Distributed Network Optimization With Heuristic Rational Agents

Ceyhun Eksin, *Student Member, IEEE*, and Alejandro Ribeiro, *Member, IEEE*

***Abstract*—A network of distributed agents wants to minimize a global cost given by a sum of local terms involving nonlinear convex functions of self and neighboring variables. Agents update their variables at random times by observing the values of neighboring agents and applying a random heuristic rule intent on minimizing the local cost with respect to their own variables. The heuristic rules are rational in that their average result is the actual optimal action with respect to the given values of neighboring variables. By identifying heuristic rational optimization with stochastic coordinate descent, it is shown that all agents visit a neighborhood of the optimal cost infinitely often with probability 1. An exponential probability bound on the worst deviation from optimality between visits to near optimal operating points is also derived. Commonly used models of consensus and opinion propagation in social networks, Markov random field estimation in wireless sensor networks, and cohesive foraging of animal herds are cast in the language of heuristic rational optimization. Numerical simulations for these three examples are presented to corroborate analytical results.**

***Index Terms*—Biological systems, convergence, distributed estimation, distributed optimization, social networks.**

I. Introduction

NETWORK optimization problems entail a group of agents with certain underlying connectivity that strive to minimize a global cost through appropriate selection of local variables. Optimal determination of local variables requires, in principle, global coordination of all agents. In *distributed* network optimization, agent coordination is further restricted to neighboring nodes. The optimization of the global objective is then achieved through iterative application of local optimization rules that update local variables based on information about the state of neighboring agents. Distributed network optimization is a common solution method for estimation and detection problems in wireless sensor networks (WSNs) [1]–[7].

Beyond its use in engineered systems, distributed network optimization is also used to model the emergence of global behavior in biological and social networks. In this context, the optimization cost models global network behavior that emerges through the application of the local optimization rules. In biological systems, network optimization models that mimic natural phenomena like bird flocking [8], [9] or animal swarming [10]–[13] have been introduced. Bird flocking models posit that individual birds try to optimize total drag by adjusting their individual positions and velocities based on the observed behavior of neighboring birds within their field of vision [8]. Similarly, the foraging behavior of animal herds and fish schools can be explained as the optimization of an objective that includes terms to account for the value of food, the value of cohesion and the cost of excessive proximity [14], [15]. As in the case of bird flocks, members of the herd or school adjust their positions with respect to the observed positions of nearby peers. Notice how these models exhibit the three hallmarks of distributed network optimization. They start from a global objective that the network agents want to optimize—like total drag for bird flocks—through the selection of local variables—birds' positions and velocities—while restricting interactions to neighboring agents—positions and velocities are updated relative to the closest neighboring birds on the field of vision. Consensus formation [15], [16] and opinion propagation [17] in social networks can also be understood in terms of distributed network optimization. In this case, network nodes represent social agents having differing opinions that they update over time based on the observed opinions of neighboring nodes. Agents determine these updates by minimizing a *local* measure of disagreement with their neighbors. As a result, the network as a whole is minimizing a *global* measure of disagreement. The difference between consensus and opinion propagation models is that in the former all nodes attempt to increase harmony, while in the latter some stubborn agents do not change their opinions.

A common feature in [8]–[17] is that agents are assumed to act rationally, in that they update their local variables in a manner that is optimal with respect to the available information. Birds in a flock choose a position and speed to minimize their own drag [8], animals in a herd choose a position to balance attraction and repulsion forces with their neighbors, [11], [12], and people update their opinions to minimize the discordance with their friends [17]–[19]. However, assuming optimal behavior in the context of natural and social, as opposed to engineered networks, limits the applicability of these models because making optimal decisions requires exceedingly high levels of awareness and cunningness.

The goal of this paper is to propose and study more realistic models whereby agents execute actions that are optimal in an average sense only. We name these rules and the agents that use them as heuristic rational, since we think of them as the application of a heuristic rule that is intent on being optimal, even though it may not be so. We show that models com-

Manuscript received October 20, 2011; revised March 09, 2012 and May 31, 2012; accepted June 09, 2012. Date of publication June 28, 2012; date of current version September 11, 2012. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Yao-Win (Peter) Hong. Work in this paper is supported by AFOSR MURI FA9550-10-1-0567, W911NF-10-1-0388, NSF CAREER CCF-0952867, and NSF CCF-1017454. Parts of the results in this paper appeared in Asilomar 2011 and ICASSP 2012.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: ceksin@seas.upenn.edu; aribeiro@seas.upenn.edu).

Digital Object Identifier 10.1109/TSP.2012.2206587

monly used to study propagation of opinions in social networks [20], foraging of animal herds [11], [21], and quantization and communication issues in field estimation using WSNs [1], [2], [22], [23] can be cast in the language of heuristic rational optimization. We also study the behavior of networks composed of heuristic rational agents and show that: (i) The global network behavior visits a neighborhood of optimality infinitely often. (ii) The probability of straying away from this neighborhood by more than a given amount is exponentially bounded. These results can be interpreted as an explanation for the emergence [cf. (i)] and sustenance [cf. (ii)] of global network behavior that is close to optimal despite imperfect decision making of individual agents in natural and social systems. We note that other efforts to lift unrealistic assumptions in distributed network optimization exist. These include the study of asynchronous updates [24], time-varying [15] or unreliable communication links [2], [3], [25], [26], and communication contaminated with random noise [2], [5], [8], [27]. Our work differs from these contributions in that we are considering the update rules themselves as being imperfect.

The paper begins by describing the induction of global behavior through the minimization of a cost given by a sum of local terms involving nonlinear functions of self and neighboring variables. At random times, agents observe current values of their neighbors' variables and apply a heuristic rule with the intent of minimizing the global cost with respect to the selection of their local variables. These heuristic rules need not be optimal but we assume that they are so in expectation (Section II). We proceed to describe how voter models [19], [20] used to study the propagation of opinions in social networks can be interpreted as a heuristic rational version of local averaging models [16], [22] (Section II-A). We further present a Markov random field (MRF) estimation problem using a WSN. In this case, heuristic rational actions can be used to model communication and quantization effects (Section II-B).

Because of the randomness associated with heuristic rational rules, we do not expect convergence to optimal global behavior. Consequently, our goal is to describe the difference in the yield of optimal variables and the values achieved by heuristic rational rules (Section III). To characterize this difference, we identify the variable updates with a stochastic coordinate descent algorithm with random activation rule. Exploiting this fact, it is possible to show that for points sufficiently far away from optimal the distance to optimality between subsequent uses of heuristic rational rules satisfies a supermartingale inequality (Lemma 1). This observation leads to the conclusion that a neighborhood of optimality is visited infinitely often with probability 1 (Theorem 1). The size of the near optimality region depends on parameters of the function being minimized and is proportional to the variance of the heuristic rational rule. We further show that between visits to optimality the probability of the gap in the yield of agents' variables exceeding a given value is bounded exponentially (Theorem 2 in Section IV). This bound shows that even though it is possible for local variables to become arbitrarily bad, significant deviations are exponentially rare.

We present numerical results for opinion propagation in social networks, MRF estimation with WSNs, and herd foraging in biological networks (Section V). In the social network case study, we use a consensus model to analyze the propagation of opinions in a social network from two stubborn agents having opposing beliefs (Section V-A). We observe that stubborn agents are influential only within their close neighborhood, their beliefs losing strength as they propagate through the network. In the MRF estimation problem, we consider a WSN deployed to estimate a spatial-varying temperature field with Markov conditional probabilities; see e.g., ([28], Ch. 1). Sensors are interested in the temperature value at their location that they estimate using a local noisy temperature measurement and cooperation with neighboring nodes (Section V-B). The biological network example concerns cohesive movement of a foraging herd. Animals in the herd move on a field looking for a food source while staying neither too close nor too far away from each other. Individual behavior is explained through attraction and repulsion forces between neighbors and an attraction force to the food source [21]. We close the paper with concluding remarks (Section VI).

II. Local Heuristic Rational Optimization

Consider a network of N agents represented by the symmetric graph $G = (V, E)$ where vertices $i \in V$ denote agents and edges $(i, j) \in E$ connections between them. Agent i can only interact with neighboring nodes $n(i) = \{j : (j, i) \in E\}$ that form an edge with it. We denote as $N_i := \#(n(i))$ the cardinality of the number of neighbors. Each of the agents $i \in V$ is associated with corresponding variable $\mathbf{x}_i \in \mathbb{R}^n$ and a convex function $f_{0i}(\mathbf{x}_i)$. Each of the edges $(i, j) \in E$ is affiliated with a convex function $f_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ that depends on the agent variables at the vertices of the given edge. To maintain symmetry, we require that functions $f_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ and $f_{ji}(\mathbf{x}_j, \mathbf{x}_i)$ be equal,

$$f_{ij}(\mathbf{x}_i, \mathbf{x}_j) = f_{ji}(\mathbf{x}_j, \mathbf{x}_i), \quad \text{for all } i, j \in n(i). \tag{1}$$

Variables $\mathbf{x}_i$ are also constrained to the convex set $\mathcal{X}_i$ in that allowable values satisfy $\mathbf{x}_i \in \mathcal{X}_i \subseteq \mathbb{R}^n$. Define the vectors $\mathbf{x} := \{\mathbf{x}_i\}_{i \in V}$ grouping all network variables and $\mathbf{x}_{n(i)} := \{\mathbf{x}_j\}_{j \in n(i)}$ containing the variables of all neighbors of i, and $\mathbf{x}_{-i} := \{\mathbf{x}_j\}_{j \neq i}$ referring to all variables except $\mathbf{x}_i$. Further introduce the set $\mathcal{X} := \prod_{i \in V} \mathcal{X}_i$ to represent the Cartesian product of sets $\mathcal{X}_i$.

The function

$$f_i\left(\mathbf{x}_i, \mathbf{x}_{n(i)}\right) := f_{0i}(\mathbf{x}_i) + \sum_{j \in n(i)} f_{ij}(\mathbf{x}_i, \mathbf{x}_j) \tag{2}$$

represents a cost that agent i would like to make as small as possible by proper selection of its variable $\mathbf{x}_i \in \mathcal{X}_i$. Since this cost depends on neighboring variables $\mathbf{x}_{n(i)}$, it follows that $\mathbf{x}_i$ and $\mathbf{x}_j$ for $j \in n(i)$ have to be jointly chosen. But these neighboring variables are jointly chosen with their respective neighbors, which depend on the values of their corresponding neighbors, and so on. It follows that as long as the network is fully connected, cost minimization requires simultaneous selection of all variables $\mathbf{x}_i$. This is not a plausible model of network behavior.

Alternatively, suppose that at random time $t \in \mathbb{R}^+$, agent i observes the values of neighboring variables $\mathbf{x}_{n(i)}(t)$. Given the interest in minimizing the local cost $f_i(\mathbf{x}_i, \mathbf{x}_{n(i)})$ in (2), a *rational* action for this agent is to update its variable by selecting the value that minimizes $f_i(\mathbf{x}_i, \mathbf{x}_{n(i)})$ given the observed values of neighboring variables,

$$\tilde{\mathbf{x}}_i(t) = \underset{\mathbf{x}_i \in \mathcal{X}_i}{\operatorname{argmin}} f_i\left(\mathbf{x}_i, \mathbf{x}_{n(i)}(t)\right). \tag{3}$$

Since the update in (3) is based on information that can be locally acquired and is unilaterally executed by i, it constitutes a possible model for network optimization, which has indeed been used to model, e.g., the propagation of opinions in a social network; see [17] and Section II-A. However, it is not always accurate to assume that agents apply optimal policies perfectly. In, e.g., social or biological systems, agents apply heuristic rules in their decision making which are prone to randomness and suboptimality. To model this type of network, we introduce the concept of heuristic rational actions as random actions that are optimal on average as we formally define next.

Definition 1: Consider network agent i associated with variable $\mathbf{x}_i$ and denote as $\mathbf{x}_{n(i)}(t)$ the values of neighboring variables at time t. We say that a probabilistic rule $\mathbf{x}_i(t) \in \mathcal{X}$ is heuristic rational if and only if its expectation is a rational action as defined in (3)

$$\mathbb{E}\left[\mathbf{x}_i(t) \,\middle|\, \mathbf{x}_{n(i)}(t)\right] = \tilde{\mathbf{x}}_i(t) = \underset{\mathbf{x}_i \in \mathcal{X}_i}{\operatorname{argmin}} f_i\left(\mathbf{x}_i, \mathbf{x}_{n(i)}(t)\right). \tag{4}$$

This paper considers network optimization models that consist of a random activation rule that determines when agents modify their variables and a heuristic rational rule that determines how the active agent updates its local values. Activations are indexed by the nonnegative integer variable $k \in \mathbb{N}$ with $k = 0$ denoting the initial state. Variable $k \neq 0$ denotes the kth activation that occurs at time t_k. Given a random activation rule, kth activation almost surely involves a unique agent $i = i_k$ modifying its local variable $\mathbf{x}_i = \mathbf{x}_{i_k}$. When an activation occurs, variables $\mathbf{x}_i(t_k)$ stay unchanged for all agents $i \neq i_k$ and are updated to $\mathbf{x}_{i_k}(t_k)$ for terminal i_k. Update rules are restricted to depend only on neighboring variables $\mathbf{x}_{n(i_k)}(t_k)$ and are assumed heuristic rational in the sense of Definition 1.

Based on the local costs in (2), we define the global cost

$$f(\mathbf{x}) := \sum_{i \in V} f_{0i}(\mathbf{x}_i) + \frac{1}{2} \sum_{i \in V, j \in n(i)} f_{ij}(\mathbf{x}_i, \mathbf{x}_j) \tag{5}$$

where the factor 1/2 is intended to account for the fact that the function $f_{ij}(\mathbf{x}_i, \mathbf{x}_j) = f_{ji}(\mathbf{x}_j, \mathbf{x}_i)$ is included twice in the sum in (5). The cost $f(\mathbf{x})$ measures the optimality of configuration $\mathbf{x} := \{\mathbf{x}_i\}_{i \in V}$ from a global perspective—as opposed to $f_i(\mathbf{x}_i, \mathbf{x}_{n(i)})$ that measures the optimality of configuration $\mathbf{x}_i$ from a local perspective. In particular, there exist globally optimal configurations $\mathbf{x}^*$ that achieve the minimum possible cost $p^* = f(\mathbf{x}^*)$ given by

$$p^* := \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}). \tag{6}$$

The goal of this paper is to compare the sequence of iterates $\mathbf{x}(t_k)$ generated by recursive application of heuristic rational rules with the optimal configuration $\mathbf{x}^*$. More to the point, we define the stochastic process $\{F_k\}_{k \in \mathbb{N}}$ of optimality gaps with elements

$$F_k := f(\mathbf{x}(t_k)) - p^*. \tag{7}$$

We will show that the optimality gap F_k achieves a small value with probability 1 infinitely often (Theorem 1 in Section III). Excursions away from this small value are possible and can be arbitrarily bad. However, we will also show that the largest value achieved in each of these excursions follows an exponential probability bound (Theorem 2 in Section IV). Before proceeding with the analysis, we discuss two examples of network optimization with heuristic rational agents.

A. Opinion in Social Networks

The propagation of opinions in a social network can be cast in the language of heuristic rational optimization. Consider a social network where some stubborn agents have fixed opinions while some other agents value agreement with friends with whom they are directly connected [20]. In this context, we interpret $x_i \in [-1, 1]$ as the opinion of a social agent. The subset $\mathcal{S}$ of stubborn agents have fixed extreme opinions $x_i \in \{-1, 1\}$ for all $i \in \mathcal{S}$. For the remaining agents, we model the desire for agreement through the penalty function $f_{ij}(x_i, x_j) = (1/2)(x_i - x_j)^2$. The resulting cost for disagreement for agent i is $f_i(x_i, \mathbf{x}_{n(i)}) = (1/2)\sum_{j \in n(i)}(x_i - x_j)^2$ as follows from (2) in which function $f_{0i}(x_i) = 0$. Through minimization of this quadratic cost, we have that the rational action, as defined by (3), for agent i at time t is

$$\tilde{x}_i(t) = \frac{1}{N_i} \sum_{j \in n(i)} x_j(t). \tag{8}$$

This action amounts to taking a local average of opinions in the network [9]. A heuristic rational rule $x_i(t)$ randomizes $\tilde{x}_i(t)$ to account for the fact that the average in (8) is not computed exactly but rather guessed. The presumption in Definition 1 is that these guesses are correct on average in that $\mathbb{E}[x_i(t)] = \tilde{x}_i(t)$.

A more interesting example of heuristic rationality stems from the observation that agents are not likely to consider opinions of all of their neighbors at each decision but rather rely on interactions with random subsets of friends. Accounting for the fact that interactions occur between a member of the network and subsets of its friends is the intent of voter models [19]. The model of opinion propagation in this case replaces the average in (8) by the average of a random sample of friends

$$x_i(t) = \frac{1}{\#(\tilde{n}_i(t))} \sum_{j \in \tilde{n}_i(t)} x_j(t) \tag{9}$$

where $\tilde{n}_i(t) \subseteq n_i$ denotes the random interaction group at time t. If all subsets of friends are equally likely to be chosen it follows that actions $\tilde{x}_i(t)$ in (8) and actions $x_i(t)$ in (9) are such that $\mathbb{E}[x_i(t)] = \tilde{x}_i(t)$. Thus, we can think of voter models [cf. (9), [19]] as heuristic rational rules for the local averaging model [cf. (8), [9], [19]].

In lieu of the quadratic cost leading to the rational action in (8), we could use the pairwise cost $f_i(x_i, x_j) = |x_i - x_j|$ given by the absolute difference between neighbor's opinions. This L_1 cost assigns more weight to small opinion discrepancies than an L_2 cost and less weight to large disagreements. In this case the local cost $f_i(x_i, x_{n(i)})$ in (2) takes the form

$$f_i\left(x_i, x_{n(i)}\right) = \sum_{j \in n(i)} |x_i - x_j|. \tag{10}$$

Given values $\mathbf{x}_{n(i)}(t)$ of the actions of neighboring terminals, a rational action $\tilde{x}_i(t)$ of agent i is to set its opinion to the median of the elements of $\mathbf{x}_{n(i)}(t)$. To be precise let $\mathbf{y}_{n(i)}(t) = [y^{(1)}_{n(i)}(t), \ldots, y^{(N_i)}_{n(i)}(t)]$ be an ordered version of the elements of $\mathbf{x}_{n(i)}(t)$ so that $y^{(1)}_{n(i)}(t) \leq \cdots \leq y^{(N_i)}_{n(i)}(t)$. Further denoting as $\text{floor}(N_i/2)$ the smallest integer not exceeding $N_i/2$ and as $\text{ceil}(N_i/2)$ the largest integer not smaller than $N_i/2$, we can write the rational action of agent i as

$$\begin{aligned}\tilde{x}_i(t) &= \text{median}\left(\mathbf{x}_{n(i)}\right) \\ &= \frac{1}{2}\left[y^{\text{floor}((N_i+1)/2)}_{n(i)}(t) + y^{\text{ceil}((N_i+1)/2)}_{n(i)}(t)\right]. \end{aligned} \tag{11}$$

When the number of neighbors N_i is odd $\text{floor}((N_i+1)/2) = \text{ceil}((N_i+1)/2) = (N_i+1)/2$ and (11) yields $\tilde{x}_i(t) = y^{((N_i+1)/2)}_{n(i)}(t)$ as the rational action. When the number of neighbors N_i is even, the rational action $\tilde{x}_i(t)$ is the average of the two actions in the middle of the ordered vector $\mathbf{y}_{n(i)}(t)$. In this latter case, the rational action $\tilde{x}_i(t)$ is not unique because we can set $\tilde{x}_i(t)$ to any value between $y^{\text{floor}((N_i+1)/2)}_{n(i)}(t)$ and $y^{\text{ceil}((N_i+1)/2)}_{n(i)}(t)$ without changing the cost $f_i(x_i, x_{n(i)})$ in (10).

A particular heuristic rational rule for the rational action in (11) is to select either $x_i(t) = y^{\text{floor}((N_i+1)/2)}_{n(i)}(t)$ or $x_i(t) = y^{\text{ceil}((N_i+1)/2)}_{n(i)}(t)$ with equal probability,

$$\begin{aligned}&\mathrm{P}\left(x_i(t) = y^{\text{floor}((N_i+1)/2)}_{n(i)}(t)\right) \\ &\qquad = \mathrm{P}\left(x_i(t) = y^{\text{ceil}((N_i+1)/2)}_{n(i)}(t)\right) = 1/2. \end{aligned} \tag{12}$$

It follows that the rational actions $\tilde{x}_i(t)$ in (11) and actions $x_i(t)$ in (12) are such that $\mathbb{E}[x_i(t)] = \tilde{x}_i(t)$. This is consistent with the definition of heuristic rationality in Definition 1. The heuristic rational rule in (12) for even N_i is similar to the voting model in (9) in that both consider random subsets of neighbors. The rules for selecting these subsets are different. In (9), neighbors are selected with equal probability irrespective of their opinions. In (12), only the neighbors with opinion closest to the median in (11) are chosen.

B. Field Estimation With Wireless Sensor Networks

Consider a WSN deployed to estimate a spatially varying field. Each sensor is interested in the value $\mathbf{x}_i$ of the field at its location which they estimate using locally collected observations $\mathbf{y}_i$. Observations $\mathbf{y}_i$ are assumed conditionally independent with probability density $\mathrm{P}(\mathbf{y}_i \,|\, \mathbf{x}_i) = e^{U(\mathbf{y}_i \,|\, \mathbf{x}_i)}$. We refer to $U(\mathbf{y}_i \,|\, \mathbf{x}_i)$ as the observation energy function of sensor i. Given the spatial correlation of field values, all observations $\mathbf{y} := \{\mathbf{y}_i\}_{i \in V}$ contain information about all field values $\mathbf{x} := \{\mathbf{x}_i\}_{i \in V}$. The network goal is consequently stipulated as the computation of maximum a posteriori (MAP) estimates $\hat{\mathbf{x}} := \arg\max P(\mathbf{x} \,|\, \mathbf{y})$ [29]. According to Bayes' rule MAP estimates $\hat{\mathbf{x}}$ can be computed as

$$\hat{\mathbf{x}} = \arg\max \mathcal{L}(\mathbf{x}) := \arg\min(-\ln \mathrm{P}(\mathbf{y} \,|\, \mathbf{x}) - \ln \mathrm{P}(\mathbf{x})) \tag{13}$$

where in the second equality, we used monotonicity of the logarithm function.

Even though according to problem definition sensor i is interested in $\hat{\mathbf{x}}_i$ only, these estimates are coupled in (13). This coupling can be handled by reducing attention to MRFs with spatial dependency coinciding with network connectivity. We therefore assume that the conditional distribution of $\mathbf{x}_i$ given all field values can be reduced to a distribution conditioned on neighboring nodes only, i.e., $\mathrm{P}(\mathbf{x}_i \,|\, \mathbf{x}_j, j \neq i) = \mathrm{P}(\mathbf{x}_i \,|\, \mathbf{x}_{n(i)})$, ([28], Ch. 1.2). A MRF with this correlation structure has an equivalent representation in terms of a Gibbs random field over the graph $G(V, E)$. In this representation, the probability distribution of the field can be written as $\mathrm{P}(\mathbf{x}) = \mathrm{e}^{-U(\mathbf{x})}/\alpha$, for some energy function $U(\mathbf{x}) = \sum_{i,j \in n(i)} u_{ij}(\mathbf{x}_i - \mathbf{x}_j)$ and normalizing constant α. Using this representation in (13), and further noting that observations are conditionally independent $\mathrm{P}(\mathbf{y} \,|\, \mathbf{x}) = \prod_i \mathrm{P}(\mathbf{y}_i \,|\, \mathbf{x}_i) = \prod_i e^{U(\mathbf{y}_i \,|\, \mathbf{x}_i)}$, it follows that

$$\hat{\mathbf{x}} = \text{argmax}\left(\sum_{i \in V} U(\mathbf{y}_i \,|\, \mathbf{x}_i) + \frac{1}{2}\sum_{i \in V, j \in n(i)} u_{ij}(\mathbf{x}_i - \mathbf{x}_j)\right). \tag{14}$$

Rearranging terms in (14), it follows that the maximand is of the form in (6) if we define the functions $f_{0i}(\mathbf{x}_i) = U(\mathbf{y}_i \,|\, \mathbf{x}_i)$ and $f_{ij}(\mathbf{x}_i, \mathbf{x}_j) = u_{ij}(\mathbf{x}_i - \mathbf{x}_j)$. The aggregate local cost in (2) would then be

$$f_i\left(\mathbf{x}_i, \mathbf{x}_{n(i)}\right) = U(\mathbf{y}_i \,|\, \mathbf{x}_i) + \sum_{j \in n(i)} u_{ij}(\mathbf{x}_i - \mathbf{x}_j). \tag{15}$$

A rational action amounts to sensors minimizing the costs $f_i(\mathbf{x}_i, \mathbf{x}_{n(i)})$ in (15) based on their individual measurement $\mathbf{y}_i$ and the neighboring estimates $\mathbf{x}_{n(i)}$. Heuristic rationality can be used to account for communication errors, quantization effects during the local signal processing or the communication stages, and model mismatch; see Section V-B.

Remark 1: The model of heuristic rationality proposed here is not to be confused with stochastic gradient descent algorithms; see e.g., [30]. These algorithms are a variation of gradient descent where the descent direction is replaced by a quantity whose expected value is a gradient of the cost function. Notice that this is fundamentally different from the stochastic heuristic rational rules of Definition 1. Agents are not descending along gradient directions but rather attempting to optimize the cost with respect to their local variables. The proper analogy is to a stochastic version of a coordinate descent algorithm, see e.g., [31]. Our motivation to study stochastic coordinate descent is that we believe it to be a better model of behavior in natural networks than stochastic gradient descent.

III. Near Optimality

The sequence of iterates $\mathbf{x}(t_k)$ generated by recursive application of heuristic rational rules is akin to a stochastic version of block coordinate descent on the function $f(\mathbf{x})$. In coordinate descent algorithms, minimization is attempted by alternation between descents on different subsets of variables chosen according to a given rule ([32], Ch. 1). The convergence properties of *deterministic* coordinate descent algorithms have been studied with cyclic and random activation rules [31], [33]–[35]. In the case of heuristic rational optimization, we can identify agents' variables as coordinate blocks and random activation as the selection rule. The structure of the local cost $f_i(\mathbf{x}_i, \mathbf{x}_{n(i)})$ in (2) allows for the distributed implementation of block coordinate descent. Given this correspondence, we expect convergence to a neighborhood of the optimal configuration $\mathbf{x}^*$, [cf. (6)] in some sense. In this section, we prove that this expectation is indeed true if the following assumptions on the cost function $f(\mathbf{x})$ and the random activation rule are satisfied.

(A1) Strong convexity. The global cost $f(\mathbf{x})$ is strongly convex in that there exists a constant $m > 0$ such that for any pair of points $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{X}$ it holds

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y}-\mathbf{x}) + \frac{m}{2}\|\mathbf{y}-\mathbf{x}\|_2^2. \tag{16}$$

(A2) Lipschitz gradients. Gradients of the global cost $f(\mathbf{x})$ are Lipschitz in that there exists a constant $M > 0$ such that for any pair of points $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{X}$ it holds

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y}-\mathbf{x}) + \frac{M}{2}\|\mathbf{y}-\mathbf{x}\|_2^2. \tag{17}$$

(A3) Random activation. At any given time t, all agents are equally likely to become active.

(A4) Bounded variance The mean squared error of the heuristic rational action $\mathbf{x}_{i_k}(t_k)$ with respect to the corresponding rational action $\tilde{\mathbf{x}}_{i_k}(t_k)$ is bounded [cf. (4)].

$$\mathbb{E}\left[\|\mathbf{x}_{i_k}(t_k) - \tilde{\mathbf{x}}_{i_k}(t_k)\|_2^2\right] \leq \sigma^2. \tag{18}$$

Assumptions (A1) and (A2) are typical in convergence analysis of coordinate descent algorithms [31]. They are satisfied by the examples discussed in Sections II-A and II-B except for opinion propagation with L_1 cost which violates Assumption (A1). Assumption (A3) states that activations occur at random times and that all agents are equally likely to become active in any given time interval. This assumption is also common; see, e.g., [36]. Among other possibilities, it can be satisfied if all agents have an activation clock based on independent exponential waiting times with equal means. This is more a matter of simplifying discussion than a fundamental requirement. It can be substituted by laxer conditions as we discuss in Remark 3. Assumption (A4) bounds the average irrationality of each agent by bounding the deviation from the rational decision (3). We emphasize that this bound holds on a mean square sense. It is possible to have isolated actions that are arbitrarily bad. Our results are parametric on the irrationality bound σ^2. As time increases, the optimality gap F_k of the global network behavior approaches a neighborhood of zero whose size is determined by the irrationality bound σ^2. Further note we are not imposing a connectivity requirement on the network. We explain why such assumption is not necessary in Remark 2.

The first result presented here considers global functional values $f(\mathbf{x}(k)) := f(\mathbf{x}(t_k))$ and $f(\mathbf{x}(k+1)) := f(\mathbf{x}(t_{k+1}))$ at subsequent update times and shows that the suboptimality of the corresponding configurations $\mathbf{x}(t_k)$ and $\mathbf{x}(t_{k+1})$ tends to be reduced in the sense specified in the following lemma.

Lemma 1: Consider a sequence of iterates $\mathbf{x}(t_k) = \{\mathbf{x}_i(t_k)\}_{i\in V}$ such that at time t_k agent i_k updates its local variables $\mathbf{x}_i(t_k)$ according to a heuristic rational update which is optimal on average as per Definition 1. If assumptions (A1)–(A4) hold, the optimality gaps F_{k+1} and F_k as defined in (7) satisfy

$$\mathbb{E}[F_{k+1} \,|\, \mathbf{x}(k)] \leq (1-\beta)F_k + \frac{\sigma^2 M}{2} \tag{19}$$

where we defined the condition number $\beta := m/(MN)$ and used the shorthand notation $\mathbf{x}(k) := \mathbf{x}(t_k)$.

Proof: At time t_{k+1}, define $\tilde{\mathbf{x}}_{i_{k+1}}(k+1)$ as the result of the application of a rational rule by agent i_{k+1}. Further denote as $\tilde{\mathbf{x}}(k+1)$ the vector with i_{k+1}th component $\tilde{\mathbf{x}}_{i_{k+1}}(k+1)$ and remaining components $\tilde{\mathbf{x}}_i(k+1) = \tilde{\mathbf{x}}_i(k)$. Consequently, the vector $\tilde{\mathbf{x}}(k+1)$ would be the network state after the application of a rational rule by node i_{k+1}, whereas the vector $\mathbf{x}(k+1)$ is the network state after a heuristic rational rule is applied. As per (3) and (4), we have that $\mathbb{E}[\mathbf{x}(k+1) \,|\, \mathbf{x}(k)] = \tilde{\mathbf{x}}(k+1)$. In the first part of the proof, we establish a relationship between the current iterate $\mathbf{x}(k)$ and the rational state $\tilde{\mathbf{x}}(k+1)$. In the second part of the proof, we use the connection between optimal $\tilde{\mathbf{x}}(k+1)$ and heuristic optimal $\mathbf{x}(k+1)$ variables to translate this result into a result relating $\mathbf{x}(k+1)$ with $\mathbf{x}(k)$.

For the first part of the proof, we use (17) to obtain a contraction between $f(\mathbf{x}(k))$ and $f(\tilde{\mathbf{x}}(k+1))$. Notice that since $\tilde{\mathbf{x}}(k+1)$ results from the application of the optimal action by node i_{k+1}, it must be

$$f(\tilde{\mathbf{x}}(k+1)) = \min_{\mathbf{x}_{i_{k+1}}} f\left(\mathbf{x}_{i_{k+1}}, \mathbf{x}_{-i_{k+1}}(k)\right). \tag{20}$$

Using the upper bound (17) restricted to the i_{k+1}th coordinate in lieu of the minimization objective in (20), we can write

$$\begin{aligned} f(\tilde{\mathbf{x}}(k+1)) \leq \min_{\mathbf{x}_{i_{k+1}}} \Big[& f(\mathbf{x}(k)) \\ & + \nabla_{i_{k+1}} f(\mathbf{x}(k))^T \left(\mathbf{x}_{i_{k+1}} - \mathbf{x}_{i_{k+1}}(k)\right) \\ & + \frac{M}{2}\left\|\mathbf{x}_{i_{k+1}} - \mathbf{x}_{i_{k+1}}(k)\right\|_2^2 \Big], \end{aligned} \tag{21}$$

where $\nabla_{i_{k+1}} f(\mathbf{x}(k))$ denotes the i_{k+1}th component of the gradient $\nabla f(\mathbf{x}(k))$. The minimization of the quadratic form in the right hand side of (21) can be performed explicitly to yield

$$f(\tilde{\mathbf{x}}(k+1)) \leq f(\mathbf{x}(k)) - \frac{1}{2M}\left\|\nabla_{i_{k+1}} f(\mathbf{x}(k))\right\|_2^2. \tag{22}$$

Subtracting p^* from both sides of (22) and using the definition $F_k := f(\mathbf{x}(t_k)) - p^*$ given in (7) we get

$$f(\tilde{\mathbf{x}}(k+1)) - p^* \leq \left(1 - \frac{1}{2M}\frac{\left\|\nabla_{i_{k+1}} f(\mathbf{x}(k))\right\|_2^2}{F_k}\right) F_k. \tag{23}$$

Equation (23) is true for any agent $i_k = 1, \ldots, N$. Notice that (23) gives us a strict inequality whenever the gradient $\nabla_{i_{k+1}} f(\mathbf{x}(k)) \neq 0$, which is to say when $\mathbf{x}_{i_{k+1}}(k)$ is not the rational action for agent i_{k+1}. However, it is possible that for some configuration, the rational updates may not result in $f(\tilde{\mathbf{x}}(k+1)) - p^*$ smaller than F_k. This may happen, e.g., if the same agent is activated twice consecutively. Nevertheless, since all agents have equal chance of becoming active as per Assumption (A3), the process will descend on average. To make this observation precise, observe that according to Assumption (A3), we have

$$\mathbb{E}\left[\|\nabla_{i_{k+1}} f(\mathbf{x}(k))\|_2^2 \mid \mathbf{x}(k)\right] = \sum_{i=1}^{N} \frac{1}{N} \|\nabla_i f(\mathbf{x}(k))\|_2^2 = \frac{1}{N} \|\nabla f(\mathbf{x}(k))\|_2^2. \quad (24)$$

Taking now expectation on both sides of (23) conditional on the network state $\mathbf{x}(k)$ and combining the result with the equality in (24) yields,

$$\mathbb{E}[f(\tilde{\mathbf{x}}(k+1)) - p^* \mid \mathbf{x}(k)] \leq \left(1 - \frac{1}{2MN} \frac{\|\nabla f(\mathbf{x}(k))\|_2^2}{F_k}\right) F_k. \quad (25)$$

Minimization of both sides of (16) with respect to $\mathbf{y}$ yields the bound $\|\nabla f(\mathbf{x})\|_2^2 \geq 2m(f(\mathbf{x}) - p^*)$; see e.g., ([37], Ch. 9). With $\mathbf{x} = \mathbf{x}(k)$, this bound takes the form $\|\nabla f(\mathbf{x}(k))\|_2^2 \geq 2m(f(\mathbf{x}(k)) - p^*) = 2mF_k$. Substituting this latter bound in (25) for the gradient norm $\|\nabla f(\mathbf{x}(k))\|_2^2$ leads to

$$\mathbb{E}[f(\tilde{\mathbf{x}}(k+1)) - p^* \mid \mathbf{x}(k)] \leq \left(1 - \frac{m}{MN}\right) F_k = (1-\beta) F_k \quad (26)$$

where we used the definition of the condition number $\beta := m/(MN)$ in the equality.

This completes the first part of the proof. To obtain a similar relation between $f(\mathbf{x}(k+1))$ and $f(\mathbf{x}(k))$, we use the definition of heuristic rationality in (4). Start using the second order mean value theorem centered at $f(\tilde{\mathbf{x}}(k+1))$ so as to write

$$\begin{aligned} f(\mathbf{x}(k+1)) &= f(\tilde{\mathbf{x}}(k+1)) \\ &+ \nabla f(\tilde{\mathbf{x}}(k+1))^T (\mathbf{x}(k+1) - \tilde{\mathbf{x}}(k+1)) \\ &+ \frac{1}{2} (\mathbf{x}(k+1) - \tilde{\mathbf{x}}(k+1))^T \nabla^2 f(\mathbf{x}) (\mathbf{x}(k+1) - \tilde{\mathbf{x}}(k+1)), \end{aligned} \quad (27)$$

which is true for some $\mathbf{x}$ lying in the segment between $\mathbf{x}(k+1)$ and $\tilde{\mathbf{x}}(k+1)$. Recall now that only the i_{k+1}th coordinate is changing at time $k+1$, which implies $\mathbf{x}_i(k+1) = \tilde{\mathbf{x}}_i(k+1)$ for all coordinates $i \neq i_{k+1}$. We can therefore simplify (27) by keeping only the i_{k+1}th coordinate components as the remaining terms are null. Using $\nabla^2_{i_{k+1}} f(\mathbf{x})$ to denote the i_{k+1}th diagonal block of the Hessian and defining the error of agent i_{k+1} as $\Delta x_{i_{k+1}}(k+1) := \mathbf{x}_{i_{k+1}}(k+1) - \tilde{\mathbf{x}}_{i_{k+1}}(k+1)$, this observation yields

$$\begin{aligned} f(\mathbf{x}(k+1)) &= f(\tilde{\mathbf{x}}(k+1)) \\ &+ \nabla^T_{i_{k+1}} f(\tilde{\mathbf{x}}(k+1)) \left(\Delta x_{i_{k+1}}(k+1)\right) \\ &+ \frac{1}{2} \left(\Delta x_{i_{k+1}}(k+1)\right)^T \nabla^2_{i_{k+1}} f(\mathbf{x}) \left(\Delta x_{i_{k+1}}(k+1)\right). \end{aligned} \quad (28)$$

But now notice that since $\tilde{\mathbf{x}}_{i_{k+1}}(k+1)$ is a rational action for agent i_{k+1}, we must have $\nabla^T_{i_{k+1}} f(\tilde{\mathbf{x}}(k+1)) = 0$. We can then further simplify (28) to

$$\begin{aligned} f(\mathbf{x}(k+1)) &= f(\tilde{\mathbf{x}}(k+1)) \\ &+ \frac{1}{2} \left(\Delta x_{i_{k+1}}(k+1)\right)^T \nabla^2_{i_{k+1}} f(\mathbf{x}) \left(\Delta x_{i_{k+1}}(k+1)\right). \end{aligned} \quad (29)$$

Now, subtract p^* from both sides of (29) so that the left hand side becomes $f(\mathbf{x}(k)) - p^* := F_k$ as defined in (7). Further, taking expectation conditioned on $\mathbf{x}(k)$ on both sides of (29) leads to

$$\begin{aligned} \mathbb{E}[F_{k+1} \mid \mathbf{x}(k)] \leq\ & \mathbb{E}[f(\tilde{\mathbf{x}}(k+1)) - p^* \mid \mathbf{x}(k)] \\ &+ \mathbb{E}\left[\frac{1}{2}\left(\Delta x_{i_{k+1}}(k+1)\right)^T \nabla^2_{i_{k+1}} f(\mathbf{x}) \left(\Delta x_{i_{k+1}}(k+1)\right) \middle| \mathbf{x}(k)\right]. \end{aligned} \quad (30)$$

We can now substitute (26) for the first term in the right hand side of (30) and use Assumption (A4) to bound the second term to obtain

$$\begin{aligned} \mathbb{E}[F_{k+1} \mid \mathbf{x}(k)] \leq\ & (1-\beta) F_k \\ &+ \frac{M}{2} \mathbb{E}\left[\left\|\Delta x_{i_{k+1}}(k+1)\right\|^2 \,\middle|\, \mathbf{x}(k)\right]. \end{aligned} \quad (31)$$

The result in (19) follows from using the mean squared error bound in (18) to bound the corresponding term in (31). ■

The result in Lemma 1 shows that the network behavior is different depending on whether $\mathbf{x}(k+1)$ is far from the optimal value or close to it. When the optimality gap is $F_k > \sigma^2 M/2\beta$, rationality dominates in the sense that the right hand side of (19) is smaller than F_k. Thus, when an update occurs, the gap F_{k+1} is expected to become smaller. In this regime, the observed behavior seems rational as the global metric $f(\mathbf{x}(k+1))$ becomes closer to the optimal p^*. When $F_k \leq \sigma^2 M/2\beta$, we do not expect a decrease on the optimality gap F_{k+1}. This is because when $f(\mathbf{x}(k))$ gets close to p^*, the network's behavior is dominated by randomness. In this regime the network behavior seems erratic or irrational in that the yield stops improving.

Considering that whenever F_k exceeds $\sigma^2 M/2\beta$ the process is attracted towards zero, we expect to see the process F_k becoming smaller than $\sigma^2 M/2\beta$ at least once. This intuition is correct as we state and prove in the following theorem.

Theorem 1: Consider the heuristic rational sequence of iterates $\mathbf{x}(k)$ [cf. Definition 1] with corresponding optimality gaps F_k [cf. (7)]. Define the best optimality gap by time t_k as $F_k^{\text{best}} := \min_{l \in [0,k]} F_l$. With the same hypotheses and definitions in Lemma 1 it holds

$$\lim_{k \to \infty} F_k^{\text{best}} \leq \frac{M\sigma^2}{2\beta}, \quad \text{a.s.} \quad (32)$$

i.e., the optimality gap becomes smaller than $M\sigma^2/2\beta$ at least once for almost all realizations.

Proof: Let $0 < \delta < 1$ be an arbitrary constant and define the sequence

$$\alpha(k) := F_k \mathbb{I}\left(F_k^{\text{best}} > \frac{\sigma^2 M}{2\beta(1-\delta)}\right). \tag{33}$$

The sequence $\alpha(k)$ tracks the optimality gap F_k until the first time at which $F_k \leq \sigma^2 M/2\beta(1-\delta)$. Thereafter, we have $F_k^{\text{best}} \leq \sigma^2 M/2\beta(1-\delta)$ and the sequence is locked at $\alpha(k) = 0$. We will show that the stochastic process with realizations $\alpha(k)$ is a supermartingale whose expectation $\mathbb{E}[\alpha(k) \,|\, \alpha(0)]$ converges to 0. The result in (32) will follow from there.

Consider the expectation of the process' value $\alpha(k+1)$ at update $k+1$, given the network state at time k, i.e., $\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k)]$. Separate this expectation into the cases $\alpha(k) \neq 0$ and $\alpha(k) = 0$ and use total probability to write

$$\begin{aligned}\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k)] &= \mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k), \alpha(k) = 0] P[\alpha(k) = 0] \\ &\quad + \mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k), \alpha(k) \neq 0] P[\alpha(k) \neq 0]. \end{aligned} \tag{34}$$

Consider the case when $\alpha(k) = 0$ in (34). It then must be that the indicator function is $\mathbb{I}(F_k^{\text{best}} > \sigma^2 M/2\beta(1-\delta)) = 0$, because it is not possible to have $F_k = 0$ without the indicator function being null. If the indicator function is null, then we must have $F_k^{\text{best}} \leq \sigma^2 M/2\beta(1-\delta)$ from which in turn it follows that $F_l^{\text{best}} \leq \sigma^2 M/2\beta(1-\delta)$ for all $l \geq k$ since $F_{k+1}^{\text{best}} \leq F_k^{\text{best}}$ by definition. This guarantees that for any update after k, the indicator function is locked to zero. Hence, the conditional expectation is zero from that point on,

$$\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k), \alpha(k) = 0] = 0. \tag{35}$$

In the other case, when $\alpha(k) \neq 0$, we can write the conditional expected value of $\alpha(k+1)$ as

$$\begin{aligned}&\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k), \alpha(k) \neq 0] \\ &\quad = \mathbb{E}\left[F_{k+1} \mathbb{I}\left(F_{k+1}^{\text{best}} > \frac{\sigma^2 M}{2\beta(1-\delta)}\right) \,\middle|\, \mathbf{x}(k), \alpha(k) \neq 0\right] \\ &\quad \leq \mathbb{E}[F_{k+1} \,|\, \mathbf{x}(k)] \\ &\quad \leq (1-\beta)F_k + \frac{\sigma^2 M}{2}. \end{aligned} \tag{36}$$

The equality in (36) follows from the definition of $\alpha(k+1)$ in (33). The first inequality is true by the fact that indicator functions cannot be greater than 1, and the second inequality follows from applying Lemma 1 to $\mathbb{E}[F_{k+1} \,|\, \mathbf{x}(k)]$.

Notice now that if $\alpha(k) \neq 0$, it must be that the indicator function is $\mathbb{I}(F_k^{\text{best}} > \sigma^2 M/2\beta(1-\delta)) = 1$. This further implies that the best optimality gap attained so far is $F_k^{\text{best}} > \sigma^2 M/2\beta(1-\delta)$. In particular, this is true for the current iterate, $F_k > \sigma^2 M/2\beta(1-\delta)$. Rearranging terms in this latter inequality, we conclude that $\sigma^2 M/2 < \beta(1-\delta)F_k$, which upon substitution in (36) yields

$$\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k), \alpha(k) \neq 0] \leq (1-\beta\delta)F_k. \tag{37}$$

Further notice that if the indicator function is $\mathbb{I}(F_k^{\text{best}} > \sigma^2 M/2\beta(1-\delta)) = 1$, it also holds that $\alpha(k) = F_k$ allowing us to rewrite (37) as

$$\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k), \alpha(k) \neq 0] \leq (1-\beta\delta)\alpha(k). \tag{38}$$

The result in (38) states that the expectation of $\alpha(k+1)$ contracts by a factor $(1-\beta\delta)$ when $\alpha(k)$ is not null, or equivalently when $F_k > \sigma^2 M/2\beta(1-\delta)$. The result in (35) states that otherwise, the expectation is null. These two results substituted in (34) combined with the observation that probabilities are smaller than one leads to

$$\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(k)] \leq (1-\beta\delta)\alpha(k). \tag{39}$$

The difference between (38) and (39) is that (38) is a statement conditional on the value of $\alpha(k)$, whereas (39) holds for any value of $\alpha(k)$. It thus follows from (39) that $\alpha(k)$ is a supermartingale and, as a consequence of the supermartingale convergence theorem, that the limit $\lim_{k\to\infty} \alpha(k)$ exists almost surely ([38], Ch. 5.2). We are left to prove that this limit is 0.

To see that this is true, observe that a second consequence of (39) is that the limit of the expectation $\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(0)]$ is null. Indeed, applying (39) recursively yields

$$\mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(0)] \leq (1-\beta\delta)^k \alpha(0). \tag{40}$$

Since the process $\alpha(k)$ is nonnegative by definition, it follows that we must have

$$\begin{aligned}0 &\leq \lim_{k\to\infty} \mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(0)] \\ &\leq \lim_{k\to\infty} (1-\beta\delta)^k \alpha(0) = 0, \end{aligned} \tag{41}$$

which establishes that $\lim_{k\to\infty} \mathbb{E}[\alpha(k+1) \,|\, \mathbf{x}(0)] = 0$.

If almost all realizations $\alpha(k)$ are nonnegative and converge—as it follows from the supermartingale convergence theorem—and their expected values converge to 0—as it follows from (41), it must be that $\lim_{k\to\infty} \alpha(k) = 0$. Recalling the definition of $\alpha(k)$ in (33), this is equivalent to the indicator function becoming null i.e., $\mathbb{I}(F_k^{\text{best}} > \sigma^2 M/2\beta(1-\delta)) = 0$, or equivalently to

$$\lim_{k\to\infty} F_k^{\text{best}} < \frac{\sigma^2 M}{2\beta(1-\delta)}, \quad \text{a.s.} \tag{42}$$

The result in (32) follows because (42) is true for arbitrary small $\delta > 0$. ∎

According to Theorem 1 it holds that, for almost all realizations, the optimality gap F_k approaches or becomes smaller than $\sigma^2 M/2\beta$ at least once as k grows. Theorem 1 also implies that this happens infinitely often. Indeed, if at given time k_0, we have $F_{k_0} > \sigma^2 M/2\beta$, a simple time shift in Theorem 1 permits concluding that there exists a future time k at which $F_k \leq \sigma^2 M/2\beta$.

For F_k to become small, we need to have the current network configuration $\mathbf{x}(k)$ close to the optimal configuration $\mathbf{x}^*$. Consequently, Theorem 1 implies that $\mathbf{x}(k)$ enters into a neighborhood of the optimal configuration infinitely often; see

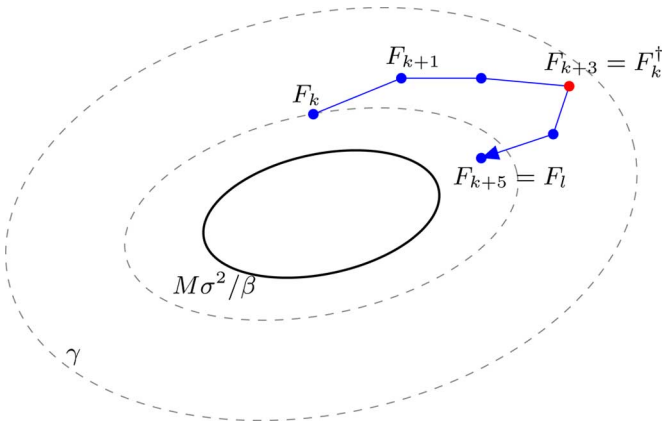

Fig. 1. Excursion from near optimality. Level sets for the optimality gap $F_k := f(\mathbf{x}(t_k)) - p^*$ are shown along with a piece of a sample path of optimality gaps. The innermost level set encloses the near optimality region and the middle level set corresponds to the value F_k taken at the current update. An excursion is defined as the path of the process until we return to a value smaller than F_k. The outermost level set is a given γ level. We want to study the probability of the worst value $F_k^\dagger = F_{k+3}$ attained during the excursion exceeding γ.

Fig. 1. The volume of this neighborhood increases with increasing mean squared error of the heuristic rule σ^2, increasing Lipschitz constant M, or decreasing condition number β. The condition number $\beta := m/(MN)$ is small for functions $f(\mathbf{x})$ having $m \ll M$ corresponding to ill conditioned functions with elongated level sets. Therefore, the dependence on β captures the difficulty of minimizing the cost $f(\mathbf{x})$. The constant M is of little consequence as it plays the role of a normalizing constant. If we multiply the function $f(\mathbf{x})$ with a constant, both, the optimality gaps F_k and the Lipschitz constant M are multiplied by the same constant. The dependence on the mean squared error σ^2 captures the increase in global suboptimality as agents' behaviors become more erratic.

If the optimality gap F_k approaches a small value infinitely often but can stray away from it, the question arises of what the process behavior is between visits to the optimality neighborhood. We answer this question in the following section after the following remarks.

Remark 2: Lemma 1 and Theorem 1 do not require a connected network and hold true if the network contains multiple connected components. If the network contains multiple connected components, each of them descends towards its optimal configuration. Given the structure of the cost function in (5), this is equivalent to the optimal configuration for the network as a whole. This is true because in a disconnected network there is no cost coupling between members of different connected components. Perhaps the most restrictive condition in Lemma 1 and Theorem 1 is the strong convexity requirement in Assumption (A1). Strong convexity is a necessary assumption in analyzing convergence properties of deterministic coordinate descent algorithms. In general, convergence fails in functions that are not strongly convex. Since variable updates in a network of heuristic rational agents are akin to a stochastic version of coordinate descent, we expect Lemma 1 and Theorem 1 to not hold in this case. There are some studies that achieve convergence under laxer conditions such as pseudo-convexity but they assume special restrictive structures for the cost function [31], [34]. A generalization of Lemma 1 and Theorem 1 to these scenarios is beyond the scope of this paper.

Remark 3: The assumption that all agents are equally likely to become active can be relaxed to the assumption that all agents have possibly different but strictly positive probabilities of becoming active. This less restrictive assumption still ensures that when the configuration $\mathbf{x}(t)$ is not optimal there is always a positive probability of the *rational* rule descending towards the optimum. Formally, when considering different activation probabilities, the expectation in (24) would be replaced by a weighted sum of block coordinate gradients. The rest of the proof of Lemma 1 would continue unchanged by redefining the gradient norm bound M. A modified version of Theorem 1 would follow with the same arguments.

IV. Excursions From Near Optimality

Although Theorem 1 shows that the network state moves within a close boundary of the optimal configuration almost surely and infinitely often, it does not claim a guarantee on staying close to the optimal value. In fact, it is easy to see that in some particular examples, the process F_k is almost sure to move out of the optimality neighborhood $F_k \leq M\sigma^2/2\beta$ and even become arbitrarily bad with small but nonzero probability. This may happen in the unlikely but not impossible situation in which the variations in the heuristic rational rule cancel out the intended drive towards optimality. In this section, we derive an exponential probability bound on these excursions from optimality. The bound shows that while arbitrarily bad excursions may be possible, they happen with exponentially small probability.

To formally define excursions away from the optimality neighborhood, suppose that at given iteration k, the optimality gap is $F_k = (1+\rho)M\sigma^2/2\beta$, i.e., larger than the neighborhood border by a factor $\rho > 0$. Further consider a given value $\gamma > F_k$. We define excursion as the trajectory $F_k, F_{k+1}, \ldots, F_{k+L}$ of the optimality gap until the process returns to a value $F_{k+L} < F_k$ smaller than the given gap F_k from which the excursion started. Notice that L is a random stopping time given by $L = \min(l > 0 : F_{k+l} < F_k)$. In particular, we are interested in the worst value $F_k^\dagger = \max(F_k, F_{k+1}, \ldots, F_{k+L})$ reached during the excursion. In formal terms, we define $F_k^\dagger$ as

$$F_k^\dagger := \max(F_{k+l} : 0 \leq l \leq L). \tag{43}$$

Our goal here is to determine the probability $\mathrm{P}(F_k^\dagger \geq \gamma)$ that the worst value attained during the excursion exceeds the given γ. This definition and corresponding goal are illustrated in Fig. 1. The innermost curve is the level set of network configurations with $M\sigma^2/2\beta$, the middle curve corresponds to states with optimality gap $F_k = (1+\rho)M\sigma^2/2\beta$, and the outermost curve is the level set for gap value γ. The process strays further away until returning to a value smaller than F_k at time $k + L = k + 5$. During this excursion the largest optimality gap is at time $k + l = k + 3$ meaning that $F_k^\dagger = F_{k+3}$. This particular realization does not exceed γ, i.e., $F_k^\dagger = F_{k+3} < \gamma$, but others may.

To bound the probability $\mathrm{P}(F_k^\dagger \geq \gamma)$, we need the following additional assumption.

(A5) Bounded Increments. The difference on optimality gaps between successive iterations is almost surely bounded by a finite constant $\kappa > 0$, i.e., for all times k we have that

$$\mathrm{P}(|F_{k+1} - F_k| \leq \kappa \,|\, F_k) = 1. \tag{44}$$

A particular case in which Assumption (A5) is satisfied is when the functions $f_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ are bounded for all feasible values $\mathbf{x}_i \in \mathcal{X}_i$ and $\mathbf{x}_j \in \mathcal{X}_j$. Assumption (A5) can be alternatively satisfied if the differences $\|\mathbf{x}_{i_k}(t_k) - \tilde{\mathbf{x}}_{i_k}(t_k)\|$ between rational and heuristic rational actions are almost surely bounded. This latter condition is more stringent than the finite variance requirement of Assumption (A4). For the opinion propagation scenario in Section II-A, the bound in (44) is the maximum number of neighbors, i.e., $\kappa = \max_i(N_i)$. This corresponds to the most connected agent flipping its opinion from -1 to 1. For the MRF example in Section II-B, Assumption (A5) is satisfied if we bound the range of values measured by sensors.

The exponential bound on $\mathrm{P}(F_k^\dagger \geq \gamma)$ is stated in the following theorem.

Theorem 2: Consider a process of heuristic rational updates $\mathbf{x}(t_k) = \{\mathbf{x}_i(t_k)\}_{i\in V}$ [cf. Definition 1] and the associated process of optimality gaps F_k [cf. (7)]. Assume that at time k the value of F_k exceeds the optimality neighborhood of Theorem 1 by a factor $\rho > 0$, i.e., $F_k = (1+\rho)M\sigma^2/2\beta$, and let $F_k^\dagger$ be the worst optimality gap achieved during the subsequent excursion as defined in (43). If assumptions (A1)–(A5) hold, then, for arbitrary given constant γ we have

$$\mathrm{P}(F_k^\dagger \geq \gamma \,|\, F_k) \leq \mathsf{e}^{-c(\gamma - F_k)} \tag{45}$$

with $c = 2\rho M\sigma^2/[(\rho M\sigma^2)^2 + \kappa^2]$.

When the optimality gap F_k is outside the neighborhood $F_k > M\sigma^2/2\beta$, it behaves like a supermartingale. It would then be easy to obtain a linear bound for the probability $\mathrm{P}(F_k^\dagger \geq \gamma \,|\, F_k)$ using Markov's inequality for martingales; see, e.g., ([38], Ch. 1). An exponential bound as the one in (45) could be obtained if we can claim that an exponential transformation of F_k yields another supermartingale-like relation. That this is indeed possible is shown in the following lemma.

Lemma 2: Let X_t be a sequence of random variables such that for some t, there are positive constants $\upsilon > 0$ and $\mu > 0$ for which $\mathrm{P}(|X_{t+1} - X_t| \leq \upsilon \,|\, X_t) = 1$ and $\mathbb{E}[X_{t+1} \,|\, X_t] \leq X_t - \mu$. Then, for any constant $c \leq 2\mu/(\mu^2 + \upsilon^2)$, the transformed sequence with values $Y_t := cX_t$ satisfies $\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t] \leq \mathsf{e}^{Y_t}$.

Proof: Observe for later reference that the conditions $\mathrm{P}(|X_{t+1} - X_t| \leq \upsilon \,|\, X_t) = 1$ and $\mathbb{E}[X_{t+1} \,|\, X_t] \leq X_t - \mu$ imposed on the sequence X_t can be translated into corresponding conditions for the sequence Y_t

$$\mathrm{P}(|Y_{t+1} - Y_t| \leq c\upsilon \,|\, Y_t) = 1, \tag{46}$$

$$\mathbb{E}[Y_{t+1} \,|\, Y_t] \leq Y_t - c\mu. \tag{47}$$

Consider the conditional expectation $\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t]$ which we want to show is smaller than e^{Y_t}. Writing $\mathsf{e}^{Y_{t+1}} = \mathsf{e}^{Y_{t+1}+Y_t-Y_t} = \mathsf{e}^{Y_t}\mathsf{e}^{Y_{t+1}-Y_t}$ and using the fact that with Y_t given, e^{Y_t} is a constant, we can write this expectation as

$$\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t] = \mathsf{e}^{Y_t}\mathbb{E}[\mathsf{e}^{Y_{t+1}-Y_t} \,|\, Y_t]. \tag{48}$$

Focus now on the function $\mathsf{e}^{Y_{t+1}-Y_t}$. Because the exponential function is convex, any point in the line that connects two points in the graph of the exponential function lies above the function. For the line segment defined by the points $(-c\upsilon, \mathsf{e}^{-c\upsilon})$ and $(c\upsilon, \mathsf{e}^{c\upsilon})$ this property can be written as

$$\mathsf{e}^{x} \leq \frac{\mathsf{e}^{-c\upsilon} + e^{c\upsilon}}{2} + \frac{e^{c\upsilon} - \mathsf{e}^{-c\upsilon}}{2c\upsilon}x \tag{49}$$

for any $x \in [-c\upsilon, c\upsilon]$, or equivalently, for any x with absolute value $|x| \leq c\upsilon$. Further notice that according to (46), the difference $Y_{t+1} - Y_t$ satisfies $|Y_{t+1} - Y_t| \leq c\upsilon$ almost surely. Therefore, we can make $x = Y_{t+1} - Y_t$ in (49) to conclude that with probability 1

$$\mathsf{e}^{Y_{t+1}-Y_t} \leq \frac{\mathsf{e}^{-c\upsilon} + e^{c\upsilon}}{2} + \frac{e^{c\upsilon} - \mathsf{e}^{-c\upsilon}}{2c\upsilon}[Y_{t+1} - Y_t]. \tag{50}$$

Since (50) holds for almost all differences $Y_{t+1} - Y_t$, it must also be true on expectation. Therefore

$$\mathbb{E}[\mathsf{e}^{Y_{t+1}-Y_t} \,|\, Y_t] \leq \frac{(\mathsf{e}^{-c\upsilon} + \mathsf{e}^{c\upsilon})}{2} + \frac{(\mathsf{e}^{c\upsilon} - \mathsf{e}^{-c\upsilon})(\mathbb{E}[Y_{t+1} \,|\, Y_t] - Y_t)}{2c\upsilon}. \tag{51}$$

Substitute now the bound in (51) for the corresponding term in (48) and subtract e^{Y_t} from both sides of the equation to obtain, after grouping terms

$$\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t] - \mathsf{e}^{Y_t} \leq \mathsf{e}^{Y_t}\left(\frac{\mathsf{e}^{-c\upsilon} + \mathsf{e}^{c\upsilon}}{2} + \frac{(\mathsf{e}^{c\upsilon} - \mathsf{e}^{-c\upsilon})(\mathbb{E}[Y_{t+1} \,|\, Y_t] - Y_t)}{2c\upsilon} - 1\right). \tag{52}$$

Using the bound for $\mathbb{E}[Y_{t+1} \,|\, Y_t]$ in (47), we can bound the term $\mathbb{E}[Y_{t+1} \,|\, Y_t] - Y_t) \leq -c\mu$. This bound substituted in (52) yields

$$\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t] - \mathsf{e}^{Y_t} \leq \mathsf{e}^{Y_t}\left(\frac{\mathsf{e}^{-c\upsilon} + \mathsf{e}^{c\upsilon}}{2} - \frac{\mu(\mathsf{e}^{c\upsilon} - \mathsf{e}^{-c\upsilon})}{2\upsilon} - 1\right). \tag{53}$$

The desired result follows if the right hand side of (53) is nonpositive. Since $\mathsf{e}^{Y_t} > 0$, this imposes the following condition in the constants μ, υ, and c

$$\frac{\mathsf{e}^{-c\upsilon} + \mathsf{e}^{c\upsilon}}{2} - \frac{\mu(\mathsf{e}^{c\upsilon} - \mathsf{e}^{-c\upsilon})}{2\upsilon} - 1 \leq 0. \tag{54}$$

We interpret μ and υ as given constants for the sequence X_t. Thus, (54) implies that for any constant c satisfying this inequality, the sequence cY_t satisfies $\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t] \leq \mathsf{e}^{Y_t}$. We will see that the hypothesis $c \leq 2\mu/(\mu^2 + \upsilon^2)$ is sufficient for (54) to be true. To do so divide both sides of (54) by $(\mathsf{e}^{-c\upsilon} + \mathsf{e}^{c\upsilon})/2$ and use the definitions of the hyperbolic tangent $\tanh(\cdot)$ and secant $\mathrm{sech}(\cdot)$ to conclude that (54) is equivalent to

$$\left(1 - \frac{\mu}{\upsilon}\tanh(c\upsilon)\right) \leq \mathrm{sech}(c\upsilon). \tag{55}$$

We square both sides of (55) which is valid because $\mathrm{sech}(x) \geq 0$ for any $x \in \mathbb{R}$, expand terms and use the relationship $\tanh^2(c\upsilon) = 1 - \mathrm{sech}^2(c\upsilon)$ to get that (54) is further equivalent to

$$\tanh(c\upsilon) \leq \frac{2\mu}{\upsilon}\left(1+\frac{\mu^2}{\upsilon^2}\right)^{-1} = \frac{2\mu\upsilon}{\mu^2+\upsilon^2}. \tag{56}$$

Notice now that the hyperbolic tangent satisfies $\tanh(c\upsilon) \leq c\upsilon$ for any nonnegative $c\upsilon$. Therefore, if we guarantee $c\upsilon \leq (2\mu\upsilon/(\mu^2+\upsilon^2))$, the inequality in (56) will be also satisfied. This latter condition is equivalent to $c \leq 2\mu/(\mu^2+\upsilon^2)$. Therefore, if the constant c satisfies $c \leq 2\mu/(\mu^2+\upsilon^2)$, the inequality in (56) is satisfied. Since (56) is equivalent to (54), $c \leq 2\mu/(\mu^2+\upsilon^2)$ also implies (54) is satisfied, which in turn implies $\mathbb{E}[\mathsf{e}^{Y_{t+1}} \,|\, Y_t] - \mathsf{e}^{Y_t} \leq 0$ because of (53). This is the result we wanted to prove. Notice that the laxer condition $c \leq \upsilon^{-1}\mathrm{arctanh}(2\mu\upsilon/(\mu^2+\upsilon^2))$ can be imposed on c. The condition $c \leq 2\mu/(\mu^2+\upsilon^2)$ suffices for our purposes and is close to the condition in (56) for small values of c. ■

Lemma 2 shows that a suitably chosen exponential transform $e^{Y_t} = e^{cX_t}$ of a sequence of random variables X_t that has almost sure finite increments and strict expected decrease adheres to a supermartingale-like inequality. From Assumption (A5), we are given that increments of the process F_k are almost surely finite. Since during the excursion, it holds $F_l \geq F_k = (1+\rho)M\sigma^2/2\beta$, it further follows from Lemma 1 that F_l is expected to decrease during the excursion. Hence, the two hypotheses of Lemma 2 are satisfied for the excursion sequence $F_k, \ldots, F_{k+L}$, from where it follows that $e^{cF_{k+l}}$ satisfies a supermartingale-like inequality. This observation is combined with a stopping time argument and the use of Markov's inequality to prove Theorem 2.

Proof of Theorem 2: Our interest is in the excursion $\{F_k, F_{k+1}, \ldots, F_{k+L}\}$ starting at F_k and finishing at time $k+L$ where $L = \min(l > 0 : F_{k+l} < F_k)$. During this proof, we consider the process of optimality gaps $\{F_{k+l}\}_{l\geq 0}$ starting at given F_k and the exponential transformation $G_{k+l} = \mathsf{e}^{cF_{k+l}}$ with $c = 2\rho M\sigma^2/[(\rho M\sigma^2)^2 + \kappa^2]$ as given in the statement of Theorem 2. We further define the stopping time

$$L^* = \min\{l > 0 : G_{k+l} < G_k, G_{k+l} > \mathsf{e}^{c\gamma}\}. \tag{57}$$

The stopping criterion $G_{k+l} < G_k$ is equivalent to $F_{k+l} \leq F_k$ and corresponds to the end of the excursion at time $k+L$ [cf. (43)]. In this case, the two stopping times L and L^* are equal, $L^* = L$. The stopping condition $G_{k+l} > \mathsf{e}^{c\gamma}$ is equivalent to $F_{k+l} > \gamma$ and corresponds to the excursion exceeding the value γ before time $k+L$. In the first stopping case, the excursion finishes before reaching γ, while in the second case the excursion reaches the value γ before finishing at $k+L$. It follows that the worst optimality gap during the excursion exceeds gamma, i.e., $F_k^\dagger > \gamma$, if and only if the value $F_{k+L^*} > \gamma$, or equivalently if $G_{k+L^*} > e^{c\gamma}$. Thus, we can write the excursion probability in (45) as

$$\mathrm{P}(F_k^\dagger \geq \gamma \,|\, F_k) = \mathrm{P}(G_{k+L^*} > e^{c\gamma} \,|\, F_k). \tag{58}$$

Also notice that L^* is a properly defined stopping time in that it is almost surely finite ([38], Ch. 4). This is true because Theorem 1 assures that for any $\rho > 0$, $F_{k+l} \leq (1+\rho)M\sigma^2/2\beta$ for some finite l with probability 1. This is the first stopping condition in (57).

Further define the stopped process $\tilde{G}_{k+l}$ that coincides with G_{k+l} before the stopping time L^* and with the stopping value G_{k+L^*} afterwards, i.e.,

$$\tilde{G}_{k+l} := \begin{cases} G_{k+l} & \text{if } l < L^*, \\ G_{k+L^*} & \text{if } l \geq L^*. \end{cases} \tag{59}$$

We will show that the stopped process $\tilde{G}_{k+l}$ is a supermartingale and conclude the proof with a simple Markov's inequality argument. Start by noticing that the bound (19) claimed by Lemma 1 can be equivalently written as

$$\mathbb{E}[F_{k+1} - F_k \,|\, \mathbf{x}(k)] \leq -\beta F_k + \frac{M\sigma^2}{2} \tag{60}$$

because F_k is set if $\mathbf{x}(k)$ is given.

Consider the excursion $\{F_k, F_{k+1}, \ldots, F_{k+L}\}$ starting at F_k and finishing at time $k+L$. Since we are assuming by hypothesis that $F_k = (1+\rho)M\sigma^2/2\beta$, it follows that for all elements of this excursion except F_{k+L}, it holds $F_{k+l} \geq (1+\rho)M\sigma^2/2\beta$. Combining this observation with the bound in (60), we have that for all $l \in [0, L-1]$, it holds

$$\mathbb{E}\left[F_{k+l+1} - F_{k+l} \,|\, \mathbf{x}(k+l)\right] \leq -\rho\frac{M\sigma^2}{2}. \tag{61}$$

Further recalling that Assumption (A5) states that $\mathrm{P}(|F_{k+1} - F_k| \leq \kappa \,|\, \mathbf{x}(k)) = 1$ it follows that during the excursion the elements of the optimality gap satisfy the conditions of Lemma 2 with $\upsilon = \kappa$ and $\mu = \rho M\sigma^2/2$. Consequently, when $l \leq L$, the process G_{k+l} satisfies the supermatingale condition

$$\mathbb{E}[G_{k+l+1} \,|\, G_{k+l}] \leq G_{k+l}. \tag{62}$$

When $l < L^*$, the stopped process (59) is equal to the transformed process G_{k+l} hence $\mathbb{E}[\tilde{G}_{k+l+1} \,|\, \tilde{G}_{k+l}] \leq \tilde{G}_{k+l}$ holds in this case. When the process has stopped $\mathbb{E}[\tilde{G}_{k+l+1} \,|\, \tilde{G}_{k+l}] = \tilde{G}_{k+l}$ for $l \geq L^*$ because all values after the stopping time are equal to $\tilde{G}_{k+L^*}$. Therefore, a relationship analogous to the one in (62) is also true for the stopped process,

$$\mathbb{E}[\tilde{G}_{k+l+1} \,|\, \tilde{G}_{k+l}] \leq \tilde{G}_{k+l}. \tag{63}$$

The difference between (62) and (63) is that the former holds for all $l \leq L$, whereas the latter is true for all times. Applying (63) recursively, it follows that a particular consequence of (63) is $\mathbb{E}[\tilde{G}_{k+l} \,|\, \tilde{G}_k] \leq \tilde{G}_k$ for all $l \geq 0$, which implies that the same is true for the value of $\tilde{G}_{k+L^*}$ at the stopping time L^* ([39], Ch. 6),

$$\mathbb{E}[\tilde{G}_{k+L^*} \,|\, \tilde{G}_k] \leq \tilde{G}_k. \tag{64}$$

If we now apply Markov's inequality to the probability $\mathrm{P}(G_{k+L^*} > e^{c\gamma} \,|\, F_k)$ in (58) we obtain

$$P(\tilde{G}_{k+L^*} \geq \mathsf{e}^{c\gamma} \,|\, \tilde{G}_k) \leq \mathbb{E}[\tilde{G}_{k+L^*} \,|\, \tilde{G}_k]e^{-c\gamma} \leq \tilde{G}_k e^{-c\gamma}. \tag{65}$$

Substituting (65) for the right hand side of (58) and using $\tilde{G}_k = G_k = e^{cF_k}$, the result follows. ■

According to Theorem 2, the probability of $F_k^\dagger$ being larger than some arbitrary constant γ decreases exponentially. This result characterizes process behavior outside the convergence region. This is a bound on the worst optimality gap attained during the process starting at a level set $F_k = (1+\rho)M\sigma^2/2\beta$ and ending at or below the starting level set F_k. The exponential bound in (45) is dependent on the coefficient $c = 2\rho M\sigma^2/[(\rho M\sigma^2)^2 + \kappa^2]$ with larger c corresponding to smaller upper bounds. Accordingly, an increase in any of the constants ρ, M, σ^2, or κ, increases the excursion probability bound in (45) because it decreases the coefficient c. This is natural. An increase in the mean squared error σ^2 implies more variation in heuristic rational actions. Likewise, an increment in κ increases the maximum possible suboptimality increase between subsequent steps. Increasing M results in functions that change faster. The constant ρ indicates how far away $F_k^\dagger$ is at the start of the excursion. Notice that since F_k in (45) is also proportional to ρ, this does not necessarily mean that the bound becomes worse for larger ρ. Rather, the bound is more or less invariant with respect to the starting point of the excursion.

Notice that since the proof of Theorem 2 uses the results in Lemma 1 and Theorem 1 suitable variations of Remarks 2 and 3 hold for Theorem 2 as well. In particular, Theorem 2 holds true if the network is not connected [cf. Remark 2]. If activation probabilities are unequal, Theorem 2 follows with a modification to the constant c [cf. Remark 3].

V. Simulations

We further study the behavior of networks composed of heuristic rational agents through numerical simulations. We use the examples in Sections II-A and II-B, namely propagation of opinions in social networks and MRF estimation with a sensor network as case studies. We add a third case study corresponding to cohesive foraging of animal herds or fish schools. In this scenario, a group of animals balance attractive forces to a food source with attractive and repulsive forces between each other.

In all three test cases, we generate network connectivity using a geometric model. We drop a group of N agents on a rectangular field of length L and width W. The coordinates $\mathbf{r}_i$ of user i are chosen inside this rectangle uniformly at random. The neighborhood set of agent i consists of all agents j positioned within a cutoff distance d of $\mathbf{r}_i$, i.e., $n(i) = \{j : \|\mathbf{r}_i - \mathbf{r}_j\| \leq d,\ j \neq i\}$. We further assume agents become active independently of each other and that times between activations of user i are exponentially distributed with parameter $\mu = 1$. This is equivalent to Poisson activations at a rate μ for individual users and at a rate $N\mu$ for the network as a whole.

A. Opinion Propagation

Consider the model of opinion propagation with stubborn agents presented in Section II-A with quadratic individual cost functions. The network consists of $N = 100$ agents forming a connected network as described above. Fig. 2 shows network structure on a 100 unit $\times$ 100 unit two dimensional field. Lines indicate connections between agents. There are two stubborn agents in the set $\mathcal{S} = \{1, 2\}$ marked with dotted squares. Agent 1 located at $\mathbf{r}_1 = (20, 3)$ has set opinion $x_1 = 1$, whereas agent 2 located at $\mathbf{r}_2 = (67, 79)$ has fixed opinion $x_2 = -1$. The remaining agents $i \in V \setminus \mathcal{S}$ are compliant. They start with a random opinion uniformly drawn from $[-1, 1]$. Opinions are updated using the rational action in (8) superimposed with zero mean noise. The noise is chosen as uniformly distributed in $[-\alpha, \alpha]$.

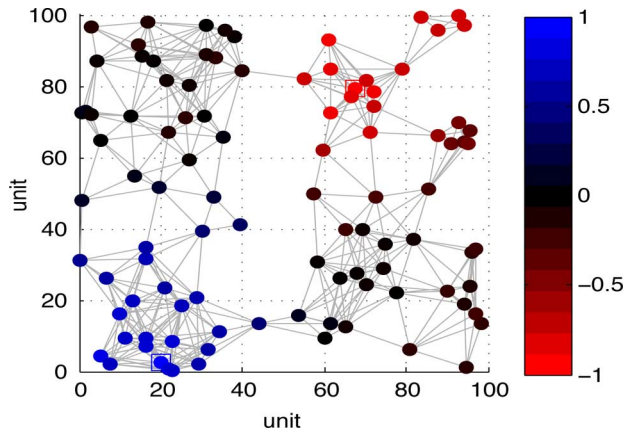

Fig. 2. Example network connectivity for opinion propagation in social networks. A total of $N = 100$ agents are randomly placed on a 100 unit $\times$ 100 unit square. Connections are drawn between agents situated less than 20 units apart. Two stubborn agents with coordinates $\mathbf{r}_1 = (20, 3)$ and $\mathbf{r}_2 = (67, 79)$ have set extreme opinions $x_1(t) = 1$ and $x_2(t) = -1$. The remaining nodes are compliant and attempt to minimize a measure of network discordance. Color encodes opinions at time $t = 50$ for rational behavior (8) superimposed with uniform additive noise in $[-0.1, 0.1]$. Stubborn agents are influential among agents within their close proximity.

The evolution of individual opinions during $t = 50$ time units is presented in Fig. 3 for $\alpha = 0$ (left), $\alpha = 0.1$ (middle), and $\alpha = 0.3$ (right). With the chosen rate of activation $\mu = 1$, an average of 50 activations per agent are observed. Further observe that $\alpha = 0$ corresponds to rational actions. In the plot for $\alpha = 0.1$, it is apparent that the actions of different agents oscillate randomly around their rational actions shown on the plot for $\alpha = 0$. This is not so clear in the plot for $\alpha = 0.3$ where the noise starts to be dominant. To further understand behavior for varying levels of randomness refer to Fig. 4 which illustrates the evolution of the normalized global optimality gap $F_k/p^* := (f(\mathbf{x}(t_k)) - p^*)/p^*$ in (7) for the noise levels $\alpha \in \{0, 0.1, 0.2, 0.3\}$. Consistent with the results of Theorem 1, the optimality gaps attained under various noise levels increase with increasing values of α. The variability around the near optimal value also become larger when the mean square error bound σ^2, which is proportional to α grows. This confirms the dependence of the excursion probability upper bound on the variance σ^2 indicated by Theorem 2.

The emergence of three opinion clusters is clear after around $t = 10$ for the plots corresponding to $\alpha = 0$ and $\alpha = 0.1$ in Fig. 3. Two of these clusters contain opinions between the intervals $0.5 < x_i(50) < 1$ and $-1 < x_i(50) < -0.5$ corresponding to strong support for the opinion of agents 1 and 2, respectively. The third cluster contains opinions between the intervals $-0.25 < x_i(50) < 0.25$ corresponding to weak support for either stubborn agent. A small number of agents stay unclustered. Clusters start loosing meaning in the plot for $\alpha = 0.3$ as

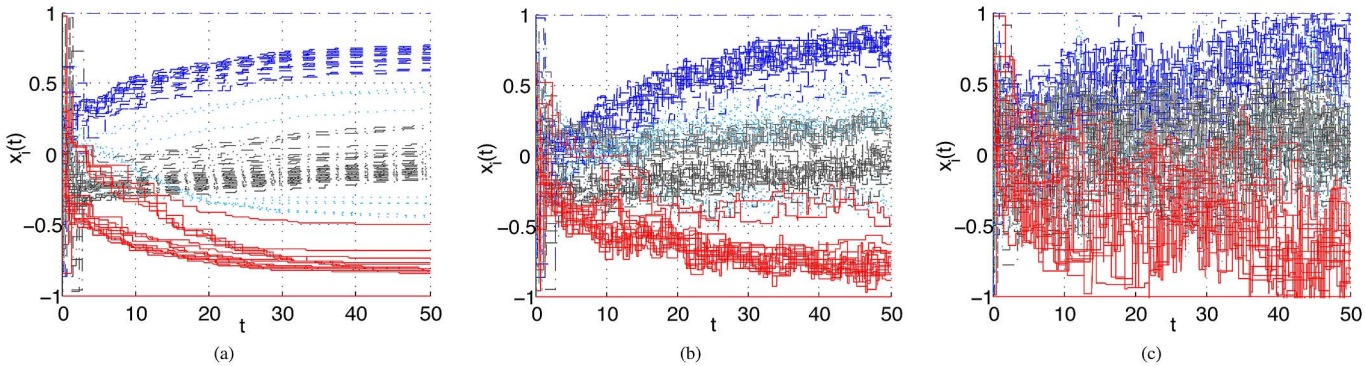

Fig. 3. Agent opinions as a function of time for the network in Fig. 2. Lines represent the path $x_i(t)$ of each agent's opinion up until time $t = 50$. Opinions are updated by rational action (8) superimposed with zero mean noise uniformly distributed in $[-\alpha, \alpha]$ with $\alpha = 0$ (left), $\alpha = 0.1$ (middle), and $\alpha = 0.3$ (right). Three main clusters are apparent when $\alpha = 0$ and $\alpha = 0.1$ with final opinions within the range $0.5 < x_i(50) < 1$, $-1 < x_i(50) < -0.5$ and $-0.25 < x_i(50) < 0.25$, corresponding to strong support for stubborn agents' opinions $x_1(t) = 1$ and $x_2(t) = -1$ and weak support for either stubborn agent.

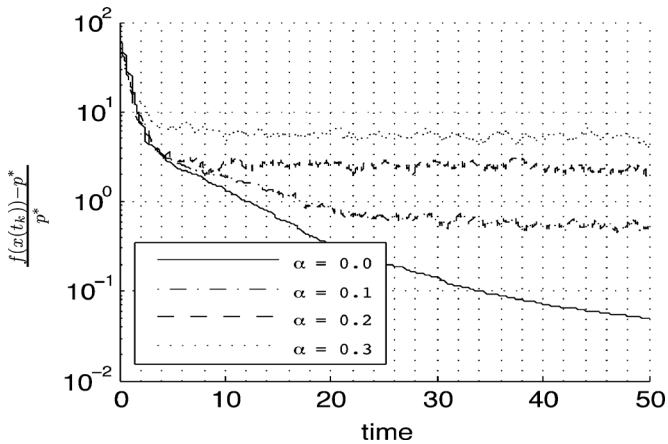

Fig. 4. Normalized global optimality gap for opinion propagation with L_2 cost. Opinions are updated by the rational action (8) superimposed with zero mean noise uniformly distributed in $[-\alpha, \alpha]$ with $\alpha \in \{0, 0.1, 0.2, 0.3\}$. The steady state optimality gap increases with α as predicted by Theorem 1.

the noise in the heuristic rational rule dominates attempts at optimality. Opinions $x_i(50)$ are also color-coded in Fig. 2 for the noise level $\alpha = 0.1$. It is noticeable that agents in the clusters with strong support for either opinion are in close proximity of the corresponding stubborn agent. Strong supporters of agent 1, i.e., those in the cluster $\{i : 0.5 < x_i(50) < 1\}$, are located in the lower-left quadrant. Strong supporters of agent 2, i.e., those in the cluster $\{i : -1 < x_i(50) < -0.5\}$, are located in the upper-right quadrant. Weak supporters of either agent are located in either the upper-left or the lower-right quadrant.

1) Small World Networks: We also consider the same opinion propagation model in a small-world network, which is thought to better capture the structure of a social network [40]. A small-world network is constructed from a geometric network by introducing a rewiring probability p_r and going through a cycle of random edge rewiring. In this cycle, an edge is removed and reconnected to a random node with probability p_r. The addition of these connections reduces the average path length of the network's graph. Three networks resulting from random rewiring of edges of the network in Fig. 2 with probability $p_r = 0.1$ are shown in Fig. 5(a)–(c). Individuals update with respect to rational action (8) superimposed with uniform noise between $[-0.1, 0.1]$. The evolution of individual opinions during $t = 50$ time units for these three networks is shown in Fig. 6(a)–(c), respectively. Different rewiring leads to different steady state opinions. However, in all three cases, the introduction of random links enhances opinion propagation among clusters decreasing the influence of stubborn agents. This is supported by the fact that the total number of strong supporters of either extreme opinion drops from 41 in the geometric network case with $\alpha = 0.1$ (cf. Fig. 3(b)) to 14 in Fig. 6(a), 12 in Fig. 6(b), and 19 in Fig. 6(c). Also note that the average opinions of individuals in the strong support clusters become less extreme.

2) Agents Using L_1 Cost: We consider the same network setup of Fig. 2 with agents using the L_1 cost function in (10) and the heuristic rational rule in (12) instead of the L_2 cost. The cost in (10) is not strongly convex and therefore theorems 1 and 2 do not apply. Fig. 7 illustrates the evolution of the normalized global optimality gap F_k/p^* when agents are rational as per (11) and when they follow the heuristic rational rule in (12). Heuristic rational agents attain a smaller global cost than rational agents. Deterministic coordinate descent fails in functions that are not strongly convex because subsequent updates tend to induce cycles that result in periodic oscillations. The randomness in the heuristic rational rule prevents this from happening and leads to better performance than the one achieved by perfectly rational agents.

B. Temperature Field Estimation

As an example of MRF estimation as presented in Section II-B, consider $N = 270$ temperature sensors placed on a 15 m $\times$ 20 m field. Sensor observations are corrupted by independent additive zero-mean Gaussian noise. The observation energy function for agent i is consequently $U(y_i \,|\, x_i) = (x_i - y_i)^2/2\sigma^2$, with σ^2 denoting noise power. The MRF is also assumed to be Gaussian distributed which results in a quadratic field energy function $U(\mathbf{x}) = \sum_{i,j \in n(i)} u_{ij}(x_i - x_j) = \sum_{i,j \in n(i)} (x_i - x_j)^2/2\lambda$, with λ representing a smoothing coefficient. Smaller λ result in smoother fields. The log-likelihood function for this estimation problem can be obtained by substituting $U(y_i \,|\, x_i) = (x_i - y_i)^2/2\sigma^2$ and $u_{ij}(x_i - x_j) = (x_i - x_j)^2/2\lambda$

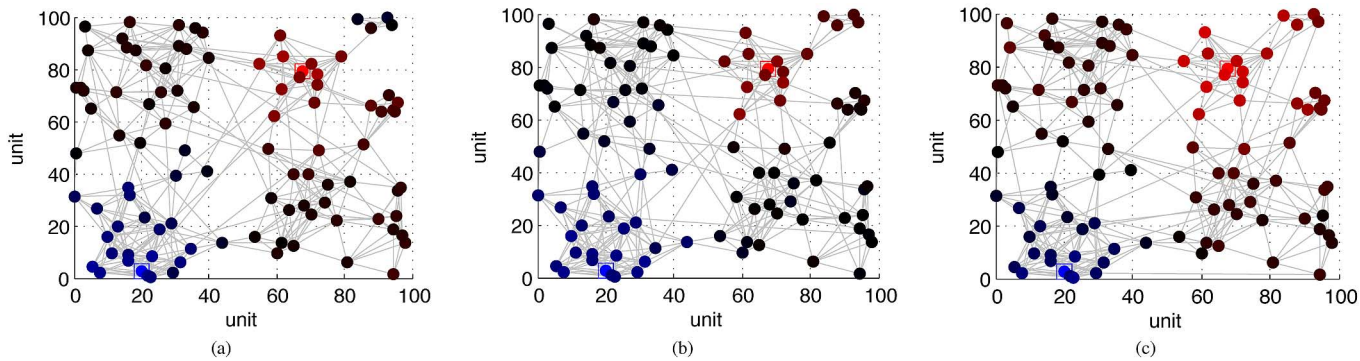

Fig. 5. Opinion propagation in small-world networks. Edges in the network in Fig. 2 are rewired to a randomly selected node with probability $p_r = 0.1$ to simulate observed long-distance connections in social networks. Networks (a)–(c) are different random outcomes of this rewiring procedure. Color encodes opinions at time $t = 50$ as in Fig. 2.

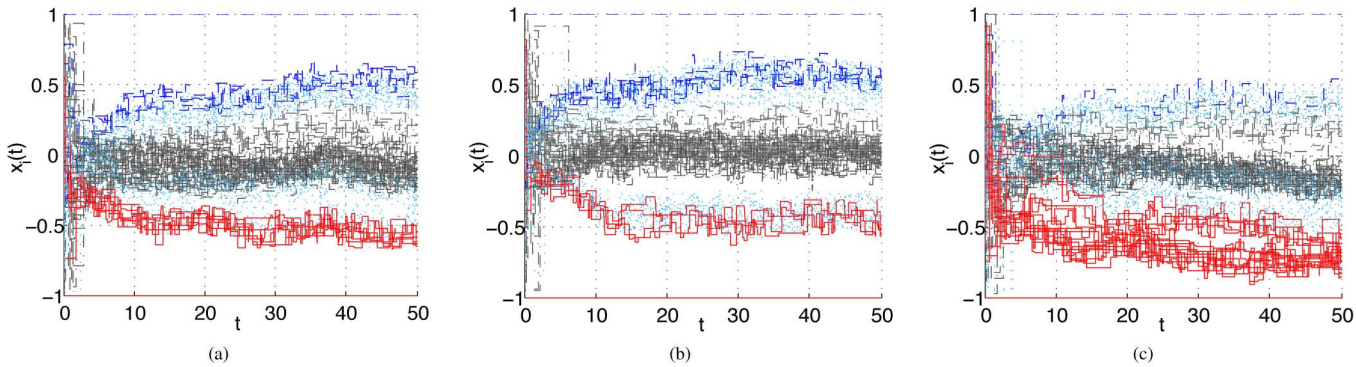

Fig. 6. Agent opinions as a function of time for the networks in Fig. 5(a)–(c). Even though different rewiring leads to different specific behavior, common features appear in small-world networks. The clusters of strong supporters in Fig. 3 become smaller and the opinion of agents in these clusters become milder. The cluster of mild supporters becomes larger.

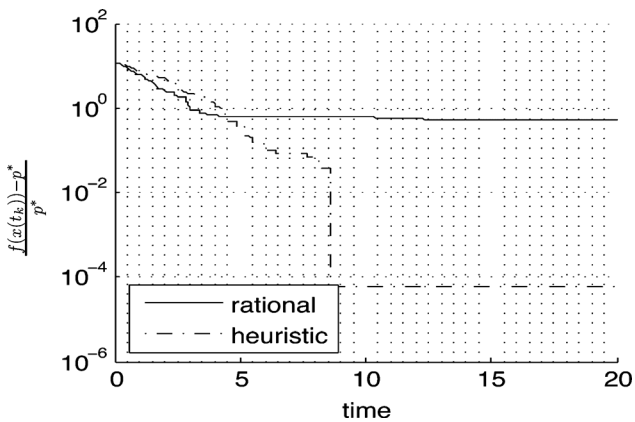

Fig. 7. Normalized global optimality gap for opinion propagation with L_1 cost. Optimality gaps for rational updates as per (11) and heuristic rational updates as per (12) are shown. A smaller optimality gap is achieved when agents follow the heuristic rational rule.

into (14). Equivalently, the aggregate local cost in (15) for agent i is

$$f_i\left(x_i, \mathbf{x}_{n(i)}\right) = \frac{1}{2\sigma^2}(x_i - y_i)^2 + \frac{1}{2\lambda}\sum_{j\in n(i)}(x_i - x_j)^2. \quad (66)$$

The resulting rational action $\tilde{x}_i(t)$ that optimizes (66) for agent i given neighboring states $x_j(t)$ at time t is

$$\tilde{x}_i(t) = \frac{\lambda y_i + \sigma^2 \sum_{j\in n(i)} x_j(t)}{\lambda + N_i\sigma^2}. \quad (67)$$

Equation (67) shows that the optimal estimate of agent i is a weighted average of its own measurement y_i and the current estimates $x_j(t)$ of neighboring sensors. The measurement is weighted by the smoothing coefficient λ and the neighboring observations by the noise variance σ^2. Thus, either smaller λ or larger σ^2 increases the value placed on neighbors' information.

Due to bandwidth limitations sensors quantize their observations before transmission. Thus, the signal received by sensor i from sensor j is not the current state $x_j(t)$ but a quantized version $x_{qj}(t)$ that we can write as $x_{qj}(t) = x_j(t) + q_j(t)$ for some quantization noise $q_j(t)$. The update in (67) is therefore not the one carried out by sensor i. Rather, the estimate of sensor i is updated to

$$x_i(t) = \frac{\lambda y_i + \sigma^2 \sum_{j\in n(i)} x_{qj}(t)}{\lambda + N_i\sigma^2}. \quad (68)$$

Assuming the quantization noise has zero mean, i.e., $\mathbb{E}[q_i(t)] = 0$, it follows that $\mathbb{E}[x_i(t)] = \tilde{x}_i(t)$. We can then think of the update in (68) as a heuristic rational version of the update in (67). We can model communication noise in identical manner.

Numerical simulations for a WSN with sensors performing the update in (68) to estimate a temperature field are shown in Figs. 8–10. The temperature field is generated as a 15 m × 20 m grid with temperature values ranging in [0°F, 255°F]. Two heat sources at locations $\mathbf{h}_1 = (7, 14)$ and $\mathbf{h}_2 = (12, 13)$ set the temperature in these points to 255°F. Temperature drops from these points at a rate of 25°F/m within an area of influence of

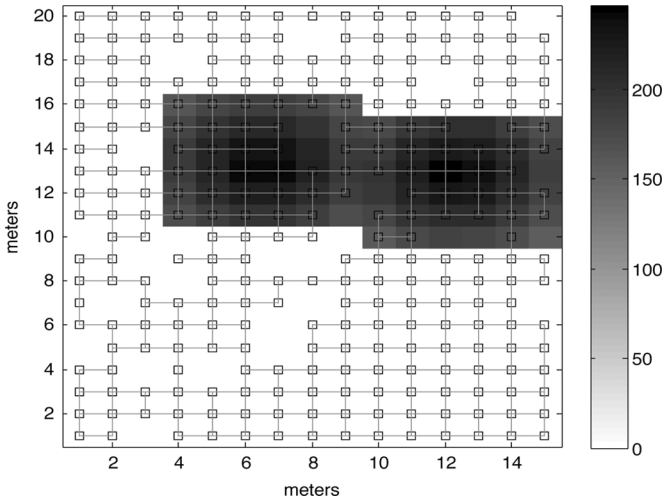

Fig. 8. Estimation of a temperature field with a sensor network. Two heat sources located at (7 m, 14 m) and (12 m, 13 m) set temperatures at their locations to 255° F. Temperature around these sources falls at a rate of 25° F/m within the radius of 3 m from the sources and is set to 0° F at other locations. The resulting temperature field is encoded according to the scale on the right. A sensor network with $N = 270$ is deployed to estimate this field with lines representing network connectivity. Sensors estimate the temperature at their location.

3 m from the sources. The sources do not influence the temperature outside this area which is therefore set to 0° F; see Fig. 8. The $N = 270$ sensors are located at random positions in a 1 m grid with communication between sensors occurring only between sensors located less than 1 m apart; see Fig. 8. Consistent with the discussion leading to (68) the temperature field is modeled as a uniform Gaussian MRF with smoothing parameter $\lambda = 10^3$ and observation noise as Gaussian with variance $\sigma^2 = 0.9 \times 10^3$. The quantization levels for temperature estimates are integers in [0, 255].

Fig. 9(a)–(c) display sensor estimates at times $t = 0$, $t = 2$, and $t = 8$, respectively. At time $t = 0$ estimates are based on local observations only and thereby show significant difference with respect to field values. By time $t = 2$ the average number of updates per sensor is just two but we already observe a significant reduction of noise effects, which becomes more substantial at time $t = 8$. Notice that the outline of the area of influence of the heat sources is better defined in the actual field than in the sensor estimates. This is because of the mismatch between the way in which the field was generated and the Gaussian MRF model presumed for its estimation. Fig. 10(a) and (b) show the evolution of the global log-likelihood function $f(\mathbf{x}) = \sum_i f_i(x_i(t), \mathbf{x}_{n(i)}(t))$ and the local cost function $f_i(x_i(t), \mathbf{x}_{n(i)}(t))$ in (66) for a selected i. The log-likelihoods tend to decrease thereby resulting on more refined estimates.

C. Cohesive Foraging in Animal Networks

As a third example, consider the problem of cohesive foraging where a group of animals search collectively for a food source. Models of swarming behavior in schools of fish, herds of animals, flocks of birds, and colonies of bacteria are explained through distance-dependent attraction and repulsion forces between individuals [10]–[12]. The attraction force between individuals at positions $\mathbf{x}_i$ and $\mathbf{x}_j$ increases with their separating distance $\|\mathbf{x}_i - \mathbf{x}_j\|$. A commonly used form for the potential associated with this force is $A(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j(t)\|^2$ [13]. The repulsion force decreases with growing distance; a frequent model being derived from the potential $R(\mathbf{x}_i, \mathbf{x}_j) = \mathrm{e}^{-\|\mathbf{x}_i - \mathbf{x}_j(t)\|^2}$ [13]. These functions are chosen such that for large distances between individuals attraction forces dominate while repulsion is dominant at short distances. A third force added to swarm foraging models results from an attractive potential of the form $\|\mathbf{x}_i - \mathbf{x}_g\|^2$, pulling individual $\mathbf{x}_i$ towards a goal destination $\mathbf{x}_g$ representing, e.g., the location of a food source [13]. Aggregating these three potentials we define the cost function for agent i as

$$f_i\left(\mathbf{x}_i, \mathbf{x}_{n(i)}\right) = \frac{h}{2}\|\mathbf{x}_i - \mathbf{x}_g\|^2 + \sum_{j \in n(i)} \frac{a}{2}\|\mathbf{x}_i - \mathbf{x}_j(t)\|^2 + \frac{bc}{2}\mathrm{e}^{-\|\mathbf{x}_i - \mathbf{x}_j(t)\|^2/c} \quad (69)$$

where h, a, b, c are positive constants. As in, e.g., [13], [21], we assume the neighborhood structure is fixed and that network agents know the relative position of their neighbors $\mathbf{x}_j$. Observe that the cost in (69) is not convex. Since convexity was required by hypotheses (A1), the results in theorems 1 and 2 are not valid in this example.

For this example, we consider a *local* heuristic rational rule where $\tilde{\mathbf{x}}_i$ is a solution to

$$0 = \nabla_{\mathbf{x}_i} f_i\left(\mathbf{x}_i, \mathbf{x}_{n(i)}\right) = h(\mathbf{x}_i - \mathbf{x}_g) + \sum_{j \in n(i)} (\mathbf{x}_i - \mathbf{x}_j)\left[a + b\mathrm{e}^{-\|\mathbf{x}_i - \mathbf{x}_j\|^2/c}\right] \quad (70)$$

given $\mathbf{x}_{n(i)}$ and $\mathbf{x}_g$. In other words, $\tilde{\mathbf{x}}_i$ is only guaranteed to be a local minimum. Now, we assume that the animal can only reach the position $\tilde{\mathbf{x}}_i$ with some error that has mean zero yielding $\mathbb{E}[\mathbf{x}_i(t)] = \tilde{\mathbf{x}}_i(t)$. This position error could be due to precision error of the animal's visual perception or other external factors such as geographical barriers.

We consider a herd of seven animals starting at random locations inside 1 m × 1 m region. The herd network is completely connected meaning each agent knows the relative position of all other animals in the herd. The environment includes a food source located at point $\mathbf{x}_g = (3, 1)$ that each agent is attracted to. We let the positive constants of (69) be equal for all animals with values $h = 2$, $a = 2$, $b = 5$, $c = 1$. Agents update their position with respect to a uniform distribution in $[\tilde{\mathbf{x}}_i(t) - \alpha, \tilde{\mathbf{x}}_i(t) + \alpha]$ with $\alpha = 0.1$ where $\tilde{\mathbf{x}}_i(t)$ is a solution to (70). Precision error in animals' visual perception is captured by α.

Network structure of seven animals at times $t = 0$, $t = 15$, $t = 30$ and $t = 45$ are shown in Fig. 11(a)–(d). Animals indicated by color coded filled circles start with random locations on a 1 m × 1 m region. Lines indicate connection between animals displaying complete network structure. Food source located at $\mathbf{x}_g = (3, 1)$ is indicated with bold circle. Observe that between times $t = 0$ and $t = 15$, the dominant forces on animal behavior are food source attraction and attraction between individuals as the individuals get closer to food source and each other. At time $t = 15$, the network structure is elongated toward food source indicating a strong attraction by the food source. As the herd gets close to food source by time $t = 30$, inter-individual attraction

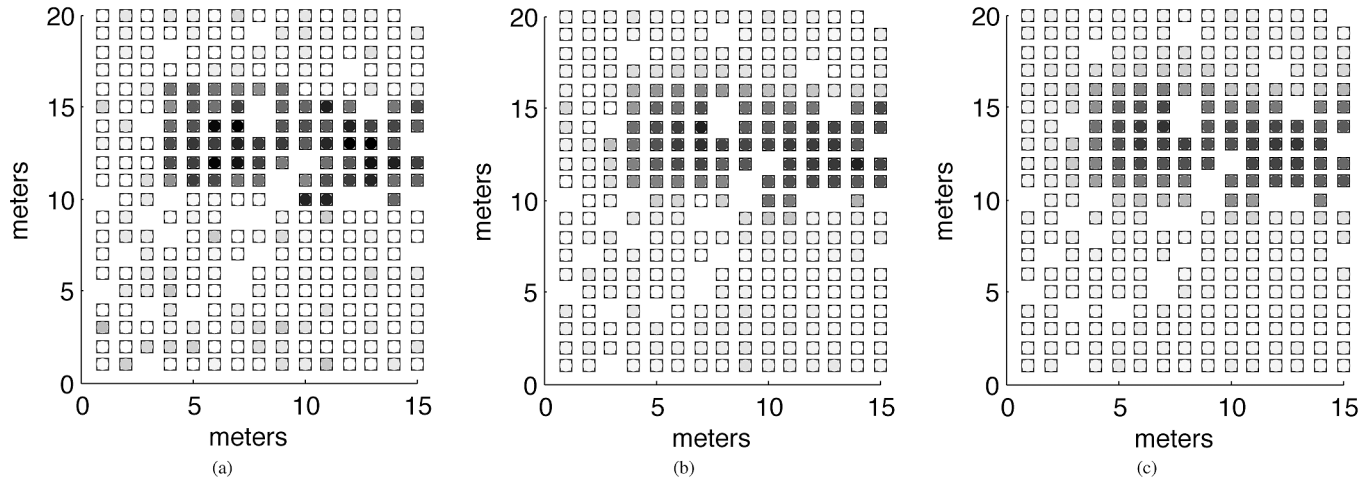

Fig. 9. Sensor temperature field estimates at times $t = 0$ (a), $t = 2$ (b), and $t = 8$ (c). Gaussian MRF has smoothing coefficient $\lambda = 10^3$. Observation noise power is equal to $\sigma^2 = 0.9 \times 10^3$. Sensors update their temperature estimates using the heuristic rational update in (68) with quantization level $\Delta = 1$. Estimates are encoded using the same scale of Fig. 8. Sensor estimates become closer to field values as information from neighboring nodes is incorporated into local estimates.

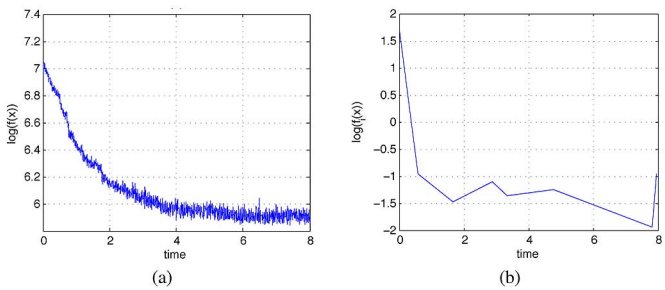

Fig. 10. Global (a) and individual (b) log-likelihoods for temperature field estimation.

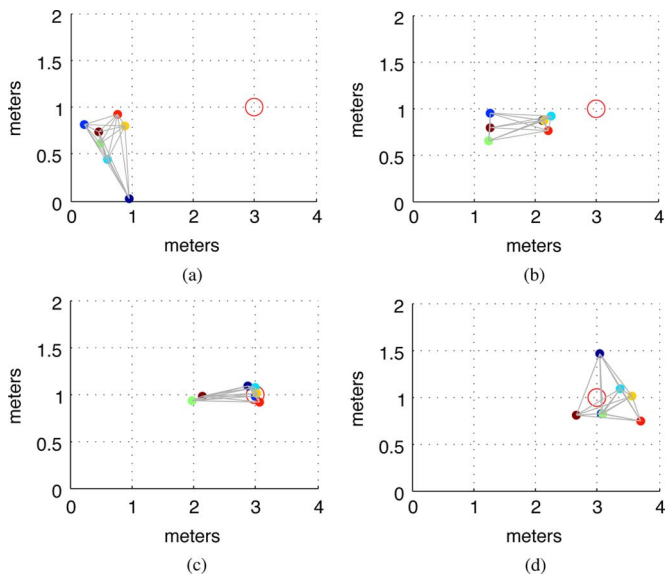

Fig. 11. Snapshots of herd network structure at times $t = 0$, $t = 15$, $t = 30$, and $t = 45$ on a 4 m $\times$ 2 m environment corresponding to (a)–(d), respectively. A herd of $N = 7$ is initially randomly located at a 1 m $\times$ 1 m region. Network is complete and fixed. Lines indicate existence of links between animals. Food is bold circled at point $\mathbf{x}_g = (3, 1)$. Filled circles are color coded to identify each individual animal. Agents move toward food source while trying to even out inter-individual attraction and repulsion forces.

and repulsion become the dominant forces guiding behavior. Observe that at time $t = 45$, herd is gathered near the food source and tries to maintain inter-individual distance that balances attraction and repulsion forces. Noise inhibits the agents from making the right decision at all times which further affects path of other agents, yet we still observe a final outcome that is close to the case when the decision making is noise free. This provides an explanation for observation of harmonic behavior of the whole herd or school even when the individual actions seem imperfect.

VI. Conclusion

We analyzed the convergence and excursion behavior of heuristic rational rules (cf. Definition 1) in distributed network optimization problems with global cost functions that are sums of local nonlinear costs. We have shown that sequences of heuristic rational actions performed individually at random activation times guarantee visits to a near optimality region infinitely often for almost all realizations of the process. This region was completely characterized in terms of cost function properties and the mean squared error of the heuristic rule with respect to the optimal action. We also studied the path of the optimality gap between visits to the near optimality region. We showed that the worst yield achieved during excursions away from optimality are exponentially bounded. An important consequence of this result is that while deviations from optimality are possible, they are rare.

We illustrated heuristic rational optimization using scenarios from social, communication, and biological networks where heuristic rational actions can explain emergent behavior. The social network example is the propagation of opinions from stubborn agents. The communication network example consisted of the estimation of a Markov random field using a wireless sensor network. The biological network example considered cohesive foraging herd behavior.

References

[1] A. Dogandžić and B. Zhang, "Distributed estimation and detection for sensor networks using hidden Markov random field models," *IEEE Trans. Signal Process.*, vol. 54, no. 8, pp. 3200–3215, 2006.

[2] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Trans. Signal Process.*, vol. 57, no. 5, pp. 355–369, 2009.
[3] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks: Quantized data and random link failures," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1383–1400, 2010.
[4] S. Kar and J. M. Moura, "Convergence rate analysis of distributed gossip (linear parameter) estimation: Fundamental limits and tradeoffs," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 4, pp. 674–690, 2011.
[5] S. Stankovic, M. Stankovic, and D. Stipanovic, "Decentralized parameter estimation by consensus based stochastic approximation," in *Proc. 46th IEEE Conf. Decision and Contr. (CDC)*, New Orleans, LA, Dec. 2007, pp. 1535–1540.
[6] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proc. 49th IEEE Conf. Decision Contr.*, Atlanta, GA, Dec. 2010.
[7] J. Chen and A. H. Sayed, Diffusion Adaptation Strategies for Distributed Optimization and Learning Over Networks 2012, Arxiv 1111. 0034v2.
[8] F. S. Cattivelli and A. H. Sayed, "Modeling bird flight formations using diffusion adaptation," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2038–2051, 2011.
[9] V. D. Blondel, J. M. Hendrickx, A. Olshevsky, and J. N. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *Proc. 44th IEEE Conf. Decision Contr.*, Seville, Spain, 2005, pp. 2996–3000.
[10] C. M. Breder, "Equations descriptive of fish schools and other animal aggregations," *Ecology*, vol. 35, no. 3, pp. 361–370, 1954.
[11] D. Grunbaum and A. Okubo, "Modeling social animal aggregations," *Frontiers in Theoret. Biol.*, vol. 100, pp. 296–325, 1994.
[12] K. Warburton and J. Lazarus, "Tendency-distance models of social cohesion in animal groups," *J. Theoret. Biol.*, vol. 150, pp. 473–488, 1991.
[13] V. Gazi and K. M. Passino, "Stability analysis of social foraging swarms," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 34, no. 1, pp. 539–557, 2004.
[14] T. Vicsek, A. Czirok, I. Ben-Jacob, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Phys. Rev. Lett.*, vol. 75, no. 6, pp. 1226–1229, 1995.
[15] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, 2003.
[16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Gossip algorithms: Design, analysis, and applications," in *Proc. IEEE INFOCOM*, Miami, FL, 2005, vol. 3, pp. 1653–1664.
[17] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar, Opinion Fluctuations and Disagreement in Social Networks 2011, arXiv: 1009.2653v3.
[18] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *Proc. Amer. Contr. Conf. (ACC)*, Portland, OR, 2005, pp. 1859–1964.
[19] E. Mossel and G. Schonenebeck, "Reaching consensus on social networks," in *Proc. Innov. Comput. Sci. (ICS)*, Beijing, China, 2010, pp. 214–229.
[20] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, Discrete Opinion Dynamics With Stubborn Agents, 2011 [Online]. Available: http://ssrn.com/abstract=1744113
[21] V. Gazi and K. M. Passino, "Stability analysis of swarms," *IEEE Trans. Autom. Control*, vol. 48, no. 4, pp. 692–697, 2003.
[22] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Trans. Autom. Control*, vol. 54, no. 11, 2009.
[23] A. Kashyap, T. Basar, and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.
[24] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. 31, no. 9, pp. 803–812, 1986.
[25] A. Nedic, A. Ozdaglar, and A. P. Parrilo, "Constrained consensus and optimization in multi-agent networks," *IEEE Trans. Autom. Control*, vol. 55, no. 4, pp. 922–938, 2010.
[26] M. Rabbat, R. Nowak, and J. Bucklew, "Generalized consensus computation in networked systems with erasure links," in *Proc. IEEE 6th Workshop on the Signal Process. Adv. Wireless Commun. (SPAWC)*, New York, NY, Jun. 2005, pp. 1088–1092.
[27] S. S. Ram, A. Nedic, and V. V. Veeravalli, "Distributed stochastic subgradient projection algorithms for convex optimization," *J. Optimiz. Theory Appl.*, vol. 147, no. 3, pp. 516–545, 2010.
[28] S. Z. Li, *Markov Random Field Modeling in Image Analysis*. Tokyo, Japan: Springer-Verlag, 2001.
[29] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 6, pp. 721–741, 1984.
[30] Y. Kushner, *Stochastic Approximation and Recursive Algorithms and Application*, 2nd ed. New York: Springer-Verlag, 2003.
[31] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *J. Optimiz. Theory Appl.*, vol. 109, no. 3, pp. 475–494, 2001.
[32] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.
[33] N. Zadeh, "A note on the cyclic coordinate ascent method," *Manage. Sci.*, vol. 16, pp. 642–644, 1970.
[34] L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear Gauss Seidel method under convex constraints," *Operat. Res. Lett.*, vol. 26, pp. 127–136, 2000.
[35] Y. Nesterov, Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems, 2010, CORE Discussion Papers.
[36] E. Yildiz, A. Ozdaglar, D. Acemoglu, and A. Scaglione, "The voter model with stubborn agents extended abstract," in *Proc. 48th Ann. Allerton Conf.*, Allerton, IL, 2010, pp. 1179–1181.
[37] S. Boyd and Vanderberghe, *Convex Programming*. New York: Wiley, 2004.
[38] R. Durrett, *Probability: Theory and Examples*, ser. Cambridge Series in Statistical and Probabilistic Mathematics, 3rd ed. Cambridge, U.K.: Duxbury Press, 2004.
[39] S. M. Ross, *Stochastic Processes*, 2nd ed. New York: Wiley, 1996.
[40] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, 1998.

Ceyhun Eksin (S'12) received the B.Sc. degree in control engineering from Istanbul Technical University, in 2005. He received the M.Sc. degree in industrial engineering from Bogazici University, Istanbul, in 2008. In this period, he also spent one semester at the Technische Universiteit Eindhoven as Erasmus Exchange student.

He joined the University of Pennsylvania, Philadelphia, in 2008 as a Ph.D. student. His research interests are in the areas of signal processing, distributed optimization, and social networks. Since 2011, he has been focusing on distributed optimization in social, biological, and technological networks.

Alejandro Ribeiro (M'07) received the B.Sc. degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, in 1998. From 2003 to 2008, he was with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, where he received the M.Sc. and Ph.D. degrees in electrical engineering.

In 2008, he joined the Department of Electrical and Systems Engineering, University of Pennsylvania (Penn), Philadelphia, an Assistant Professor. From 1998 to 2003, he was a Member of the Technical Staff at Bellsouth Montevideo. His research interests lie in the areas of communication, signal processing, and networking. His current research focuses on network and wireless communication theory.

Dr. Ribeiro received the 2012 S. Reid Warren, Jr. Award presented by Penn's undergraduate student body for outstanding teaching and the NSF CAREER award in 2010. He is also a Fulbright scholar and the recipient of student paper awards at ICASSP 2005 and ICASSP 2006.