical barriers may come to exist during a natural disaster or a sudden weather change, preventing or delaying robot navigation, or situations where resources may deplete due to the action of the robots themselves or of external agency.

In this paper, we extend our work on the best-of- n in the dynamic environment first introduced by Prasetyo et al. (2018b). As in that preliminary study, we consider a problem where the environment is static and symmetric with respect to option distribution, but the quality is asymmetric and abruptly changes at a given moment in time. In particular, the environmental change is modeled by swapping the quality of the two options, a choice that allows us to model abrupt changes while keeping constant the quality ratio between the two options. The goal of the swarm is to collectively chase the best option: The swarm must achieve consensus to the option corresponding to the best quality and change the consensus state when the best option changes. We consider the voter model as the main decision mechanism. We couple it with the positive feedback modulation mechanism first proposed by Valentini et al. (2014), whereby robots advertise their opinion to other robots within their range of sight for a time that is proportional, on average, to the quality of the option corresponding to their opinion, an idea that is inspired by honeybees waggle dance behavior (Seeley 2010).

We use multi-agent simulations where the spatial dimension is taken into account, and each robot is abstracted by an agent. We extend the study done by Prasetyo et al. (2018b) in several directions. First, we consider two mechanisms (rather than one) to tackle the problem. The first one, already introduced by Prasetyo et al. (2018b), considers *stubborn* agents, that is, agents that never change their opinion. Stubborn agents can be seen as scouts, constantly exploring their favorite opinion, irrespective of the opinion of others and of the consensus state of the swarm. The second mechanism, that is new to this paper, is the spontaneous opinion switching: After applying the decision rule, each agent in the swarm has a small probability to randomly switch its opinion to a different one. The spontaneous switching mechanism acts as negative feedback and again has the effect to add exploratory capabilities to the swarm, analogously to the abandonment component already seen in different collective decision-making models (Reina et al. 2015b, 2017). We study the first mechanism in more detail compared to the work of Prasetyo et al. (2018b): After confirming the effect of swarm size, of the proportion of stubborn individuals, and of the ratio between the option qualities in larger swarms than those considered by Prasetyo et al. (2018b), we now also confirm that the decision making is indeed affected by the swarm size alone and not by increased agent density, as we find that the agent density does not play a role unless it is below a critical threshold. We also discover that in large swarms and when options are difficult to discern, increasing the overall number of stubborn individuals has a detrimental effect. Therefore, we perform a study in which the number of stubborn individuals is kept fixed and does not change with the swarm size. Additionally, we study the new mechanism, the spontaneous opinion switching mechanism, with respect to its key parameter—the switching probability—and with respect to the swarm size.

As an additional novel contribution, our simulation results are complemented with a study performed with two ordinary differential equations models (ODEs). Both models are extensions of the ODEs defined by Valentini et al. (2014): In the first, we introduce new equations and state variables to represent subpopulations of stubborn individuals, while in the second, we extend the previous equations in order to include the spontaneous switching parameter. We study the asymptotic stability of both models with respect to their characteristic parameter (the proportion of stubborn individuals in the first, and the switching probability in the second), we qualitatively validate their prediction with respect to the simulation results, and we analyze more in general how the characteristic parameter affects the collective decision-making dynamics.

The rest of the paper is organized as follows. In Sect. 2, we relate our work to the collective decision-making literature. In Sect. 3, we introduce the dynamic best-of- n problem, the collective decision-making method, the definition of stubborn individuals, and the spontaneous opinion switching mechanism. In Sect. 4, we present our experimental setup in terms of environment and parameter settings that have been studied, and the metric of evaluations. In Sect. 5, we present the results. In Sect. 6, we present the mathematical model, its study, and its validation against simulation results. Finally, in Sect. 7, we give a conclusion and discuss our future research agenda on this topic.

2 Related work

The best-of- n problem and the particular scenario we consider have biological inspiration coming from the collective behaviors of social insects such as ants (Franks et al. 2002) and more specifically bees (Marshall et al. 2009; Seeley 2010). We review the literature

on the best-of- n problem in swarm robotics by considering the two categories introduced by Valentini et al. (2017). We also analyze some work done on the best-of- n in settings that can be considered as a dynamic environments, and we discuss work related to the notions of stubborn individuals and spontaneous opinion switching.

In the first category, we place work whereby the quality of the different options cannot be measured directly by the robots. Instead, asymmetries in the environment can bias the collective decision toward one of the n options. For example, Garnier et al. (2009) and Campo et al. (2010) presented a classical aggregation task inspired by cockroaches. In this work, size differences in aggregation sites induce asymmetries in the environment; however, robots do not have the ability to discern the sites. Thanks to these asymmetries, robots are able to aggregate only in one shelter, which in the study by Campo et al. (2010) is a specific one (the one that has the right size to host all the robots, but not bigger). Another example of environmental asymmetry is shown in the work by Montes de Oca et al. (2011), Valentini et al. (2013) and Brutschy et al. (2012), whereby robots move in a classical double-bridge environment (Deneubourg and Goss 1989) and have to find the shortest path between two bridges connecting the nest to the food source. The asymmetries between the two paths induce agents to select the shortest path to appear more frequently in the nest, and therefore, biasing the process toward that path. Montes de Oca et al. (2011) used the majority rule as decision mechanism, whereby agents switch opinion to the opinion held by the majority of a group of neighbors with predefined size. In a subsequent work, Scheidler et al. (2016) studied the same scenario but applied another mechanism called the k -unanimity rule: The agent switches opinion only after observing the same option k times in a row, where each time the agent observes the opinion of a random neighbor.

In the second category, we place work in which the quality can be directly measured, as per our case. The baseline studies on direct modulation of positive feedback through quality were performed by Valentini et al. (2014), Valentini et al. (2015), and Valentini et al. (2016). In this articles, the authors thoroughly analyzed the voter model and the majority rule through real-robot experiments, simulations, ordinary differential equations, and chemical reaction network models and studied the speed versus accuracy trade-off. Reina et al. (2015a), Reina et al. (2015b), and Reina et al. (2017) developed a decision-making strategy that, differently from our work, includes also an uncommitted opinion (neither of the n alternatives), a recruitment mechanism, an inhibition mechanism (as in honeybees studied by Seeley et al. (2012)), and an abandonment or decay mechanism, which is analogous to our spontaneous opinion switching. Experiments using real robots with this mechanism have been done by Reina et al. (2015a, 2018a). In a recent follow-up study by Reina et al. (2018b), they have shown how this model can be generalized to encompass not only decision making in social insects but also in the human brain (Marshall et al. 2009). Finally, Parker and Zhang (2009) considered the best-of- n problem in an aggregation task, whereby agents use a direct recruitment mechanism and are able to commit by using a quorum-based mechanism that makes the swarm aware of the consensus level reached.

In the context of dynamic environment, relatively little research has been done. Among the exceptions, Parker and Zhang (2010) considered a task-sequencing problem that can be seen as a best-of-2 with two options: “task complete” and “task incomplete.” The two options have dynamic qualities because the task completion level changes over time. Arvin et al. (2014) studied a dynamic version of aggregation. Here, each shelter emits a different sound that varies over time, and the swarm has to aggregate in the shelter with the loudest sound. The method is based on a fuzzy version of the original BEECLUST algorithm (Kernbach et al. 2009; Schmickl et al. 2009). In the original BEECLUST, after a waiting period, each agent chooses a new direction of motion at random, while Arvin et al. (2014) use a fuzzy

controller that maps the loudness and the bearing of the sound to the new direction of motion. Differently from all these works that focused on specific application scenarios, in this paper, we perform a systematic study of a minimal model of the dynamic best-of- n problem, in order to understand better the effect of the most important parameters.

The idea of having the swarm not converging to a full unanimity when seeking consensus is not new to this paper. For example, biological studies have found that having only a large majority committing to an option rather than the unanimity allows fish schools to swiftly adapt to perturbations (Calovi et al. 2015). Stubborn individuals and spontaneous opinion switching are two ways to achieve this. Concerning studies on stubborn individuals in a population, an increased interest is emerging in social dynamics literature. While the introduction of stubborn individuals can be a way to increase the realism of opinion dynamics models applied to social systems, the topic is nowadays more and more relevant for national security issues, such as the risk of election and referendum manipulation as reported in the USA and in Europe. For example, Hunter and Zaman (2018) showed that only few stubborn individuals can strongly impact the overall opinion of other agents. They study also the role of different placements of stubborn individuals to maximally shift the average opinion of the others. Mukhopadhyay and Mazumdar (2016) showed that with the majority rule, the presence of stubborn individuals introduces metastability, that is, fluctuations between different equilibrium points. Also according to a study by Yildiz et al. (2013), the presence of stubborn individuals prevents the formation of consensus, introducing instabilities and fluctuations. While the presence and role of stubborn individuals have been confirmed and evaluated in groups of humans, it is much harder to find evidence of such individuals in social insects. A recent paper has detected “contrarian effects” in a collective decision-making system where the well-mixed assumption fails due to spatial correlations (Hamann 2018). Those effects could potentially be similar to those exhibited by stubborn individuals.

Mechanisms analogous to spontaneous opinion switching have been studied by many authors, such as Pratt et al. (2002) and Britton et al. (2002). Marshall et al. (2009) provided an interesting discussion across collective decision-making models, pointing out that different theoretical models of collective decision making in social insects include two main types of switching: indirect switching, whereby agents committed to an option spontaneously become uncommitted before recommitting to another option, and direct switching between two options, which can only occur through recruitment and therefore cannot be spontaneous. Therefore, to the best of our knowledge, both stubborn and spontaneous opinion switching mechanism in the way considered in this paper are not featured in social insects and therefore can be seen as engineering mechanisms to tackle the best-of- n problem in artificial collective systems.

3 The model

In this section, we define the dynamic best-of- n problem (Sect. 3.1) and the collective decision making model (Sect. 3.2).

3.1 The dynamic best-of- n problem

The best-of- n problem requires a swarm of agents to make a collective decision among n possible alternatives toward the choice that has the best quality. A typical example is the choice of best location for honeybees’ swarm foraging. Each of the n options has an intrinsic

quality ρ_i with $i \in 1, \dots, n$. A best-of- n problem reaches the optimal solution when the collective decision for a swarm composed of N individuals is for the option with maximum quality. That means that a large majority $M \geq N(1 - \delta)$ of agents agrees on the same option, where δ is a small number chosen by the designer. In the case where $\delta = 0$, there is *perfect consensus* or *unanimity*.

In this paper, as for the majority of the studies (Valentini et al. 2016), we restrict n to 2 options, labeled A and B , having intrinsic quality ρ_A and ρ_B . To reduce the number of parameters to study, one option quality ρ_a is set to 1 while $\rho_b > 1$. No cost is included in the current model, which means that the time needed to explore and assess the quality of both options is symmetric (Valentini et al. 2017). Each agent can measure the quality of different options and can only advertise that option using local communication (see Sect. 3.2). In dynamic environments, qualities can change over time: $\rho_A = \rho_A(t)$ and $\rho_B = \rho_B(t)$. In this study, we only consider qualities that are piece-wise constant: At a given time T_C , the two qualities are swapped. Namely $\rho_A(t)$ and $\rho_B(t)$ remain constant for $t < T_C$, they are swapped at T_C ($\rho_A(T_C) = \rho_B(T_C - 1)$, $\rho_B(T_C) = \rho_A(T_C - 1)$), and again remain constant afterward.

3.2 The decision mechanism in its vanilla form

The agents controller is represented by the finite-state machine in Fig. 1a. Accordingly, agents can have one of the following 4 possible states: dissemination state of opinion A (d_A), dissemination state of opinion B (d_B), exploration state of opinion A (e_A), exploration state of opinion B (e_B). In Fig. 1a, solid lines represent deterministic transitions, while dotted lines stochastic transitions. The symbol VM indicates that the voter model is used at the end of the dissemination state.

As initial conditions, agents are initialized inside the nest. Half of the agents are initialized with the e_A state, the other half with the e_B state, and they move toward the site associated with their opinion to explore that option. Once they reach the site, they explore it for an exponentially distributed amount of time (sampled independently per agent) that does not depend on the option or option quality. During this time, agents measure the quality of that site. Subsequently, they switch to the dissemination state associated with their current opinion (d_A if they were in e_A , d_B if they were in e_B), travel back to the nest, each at a different time




Fig. 1 **a** Probabilistic finite-state machine. d_A , d_B , e_A , and e_B represent the dissemination and exploration state. Solid lines represent deterministic transitions, while dotted lines stochastic transitions. The symbol VM indicates that the voter model is used at the end of the dissemination state. Note that stochastic transition may be the results of either the application of the decision rule, or of the spontaneous opinion switching mechanism, if enabled. **b** Screenshot of the simulation arena. This image is taken from NetLogo software

due to independent sampling, where they initiate opinion dissemination. While at the nest, we aim at having agents that are well mixed with respect to their opinion and to which site they come from, to avoid agents with same opinion clustering near each other and create spatial correlations (Hamann 2018). To meet this criterion as much as possible, agents perform a correlated random walk while disseminating and before applying the decision mechanism.

In the dissemination state, each agent locally broadcasts his opinion continuously, and this message is sensed by other agents that are also in the dissemination state and situated within a limited range from the broadcasting agent. The time spent by the agent disseminating its opinion is randomly sampled from an exponential distribution characterized by a parameter proportional to site quality they have last visited. As a consequence, it is more probable to meet neighbors with the best opinion than meeting those with the worst one, because the former will disseminate longer than the latter. This mechanism is called *modulation of positive feedback*, and it is the driving mechanism to make the group converge on the option with the best quality. At the end of dissemination, each agent can change its opinion based on the opinions of other agents and using the voter model. The result of the voter model depends on the neighbors' opinion, that is, the agents within a specified spatial radius (in our simulations set to 10 units): The agent switches its opinion to the one of a random neighbors within the interaction radius.

In the following, we explain the two mechanisms we introduced in order to tackle the dynamic version of the best-of- n problem: stubborn agents and the spontaneous opinion switching mechanism.

The stubborn agents In simulations with stubborn agents, we consider two kinds of agents: *normal* and *stubborn*. Each agent has an initial opinion, which consists in one of the two options A or B . Normal agents are able to change their opinion by applying a decision mechanism that relies on the observation of other agents in local proximity. Stubborn agents instead never change their opinion and keep the one they have at the very beginning, either A or B . In Sect. 5, we will show the effect of introducing a number of stubborn individuals in the swarm that can either scale with the swarm size or remain fixed and independent of it. *The spontaneous opinion switching mechanism* Spontaneous opinion switching is an alternative mechanism to the one represented by stubborn agents. Here, every agent is considered as *normal* (i.e., not *stubborn*), in the sense that every agent is allowed to change its opinion using one of the decision rules. However, right after applying the decision rule, each agent can spontaneously change its opinion: With a probability p , an agent will switch to B if its opinion after the application of the decision rule was A , and to A if its opinion after the application of the decision rule was B . With probability $1 - p$, the agent will keep the opinion resulting from the application of the decision rule. After this opinion has been determined (either via switching or not), the agent will transition to the corresponding exploration e_A or e_B as normal.

4 Experimental setup

We conducted systematic simulations using the simulator developed by Valentini et al. (2016). Agents move on a two-dimensional arena. Space is explicitly modeled, but collisions between agents are not taken into account: Despite this, our previous study (Valentini et al. 2016) showed that these types of simulations can reproduce real-robot collective decision-making dynamics quite well. We considered two types of simulations: with variable and with constant agent density, measured as agents per square units. In simulations with variable density, the

Table 1 Model parameters used in simulations

Parameter	Description	Values
N	Swarm size	{100; 1000; 10,000}
ρ_A (ρ_B after change)	Site A (B) quality	1
ρ_B (ρ_A after change)	Site B (A) quality	{1.05, 3}
x_S	Prop. of stubborn individuals	{0.05, 0.2}
S	Num. of stubborn individuals	10
p	Switching probability	{0.0001, 0.001, 0.005, 0.01, 0.02}
D	Density (agents per square unit)	$\{5 \times 10^{-i}\}, i \in [1, 7]$

arena size is kept fixed to a nominal size of 200 (width) $\times$ 100 (height) units, while the swarm size is varied. In simulations with constant agent density, the arena size was rescaled when the swarm size varied in order to meet the target agent density. We considered the following agent densities that varied across 6 scales $D \in \{5 \times 10^{-i}\}, i \in [1, 6]$ (i.e., D varied from 0.5 to 0.0000005 in the $\log_{10}$ scale). When only one density was studied, we considered the nominal density $D = 0.005$. Figure 1 depicts a screenshot within NetLogo that was used only for fast prototyping and visualization. The arena comprises a central region called the *nest*, where we initialize all agents and where they subsequently meet to perform the decision-making process. The two external areas are the *sites* and represent the two options: option A on the left and option B on the right.

In order to test the robustness of the model, some key parameters have been studied. As evident from Table 1, we study three different values for swarm sizes: 100, 1000, 10,000. Without loss of generality, the interplay between ρ_A and ρ_B can be studied simply by keeping one of them fixed (ρ_A before the environment changes, and ρ_B after it changes) to a value of 1 and by changing the other one. The values of the second quality are 1.05 and 3, indicating small and large difference in quality, respectively. To study the effect of stubborn individuals, we considered two cases: fixed proportion of stubborn individuals, indicated with x_S , and fixed number of stubborn individuals, indicated with S : In the first case, the number of stubborn individuals scales up with the swarm size N , while in the second case, it is kept fixed and independent of N . We considered $x_S \in \{0.05, 0.2\}$ and $S = 10$, and in both cases, stubborn individuals are equally distributed between the two opinions. Finally, when studying the new mechanism based on probabilistic switching, we studied a wide range of values for the parameter $p \in \{0.0001, 0.001, 0.005, 0.01, 0.02\}$. As initial conditions of each run, $\frac{N}{2}$ agents are initialized with opinion A and $\frac{N}{2}$ agents are initialized with opinion B .

The dissemination time is exponentially distributed. The parameter of the distribution is $\tau_D = g \cdot \rho_i, i \in \{A, B\}$ with $g = 100$. The time of exploration is also exponentially distributed, with parameter set to $\tau_E = 10$, therefore independent of the site. These stochastic times have been modeled through exponential distributions because their lack of memory enhances the predictability of mathematical models (Valentini et al. 2014), such as the one we introduce in Sect. 3. The main fundamental difference between the dissemination and the exploration time is that the former is a design parameter which needs to be chosen to achieve a good trade-off between accuracy and speed (Valentini et al. 2016), while the latter depends on the experimental conditions. The value chosen in this paper is consistent with those used in the previous study on the voter model (Valentini et al. 2014).

The total duration of one simulation run is $T = 40,000$ simulated seconds. In the dynamic environment considered in this paper, a new time parameter T_C is introduced: the time when the values of ρ_A and ρ_B are abruptly changed by swapping their values. In this study $T_C = 12,000$, a value empirically chosen as a compromise between reaching consensus to the best option prior to change and reasonably short runs, in the most challenging settings in terms of speed (large swarms and low-quality ratios). For each configuration of parameters, an ensemble of simulation has been realized, consisting of $R = 50$ runs.

For the study in Sect. 5.4, we also calculate two metrics that we formally define here. As a first metric, in order to evaluate the accuracy of the decision-making process, we calculate the square root of the mean square error ($\sqrt{\text{MSE}}$) of the process as:

$$\sqrt{\text{MSE}} = \sqrt{\frac{1}{R} \sum_{i=1}^R (\hat{x}_A - x_{A,i})^2},$$

where $\hat{x}_A$ is the target value of the consensus state, which is equal to 0 before T_C (where B is the best opinion) and to 1 after T_C (where A becomes the best opinion), and $x_{A,i}$ is the proportion of agents with opinion A in run i . The square-root operator is applied in order to bring the error measure to a scale that can be easily related to the original scale of the x_A quantity. As a second metric, in order to evaluate the quality of the response to the environmental change, we calculate the standard deviation of the response time of the system to the change. To do this, we first determine what is the time at which the system switches opinion, $T_{s,r}$ for each run $r \in R$: $T_{s,r}$ is set to the last time at which the average opinion x_A crossed the value 0.5 while increasing, or it is set to $T = 40,000$, the highest possible value, in case the system did not converge to opinion A , which is the best option after the environmental change. Once $T_{s,r}$ is determined for each run, the metric of interest is the standard deviation of $T_{s,r}$ across the R runs.

5 Results

We analyze the different parameter configurations by reporting the temporal evolution of opinions. Only the proportion of agents with opinion A (x_A) is reported, as the percentage of agents with opinion B (x_B) is simply given by $x_B = 1 - x_A$. These plots report all the trajectories of x_A over time (in simulated seconds, sampled every $\Delta t = 0.1$ steps) for all runs. We report in the main text only the plots that are most relevant for our discussion. The full set of results is available as Supplementary Material (Prasetyo et al. 2018a).

5.1 Preliminary analysis on the effect of swarm size and of the proportion of stubborn individuals

We start our analysis by summarizing the results that were obtained in our earlier study (Prasetyo et al. 2018b). The first outcome of this study was that the vanilla voter model without stubborn individual produced consensus dynamics that did not adapt to dynamic environments. We reproduced these results in Fig. 2, where we compare two different values of quality ratio: 1.05 (low) and 3 (high). For lower values of quality ratio, we observed low convergence speed and the consensus state was one of the two sites at random: The agents did not have the capability to discern between the two opinions. For high value of quality




Fig. 2 Opinion evolution for a voter model with no stubborn with $N = 100$, for two different values of quality ratio: 1.05 (a) and 3 (b). For low quality ratio there is no convergence. For high quality ratio the convergence to one option is reached but there is no adaptation to the change of opinion quality




Fig. 3 Different cases of systems of $N = 100$ agents. **a** $S = 5\%$ and $\rho_A/\rho_B = 1.05$, **b** $S = 20\%$ and $\rho_A/\rho_B = 1.05$, **c** $S = 5\%$ and $\rho_A/\rho_B = 3$, and **d** $S = 20\%$ and $\rho_A/\rho_B = 3$. It shows that quality ratio has a stronger effect than the percentage of stubborn

ratio, the swarm converged to whichever site had the optimal value at the beginning of the simulation run, and agents could not adapt to changes in the environment.

Figure 3 reports the results of runs for four different cases of systems of 100 agents, as shown by Prasetyo et al. (2018b), but with new runs that lasted four times longer than in the original study. Across rows, we vary the ratio ρ_A/ρ_B from low (1.05) to high (3). Across columns, we vary the stubborn percentage from 5% ($x_S = 0.05$) to 20% ($x_S = 0.2$). As in our previous work (Prasetyo et al. 2018b), we obtain that the mere presence of stubborn




Fig. 4 The effect of the swarm size 1000 and 10,000 with $S = 0.05$ for the two quality differences 1.05 and 3: **a** $N = 1000$ and $\rho_A/\rho_B = 1.05$, **b** $N = 10,000$ and $\rho_A/\rho_B = 1.05$, **c** $N = 1000$ and $\rho_A/\rho_B = 3$, and **d** $N = 10,000$ and $\rho_A/\rho_B = 3$. In the case of low quality ratio, increasing the size of the swarm shows a certain tendency to convergence. In the other case (high quality ratio), increasing the swarm size reduces the variance of adaptation time

individuals is enough to achieve adaptability when the quality ratio ρ_A/ρ_B is high, while the proportion of stubborn individuals does not play a significant role for smaller swarms, by only affecting the final value of the consensus state in a way that is decreasingly proportional to the proportion of stubborn individuals employed. In the case where the quality ratio is low, convergence of opinions and adaptation are very poor.

We analyzed the effect of the swarm size in our previous study (Prasetyo et al. 2018b); however, the largest swarm considered in that study was $N = 500$. Here, we perform scalability analysis up to $N = 10,000$ and we consider also longer runs. Keeping constant the percentage of stubborn individuals, the big role of the swarm size is disclosed in Fig. 4. (The quality ratio varies across rows, while the swarm size varies across columns). This figure should be analyzed by also comparing it with the first column in Fig. 3. The three Fig. 5b–d show three swarm sizes: $N = 100$, $N = 1000$, and $N = 10,000$. Increasing the population size decreases the variance of fraction of agents following a certain opinion (here A), while the convergence or non-convergence is determined by the value of the quality ratio. In the case of low quality ratio, the decrease in variance allows us to see a certain pattern of convergence; however, the final value of the convergence state seems too far from the ideal one ($x_A = 1$ or $x_A = 0$). In principle, the presence of stubborn individuals has the natural effect of modifying the consensus state, as the highest (respectively, the lowest) possible consensus state is $x_A = 1$ (respectively, $x_A = 0$) minus the proportion of stubborn individuals divided by 2, which correspond to the individuals committed to the other option that do not contribute

to the consensus. However, in Fig. 4a, b, we observe that the deviation from the consensus state is much larger than that. This fact will be investigated in details in Sect. 5.2.

We conclude this section with an analysis of the response times, that is, the time the system takes to adapt to the environmental changes (for how it is estimated refer to Sect. 4). In Fig. 6, we report the distribution of response times as a function of the swarm size (Fig. 6a) and of the proportion of stubborn individuals (Fig. 6b). As we can see, larger swarm sizes result in larger response times, which is to be expected as larger swarms take longer to reach consensus (Valentini et al. 2014; Montes de Oca et al. 2011). Additionally, increased proportion of stubborn individuals x_S has an effect to reduce the response times; however, this effect is nonlinear and quickly saturates. We will further analyze response times more in general in Sect. 6 while studying the ODE model.

5.2 Results with fixed number of stubborn individuals

Here, we analyze why in large swarms and with low quality ratio, the swarm achieves consensus and adaptation with a deviation from the ideal consensus state that is much larger than what can be produced by the stubborn individuals alone. For example, in Fig. 4a, b, we considered a swarm of 1000 and 10,000 individuals, with only 5% of the individuals stubborn: Here, the deviation from the consensus state is above 0.2, which is ten times larger than the expected deviation of approximately 0.025 (because 2.5% of the individuals are stubborn to the opinion opposite to the one of the consensus state reached at any point). This “ideal” deviation from the consensus state is indeed observed when the quality difference is high (e.g., in Fig. 4c, d). We hypothesize that when the quality ratio is low, increasing the overall number of stubborn individuals has a detrimental effect, and this is why this is especially noticeable in larger swarms. This hypothesis is further supported when we increase the percentage of stubborn agents even further to 20%, whose results we report in Fig. 5a for a large swarm of 10,000 and low quality difference $\rho_A/\rho_B = 1.05$. Here, the convergence dynamics are almost entirely flat, with both consensus states very close to $x_A = 0.5$.

So far, we have therefore evidence that stubborn individuals are needed to achieve adaptability but that larger numbers have either no effect or have a detrimental effect. It appears therefore that stubborn individuals have to be included in the swarm, but their number has to be kept to a bare minimum. To confirm whether this hypothesis is supported, we have decided to run another set of simulations with stubborn individuals, by keeping their number fixed to 10 individuals in total, 5 per option, independently of the swarm size. The other parameters have been fixed as follows: We considered swarms of 100, 1000, and 10,000 individuals, both with low ($\rho_A/\rho_B = 1.05$) and high ($\rho_A/\rho_B = 3$) quality ratio.

Results are shown in panels (b-f) of Fig. 5. As we can see, results with a small constant number of stubborn individuals are very good. (This is also supported by the complete analysis which is available at our supplementary material webpage (Prasetyo et al. 2018a).) For high quality difference (Fig. 5e, f), the mechanism still performs very well in terms of consensus dynamics and adaptation. For low quality difference, a small constant number of stubborn individuals achieves good levels or a small constant number of stubborn individuals corresponds to good levels of consensus and adaptation as long as the swarm size is large enough: In our case, the system does not converge for $N = 100$ (Fig. 5b), but converges and adapts well for $N = 1000$ (Fig. 5c) and $N = 10,000$ (Fig. 5d).




Fig. 5 Simulations with fixed number of stubborn individuals. **a** Result with 10,000 agents, low quality ratio $\rho_A/\rho_B = 1.05$, and high percentage of stubborn individuals $S = 20\%$. **b–d** Results obtained by using only 5 stubborn individuals per site and low quality ratio, over different swarm sizes: **b** 100, **c** 1000, and **d** 10,000. The last row shows results obtained with high quality ratio $\rho_A/\rho_B = 3$ with **e** small $N = 100$ swarms and **f** large $N = 10,000$. Using only *ten* stubborn individuals produces the best results in all these settings

5.3 Disentangling the effect of swarm size and density

In our previous work (Prasetyo et al. 2018b), we were not conclusive in determining whether performance increased merely as a result of increased swarm size alone or not. This is because the arena size was kept fixed, and therefore, it could have been that the swarm density, rather than the swarm size, played a role in improving the consensus dynamics. Here, we shed light on this issue. Figure 7 shows what happens if we keep the density fixed. The swarm size varies across rows: $N = 100$ in the first row and $N = 10,000$ in the second row. The density




Fig. 6 Convergence times as a function of **a** swarm size N (fixed 5 stubborn individuals per site) and **b** proportion of stubborn individuals x_S ($N = 10,000$). In both figures, $\rho_A/\rho_B = 3$. The calculation of the response time is explained in Sect. 4




Fig. 7 We disambiguate the effect of swarm size and density. In all four plots, $S = 10$ (5 per side). **a** $D = 0.005$ and $N = 100$, **b** $D = 0.5$ and $N = 100$, **c** $D = 0.005$ and $N = 10,000$, and **d** $D = 0.5$ and $N = 10,000$. For the values chosen, the swarm size only and not the density has an effect on the dynamics

varies across columns: In the left column, the density is fixed at $D = 0.005$ agents per square unit (this is the density that 100 agents had in the original 200×100 arena, while for 10,000 agents, we now consider an arena that is 100 times bigger), while in the right column, the density is fixed at $D = 0.5$ agents per square unit. (This is the density that 10,000 agents had in the original 200×100 arena, while for 100 agents, we now consider an arena that is 100 times smaller.) The number of stubborn agents is 5 per site. As we can see, we can clearly state that the swarm size has an effect on consensus dynamics, while the density seems to have no effect, at least in the range considered so far.




Fig. 8 Results obtained with large swarms ($N = 10,000$), low quality ratio ($\rho_A/\rho_B = 1.05$, and fixed number of stubborn individuals ($S = 10$, 5 per site) in very low density environments. **a** Reports the distribution of the $N0$ statistics for different values of density. The remaining three panels show the time dynamics with **b** $D = 5 \times 10^{-4}$, **c** $D = 5 \times 10^{-5}$, and **d** $D = 5 \times 10^{-6}$. The systems dynamics are affected by density only from $D = 5 \times 10^{-5}$ onward

We found the result above surprising and we decided to investigate further. We hypothesized that the reason why dynamics are not affected by densities, even if they are different by two orders of magnitude, could be the intrinsic resilience of the voter model to density changes. In fact, while using the voter model, each agent needs to interact only with one other agent. Therefore, one agent in range of sight at each application of the voter model is enough for the dynamics to be unperturbed, as this agent will be a random agent when we assume a well-mixed distribution of agents in the nest. This hypothesis therefore suggests that if we progressively decrease the swarm density even more, we will eventually encounter a situation where agents are no longer guaranteed to have at least one neighbor when applying the voter mechanism. To investigate this, we performed simulations with even lower values of densities ($D \in \{5 \times 10^{-i}\}, i \in [1, 6]$), for both swarm sizes $N = 100$ and $N = 10,000$, and we considered a new statistics that is the $N0$ = “number of times agents do not interact with any other agent when applying the voter model.” Results are shown in Fig. 8. Figure 8a shows the violin plot of the distribution of the $N0$ statistics for the different values of densities. The first thing we notice is that $N0$ is always 0 at all times (confirmed also by inspecting the data) when $D = 0.05$ and $D = 0.5$. It then only starts to assume values greater than 0 with $D = 0.0005$. However, having about *ten* failed applications of the voter model per time-step in a swarm of $N = 10,000$ seems to be negligible, and as we saw, it does not affect the time dynamics. The dynamics start to be severely affected only at very low densities, that is,




Fig. 9 Spontaneous opinion switching mechanism. Dynamics in **a** medium ($N = 1000$ agents) and **b–d** large ($N = 10,000$) systems for several values of p ($p = 0.0001$ in the first row, $p = 0.001$ in (c), and $p = 0.02$ in (d)). Depending on the parameters, we observe either randomly delayed response (a), good response and good accuracy (b, c), or good response and low accuracy (d)

roughly for density equal and below 5×10^{-5} (see Fig. 8c, d). For these densities, also the $N0$ statistics undergo a significant qualitative and quantitative change, with its distribution increasing about *three* times or more in terms of median and even more in terms of spread.

As a summary, we have demonstrated that the introduction of a fixed and low number of stubborn individuals can achieve adaptation to dynamic environments in different environmental settings: small to large swarms, small to large difference in quality, and different swarm densities except for very low values.

5.4 Results with the spontaneous opinion switching

As an alternative to the introduction of stubborn individuals, we introduced in Sect. 3 and we study here the spontaneous opinion switching mechanism, whereby the swarm is composed of homogeneous individuals where each of them has a probability p to switch opinion, after and independently of the application of the decision rule.

In Fig. 9, we report the time dynamics of four interesting cases, one with medium swarm size ($N = 1000$) and three with large swarm size ($N = 10,000$), all executed with low quality ratio $\rho_A/\rho_B = 1.05$. These simulations are all executed with constant density $D = 0.05$. We observe interesting dependency of the system from both parameters, swarm size and probability. Smaller swarm sizes with smaller values of p (see Fig. 9a) exhibit randomly delayed




Fig. 10 Spontaneous opinion switching mechanism. In **a**, a heatmap shows the result of a systematic study as a function of p and N evaluating the square root of the mean square error ($\sqrt{MSE}$) between the best opinion and the average swarm opinion across the different runs. In **b**, a heatmap shows the result of the same systematic study but evaluating the standard deviation between the times at which the swarm switches opinion, across the different runs. In the second row, the MSE is plotted over time for different values of p and for the 10 stubborn individuals case, **c** for $N = 1000$ and **d** for $N = 10,000$

switching dynamics, whereby the system switches its consensus state but with a response time that is delayed with respect to when the environment changes, and furthermore this delay has a high standard deviation. Interestingly, in large systems with otherwise identical parameters, the system still exhibits variation in the response but this time with a much smaller standard deviation (see Fig. 9b). This trend is confirmed when analyzing smaller systems, in which the standard deviation of the response time is even higher than with $N = 1000$ (results available in our supplementary materials page (Prasetyo et al. 2018a)). The standard deviation in the response time can be lowered by increasing the value of p . Figure 9c shows the results obtained with the same swarm in Fig. 9b but with $p = 0.001$. Here, we observe a quite ideal response, comparable with the one we had obtained with *ten* stubborn individuals in Sect. 5.2. If we increase p even further, we observe that now the consensus states move toward 0.5 and away from the ideal states 0 and 1, analogously to what we observed with higher numbers of stubborn individuals in Sect. 5.2.

We decided to do a systematic analysis to confirm or deny the trends identified above. We launched 50 simulation runs for the following parameter configurations: $N \times p \in \{40; 100; 1000; 10,000\} \times \{0.0001, 0.001, 0.005, 0.01, 0.02\}$ (i.e., we executed all combinations between these listed values of N and p). Results are shown in Fig. 10. In the first row, in Fig. 10a, we report a heatmap showing the value of the square root of the mean square error ($\sqrt{MSE}$) between the ideal consensus state and x_A , while in Fig. 10b, we report a heatmap of the standard deviation of the response times. (Both metrics are defined in Sect. 4.) In our color coding, lower values are represented with darker colors, and both metrics need to be minimized; therefore, it is easy to identify visually what is the best region of the parameter

space. As observed from Fig. 10b, in case of large swarms, response times show low variance for all the values of p that we studied, and thus, large swarms alone are able to reduce the variation in the response time of the system. In Sect. 6, we will show how the analytical model, which assumes infinite swarm size, also supports this results. Additionally, Fig. 10a seems to suggest that there is an interplay between swarm size and value of the p parameter, with intermediate values of p performing better irrespective of the swarm size, and best parameters found for large swarms and intermediate values of p .

It is interesting to compare how the system performs over time with respect to different values of p and also relate this to the performance with the best identified case for the stubborn individuals. In the second row of Fig. 10, we report the evolution over time of the $\sqrt{\text{MSE}}$ for different values of p and also of the stubborn individuals mechanism (with ten stubborn individuals). We report these results for $N = 1000$ (see Fig. 10c) and for $N = 10,000$ (see Fig. 10d). In both swarm sizes, we observe an interesting trade-off between accuracy (lowest value reached by $\sqrt{\text{MSE}}$) and speed of adaptation (the rate at which the $\sqrt{\text{MSE}}$ goes down). Lower values of p produce slower but more accurate systems. Interestingly, the performance of the system with stubborn individuals (denoted by the thick black line) performs analogously to one of the parameters ($p = 0.005$) for $N = 1000$ and has a performance that is in between $p = 0.001$ and $p = 0.0001$.

The analysis of the spontaneous opinion switching mechanism and its comparison with the stubborn individuals reveals strengths and weaknesses of both: On the one hand, the spontaneous opinion switching mechanism allows the designer to tune the desired level of accuracy and speed, depending on the relative importance of the two in the application scenario where this method is to be applied. On the other hand, parameter tuning implies that either optimization or trial and error is required to find good parameters, which implies extra simulations or physical robot experiments. With stubborn individuals, the “recipe” is much simpler: Stubborn individuals must be included in small numbers, where this number should be enough just to guarantee the desired level of redundancy and fault tolerance.

6 The ordinary differential equation models

In this section, two ordinary differential equation (ODE) models are introduced to study how the collective decision-making dynamics are influenced by the introduction of the two new mechanisms: the stubborn individuals and the spontaneous opinion switching. All these models assume a continuum of agents ($N \rightarrow \infty$). The focus is on the time evolution of two subpopulations, one with opinion A and one with opinion B . Furthermore, the model is compartmentalized in a way to reflect the probabilistic finite-state machine introduced in Sect. 3.2 that models the individual behavior of the agents: The four state variables e_A , e_B , d_A , and d_B are considered, where e_A is the proportion of agents with opinion A in the exploration state, e_B is the proportion of agents with opinion B in the exploration state, d_A is the proportion of agents with opinion A in the dissemination state, and d_B is the proportion of agents with opinion B in the dissemination state. In the model with stubborn individuals (Sect. 6.1), we further compartmentalize each subpopulation into two (normal and stubborn), resulting in a total of eight state variables.

With ODEs, it is possible to monitor the deterministic evolution of the system, while stochastic fluctuations and potential effects of finite population sizes are neglected in these models. Using compartmentalized ODEs, we can study the dynamics at two scales: mesoscopic if we focus on subpopulations e_A , e_B , d_A , and d_B and macroscopic if we focus on the

total number of agents with opinion A (i.e., $(d_A + e_A)$) and on the total number of agents with opinion B (i.e., $(d_B + e_B)$). The model is solved at the mesoscopic scale, whereas the results will be reported at a macroscopic scale to enhance interpretability. Analytical methods from dynamical systems theory are applied to find and study the equilibria of the system, and integration is used to calculate some of the trajectories.

6.1 ODE model with stubborn agents

To model stubborn individuals as studied in Sect. 5.2, we extended the ODE model by Valentini et al. (2014) by introducing new subpopulations of stubborn agents, e_{AS} , e_{BS} , d_{AS} , and d_{BS} . Their sum is constant and equal to the x_S , the proportion of stubborn individuals in the population: $e_{AS} + e_{BS} + d_{AS} + d_{BS} = x_S$. The total number of agents is conserved $e_A + e_B + d_A + d_B + e_{AS} + e_{BS} + d_{AS} + d_{BS} = 1$, and each individual subpopulation must be $0 \leq e_A, e_B, d_A, d_B, e_{AS}, e_{BS}, d_{AS}, d_{BS} \leq 1$.

The system of ODEs is given by:

$$\dot{d}_A = -\frac{1}{\rho_{Ag}}d_A + \frac{1}{q}e_A \quad (1)$$

$$\dot{d}_B = -\frac{1}{\rho_{Bg}}d_B + \frac{1}{q}e_B \quad (2)$$

$$\dot{e}_A = -\frac{1}{q}e_A + \frac{\sigma_{AS}}{\rho_{Ag}}d_A + \frac{\sigma_{AS}}{\rho_{Bg}}d_B \quad (3)$$

$$\dot{e}_B = -\frac{1}{q}e_B + \frac{1 - \sigma_{AS}}{\rho_{Ag}}d_A + \frac{1 - \sigma_{AS}}{\rho_{Bg}}d_B \quad (4)$$

$$\dot{d}_{AS} = -\frac{1}{\rho_{Ag}}d_{AS} + \frac{1}{q}e_{AS} \quad (5)$$

$$\dot{d}_{BS} = -\frac{1}{\rho_{Bg}}d_{BS} + \frac{1}{q}e_{BS} \quad (6)$$

$$\dot{e}_{AS} = -\frac{1}{q}e_{AS} + \frac{1}{\rho_{Ag}}d_{AS} \quad (7)$$

$$\dot{e}_{BS} = -\frac{1}{q}e_{BS} + \frac{1}{\rho_{Bg}}d_{BS} \quad (8)$$

In the above model, Eqs. 1–4 model the evolution of non-stubborn agents and are very similar to those of the original model by Valentini et al. (2014). In Eqs. 1 and 2, d_A (resp. d_B), the proportion of non-stubborn agents disseminating A (resp. B) increases at a rate q^{-1} due to agents returning from the exploration of the sites and decreases at a rate $(\rho_{Ag})^{-1}$ (resp. $(\rho_{Bg})^{-1}$) due to agents leaving the dissemination state with a rate proportional to the quality of the sites. In Eqs. 3 and 4, e_A (resp. e_B), the proportion of non-stubborn agents exploring site A (resp. B) decreases at a rate q^{-1} due to agents finishing exploring site A (resp. B), while it increases at a rate which depends on the application of the voter model. In particular, the result of the application of the voter model will depend on the probability of observing opinion A or B as random neighbor opinion and therefore depends on the current state of the swarm. In this model with stubborn individuals, we define the voter model probability as:

$$\sigma_{AS} = \frac{d_A + d_{AS}}{d_A + d_{AS} + d_B + d_{BS}},$$

that is the probability to observe A is defined as the proportion of individuals disseminating A (normal and stubborn) divided by the total proportion of agents in the dissemination state. The probability to observe B can be simply defined as $\sigma_B = 1 - \sigma_A$. The definition of σ_A is the only deviation between these four equations and the model of Valentini et al. (2014), where the voter model probability did not include stubborn individuals but was otherwise defined in the same way. After having defined the voter model probability, the rate at which agents exploring A increase can be defined as proportional to the voter model probability and to $(\rho_{Ag})^{-1}$ for agents that were already of opinion A , or to $(\rho_{Bg})^{-1}$ for agents that were of opinion B and switch to opinion A after the application of the voter model. A similar reasoning can be applied for the rate of increase in agents exploring site B .

Equations 5–8 model the evolution of stubborn agents. Equations 5 and 6 model the increase and decrease in agents in the dissemination state and are similar to Eqs. 1 and 2, with variables modeling stubborn agents replacing variables modeling non-stubborn agents. Equations 7 and 8 model the increase and decrease in agents in the exploration state. The term indicating agent decrease is the same as the one in Eqs. 3 and 4, with variables modeling stubborn individuals replacing variables modeling non-stubborn individuals. To express the term indicating increase, note that stubborn individuals do not change opinions; therefore, all agents disseminating opinion A (resp. B) will switch to exploration at a rate $(\rho_{Ag})^{-1}$ (resp. $(\rho_{Bg})^{-1}$). Note also that the equations modeling the evolution of stubborn agents (Eqs. 5–8) are independent and not coupled with non-stubborn agents state variables, consistently with the fact that stubborn individuals are not influenced by other agents except themselves. Equations modeling the evolution of non-stubborn agents (Eqs. 1–4) are coupled to the stubborn agents state variables only through the voter model probability σ_A , consistently with the fact that stubborn individuals influence the behavior of non-stubborn individuals only during dissemination and voting. Note that by setting $x_S = 0$ and by observing the constraints on the variables defined above, we can recover the model by Valentini et al. (2014).

The parameters of the model have been set consistently with Valentini et al. (2014) and with the parameters used in Sect. 5. The exploration time is set to $q = 10$. The dissemination times are proportional to the quality of sites A and B , and we set the coefficient $g = 100$. Continuing from Sect. 5, and to keep this section concise, we consider here only the more interesting case with low quality ratio. Therefore, we set $\rho_A = 1$ and $\rho_B = 1.05$.

6.2 Dynamics of the ODE model with stubborn agents

We analytically found the equilibria of the ODEs for different values of the x_S parameter. The analysis is performed by projecting the system in two dimensions, $x_A = d_A + e_A + d_{AS} + e_{AS}$ and $x_B = d_B + e_B + d_{BS} + e_{BS}$. The equilibria are plotted in Fig. 11a. Asymptotically stable equilibria are plotted as two continuous lines, indicating the coordinates of x_A and x_B for each value of x_S , while unstable equilibria are plotted as pairs of empty circles. For $x_S = 0$, the system presents two equilibria, $\{x_A, x_B\} = \{0, 1\}$ and $\{x_A, x_B\} = \{1, 0\}$, that correspond to the two consensus states, the former being stable and the latter being unstable. These results are consistent with the study of Valentini et al. (2014), where $\{x_A, x_B\} = \{1, 0\}$ is the stable equilibrium whenever $\rho_A > \rho_B$. This does not necessarily reflect the behavior of a real system, due to the infinite system size approximation and the neglecting of stochastic fluctuations. For $x_S > 0$, the unstable equilibrium disappears and only the stable one survives. This stable equilibrium is characterized by a decay of the value of x_B asymptotically toward 0.5 and an increase in the value of x_B also asymptotically toward 0.5. This result is consistent

with those obtained in simulations (see Fig. 5d for $x_S = 0.001$, Fig. 4b for $x_S = 0.05$, and Fig. 5a for $x_S = 0.2$), where we observed a progressive tendency of the consensus state to move toward 0.5 for increasing values of x_S . The study with ODEs confirms that only small values of x_S are able to induce a consensus state close but not exactly equal to full unanimity, in order to achieve adaptability.

In Fig. 11b, we report the dynamics obtained by numerically integrating the ODEs, by starting with initial conditions $d_A = d_B = 0.1$ and $e_A = e_B = 0.49$, which means almost all agents are in the exploration state (similarly to the simulations) but split with respect to their opinions. (We initialize d_A and d_B to a small value in order to avoid zero denominators in σ_A in the ODEs.) The initial conditions for the stubborn individuals state variables are $d_{AS} = 0.01 \cdot x_S$, $d_{BS} = 0.01 \cdot x_S$, $e_{AS} = 0.49 \cdot x_S$, $e_{BS} = 0.49 \cdot x_S$. We report the value of x_A over time for different values of x_S , which include those used in the simulations and few more to have a more complete picture. At $t = T_C = 12,000$, we stop the process, we record the value of the state variables, we swap the values of the quality parameters ρ_A and ρ_B , and we integrate the system again with the new initial conditions given by these recorded state variables, in order to reproduce the dynamic environment. As we can see, the trend detected in Fig. 11a is confirmed here, with the value of the consensus state flattening toward $x_A = 0.5$ for increasing values of x_S . This new figure also gives us additional information about the behavior of the convergence times. We observe the typical speed vs. accuracy trade-off, with lower values of x_S corresponding both to higher consensus state as well as longer convergence times. A potentially disturbing result is represented by the curve corresponding to $x_S = 0$ shown in Fig. 11b, which shows the system achieving adaptability also in this case. This is, however, simply explained by the fact that dynamics of ODE models only reach the steady states for $t \rightarrow \infty$. Therefore, for any finite time t , the trajectories of the ODEs have not reached unanimity, and therefore, ODEs would predict that adaptability is always possible. However, in finite systems, the consensus state is reached in finite time, and therefore, mechanisms to prevent unanimity like those proposed in this paper are needed.

6.3 ODE model with spontaneous opinion switching

In the case of the spontaneous opinion switching mechanism, the original four state variables in Valentini et al. (2014) are sufficient in the corresponding ODE model. The opinions switching probability, however, introduces new terms to the four original equations. The extended model is the following:

$$\dot{d}_A = -\frac{1}{\rho_{Ag}}d_A + \frac{1}{q}e_A \quad (9)$$

$$\dot{d}_B = -\frac{1}{\rho_{Bg}}d_B + \frac{1}{q}e_B \quad (10)$$

$$\dot{e}_A = -\frac{1}{q}e_A + \frac{\sigma_A}{\rho_{Ag}}(1-p)d_A + \frac{\sigma_A}{\rho_{Bg}}(1-p)d_B + \frac{1-\sigma_A}{\rho_{Ag}}pd_A + \frac{1-\sigma_A}{\rho_{Bg}}pd_B \quad (11)$$

$$\dot{e}_B = -\frac{1}{q}e_B + \frac{1-\sigma_A}{\rho_{Ag}}(1-p)d_A + \frac{1-\sigma_A}{\rho_{Bg}}(1-p)d_B + \frac{\sigma_A}{\rho_{Ag}}pd_A + \frac{\sigma_A}{\rho_{Bg}}pd_B \quad (12)$$

Conservation and proportion constraints are also defined in this case: $e_A + e_B + d_A + d_B = 1$ and $0 \leq e_A, e_B, d_A, d_B \leq 1$. Equations 9 and 10 modeling the evolution of agents in the dissemination states are identical to those of the original model (Valentini et al. 2014) and to those of the model for stubborn individuals (see Sect. 6.1). Equations 11 and 12 modeling

the evolution of agents in the exploration state instead are different. The only unmodified component is the rate of decrease which is still proportional to q^{-1} . Conversely, agents in the exploration state can increase in four possible ways. This is because we apply the opinion switching model *after* the application of the voter model. Therefore, to explain Eq. 11, agents exploring site A can increase in four possible ways: via agents disseminating A that remain of opinion A after the application of the voter model (proportionally to σ_A) and the (non-) application of the opinion switching mechanism (proportionally to $(1 - p)$); via agents disseminating B that switch to A (proportionally to σ_A) and that remain in A (proportionally to $1 - p$); via agents disseminating A that switch to B after the application of the voter model (proportionally to $1 - \sigma_A$) but that again switch to A after the application of the spontaneous opinion switching (proportionally to p); via agents disseminating B that remain in B after the voter model (proportionally to $1 - \sigma_A$) but that switch to A (proportionally to p). Equation 12 can be explained using an analogous reasoning. The expression of σ_A in this model is the same as the one in (Valentini et al. 2014):

$$\sigma_A = \frac{d_A}{d_A + d_B}.$$

By setting $p = 0$, we recover the original model by Valentini et al. (2014). Concerning the value of the parameters, we use the same as in Sects. 6.1 and 5: $q = 10$, $g = 100$, $\rho_A = 1$, and $\rho_B = 1.05$.

6.4 Dynamics of the ODE model with spontaneous opinion switching

We analytically found the equilibria of the ODEs for different values of the p parameter. The equilibria are plotted in Fig. 12a. As for the stubborn agents' case, asymptotically stable equilibria are plotted as two continuous lines, indicating the coordinates of x_A and x_B for each value of p , while unstable equilibria as pairs of empty circles. Also similarly to the stubborn agents' case, for $p = 0$, the system presents two equilibria, $\{x_A, x_B\} = \{0, 1\}$ and $\{x_A, x_B\} = \{1, 0\}$, that correspond to the two consensus states, the former being stable and the latter being unstable. This is to be expected as, for $p = 0$, we recover the system in Valentini et al. (2014) which had the same equilibria. For $p > 0$, the unstable equilibrium disappears and only the stable one survives. This stable equilibrium is characterized by a decay of the value of x_B asymptotically toward 0.5 and an increase in the value of x_B also asymptotically toward 0.5. This result is consistent with those obtained in simulation (see Fig. 9), where we observed a flattening of the consensus state toward 0.5 for increasing values of p . The study with ODEs confirms that only small values of p are able to induce a consensus state that is close but not exactly equal to full unanimity, required for adaptability.

In Fig. 12b, we report the dynamics obtained by numerically integrating the ODEs, by starting with initial conditions $d_A = d_B = 0.1$ and $e_A = e_B = 0.49$. We report the value of x_A over time for different values of p , which include those used in the simulations and few more to have a more complete picture. To model dynamic environments, we use the same protocol explained in Sect. 6.2. As we can see, the trend detected in Fig. 12a is confirmed here, with the value of the consensus state flattening toward $x_A = 0.5$ for increasing values of p . Concerning the convergence times, also here we observe the typical speed vs. accuracy trade-off, with lower values of p corresponding both to higher consensus state as well as longer convergence times. This trend was similarly observed also in our simulations, such as in Fig. 10 (second row): Although the trend is confirmed in both Fig. 10c, d, similarly to the case with stubborn individuals, also in this case, the predictions of the mathematical




Fig. 11 Analysis of the ODE model in the presence of stubborn individuals. In **a**, we report the stable equilibria as a function of the proportion of stubborn individuals x_S . For $x_S > 0$, the system has only one equilibrium that is stable, that is, reported as a line (with different styles as explained in the legend). For $x_S = 0$, the system has a stable equilibrium and an unstable equilibrium, and the latter reported as an empty circle. In **b**, we report the time evolution of the consensus dynamics (via x_A , the proportion of agents choosing A) over time, for several values of x_S . Continuous lines correspond to parameter values that were also studied via numerical simulations in Sect. 5, while dashed lines correspond to additional parameter values studied here to give a broader picture of the dynamics




Fig. 12 Analysis of the ODE model with spontaneous opinion switching. In **a**, we report the stable equilibria as a function of the spontaneous opinion switching probability p . For $p > 0$, the system has only one equilibrium that is stable, that is, reported as a line (with different styles as explained in the legend). For $p = 0$, the system has a stable equilibrium and an unstable equilibrium, and the latter reported as an empty circle. In **b**, we report the time evolution of the consensus dynamics over time (via x_A , the proportion of agents choosing A), for several values of the opinion switching probability p . Continuous lines correspond to parameter values that were also studied via numerical simulations in Sect. 5, while dashed lines correspond to additional parameter values studied here to give a broader picture of the dynamics

model becomes quantitatively more accurate as the system size increases. As for the case of stubborn individuals, the fact that curve corresponding to $p = 0$ in Fig. 12b shows the system achieving adaptability against evidences from simulations can be explained considering the difference between ODE models and finite time simulations.

6.5 Relating the two models between each other and with simulations

We observe a striking duality between the two models and the two adaptation mechanism, namely the dynamics of the two systems are very similar both in terms of how equilibria vary as a function of the respective parameter (x_S or p in Figs. 11a, 12a), as well as in terms of the trajectories over time (see Fig. 11b compared to Fig. 12b). In particular, the first of

the two types of plots suggests that the best values for both x_S and p in terms of accuracy of the consensus state are infinitesimally small nonzero values, while the four plots altogether suggest that if seeking a compromise between speed and accuracy, the best values for both x_S and p seem to be around 0.001. Despite ODE dynamics of both models seem to be equivalent, these are a good predictor of the real system only in case of very large populations, as observed by comparing the results of this section with those in Sect. 5. However, for finite population size and in particular for small populations, ODE models are not sufficient to give an accurate prediction. For example, results in Sect. 5 suggest that for small swarm, stubborn individuals achieve better results in terms of fluctuations around average performance compared to the spontaneous opinion switching mechanism, as shown specifically in Fig. 10 that showed very high values for the standard deviation of the response times, which were not observed in experiments with stubborn individuals.

7 Conclusion, discussion, and future work

In this work, we have introduced the dynamic best-of- n problem, in the presence of dynamic option qualities that can abruptly change over time. The traditional voter model is not suitable to ensure adaptability of the swarm in case the best option dynamically changes after consensus is reached. To achieve adaptability, we have proposed two mechanisms. Both are applied in the context of a decision-making mechanism based on direct modulation of positive feedback and on the voter model. The first solution mechanism is represented by stubborn agents, that is, agents that do not change their opinion and stay committed to their initial option. As a second solution mechanism, we introduce spontaneous opinion switching, whereby all agents are identical and can probabilistically change their opinion after and independently of the application of the decision mechanism. Both mechanisms are artificial and do not have a direct counterpart within natural biological systems, and thus, they represent an engineering artificial mechanism to adapt the voter model to dynamic environments.

Through computer simulations, we have shown that the voter model alone (i.e., without the stubborn agents) cannot make the swarm adapt to abrupt changes in the option qualities. We thoroughly extended the study performed by Prasetyo et al. (2018b), where we found that, consistently with the previous work (Montes de Oca et al. 2011), the difference in site quality plays a crucial role, whereby higher level of adaptability is observed with increasing ratio between the qualities. We extended the study to larger swarms, where we found that increasing the ratio of stubborn individuals has a detrimental effect on accuracy and on adaptability when the ratio between the qualities is low. We further confirmed that by increasing the swarm size, both accuracy and adaptability are beneficially affected. We disambiguated the effect of the swarm size from the effect of swarm density, and we found that only the swarm size affects positively the performance, while the density has no effect unless it is below a very low critical threshold. Finally, we studied the spontaneous opinion switching mechanism with respect to swarm size and of its key parameter, the switching probability p . Once again, we confirmed that larger swarm sizes result in improved performance, this time with respect to the response time of the system which becomes more reliable in terms of its variation across runs. We also found that by regulating the parameter p , it is possible to regulate the trade-off between the accuracy of the decision making and the variation in the response time of the system. It is worth to make a comparison between the two models: Using the spontaneous opinion switching mechanism, the designer is able to tune the level of accuracy and variability of response speed to the task at hand, by paying the cost of parameter tuning. On the other

hand, the utilization of stubborn individuals achieves a given trade-off between accuracy and response speed variation, while avoiding expensive parameter tuning.

One of the main contributions of this work has been the design of a collective system able to exhibit collective response to environmental changes, in a way that is not only scale-invariant (Khaluf et al. 2017) but that had superior performance as the system scale increased. There are many possible directions for future work. First, mathematical models that allow a richer study compared to the ODEs considered here, such as chemical reaction networks, can be developed to study the effect of finite sizes and of fluctuations. We also plan to use novel analysis methods such as those based on information transfer (Valentini et al. 2018) in order to quantify the system response to the environmental change. Secondly, in our previous work (Prasetyo et al. 2018b), we performed a preliminary study of the majority rule model, where we showed that this model is ineffective in reaching consensus to the right option and at adapting to environmental changes, due to the effect of spatiality, as stubborn individuals committed to the same options are very unlikely to appear next to each other. We completely neglected the majority rule model in this paper as preliminary results were not promising and therefore deserved a much deeper study, which we plan to do in the near future. Thirdly, in this work, we mainly considered abrupt environmental changes, but future work may focus on different dynamic environments, such as non-abrupt changes following different types of dynamics. Another possible direction for future work is to study whether the decision-making process and the adaptability are sensitive not only to the relative ratio between the qualities but also to their absolute value (Pais et al. 2013; Reina et al. 2018a). Finally, provided enough resources, we plan to perform experiments on real robots, likely kilobots (Rubenstein et al. 2014), in order to have a proof of concept in the real world and potentially discover new factors influencing adaptability.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Arvin, F., Turgut, A. E., Bazyari, F., Arıkan, K. B., Bellotto, N., & Yue, S. (2014). Cue-based aggregation with a mobile robot swarm: A novel fuzzy-based method. *Adaptive Behavior*, 22(3), 189–206.
- Baronchelli, A., & Díaz-Guilera, A. (2012). Consensus in networks of mobile communicating agents. *Physical Review E*, 85, 016113.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. New York: Oxford University Press.
- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.
- Britton, N. F., Franks, N. R., Pratt, S. C., & Seeley, T. D. (2002). Deciding on a new home: How do honeybees agree? *Proceedings Biological Sciences*, 269(1498), 1383–8.
- Brutschy, A., Scheidler, A., Ferrante, E., Dorigo, M., & Birattari, M. (2012). Can ants inspire robots? Self-organized decision making in robotic swarms. In *Proceedings of the 2012 IEEE/RSJ international conference on intelligent robots and systems (IROS'12)* (pp. 4272–4273). IEEE Computer Society Press, Los Alamitos, CA.
- Calovi, D. S., Lopez, U., Schuhmacher, P., Chaté, H., Sire, C., & Theraulaz, G. (2015). Collective response to perturbations in a data-driven fish school model. *Journal of the Royal Society Interface*, 12(104), 20141362.
- Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-organization in biological systems*. Princeton, NJ: Princeton University Press.

- Campo, A., Garnier, S., Dédriche, O., Zekkri, M., & Dorigo, M. (2010). Self-organized discrimination of resources. *PLoS ONE*, 6(5), e19888.
- Correll, N., & Martinoli, A. (2011). Modeling and designing self-organized aggregation in a swarm of miniature robots. *The International Journal of Robotics Research*, 30(5), 615–626.
- Deneubourg, J. L., & Goss, S. (1989). Collective patterns and decision-making. *Ethology Ecology & Evolution*, 1(4), 295–311.
- Ferrante, E., Turgut, A. E., Huepe, C., Stranieri, A., Pinciroli, C., & Dorigo, M. (2012). Self-organized flocking with a mobile robot swarm: A novel motion control method. *Adaptive Behavior*, 20(6), 460–477.
- Font Llenas, A., Talamali, M. S., Xu, X., Marshall, J. A. R., & Reina, A. (2018). Quality-sensitive foraging by a robot swarm through virtual pheromone trails. In M. Dorigo, M. Birattari, C. Blum, A. L. Christensen, A. Reina, & V. Trianni (Eds.), *Swarm intelligence (ANTS 2018)*, LNCS (Vol. 11172, pp. 135–149). Berlin: Springer.
- Franks, N. R., Pratt, S. C., Mallon, E. B., Britton, N. F., & Sumpter, D. J. T. (2002). Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357(1427), 1567–1583.
- Garnier, S., Gautrais, J., Asadpour, M., Jost, C., & Theraulaz, G. (2009). Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, 17(2), 109–133.
- Garnier, S., Gautrais, J., & Theraulaz, G. (2007). The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1), 3–31.
- Hamann, H. (2018). Opinion dynamics with mobile agents: Contrarian effects by spatial correlations. *Frontiers in Robotics and AI*, 5, 63.
- Hunter, D. S. & Zaman, T. (2018). *Opinion dynamics with stubborn agents*. [arXiv:1806.11253](https://arxiv.org/abs/1806.11253).
- Kernbach, S., Thenius, R., Kernbach, O., & Schmickl, T. (2009). Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic system. *Adaptive Behavior*, 17(3), 237–259.
- Khaluf, Y., Ferrante, E., Pieter, S., & Huepe, C. (2017). Scale invariance in natural and artificial collective systems: A review. *Journal of the Royal Society Interface*, 14(136), 1–20.
- Marshall, J. A. R., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., & Franks, N. R. (2009). On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society Interface*, 6(40), 1065–1074.
- Mukhopadhyay, A., & Mazumdar, R. R. (2016). Binary opinion dynamics with biased agents and agents with different degrees of stubbornness. *IEEE*, 01, 261–269.
- Montes de Oca, M. A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., & Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5, 305–327.
- Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., & Marshall, J. A. R. (2013). A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9), 1–9.
- Parker, C. A. C., & Zhang, H. (2009). Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2), 240–251.
- Parker, C. A. C., & Zhang, H. (2010). Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, 4, 199–220.
- Prasetyo, J., De Masi, G., & Ferrante, E. (2018a). *The best-of-n problem in dynamic environments*. <http://swarm.live/sispecial2018/>, Supplementary material. Accessed 30 November 2018.
- Prasetyo, J., De Masi, G., Ranjan, P., & Ferrante, E. (2018b). The best-of-n problem with dynamic site qualities: Achieving adaptability with stubborn individuals. In M. Dorigo, M. Birattari, C. Blum, A. L. Christensen, A. Reina, & V. Trianni (Eds.), *Swarm intelligence (ANTS 2018)*, LNCS (Vol. 11172, pp. 239–251). Berlin: Springer.
- Pratt, S. C., Mallon, E. B., Sumpter, D. J., & Franks, N. R. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *leptothorax albigipennis*. *Behavioral Ecology and Sociobiology*, 52(2), 117–127.
- Reina, A., Bose, T., Trianni, V., & Marshall, J. A. R. (2018a). Effects of spatiality on value-sensitive decisions made by robot swarms. In *Distributed autonomous robotic systems (DARS 2016): The 13th international symposium, SPAR* (Vol. 6, pp. 461–473). Cham, Switzerland: Springer.
- Reina, A., Bose, T., Trianni, V., & Marshall, J. A. R. (2018b). Psychophysical laws and the superorganism. *Scientific Reports*, 8, 4387.
- Reina, A., Marshall, J. A. R., Trianni, V., & Bose, T. (2017). Model of the best-of-N nest-site selection process in honeybees. *Physical Review E*, 95(5), 052411.
- Reina, A., Miletitch, R., Dorigo, M., & Trianni, V. (2015a). A quantitative micro-macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*, 9(2–3), 75–102.
- Reina, A., Valentini, G., Fernández-Oto, C., Dorigo, M., & Trianni, V. (2015b). A design pattern for decentralised decision making. *PLoS ONE*, 10(10), e0140950.

- Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A., & Nagpal, R. (2014). Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robotics and Autonomous Systems*, 62(7), 966–975.
- Scheidler, A., Brutschy, A., Ferrante, E., & Dorigo, M. (2016). The k -unanimity rule for self-organized decision-making in swarms of robots. *IEEE Transactions on Cybernetics*, 46(5), 1175–1188.
- Schmickl, T., Thenius, R., Moeslinger, C., Radspieler, G., Kernbach, S., Szymanski, M., et al. (2009). Get in touch: Cooperative decision making based on robot-to-robot collisions. *Autonomous Agents and Multi-Agent Systems*, 18(1), 133–155.
- Seeley, T. D. (2010). *Honeybee democracy*. Princeton: Princeton University Press.
- Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. R. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064), 108–111.
- Valentini, G., Birattari, M., & Dorigo, M. (2013). Majority rule with differential latency: An absorbing Markov chain to model consensus. In: T. Gilbert, M. Kirkilionis, & G. Nicolis (Eds.), *Proceedings of the European conference on complex systems 2012, Springer proceedings in complexity* (pp. 651–658). Springer
- Valentini, G., Ferrante, E., & Dorigo, M. (2017). The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4, 9.
- Valentini, G., Ferrante, E., Hamann, H., & Dorigo, M. (2016). Collective decision with 100 kilobots: Speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 30(3), 553–580.
- Valentini, G., Hamann, H., & Dorigo, M. (2014). Self-organized collective decision making: The weighted voter model. In: A. Lomuscio, P. Scerri, A. Bazzan, & M. Huhns (Eds.), *Proceedings of the 13th international conference on autonomous agents and multiagent systems, IFAAMAS, AAMAS '14* (pp. 45–52)
- Valentini, G., Hamann, H., & Dorigo, M. (2015). Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In: R. Bordini, E. Elkind, G. Weiss, & P. Yolum (Eds.) *Proceedings of the 14th international conference on autonomous agents and multiagent systems, IFAAMAS, AAMAS '15* (pp. 1305–1314)
- Valentini, G., Moore, D. G., Hanson, J. R., Pavlic, T. P., Pratt, S. C., & Walker, S. I. (2018). *Transfer of information in collective decisions by artificial agents* (pp. 641–648). Cambridge: MIT Press.
- Yildiz, E., Ozdaglar, A., Acemoglu, D., Saberi, A., & Scaglione, A. (2013). Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*, 1(4), 19:1–19:30.