



Opinion Dynamics with Limited Information

Dimitris Fotakis¹ · Vardis Kandiros²  · Vasilis Kontonis³ · Stratis Skoulakis⁴

Received: 31 December 2020 / Accepted: 9 June 2023 / Published online: 4 September 2023
© The Author(s) 2023

Abstract

We study opinion formation games based on the famous model proposed by Friedkin and Johsen (FJ model). In today’s huge social networks the assumption that in each round agents update their opinions by taking into account the opinions of *all* their friends is unrealistic. So, we are interested in the convergence properties of simple and natural variants of the FJ model that use limited information exchange in each round and converge to the same stable point. As in the FJ model, we assume that each agent i has an intrinsic opinion $s_i \in [0, 1]$ and maintains an expressed opinion $x_i(t) \in [0, 1]$ in each round t . To model limited information exchange, we consider an opinion formation process where each agent i meets with one random friend j at each round t and learns only her current opinion $x_j(t)$. The amount of influence j imposes on i is reflected by the probability p_{ij} with which i meets j . Then, agent i suffers a disagreement cost that is a convex combination of $(x_i(t) - s_i)^2$ and $(x_i(t) - x_j(t))^2$. An important class of dynamics in this setting are *no regret* dynamics, i.e. dynamics that ensure vanishing regret against the experienced disagreement cost to the agents. We show an exponential gap between the convergence rate of no regret dynamics and of more general dynamics that do not ensure no regret. We prove that no regret dynamics require roughly $\Omega(1/\varepsilon)$ rounds to be within distance ε from the stable point of the FJ model. On the other hand, we provide an opinion update rule that does not ensure no

✉ Vardis Kandiros
kandiros@mit.edu

Dimitris Fotakis
fotakis@cs.ntua.gr

Vasilis Kontonis
kontonis@wisc.edu

Stratis Skoulakis
efstratios@sutd.edu.sg

¹ National Technical University of Athens, Athens, Greece

² Massachusetts Institute of Technology, Cambridge, MA, USA

³ University of Wisconsin-Madison, Madison, WI, USA

⁴ Singapore University of Technology and Design, Singapore, Singapore

regret and converges to x^* in $\tilde{O}(\log^2(1/\varepsilon))$ rounds. Finally, in our variant of the FJ model, we show that the agents can adopt a simple opinion update rule that ensures no regret to the experienced disagreement cost and results in an opinion vector that converges to the stable point x^* of the FJ model within distance ε in $\text{poly}(1/\varepsilon)$ rounds. In view of our lower bound for no regret dynamics this rate of convergence is close to best possible.

Keywords Opinion dynamics · Opinion formation games · No-regret dynamics · Game theory

1 Introduction

The study of *Opinion Formation* has a long history (see e.g. [29]). Opinion Formation is a *dynamic process* in the sense that socially connected people (e.g. family, friends, colleagues) exchange information and this leads to changes in their expressed opinions over time. Today, the advent of the internet and social media makes the study of opinion formation in large social networks even more important; realistic models of how people form their opinions by interacting with each other are of great practical interest for prediction, advertisement etc. In an attempt to formalize the process of opinion formation, several models have been proposed over the years (see e.g., [13, 14, 18, 28]). The common assumption underlying all these models, which dates back to DeGroot [13], is that opinions evolve through a form of repeated averaging of information collected from the agents' social neighborhoods.

Our work builds on the model proposed by Friedkin and Johnsen [18]. The FJ model is a variation on the DeGroot model capturing the fact that consensus on the opinions is rarely reached. According to FJ model each person i has a public opinion $x_i \in [0, 1]$ and an internal opinion $s_i \in [0, 1]$, which is private and invariant over time. There also exists a weighted graph $G(V, E)$ representing a social network where V stands for the persons ($|V| = n$) and E their social relations. Initially, all nodes start with their internal opinion and at each round t , update their public opinion $x_i(t)$ to a weighted average of the public opinions of their neighbors and their internal opinion,

$$x_i(t) = \frac{\sum_{j \in N_i} w_{ij} x_j(t-1) + w_{ii} s_i}{\sum_{j \in N_i} w_{ij} + w_{ii}}, \quad (1)$$

where $N_i = \{j \in V : (i, j) \in E\}$ is the set of i 's neighbors, the weight w_{ij} associated with the edge $(i, j) \in E$ measures the extent of the influence that j poses on i and the weight $w_{ii} > 0$ quantifies how susceptible i is in adopting opinions that differ from her internal opinion s_i .

The FJ model is one of most influential models for opinion formation. It has a very simple update rule, making it plausible for modeling natural behavior and its basic assumptions are aligned with empirical findings on the way opinions are formed [1, 32]. At the same time, it admits a unique stable point $x^* \in [0, 1]^n$ to which it converges with a *linear* rate [23]. The FJ model has also been studied under a game theoretic viewpoint. Bindel et al. considered its update rule as the minimizer

of a quadratic disagreement cost function and based on it they defined the following opinion formation game [5]. Each node i is a selfish agent whose strategy is the public opinion x_i that she expresses incurring her a disagreement cost

$$C_i(x_i, x_{-i}) = \sum_{j \in N_i} w_{ij}(x_i - x_j)^2 + w_{ii}(x_i - s_i)^2 \quad (2)$$

Note that the FJ model is the *simultaneous best response dynamics* and its stable point x^* is the unique Nash equilibrium of the above game. In [5] they quantified its inefficiency with respect to the total disagreement cost. They proved that the *Price of Anarchy* (PoA) is $9/8$ in case G is undirected and $w_{ij} = w_{ji}$. They also provided PoA bounds in the case of unweighted Eulerian directed graphs. We remark that in [5] an alternative framework for studying the way opinions evolve was introduced. The opinion formation process can be described as the *dynamics* of an opinion formation game. This framework is much more comprehensive since different aspects of the opinion formation process can be easily captured by defining suitable games. Subsequent works [3, 4, 15] considered variants of the above game and studied the convergence properties of the *best response dynamics*.

1.1 Motivation and Our Setting

Many recent works study the Nash equilibrium x^* of the opinion formation game defined in [5] under various perspectives. In [10] they extended the bounds for PoA in more general classes of directed graphs, while many recently introduced influence maximization problems [2, 24, 33], which are defined with respect to x^* . The reason for this scientific interest is evident: the equilibrium x^* is considered as an appropriate way to model the final opinions formed in a social network, since the *well established* FJ model converges to it.

Our work is motivated by the fact that there are notable cases in which the FJ model is not an appropriate model for the dynamic of the opinions, due to the large amount of information exchange that it implies. More precisely, at each round its update rule (1) requires that every agent learns all the opinions of her social neighbors. In today's large social networks where users usually have several hundreds of friends it is highly unlikely that, each day, they learn the opinions of all their social neighbors. In such environments it is far more reasonable to assume that individuals randomly meet a small subset of their acquaintances and these are the only opinions that they learn. Such information exchange constraints render the FJ model unsuitable for modeling the opinion formation process in such large networks and therefore, it is not clear whether x^* captures the limiting behavior of the opinions. In this work we ask:

Question 1 Is the equilibrium x^* an adequate way to model the final formed opinions in large social networks? Namely, are there simple variants of the FJ model that require limited information exchange and converge fast to x^* ? Can they be justified as natural behavior for selfish agents under a game-theoretic solution concept?

To address these questions, one could define precise dynamical processes whose update rules require limited information exchange between the agents and study their

convergence properties. Instead of doing so, we describe the opinion formation process in such large networks as *dynamics* of a suitable opinion formation game that captures these information exchange constraints. This way we can precisely define which *dynamics* are *natural* and, more importantly, to study general classes of *dynamics* (e.g. no regret dynamics) without explicitly defining their update rule. The opinion formation game that we consider is a variant of the game in [5] based on interpreting the weight w_{ij} as a measure of how frequently i meets j .

Definition 1 For a given opinion vector $x \in [0, 1]^n$, the disagreement cost of agent i is the random variable $C_i(x_i, x_{-i})$ defined as follows:

- Agent i meets one of her neighbors j with probability $p_{ij} = w_{ij} / \sum_{j \in N_i} w_{ij}$.
- Agent i suffers cost $C_i(x_i, x_{-i}) = (1 - a_i)(x_i - x_j)^2 + a_i(x_i - s_j)^2$, where

$$\alpha_i = \frac{w_{ii}}{\sum_{j \in N_i} w_{ij} + w_{ii}}.$$

Note that the expected disagreement cost of each agent in the above game is the same as the disagreement cost in [5] (scaled by $\sum_{j \in N_i} w_{ij} + w_{ii}$). Moreover its Nash equilibrium, with respect to the expected disagreement cost, is x^* . This game provides us with a general template of all the *dynamics* examined in this paper. At round t , each agent i selects an opinion $x_i(t)$ and suffers a disagreement cost based on the opinion of the neighbor that she randomly met. At the end of round t , she is informed only about the opinion and the index of this neighbor and may use this information to update her opinion in the next round. Obviously different update rules lead to different *dynamics*, however all of these respect the information exchange constraints: at every round each agent learns the opinion of *just one* of her neighbors. Question 1 now takes the following more concrete form.

Question 2 Can the agents update their opinions according to the limited information that they receive such that the produced opinion vector $x(t)$ converges to the equilibrium x^* ? How is the convergence rate affected by the limited information exchange? Are there dynamics that ensure that the cost that the agents experience is minimal?

In what follows, we are mostly concerned about the dependence of the rate of convergence on the distance ε from the equilibrium x^* . Thus, we shall suppress the dependence on other parameters such as the size of the graph, n . We remark that the dependence of our dynamics on these constants is in fact rather good (see Sect. 2), and we do this only for clarity of exposition.

Definition 2 (Informal) We say that a dynamics converges slowly resp. fast to the equilibrium x^* if it requires $\text{poly}(1/\varepsilon)$ resp. $\text{poly}(\log(1/\varepsilon))$ rounds to be within (expected) error ε of x^* .

1.2 Contribution

The major contribution of the paper is proving an exponential separation on the convergence rate of *no regret dynamics* and the convergence rate of more general dynamics produced by update rules that do not ensure no regret.

No regret dynamics are produced by update rules that ensure no regret to any agent that adopts them. Namely, the total disagreement cost of an agent that follows such a rule is close to the total disagreement cost that she would experience by selecting the best fixed opinion in hindsight. The latter must hold regardless of the way the other agents update their opinions and of the neighbors that the agent gets to meet. This powerful property renders no regret dynamics *natural dynamics* for describing the behavior of agents [8, 16, 31, 37]. We prove that if all the agents adopt an update rule that ensures no regret, then there exists an instance of the game such that the produced opinion vector $x(t)$ requires roughly $\Omega(1/\varepsilon)$ rounds to be ε -close to x^* . No regret comes at the price of slow convergence because it provides robust guarantees. Agents who adopt no regret update rules suffer minimal total disagreement cost even if the other agents play irrationally or adversarially. In order to provide such strong guarantees, no regret rules must only depend on the opinions that the agent observes and not take into account the weights w_{ij} of the outgoing edges (see Sect. 5). We call the update rules with the latter property, *graph oblivious*. In Sect. 5 we use a novel information theoretic argument to prove the aforementioned lower bound for this more general class.

In Sect. 6, we present a simple update rule whose resulting dynamics converges fast, i.e. the opinion vector $x(t)$ is ε -close to x^* in $O(\log^2(1/\varepsilon))$ rounds. The reason that the previous lower bound doesn't apply is that this rule does not ensure *no regret* to the agents that adopt it. In fact there is a very simple example with two agents, in which the first follows the rule while the second selects her opinions adversarially, where the first agent experiences regret (see Example 1 in Sect. 6).

We introduce an intuitive *no regret* update rule and we show that if all agents adopt it, the resulting opinion vector $x(t)$ converges to x^* . Our rule is a *Follow the Leader algorithm*, meaning that at round t , each agent updates her opinion to the minimizer of total disagreement cost that she experienced until round $t - 1$. It also has a very simple form: it is roughly the time average of the opinions that the agent observes. In Sect. 3, we bound its convergence rate and show that in order to achieve ε distance from x^* , $\text{poly}(1/\varepsilon)$ rounds are sufficient. In view of our lower bound this rate is close to best possible. In Sect. 4, we prove its *no regret* property. This can be derived by the more general results in [25]. However, we give a short and simple proof that may be of interest.

In conclusion, our results reveal that the equilibrium x^* is a robust choice for modeling the limiting behavior of the opinions of agents since, even in our limited information setting, there exist simple and natural dynamics that converge to it. The convergence rate crucially depends on whether the agents act selfishly, i.e. they are only concerned about their individual disagreement cost. We present an update rule that selfish agents can adopt (no regret update rule) and show that the resulting opinion vector converges to x^* but with a slow rate, while, for non selfish agents, the update rule in Sect. 6 leads to a *dynamics* with fast convergence rate.

1.3 Related Work

There exists a large amount of literature concerning the FJ model. Many recent works [3, 4, 12, 15] bound the inefficiency of equilibrium in variants of opinion formation game defined in [5]. In [23] they bound the convergence time of the FJ model in special graph topologies. In [3], a variant of the opinion formation game, in which social relations depend on the expressed opinions, is studied. They prove that the discretized version of the above game admits a potential function and thus best-response converges to the Nash equilibrium. Convergence results in other discretized variants of the FJ model can be found in [17, 40]. In [19] they provide convergence results for limited information variants of the Heglesmann-Krause model [28] and the FJ model. Although the considered limited information variant of the FJ model is very similar to ours, their convergence results are much weaker, since they concern the expected value of the opinion vector.

Other works that relate to ours concern the convergence properties of dynamics based on no regret learning algorithms. In [20, 21, 36, 37] it is proved that in a finite n -person game, if each agent updates her mixed strategy according to a no regret algorithm, the resulting *time-averaged* strategy vector converges to Coarse Correlated Equilibrium. The convergence properties of no regret dynamics for games with infinite strategy spaces were considered in [16]. They proved that for a large class of games with concave utility functions (socially concave games), the time-averaged strategy vector converges to Pure Nash Equilibrium (PNE). More recent work investigates a stronger notion of convergence of no regret dynamics. In [11] they show that, in n -person finite generic games that admit unique Nash equilibrium, the strategy vector converges *locally* and fast to it. They also provide conditions for *global* convergence. Our results fit in this line of research since we show that for a game with *infinite* strategy space, the strategy vector (and not the time-averaged) converges to the Nash equilibrium x^* .

No regret dynamics in limited information settings have recently received substantial attention from the scientific community since they provide realistic models for the practical applications of game theory. Perfect payoff information is rare in practice; agents act based on random or noisy past-payoff observations. Kleinberg et al. in [30] treated load-balancing in distributed systems as a repeated game and analyzed the convergence properties of no regret learning algorithms under the *full information assumption* that each agent learns the load of every machine. In a subsequent work [31], the same authors consider the same problem in a *limited information setting* (“bulletin board model”), in which each agent learns the load of just the machine that served him. Most relevant to ours are the works [6, 11, 27, 35], where they examine the convergence properties of no regret learning algorithms when the agents observe their payoffs with some additive zero-mean random noise. In our limited information setting the agents experience random disagreement cost with expected value equal to the actual cost. The main difference is that our *noise* is not additive but due to a sampling process.

2 Our Results and Techniques

We have adopted the convention of using \ln to denote the natural logarithm. We will also use the notation \log freely without specifying a base when inside the big- O notation or when we have a constant C that is arbitrary. As previously mentioned, an instance of the game in [5] is also an instance of the game of Definition 1. Following the notation introduced earlier we have that $p_{ij} = w_{ij} / \sum_{j \in N_i} w_{ij}$ if $j \in N_i$ and 0 otherwise. Moreover, $\alpha_i = w_{ii} / (\sum_{j \in N_i} w_{ij} + w_{ii}) > 0$ since $w_{ii} > 0$ by the definition of the game in [5]. If an agent i does not have outgoing edges ($N_i = \emptyset$) then $p_{ij} = 0$ for all j . Therefore $\sum_{j=1}^n p_{ij} = 0$, $\alpha_i = 1$ if $N_i = \emptyset$ and $\sum_{j=1}^n p_{ij} = 1$, $\alpha_i \in (0, 1)$ otherwise. For simplicity we adopt the following notation for an instance of the game of Definition 1.

Definition 3 We denote an instance of the opinion formation game of Definition 1 as $I = (P, s, \alpha)$, where P is a $n \times n$ matrix with non-negative elements p_{ij} , with $p_{ii} = 0$ and $\sum_{j=1}^n p_{ij}$ is either 0 or 1, $s \in [0, 1]^n$ is the internal opinion vector, $\alpha \in (0, 1]^n$ the self confidence vector.

An instance $I = (P, s, \alpha)$ is also an instance of the FJ model, since by the update rule (1) $x_i(t) = (1 - \alpha_i) \sum_{j \in N_i} p_{ij} x_j(t - 1) + \alpha_i s_i$. It also defines the opinion vector $x^* \in [0, 1]^n$ which is the stable point of the FJ model and the Nash equilibrium of the game in [5].

Definition 4 For a given instance $I = (P, s, \alpha)$ the equilibrium $x^* \in [0, 1]^n$ is the unique solution of the following linear system, for every $i \in V$:

$$x_i^* = (1 - \alpha_i) \sum_{j \in N_i} p_{ij} x_j^* + \alpha_i s_i.$$

The fact that the above linear system always admits a solution follows by matrix norm properties. Throughout the paper we study *dynamics* of the game of Definition 1. We denote as W_i^t the neighbor that agent i met at round t , which is a random variable whose probability distribution is determined by the instance $I = (P, s, \alpha)$ of the game, $\mathbf{P}[W_i^t = j] = p_{ij}$. Another parameter of an instance I that we often use is $\rho = \min_{i \in V} \alpha_i$.

In Sect. 3, we examine the convergence properties of the opinion vector $x(t)$ when all agents update their opinions according to the *Follow the Leader* principle. Since each agent i must select $x_i(t)$, before knowing which of her neighbors she will meet and what opinion her neighbor will express, this update rule says “*play the best according to what you have observed*”. For a given instance (P, s, a) of the game the Follow the Leader dynamics $x(t)$ is defined in Dynamics 1 and Theorem 1 shows its convergence rate to x^* .

Theorem 1 *Let $I = (P, s, \alpha)$ be an instance of the opinion formation game of Definition 1 with equilibrium $x^* \in [0, 1]^n$. The opinion vector $x(t) \in [0, 1]^n$ produced by update rule (3) after t rounds satisfies*

$$\mathbf{E} [\|x(t) - x^*\|_\infty] \leq C \sqrt{\log n} \frac{(\ln t)^{3/2}}{t^{\min(1/2, \rho)}},$$

Algorithm 1 Follow the Leader dynamics

- 1: Initially $x_i(0) = s_i$ for all agents i .
- 2: At round $t \geq 0$ each agent i :
 - 3: Meets neighbor with index W_i^t , $\mathbf{P}[W_i^t = j] = p_{ij}$.
 - 4: Suffers cost $(1 - \alpha_i)(x_i(t) - x_{W_i^t}(t))^2 + \alpha_i(x_i(t) - s_i)^2$ and learns the opinion $x_{W_i^t}(t)$.
 - 5: Updates her opinion

$$x_i(t+1) = \operatorname{argmin}_{x \in [0,1]} \sum_{\tau=0}^t (1 - \alpha_i)(x - x_{W_i^\tau}(\tau))^2 + \alpha_i(x - s_i)^2 \quad (3)$$

where $\rho = \min_{i \in V} \alpha_i$ and C is a universal constant and $t \geq 6$.

In Sect. 4 we argue that, apart from its simplicity, update rule (3) ensures no regret to any agent that adopts it and therefore the FTL dynamics can be considered as *natural dynamics* for selfish agents. Since each agent i selfishly wants to minimize the disagreement cost that she experiences, it is natural to assume that she selects $x_i(t)$ according to a *no regret algorithm* for the *online convex optimization problem* where the adversary chooses a function $f_t(x) = (1 - \alpha_i)(x - b_t)^2 + \alpha_i(x - s_i)^2$ at each round t . In Theorem 2 we prove that *Follow the Leader* is a no regret algorithm for the above OCO problem. We remark that this does not hold, if the adversary can pick functions from a different class (see e.g. chapter 5 in [26]).

Theorem 2 Consider the function $f : [0, 1]^2 \mapsto [0, 1]$ with

$$f(x, b) = (1 - \alpha)(x - b)^2 + \alpha(x - s)^2,$$

for some constants $s, \alpha \in [0, 1]$. Let $(b_t)_{t=0}^\infty$ be an arbitrary sequence with $b_t \in [0, 1]$. If

$$x_t = \operatorname{argmin}_{x \in [0,1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$$

then for all t ,

$$\sum_{\tau=0}^t f(x_\tau, b_\tau) \leq \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau) + O(\log t).$$

On the positive side, the FTL dynamics converges to x^* and its update rule is simple and ensures no regret to the agents. On the negative side, its convergence rate is outperformed by the rate of FJ model. For a fixed instance $I = (P, s, \alpha)$, the FTL dynamics converges with rate $\tilde{O}(1/t^{\min(\rho, 1/2)})$ while FJ model converges with rate $O(e^{-\rho t})$ [23].

Question 3 Can the agents adopt other no regret update rules such that the resulting dynamics converges fast to x^* ?

The answer is no. In Sect. 5, we prove that fast convergence cannot be established for any *no regret dynamics*. The reason that FTL dynamics converges slowly is that rule (3) only depends on the opinions of the neighbors that agent i meets, α_i , and s_i . This is also true for any update rule that ensures no regret to the agents (see Sect. 5). As already mentioned, we call this larger class of update rules *graph oblivious*, and we prove that fast convergence cannot be established for *graph oblivious dynamics*.

Definition 5 (graph oblivious update rule) A graph oblivious update rule A is a sequence of functions $(A_t)_{t=0}^\infty$ where $A_t : [0, 1]^{t+2} \mapsto [0, 1]$.

Definition 6 (graph oblivious dynamics) Let A be a graph oblivious update rule. For a given instance $I = (P, s, \alpha)$ the rule A produces a graph oblivious dynamics $x_A(t)$ defined as follows:

- Initially each agent i selects her opinion $x_i^A(0) = A_0(s_i, \alpha_i)$
- At round $t \geq 1$, each agent i selects her opinion

$$x_i^A(t) = A_t \left(x_{W_i^0}(0), \dots, x_{W_i^{t-1}}(t-1), s_i, \alpha_i \right),$$

where W_i^t is the neighbor that i meets at round t .

Theorem 3 states that for any graph oblivious dynamics there exists an instance $I = (P, s, \alpha)$, where roughly $\Omega(1/\varepsilon)$ rounds are required to achieve convergence within error ε .

Theorem 3 *Let A be a graph oblivious update rule, which all agents use to update their opinions. For any $c > 0$ there exists an instance $I = (P, s, \alpha)$ such that*

$$\mathbf{E} \left[\|x_A(t) - x^*\|_\infty \right] = \Omega(1/t^{1+c}),$$

where $x_A(t)$ denotes the opinion vector produced by A for the instance $I = (P, s, \alpha)$.

To prove Theorem 3, we show that graph oblivious rules whose dynamics converge fast imply the existence of estimators for Bernoulli distributions with “small” sample complexity. The key part of the proof lies in Lemma 6, in which it is proven that such estimators cannot exist. We also briefly discuss two well-known sample complexity lower bounds from the statistics literature and explain why they do not work in our case.

In Sect. 6, we present a simple update rule that achieves error rate $e^{-\tilde{O}(\sqrt{t})}$. This update rule is a function of the opinions and the indices of the neighbors that i met, s_i, α_i and the i -th row of the matrix P . Obviously this rule is not *graph oblivious*, due to its dependency on the i -th row and the indices, and thus does not ensure no regret to an agent that adopts it (see Example 1 in Sect. 6). However it reveals that slow convergence is not a generic property of the limited information dynamics, but comes with the assumption that agents act selfishly.

3 Convergence Rate of FTL Dynamics

In this section we prove Theorem 1 which bounds the convergence time of FTL dynamics to the unique equilibrium point x^* . Notice that for an instance $I = (P, s, \alpha)$, the opinion vector $x(t) \in [0, 1]^n$ of the FTL dynamics (see Dynamics 1) can be written equivalently as follows:

- Initially all agents adopt their internal opinion, $x_i(0) = s_i$.
- At round $t \geq 1$, each agent i updates her opinion

$$x_i(t) = (1 - \alpha_i) \sum_{\tau=0}^{t-1} \frac{x_{W_i^\tau}(\tau)}{t} + \alpha_i s_i,$$

where W_i^τ is the neighbor that i met at round t .

Since the opinion vector $x(t)$ is a random vector, the convergence metric used in Theorem 1 is $\mathbf{E} [\|x(t) - x^*\|_\infty]$ where the expectation is taken over the random meeting of the agents. At first we present a high level idea of the proof. We remind that the unique equilibrium $x^* \in [0, 1]^n$ of the instance $I = (P, s, \alpha)$ satisfies the following equations for each agent $i \in V$,

$$x_i^* = (1 - \alpha_i) \sum_{j \in N_i} p_{ij} x_j^* + \alpha_i s_i$$

Since our metric is $\mathbf{E} [\|x(t) - x^*\|_\infty]$, we can use the above equations to bound $|x_i(t) - x_i^*|$.

$$\begin{aligned} |x_i(t) - x_i^*| &= (1 - \alpha_i) \left| \frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}(\tau)}{t} - \sum_{j \in N_i} p_{ij} x_j^* \right| \\ &= (1 - \alpha_i) \left| \sum_{j \in N_i} \frac{\sum_{\tau=0}^{t-1} \mathbf{1}[W_i^\tau = j] x_j(\tau)}{t} - \sum_{j \in N_i} p_{ij} x_j^* \right| \\ &\leq (1 - \alpha_i) \sum_{j \in N_i} \left| \frac{\sum_{\tau=0}^{t-1} \mathbf{1}[W_i^\tau = j] x_j(\tau)}{t} - p_{ij} x_j^* \right| \end{aligned}$$

Now assume that

$$\left| \frac{\sum_{\tau=0}^{t-1} \mathbf{1}[W_i^\tau = j]}{t} - p_{ij} \right| = 0$$

for all $t \geq 1$, then with simple algebraic manipulations one can prove that $\|x(t) - x^*\|_\infty \leq e(t)$ where $e(t)$ satisfies the recursive equation

$$e(t) = (1 - \rho) \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t},$$

where $\rho = \min a_i$. It follows that $\|x(t) - x^*\|_\infty \leq 1/t^\rho$ meaning that $x(t)$ converges to x^* . Obviously the latter assumption does not hold, however since W_i^τ are independent random variables with $\mathbf{P}[W_i^\tau = j] = p_{ij}$, the quantity

$$\left| \frac{\sum_{\tau=0}^{t-1} \mathbf{1}[W_i^\tau = j]}{t} - p_{ij} \right|$$

tends to 0 with probability 1. In Lemma 1 we use this fact to obtain a similar recursive equation for $e(t)$ and then in Lemma 2 we upper bound its solution.

Lemma 1 *Let $e(t)$ be the solution of the recursion*

$$e(t) = \delta(t) + (1 - \rho) \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t},$$

where $e(0) = \|x(0) - x^*\|_\infty$, $\delta(t) = \sqrt{\ln(\pi^2 n t^2 / 6p) / t}$ and $\rho = \min_{i \in V} \alpha_i$. Then,

$$\mathbf{P}[\text{for all } t \geq 1, \|x(t) - x^*\|_\infty \leq e(t)] \geq 1 - p$$

Proof At first we prove that with probability at least $1 - p$, for all $t \geq 1$ and all agents i :

$$\left| \frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}^*}{t} - \sum_{j \in N_i} p_{ij} x_j^* \right| \leq \sqrt{\frac{\ln(\pi^2 n t^2 / (6p))}{t}} := \delta(t). \tag{4}$$

Since W_i^τ are independent random variables with $\mathbf{P}[W_i^\tau = j] = p_{ij}$ and $\mathbf{E}[x_{W_i^\tau}^*] = \sum_{j \in N_i} p_{ij} x_j^*$. By the Hoeffding’s inequality we get

$$\mathbf{P} \left[\left| \frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}^*}{t} - \sum_{j \in N_i} p_{ij} x_j^* \right| > \delta(t) \right] < 6p / (\pi^2 n t^2).$$

To bound the probability of error for all rounds $t \geq 1$ and all agents i , we apply the union bound

$$\sum_{t=1}^\infty \mathbf{P} \left[\max_i \left| \frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}^*}{t} - \sum_{j \in N_i} p_{ij} x_j^* \right| > \delta(t) \right] \leq \sum_{t=1}^\infty \frac{6}{\pi^2} \frac{1}{t^2} \sum_{i=1}^n \frac{p}{n} = p$$

As a result with probability at least $1 - p$ we have that inequality (4) holds for all $t \geq 1$ and all agents i . We now prove our claim by induction. Let $\|x(\tau) - x^*\|_\infty \leq e(\tau)$ for all $\tau \leq t - 1$. Then

$$\begin{aligned}
 x_i(t) &= (1 - \alpha_i) \frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}(\tau)}{t} + \alpha_i s_i \\
 &\leq (1 - \alpha_i) \frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}^* + \sum_{\tau=0}^{t-1} e(\tau)}{t} + \alpha_i s_i \\
 &\leq (1 - \alpha_i) \left(\frac{\sum_{\tau=0}^{t-1} x_{W_i^\tau}^*}{t} + \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \right) + \alpha_i s_i \tag{5}
 \end{aligned}$$

$$\begin{aligned}
 &\leq (1 - \alpha_i) \left(\sum_{j \in N_i} p_{ij} x_j^* + \delta(t) + \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \right) + \alpha_i s_i \\
 &\leq x_i^* + \delta(t) + (1 - \rho) \left(\frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \right) \tag{6}
 \end{aligned}$$

We get (5) from the induction step and (6) from inequality (4). Similarly, we can prove that

$$x_i(t) \geq x_i^* - \delta(t) - (1 - \rho) \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t}.$$

As a result $\|x(t) - x^*\|_\infty \leq e(t)$ and the induction is complete. Therefore, we have that with probability at least $1 - p$, $\|x(t) - x^*\|_\infty \leq e(t)$ for all $t \geq 1$. \square

Now that we have obtained the recursive equation for the error, we can solve it using straightforward computation. The idea is to express the term $e(t + 1)$ in terms of the previous term $e(t)$ and to apply this expression repeatedly to obtain a formula for $e(t)$. The main technical difficulty is upper bounding the sums that arise during this computation. This is done in the following lemma.

Lemma 2 *Let $e(t)$ be a function satisfying the recursion*

$$e(t) = \delta(t) + (1 - \rho) \sum_{\tau=0}^{t-1} e(\tau)/t \text{ and } e(0) = \|x(0) - x^*\|_\infty,$$

where $\delta(t) = \sqrt{\ln(Dt^{5/2})}/t$, $\delta(0) = 0$, and $D > e^{5/2}$ is a positive constant. Then

$$e(t) \leq 2\sqrt{5} \frac{(\ln(Dt))^{3/2}}{t^{\min(\rho, 1/2)}}.$$

Therefore, for all $t \geq 6$:

$$e(t) \leq 2\sqrt{5(\ln D)^3} \frac{(\ln(t))^{3/2}}{t^{\min(\rho, 1/2)}}.$$

Proof Observe that for all $t \geq 0$ the function $e(t)$ satisfies the following recursive relation

$$e(t + 1) = e(t) \left(1 - \frac{\rho}{t + 1} \right) + \delta(t + 1) - \delta(t) + \frac{\delta(t)}{t + 1}. \tag{7}$$

For $t = 0$ we have that

$$e(1) = (1 - \rho)e(0) + \delta(1) = (1 - \rho)e(0) + \sqrt{\ln D}. \tag{8}$$

Observe that for $D > e^{2.5}$, $\delta(t)$ is decreasing for all $t \geq 1$. Therefore,

$$\delta(t + 1) - \delta(t) + \frac{\delta(t)}{t + 1} \leq \frac{\delta(t)}{t + 1}.$$

Also, note that

$$\frac{\delta(t)}{t + 1} = \frac{\sqrt{\ln(Dt^{2.5})}}{t^{1/2}(t + 1)} \leq \frac{\sqrt{5 \ln(D(t + 1))}}{(t + 1)^{3/2}}$$

where in the last inequality we used the fact that $(t + 1)/t \leq 2$ for all t and $\ln t \leq \ln(t + 1)$. Thus, from equations (7) and (8) we get that for all $t \geq 0$

$$e(t + 1) \leq e(t) \left(1 - \frac{\rho}{t + 1} \right) + \frac{\sqrt{5 \ln(D(t + 1))}}{(t + 1)^{3/2}}. \tag{9}$$

Now that we have expressed $e(t + 1)$ in terms of $e(t)$, we can apply this expression to obtain a formula for $e(t)$. We denote by H_t the t -th partial sum of the harmonic series. To simplify notation, we define

$$g(t) = \frac{\sqrt{5 \ln(Dt)}}{t^{3/2}}.$$

In the following, we will make heavy use of the following elementary inequality, which holds for all $p > 0$:

$$(1 - p) \left(1 - \frac{p}{2} \right) \dots \left(1 - \frac{p}{t} \right) \leq e^{-pH_t} \leq \frac{1}{t^p} \tag{10}$$

By “unrolling” the recurrence of Eq. 9 we obtain:

$$\begin{aligned} e(t) &\leq \left(1 - \frac{\rho}{t} \right) e(t - 1) + g(t) \\ &\leq \left(1 - \frac{\rho}{t} \right) \left(1 - \frac{\rho}{t - 1} \right) e(t - 2) + \left(1 - \frac{\rho}{t} \right) g(t - 1) + g(t) \end{aligned}$$

$$= \left(1 - \frac{\rho}{t}\right) \cdots (1 - \rho)e(0) + \sum_{\tau=1}^t g(\tau) \prod_{i=\tau+1}^t \left(1 - \frac{\rho}{i}\right).$$

Next, by using (10) we obtain that

$$e(t) \leq \frac{e(0)}{t^\rho} + \sum_{\tau=1}^t g(\tau)e^{-\rho \sum_{i=\tau+1}^t \frac{1}{i}} = \frac{e(0)}{t^\rho} + \sum_{\tau=1}^t g(\tau)e^{-\rho(H_t - H_\tau)}.$$

We now use the following well known upper and lower bounds for H_t , which holds for all t and can be found in page 2 of [22]

$$\frac{1}{2(n+1)} < H_n - \ln n - \gamma < \frac{1}{2n} \tag{11}$$

for all n , where γ is the Euler number. This immediately gives for $n = t$

$$\gamma + \ln t \leq H_t$$

Also, by (11) for $n = t + 1$ we have

$$H_{t+1} = H_t + \frac{1}{t+1} \leq \ln(t+1) + \gamma + \frac{1}{2(t+1)}$$

which implies that

$$H_t \leq \ln(t+1) + \gamma + \frac{1}{2(t+1)} - \frac{1}{t+1} < \ln(t+1) + \gamma$$

Putting everything together, we have obtained the following for all t

$$\gamma + \ln t \leq H_t \leq \gamma + \ln(t+1),$$

This implies that for all t and $\tau \leq t$, we have

$$e^{\rho H_\tau} \leq e^{\gamma \rho} (\tau + 1)^\rho, \quad e^{-\rho H_t} \leq \frac{e^{-\gamma \rho}}{t^\rho}$$

Thus, we obtain:

$$e(t) \leq \frac{e(0)}{t^\rho} + e^{-\rho H_t} \sum_{\tau=1}^t g(\tau)e^{\rho H_\tau} \leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} \sum_{\tau=1}^t (\tau + 1)^\rho \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2}}$$

Now observe that

$$(\tau + 1)^\rho = \frac{(\tau + 1)^\rho}{\tau^\rho} \tau^\rho \leq 2^\rho \tau^\rho \leq 2\tau^\rho$$

Putting all these together, we obtain

$$e(t) \leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} \sum_{\tau=1}^t (\tau + 1)^\rho \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2}} \leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} \sum_{\tau=1}^t \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2-\rho}}.$$

Now the remaining task is to bound the sum on the right hand side. A standard way of bounding a sum of decreasing terms is with the corresponding Riemann integral. Indeed, we observe that

$$\sum_{\tau=1}^t \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2-\rho}} \leq \int_{\tau=1}^t \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2-\rho}} d\tau \tag{12}$$

since, $\tau \mapsto \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2-\rho}}$ is a decreasing function of τ for all $\rho \in [0, 1]$. To see that, notice that the derivative of this function is

$$\frac{\frac{1}{2\tau\sqrt{\ln(D\tau)}}\tau^{3/2-\rho} - (\frac{3}{2} - \rho)\tau^{1/2-\rho}\sqrt{\ln(D\tau)}}{\tau^{3-2\rho}} < \frac{\tau^{1/2-\rho}}{2\sqrt{\ln(D\tau)}} - \frac{\tau^{1/2-\rho}}{2}\sqrt{\ln(D\tau)}}{\tau^{3-2\rho}}$$

where the inequality holds since $\rho < 1$. Now, for any $\tau \geq 1$ we have that

$$\ln(D\tau) = \ln D + \ln \tau \geq \ln D > \frac{5}{2} > 1$$

which implies that

$$\frac{\tau^{1/2-\rho}}{2\sqrt{\ln(D\tau)}} - \frac{\tau^{1/2-\rho}}{2}\sqrt{\ln(D\tau)} = \frac{\tau^{1/2-\rho}}{2} \left(\frac{1}{\sqrt{\ln(D\tau)}} - \sqrt{\ln(D\tau)} \right) < 0,$$

meaning that the function is indeed decreasing for all $\tau \geq 1$. To bound the integral in (12), we have to distinguish cases for ρ . Intuitively, if ρ is small, then the fraction decays faster than $1/t$, which translates to the overall integral being polylogarithmic. If ρ is large, then a polynomial term with a small exponent might arise in the calculation.

– If $\rho \leq 1/2$ then

$$\int_{\tau=1}^t \tau^\rho \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2}} d\tau \leq \sqrt{\ln(Dt)} \int_{\tau=1}^t \frac{1}{\tau} d\tau = \sqrt{\ln(Dt)} \ln t \leq (\ln(Dt))^{3/2},$$

since $\ln(Dt) > \ln t$ for all $t \geq 1$. Hence

$$\begin{aligned} e(t) &\leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} \sum_{\tau=1}^t \frac{\sqrt{\ln \tau}}{\tau^{3/2-\rho}} \\ &\leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} (\ln(Dt))^{3/2} \leq 2\sqrt{5} \frac{(\ln(Dt))^{3/2}}{t^\rho}, \end{aligned} \tag{13}$$

where we used the fact that $e(0) \leq 1$ and $\sqrt{5}(\ln(Dt))^{3/2} \geq \sqrt{5}(\ln(D))^{3/2} > 1$ for all $t \geq 1$.

– If $\rho > 1/2$ then

$$\begin{aligned} \int_{\tau=1}^t \tau^\rho \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2}} d\tau &= \int_{\tau=1}^t \tau^{\rho-1/2} \frac{\sqrt{\ln(D\tau)}}{\tau} d\tau \\ &= \frac{2}{3} \int_{\tau=1}^t \tau^{\rho-1/2} ((\ln(D\tau))^{3/2})' d\tau \\ &= \frac{2}{3} t^{\rho-1/2} (\ln(Dt))^{3/2} - (\rho - 1/2) \frac{2}{3} \int_{\tau=1}^t \tau^{\rho-3/2} (\ln(D\tau))^{3/2} d\tau \\ &\leq \frac{2}{3} t^{\rho-1/2} (\ln(Dt))^{3/2}. \end{aligned}$$

Hence

$$\begin{aligned} e(t) &\leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} \sum_{\tau=1}^t \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2-\rho}} \\ &\leq \frac{e(0)}{t^\rho} + \frac{\sqrt{5}}{t^\rho} \frac{2}{3} t^{\rho-1/2} (\ln(Dt))^{3/2} \leq \frac{4\sqrt{5}}{3} \frac{(\ln(Dt))^{3/2}}{t^{1/2}}. \end{aligned} \tag{14}$$

For the last inequality, we used the fact that $\ln D > 0$ to conclude that

$$e(0) \leq 1 < \frac{2\sqrt{5}}{3} \leq \frac{2\sqrt{5}}{3} (\ln(Dt))^{3/2}$$

for all $t \geq 1$. Combining inequalities 13 and 14 yields that for all $t \geq 1$

$$e(t) \leq 2\sqrt{5} \frac{(\ln(Dt))^{3/2}}{t^{\min(\rho, 1/2)}},$$

which proves the first claim of the lemma. We would like the following inequality to be satisfied:

$$\ln(Dt) \leq \ln D \ln t,$$

which is equivalent to

$$\ln t \geq \frac{\ln D}{\ln D - 1}$$

If $\ln D > 2.5$, the right hand side is at most $1 + 2/3$. Numerically, we observe that for $t \geq 6$, $\ln t \geq 1 + 2/3$. Thus, the second inequality of the lemma follows for $t \geq 6$. □

An interesting consequence of Lemma 2 is that the rate of convergence is never better than $1/\sqrt{t}$ regardless of the value of ρ . In Sect. 5 we provide evidence that no reasonable protocol can achieve a better convergence rate.

We are now ready to prove Theorem 1.

Theorem 1 *Let $I = (P, s, \alpha)$ be an instance of the opinion formation game of Definition 1 with equilibrium $x^* \in [0, 1]^n$. The opinion vector $x(t) \in [0, 1]^n$ produced by update rule (3) after t rounds satisfies*

$$\mathbf{E} [\|x(t) - x^*\|_\infty] \leq C\sqrt{\log n} \frac{(\ln t)^{3/2}}{t^{\min(1/2, \rho)}},$$

where $\rho = \min_{i \in V} a_i$ and C is a universal constant and $t \geq 6$.

Proof By Lemma 1 we have that for all $t \geq 1$ and $p \in [0, 1]$,

$$\mathbf{P} [\|x(t) - x^*\|_\infty \leq e_p(t)] \geq 1 - p$$

where $e_p(t)$ is the solution of the recursion

$$e_p(t) = \delta(t) + (1 - \rho) \frac{\sum_{\tau=0}^{t-1} e_p(\tau)}{t}$$

with $\delta(t) = \sqrt{\frac{\ln(\pi^2 nt^2/(6p))}{t}}$. Setting $p = \frac{1}{12\sqrt{t}}$ we have that

$$\mathbf{P} [\|x(t) - x^*\|_\infty \leq e(t)] \geq 1 - \frac{1}{12\sqrt{t}}$$

where $e(t)$ is the solution of the recursion

$$e(t) = \delta(t) + (1 - \rho) \frac{\sum_{\tau=0}^{t-1} e_p(\tau)}{t}$$

with $\delta(t) = \sqrt{\frac{\ln(2\pi^2 nt^{2.5})}{t}}$. Since $2\pi^2 \geq e^{2.5}$, Lemma 2 applies and

$$e(t) \leq C\sqrt{\log n} \frac{\log t^{3/2}}{t^{\min(\rho, 1/2)}},$$

for some universal constant C and for all $t \geq 6$. Finally,

$$\begin{aligned} \mathbf{E} [\|x(t) - x^*\|_\infty] &\leq \frac{1}{12\sqrt{t}} + \left(1 - \frac{1}{12\sqrt{t}}\right) C\sqrt{\log n} \frac{(\log t)^{3/2}}{t^{\min(\rho, 1/2)}} \\ &\leq \left(C + \frac{1}{12}\right) \sqrt{\log n} \frac{(\log t)^{3/2}}{t^{\min(\rho, 1/2)}} \end{aligned}$$

□

Hence, FTL dynamics converges to the same equilibrium point as the original FJ-model, albeit slower. In the next section we provide justification about why this strategy is a natural one for players to adopt, given that they operate in an adversarial environment.

4 Follow the Leader Ensures No Regret

In this section we provide rigorous definitions of *no regret* algorithms and explain why update rule (3) ensures no regret to any agent that repeatedly plays the game of Definition 1. Based on the cost that the agents experience, we consider an appropriate *Online Convex Optimization* problem. This problem can be viewed as a “game” played between an adversary and a player. At round $t \geq 0$,

1. the player selects a value $x_t \in [0, 1]$.
2. the adversary observes the x_t and selects a $b_t \in [0, 1]$
3. the player receives cost $f(x_t, b_t) = (1 - \alpha)(x_t - b_t)^2 + \alpha(x_t - s)^2$.

where s, α are constants in $[0, 1]$. The goal of the player is to pick x_t based on the history (b_0, \dots, b_{t-1}) in a way that minimizes her total cost. Generally, different OCO problems can be defined by a set of functions \mathcal{F} that the adversary chooses from and a feasibility set \mathcal{K} from which the player picks her value (see [26] for an introduction to the OCO framework). In our case the feasibility set is $\mathcal{K} = [0, 1]$ and the set of functions is

$$\mathcal{F}_{s,\alpha} = \{x \mapsto (1 - \alpha)(x - b)^2 + \alpha(x - s)^2 : b \in [0, 1]\}.$$

As a result, each selection of the constants s, α leads to a different OCO problem.

Definition 7 An algorithm A for the OCO problem with $\mathcal{F}_{s,\alpha}$ and $\mathcal{K} = [0, 1]$ is a sequence of functions $(A_t)_{t=0}^\infty$ where $A_t : [0, 1]^t \mapsto [0, 1]$.

Definition 8 An algorithm A is no regret for the OCO problem with $\mathcal{F}_{s,\alpha}$ and $\mathcal{K} = [0, 1]$ if and only if for all sequences $(b_t)_{t=0}^\infty$ that the adversary may choose, if $x_t = A_t(b_0, \dots, b_{t-1})$ then for all t ,

$$\sum_{\tau=0}^t f(x_\tau, b_\tau) \leq \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau) + o(t).$$

Informally speaking, if the player selects the value x_t according to a *no regret algorithm* then she does not regret not playing any fixed value no matter what the choices of the adversary are. Theorem 2 states that *Follow the Leader* i.e.

$$x_t = \operatorname{argmin}_{x \in [0,1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$$

is a no regret algorithm for all the OCO problems with $\mathcal{F}_{s,\alpha}$.

Returning to the dynamics of the game in Definition 1, it is reasonable to assume that each agent i selects $x_i(t)$ according to no regret algorithm A_i for the OCO problem with $\mathcal{F}_{s_i, \alpha_i}$, since by Definition 8,

$$\frac{1}{t} \sum_{\tau=0}^t f_i(x_i(\tau), x_{W_i^\tau}(\tau)) \leq \frac{1}{t} \min_{x \in [0,1]} \sum_{\tau=0}^t f_i(x, x_{W_i^\tau}(\tau)) + \frac{o(t)}{t}$$

The latter means that the time averaged total disagreement cost that she suffers is close to the time averaged cost by expressing the best fixed opinion and this holds regardless of the opinions of the neighbors that i meets. Meaning that even if the other agents selected their opinions maliciously, her total experienced cost would still be in a sense minimal. Under this perspective update rule (3) is a rational choice for selfish agents and as a result FTL dynamics is a natural limited information variant of the FJ model. We would like to prove the following.

Theorem 2 Consider the function $f : [0, 1]^2 \mapsto [0, 1]$ with

$$f(x, b) = (1 - \alpha)(x - b)^2 + \alpha(x - s)^2,$$

for some constants $s, \alpha \in [0, 1]$. Let $(b_t)_{t=0}^\infty$ be an arbitrary sequence with $b_t \in [0, 1]$. If

$$x_t = \operatorname{argmin}_{x \in [0,1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$$

then for all t ,

$$\sum_{\tau=0}^t f(x_\tau, b_\tau) \leq \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau) + O(\log t).$$

We now present the key steps for proving Theorem 2. We first prove that a similar strategy that also takes into account the value b_t admits no regret (Lemma 3). Obviously, knowing the value b_t before selecting x_t is in direct contrast with the OCO framework, however proving the no regret property for this algorithm easily extends to establishing the no regret property of *Follow the Leader*. Theorem 2 follows by direct application of Lemma 4.

Lemma 3 Let $(b_t)_{t=0}^\infty$ be an arbitrary sequence with $b_t \in [0, 1]$. Let $y_t = \operatorname{argmin}_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau)$ then for all t ,

$$\sum_{\tau=0}^t f(y_\tau, b_\tau) \leq \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau).$$

Proof By definition of y_t , $\sum_{\tau=0}^t f(y_t, b_\tau) = \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau)$, so

$$\begin{aligned} \sum_{\tau=0}^t f(y_\tau, b_\tau) - \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau) &= \sum_{\tau=0}^t f(y_\tau, b_\tau) - \sum_{\tau=0}^t f(y_t, b_\tau) \\ &= \sum_{\tau=0}^{t-1} f(y_\tau, b_\tau) - \sum_{\tau=0}^{t-1} f(y_t, b_\tau) \leq \sum_{\tau=0}^{t-1} f(y_\tau, b_\tau) - \sum_{\tau=0}^{t-1} f(y_{t-1}, b_\tau) \end{aligned}$$

The last inequality follows by the fact that $y_{t-1} = \operatorname{argmin}_{x \in [0,1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$. Inductively, we prove that $\sum_{\tau=0}^t f(y_\tau, b_\tau) \leq \min_{x \in [0,1]} \sum_{\tau=0}^t f(x, b_\tau)$. □

Now we can understand why *Follow the Leader* admits no regret. Since the cost incurred by the sequence y_t is at most that of the best fixed value, we can compare the cost incurred by x_t with that of y_t . Since the functions in $\mathcal{F}_{s,\alpha}$ are quadratic, the extra term $f(x, b_t)$ that y_t takes into account doesn't change dramatically the minimum of the total sum. Namely, x_t, y_t are relatively close. Hence, the costs incurred by the two sequences are not very different.

Lemma 4 For all $t \geq 0$, $f(x_t, b_t) \leq f(y_t, b_t) + 2\frac{1-\alpha}{t+1} + \frac{(1-\alpha)^2}{(t+1)^2}$.

Proof We first prove that the two sequences are close. Namely, for all t ,

$$|x_t - y_t| \leq \frac{1 - \alpha}{t + 1}. \tag{15}$$

By definition $x_t = \alpha s + (1 - \alpha) \frac{\sum_{\tau=0}^{t-1} b_\tau}{t}$ and $y_t = \alpha s + (1 - \alpha) \frac{\sum_{\tau=0}^t b_\tau}{t+1}$.

$$\begin{aligned} |x_t - y_t| &= (1 - \alpha) \left| \frac{\sum_{\tau=0}^{t-1} b_\tau}{t} - \frac{\sum_{\tau=0}^t b_\tau}{t + 1} \right| \\ &= (1 - \alpha) \left| \frac{\sum_{\tau=0}^{t-1} b_\tau - t b_t}{t(t + 1)} \right| \\ &\leq \frac{1 - \alpha}{t + 1} \end{aligned}$$

The last inequality follows from the fact that $b_\tau \in [0, 1]$. We now use inequality (15) to bound the difference $f(x_t, b_t) - f(y_t, b_t)$. Since f is a quadratic function, the bound follows easily from calculations.

$$\begin{aligned} f(x_t, b_t) &= \alpha(x_t - s)^2 + (1 - \alpha)(x_t - b_t)^2 \\ &\leq \alpha(y_t - s)^2 + 2\alpha |y_t - s| |x_t - y_t| + \alpha |x_t - y_t|^2 \\ &\quad + (1 - \alpha)(y_t - b_t)^2 + 2(1 - \alpha) |y_t - b_t| |x_t - y_t| + (1 - \alpha) |x_t - y_t|^2 \end{aligned}$$

$$\begin{aligned} &\leq f(y_t, b_t) + 2|x_t - y_t| + |y_t - x_t|^2 \\ &\leq f(y_t, b_t) + 2\frac{1 - \alpha}{t + 1} + \frac{(1 - \alpha)^2}{(t + 1)^2} \end{aligned}$$

□

We are now ready to prove that FTL-dynamics has the no regret property.

Theorem 2 Consider the function $f : [0, 1]^2 \mapsto [0, 1]$ with

$$f(x, b) = (1 - \alpha)(x - b)^2 + \alpha(x - s)^2,$$

for some constants $s, \alpha \in [0, 1]$. Let $(b_t)_{t=0}^\infty$ be an arbitrary sequence with $b_t \in [0, 1]$. If

$$x_t = \operatorname{argmin}_{x \in [0, 1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$$

then for all t ,

$$\sum_{\tau=0}^t f(x_\tau, b_\tau) \leq \min_{x \in [0, 1]} \sum_{\tau=0}^t f(x, b_\tau) + O(\log t).$$

Proof Theorem 2 easily follows by Lemma 3

$$\begin{aligned} \sum_{\tau=0}^t f(x_\tau, b_\tau) &\leq \sum_{\tau=0}^t f(y_\tau, b_\tau) + \sum_{\tau=0}^t 2\frac{1 - \alpha}{\tau + 1} + \sum_{\tau=0}^t \frac{(1 - \alpha)^2}{(\tau + 1)^2} \\ &\leq \min_{x \in [0, 1]} \sum_{\tau=0}^t f(x, b_\tau) + 2(1 - \alpha)(\ln t + 1) + (1 - \alpha)\frac{\pi^2}{6} \\ &\leq \min_{x \in [0, 1]} \sum_{\tau=0}^t f(x, b_\tau) + O(\log t) \end{aligned}$$

□

In the next section, we are going to prove that FTL dynamics is the fastest possible no regret protocol to solve the problem of opinion formation.

5 Lower Bound for Graph Oblivious Dynamics

In this section we prove that any no regret dynamics cannot converge much faster than FTL dynamics (Dynamics 1).

Definition 9 (no regret dynamics) Consider a collection of no regret algorithms such that for each $(s, \alpha) \in [0, 1]^2$ a no regret algorithm $A_{s,\alpha}$ ¹ for the OCO problem with $\mathcal{F}_{s,\alpha}$ and $\mathcal{K} = [0, 1]$, is selected. For a given instance $I = (P, s, \alpha)$ this selection produces the no regret dynamics $x(t)$ defined as follows:

- Initially each agent i selects her opinion $x_i(0) = A_0^{s_i, \alpha_i}(s_i, \alpha_i)$
- At round $t \geq 1$, each agent i selects her opinion

$$x_i(t) = A_t^{s_i, \alpha_i} \left(x_{W_i^0}(0), \dots, x_{W_i^{t-1}}(t-1), s_i, \alpha_i \right)$$

where W_i^t is the neighbor that i meets at round t .

Such a selection of no regret algorithms can be encoded as a graph oblivious update rule. Specifically, the function $A_t : \{0, 1\}^{t+2} \mapsto [0, 1]$ is defined as $A_t(b_0, \dots, b_{t-1}, s, \alpha) = A_{s,\alpha}^t(b_0, \dots, b_{t-1})$. Theorem 3 applies and establishes the existence of an instance $I = (P, s, \alpha)$ such that the produced $x(t)$ converges at best slowly to x^* .

The rest of the section is dedicated to prove Theorem 3. In Lemma 5 we show that any graph oblivious update rule A can be used as an estimator of the parameter $p \in [0, 1]$ of a Bernoulli random variable. Since we prove Theorem 3 using a reduction to an estimation problem, we shall first briefly introduce some definitions and notation. For simplicity we will restrict the following definitions of estimators and risk to the case of estimating the parameter p of Bernoulli random variables. Given t independent samples from a Bernoulli random variable $B(p)$, an estimator is an algorithm that takes these samples as input and outputs an answer in $[0, 1]$.

Definition 10 An estimator $\theta = (\theta_t)_{t=1}^\infty$ is a sequence of functions, $\theta_t : \{0, 1\}^t \mapsto [0, 1]$.

Perhaps the first estimator that comes to one’s mind is the *sample mean*, that is $\theta_t = \sum_{i=1}^t X_i/t$. To measure the efficiency of an estimator we define the *risk*, which corresponds to the expected error of an estimator.

Definition 11 Let P be a Bernoulli distribution with mean p and P^t be the corresponding t -fold product distribution. The risk of an estimator $\theta = (\theta_t)_{t=1}^\infty$ is $\mathbf{E}_{(X_1, \dots, X_t) \sim P^t} [|\theta_t(X_1, \dots, X_t) - p|]$, which we will denote by

$$\mathbf{E}_p [|\theta_t(X_1, \dots, X_t) - p|] \text{ or } \mathbf{E}_p [|\theta_t - p|]$$

for brevity.

The risk $\mathbf{E}_p [|\theta_t - p|]$ quantifies the error rate of the estimated value $\hat{p} = \theta_t(Y_1, \dots, Y_t)$ to the real parameter p as the number of samples t grows. Since p is unknown, any meaningful estimator $\theta = (\theta_t)_{t=1}^\infty$ must guarantee that $\lim_{t \rightarrow \infty} \mathbf{E}_p [|\theta_t - p|] = 0$ for all p . For example, *sample mean* has error rate

$$\mathbf{E}_p [|\theta_t - p|] \leq \frac{1}{2\sqrt{t}}.$$

¹ These s, α are scalars in $[0, 1]$ and should not be confused with the internal opinion vector s and the self confidence vector α of an instance $I = (P, s, \alpha)$.




Fig. 1 This is an instance of an opinion formation game, where any algorithm for approximating the equilibrium can be used to construct an estimator for the mean of a Bernoulli random variable

Lemma 5 *Let A be a graph oblivious update rule such that for all instances $I = (P, s, \alpha)$,*

$$\lim_{t \rightarrow \infty} t^{1+c} \mathbf{E} [\|x_A(t) - x^*\|_\infty] = 0.$$

Then there exists an estimator $\theta_A = (\theta_t^A)_{t=1}^\infty$ such that for all $p \in [0, 1]$,

$$\lim_{t \rightarrow \infty} t^{1+c} \mathbf{E}_p [|\theta_t^A - p|] = 0.$$

Proof We construct an estimator $\theta_A = (\theta_t^A)_{t=1}^\infty$ using the update rule A . Consider the instance I_p described in Fig. 1. By straightforward computation, we get that the equilibrium point of the graph is $x_c^* = p/3, x_1^* = p/6 + 1/2, x_0^* = p/6$. Now consider the opinion vector $x_A(t)$ produced by the update rule A for the instance I_p . Note that for $t \geq 1$,

- $x_1^A(t) = A_t(x_c(0), \dots, x_c(t-1), 1, 1/2)$
- $x_0^A(t) = A_t(x_c(0), \dots, x_c(t-1), 0, 1/2)$
- $x_c^A(t) = A_t(x_{W_c^0}(0), \dots, x_{W_c^{t-1}}(t-1), 0, 1/2)$

The key observation is that the opinion vector $x_A(t)$ is a deterministic function of the index sequence W_c^0, \dots, W_c^{t-1} and does not depend on p . Thus, we can construct the estimator θ_A with $\theta_t^A(W_c^0, \dots, W_c^{t-1}) = 3x_c^A(t)$. For a given instance I_p the choice of neighbor W_c^t is given by the value of the Bernoulli random variable with parameter p ($\mathbf{P}[W_c^t = 1] = p$). As a result,

$$\mathbf{E}_p [|\theta_t^A - p|] = 3\mathbf{E} [|x_c^A(t) - p/3|] \leq 3\mathbf{E} [\|x_A(t) - x^*\|_\infty].$$

Since for any instance I_p , we have that

$$\lim_{t \rightarrow \infty} t^{1+c} \mathbf{E} [\|x_A(t) - x^*\|_\infty] = 0,$$

it follows that

$$\lim_{t \rightarrow \infty} t^{1+c} \mathbf{E}_p [|\theta_t^A - p|] = 0$$

for all $p \in [0, 1]$. □

In order to prove Theorem 3 we just need to prove the following claim.

Claim For any estimator $\theta = (\theta_t)_{t=1}^\infty$ there exists a $p \in [0, 1]$ such that $\lim_{t \rightarrow \infty} t^{1+c} \mathbf{E}_p [|\theta_t - p|] > 0$.

The above claim states that for any estimator $\theta = (\theta_t)_{t=1}^\infty$, we can inspect the functions $\theta_t : \{0, 1\}^t \mapsto [0, 1]$ and then choose a $p \in [0, 1]$ such that the function $\mathbf{E}_p [|\theta_t - p|] = \Omega(1/t^{1+c})$. As a result, we have reduced the construction of a lower bound concerning the round complexity of a dynamical process to a lower bound concerning the sample complexity of estimating the parameter p of a Bernoulli distribution. The claim follows by Lemma 6, which we present at the end of the section.

At this point we should mention that it is known that $\Omega(1/\varepsilon^2)$ samples are needed to estimate the parameter p of a Bernoulli random variable within additive error ε . Another well-known result is that taking the average of the samples is the *best* way to estimate the mean of a Bernoulli random variable. These results would indicate that the best possible rate of convergence for an *graph oblivious dynamics* would be $O(1/\sqrt{t})$. However, there is some fine print in these results which does not allow us to use them. In order to explain the various limitations of these methods and results we will briefly discuss some of them. We remark that this discussion is not needed to understand the proof of Lemma 6.

The oldest sample complexity lower bound for estimation problems is the well-known Cramer-Rao inequality. Let $\theta_t : \{0, 1\}^t \mapsto [0, 1]$ be a function such that $\mathbf{E}_p [\theta_t] = p$ for all $p \in [0, 1]$, then

$$\mathbf{E}_p [(\theta_t - p)^2] \geq \frac{p(1-p)}{t}. \quad (16)$$

Since $\mathbf{E}_p [|\theta_t - p|]$ can be lower bounded by $\mathbf{E}_p [(\theta_t - p)^2]$ we can apply the Cramer-Rao inequality and prove our claim in the case of *unbiased* estimators, $\mathbf{E}_p [\theta_t] = p$ for all t . Obviously, we need to prove it for any estimator θ , however this is a first indication that our claim holds.

Sample complexity lower bounds without assumptions about the estimator are usually given as lower bounds for the *minimax risk*, which was defined² by Wald in [39] as

$$\min_{\theta_t} \max_{p \in [0, 1]} \mathbf{E}_p [|\theta_t - p|].$$

Minimax risk captures the idea that after we pick the best possible algorithm, an adversary inspects it and picks the worst possible $p \in [0, 1]$ to generate the samples that our algorithm will get as input. The methods of Le’Cam, Fano, and Assouad are well-known information-theoretic methods to establish lower bounds for the minimax risk. For more on these methods see [38, 41]. As we stated before, it is well known

² Although the minimax risk is defined for any estimation problem and loss function, for simplicity, we write the minimax risk for estimating the mean of a Bernoulli random variable.

that the minimax risk for the case of estimating the mean of a Bernoulli is lower bounded by $\Omega(1/\sqrt{t})$ and this lower bound can be established by Le Cam’s method. In order to show why such results do not work for our purposes we shall sketch how one would apply Le Cam’s method to get this lower bound. To apply Le Cam’s method, one typically chooses two Bernoulli distributions whose means are far but their total variation distance is small. Le Cam showed that when two distributions are close in total variation then given a sequence of samples X_1, \dots, X_t it is hard to tell whether these samples were produced by P_1 or P_2 . The hardness of this *testing* problem implies the hardness of *estimating* the parameters of a family of distributions. For our problem the two distributions would be $B(1/2 - 1/\sqrt{t})$ and $B(1/2 + 1/\sqrt{t})$. It is not hard to see that their total variation distance is at most $O(1/t)$, which implies a lower bound $\Omega(1/\sqrt{t})$ for the minimax risk. The problem here is that the parameters of the two distributions depend on the number of samples t . The more samples the algorithm gets to see, the closer the adversary takes the 2 distributions to be. For our problem we would like to *fix* an instance and then argue about the rate of convergence of any algorithm on this instance. Namely, having an instance that depends on t does not work for us.

Trying to get a lower bound without assumptions about the estimators while respecting our need for a fixed (independent of t) p we prove Lemma 6. In fact, we show something stronger: for *almost all* $p \in [0, 1]$, any estimator θ cannot achieve rate $o(1/t^{1+c})$.

Lemma 6 *Let $\theta = (\theta_t)_{t=1}^\infty$ be a Bernoulli estimator with error rate $\mathbf{E}_p [|\theta_t - p|]$. For any $c > 0$, if we select p uniformly at random in $[0, 1]$ then*

$$\lim_{t \rightarrow \infty} t^{1+c} \mathbf{E}_p [|\theta_t - p|] > 0$$

with probability 1.

Proof Since θ_t is a function from $\{0, 1\}^t$ to $[0, 1]$, θ_t can have at most 2^t different values. Without loss of generality, we assume that θ_t takes the same value $\theta_t(x)$ for all $x \in \{0, 1\}^t$ with the same number of 1’s. For example, $\theta_3(\{1, 0, 0\}) = \theta_3(\{0, 1, 0\}) = \theta_3(\{0, 0, 1\})$. This is due to the fact that for any $p \in [0, 1]$, by Jensen’s inequality, we have

$$\sum_{0 \leq i \leq t} \sum_{\|x\|_1=i} |\theta_t(x) - p| p^i (1 - p)^{t-i} \geq \sum_{0 \leq i \leq t} \binom{t}{i} \left| \frac{\sum_{\|x\|_1=i} \theta_t(x)}{\binom{t}{i}} - p \right| p^i (1 - p)^{t-i}.$$

Therefore, for any estimator θ with error rate $\mathbf{E}_p [|\theta_t - p|]$ there exists another estimator θ' that satisfies the above property and

$$\mathbf{E}_p [|\theta'_t - p|] \leq \mathbf{E}_p [|\theta_t - p|]$$

for all $p \in [0, 1]$. Thus, we can assume that θ_t takes at most $t + 1$ different values.

Let A denote the set of p for which the estimator has error rate $o(1/t^{1+c})$, that is

$$A = \{p \in [0, 1] : \lim_{t \rightarrow \infty} t^{1+c} \mathbf{E}_p [|\theta_t - p|] = 0\}.$$

We show that if we select p uniformly at random in $[0, 1]$ then $\mathbf{P}[p \in A] = 0$. We also define the set

$$A_k = \{p \in [0, 1] : \text{for all } t \geq k, t^{1+c} \mathbf{E}_p [|\theta_t - p|] \leq 1\}.$$

Observe that if $p \in A$ then there exists t_p such that $p \in A_{t_p}$, meaning that $A \subseteq \bigcup_{k=1}^{\infty} A_k$. As a result,

$$\mathbf{P}[p \in A] \leq \mathbf{P}\left[p \in \bigcup_{k=1}^{\infty} A_k\right] \leq \sum_{k=1}^{\infty} \mathbf{P}[p \in A_k].$$

To complete the proof we show that $\mathbf{P}[p \in A_k] = 0$ for all k . Notice that $p \in A_k$ implies that for $t \geq k$, the estimator θ must always have a value $\theta_t(i)$ close to p . Using this intuition we define the set

$$B_k = \{p \in [0, 1] : \text{for all } t \geq k, t^{1+c} \min_{0 \leq i \leq t} |\theta_t(i) - p| \leq 1\}.$$

We now show that $A_k \subseteq B_k$. Since $p \in A_k$ we have that for all $t \geq k$

$$\begin{aligned} t^{1+c} \min_{0 \leq i \leq t} |\theta_t(i) - p| \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} &\leq t^{1+c} \sum_{i=0}^t \binom{t}{i} |\theta_t(i) - p| p^i (1-p)^{t-i} \\ &= t^{1+c} \mathbf{E}_p [|\theta_t - p|] \leq 1. \end{aligned}$$

Thus, $\mathbf{P}[p \in A_k] \leq \mathbf{P}[p \in B_k]$. We write the set B_k as

$$B_k = \bigcap_{t=k}^{\infty} \{p \in [0, 1] : \min_{0 \leq i \leq t} |\theta_t(i) - p| \leq 1/t^{1+c}\}.$$

As a result,

$$\mathbf{P}[p \in B_k] \leq \mathbf{P}\left[\min_{0 \leq i \leq t} |\theta_t(i) - p| \leq 1/t^{1+c}\right], \text{ for all } t \geq k.$$

Each value $\theta_t(i)$ ‘‘covers’’ length $1/t^{1+c}$ from its left and right, as shown in Fig. 2, and since there are at most $t + 1$ such values,

by the union bound we get $\mathbf{P}[p \in B_k] \leq 2(t+1)/t^{1+c}$, for all $t \geq k$. More formally, for a fixed i we get

$$\mathbf{P}\left[|\theta_t(i) - p| \leq \frac{1}{t^{1+c}}\right] \leq \frac{2}{t^{1+c}},$$




Fig. 2 Estimator output at time t

since p is picked uniformly at random. By the union bound, we have that

$$\begin{aligned} \mathbf{P} \left[\min_{0 \leq i \leq t} |\theta_t(i) - p| \leq 1/t^{1+c} \right] &= \mathbf{P} \left[\cup_{0 \leq i \leq t} \{|\theta_t(i) - p| \leq 1/t^{1+c}\} \right] \\ &\leq \sum_{i=0}^t \mathbf{P} \left[|\theta_t(i) - p| \leq 1/t^{1+c} \right] \leq \frac{2(t+1)}{t^{1+c}} \end{aligned}$$

We conclude that $\mathbf{P}[p \in B_k] = 0$. □

Lemma 6 essentially shows that we cannot construct a protocol that is graph-oblivious and converges exponentially fast to the equilibrium, as the dynamics of the original FJ model does. However, as we show in the next section, even a small amount of information about the topology of the graph results in faster protocols.

6 Limited Information Dynamics with Fast Convergence

We already discussed that the reason that graph oblivious dynamics suffer slow convergence is that the update rule depends only on the observed opinions. Based on works for asynchronous distributed minimization algorithms [7, 9], we provide an update rule showing that information about the graph G combined with agents that do not act selfishly can restore the fast convergence rate. Our update rule depends not only on the expressed opinions of the neighbors that an agent i meets, but also on the i -th row of matrix P .

In update rule (6), each agent stores the *most recent* opinions of the random neighbors that she meets in an array and then updates her opinion according to their weighted sum (each agent knows row i of P). For a given instance $I = (P, s, \alpha)$ we call the produced dynamics *Row Dependent dynamics* (Dynamics 2). We have already mentioned that while this update rule guarantees fast convergence it does not guarantee no regret to the agents. To make this concrete we include a simple example.

Example 1 The purpose of this example is to illustrate that the update rule (6) does not ensure the no regret property. If some agents for various reasons exhibit irrational or adversarial behavior, agents that adopt update rule (6) may experience regret. That is the reason that *Row Dependent dynamics* converge exponentially faster than any *no regret dynamics* including the *FTL dynamics*.

Consider the instance of the game of Definition 1 consisting of two agents. Agent 1 adopts update rule (6) and has $s_1 = 0, \alpha_1 = 1/2, p_{12} = 1$ and agent 2 plays adversarially. Thus, s_2, α_2, p_{21} don't need to be specified. By update rule (6), $x_1(t) =$

$x_2(t - 1)/2$ and thus total disagreement cost that agent 1 experiences until round t is

$$\sum_{\tau=0}^t \frac{1}{2}x_1(\tau)^2 + \frac{1}{2}(x_1(\tau) - x_2(\tau))^2 = \sum_{\tau=0}^t \frac{1}{8}x_2(\tau - 1)^2 + \frac{1}{2} \left(\frac{1}{2}x_2(\tau - 1) - x_2(\tau) \right)^2.$$

Since agent 2 plays adversarially, she selects $x_2(t) = 0$ if t is even and 1 otherwise. As a result, the total cost that agent 1 experiences is $\sum_{\tau=0}^t \frac{1}{2}x_1(\tau)^2 + \frac{1}{2}(x_1(\tau) - x_2(\tau))^2 \simeq 3t/8$. Now agent 1 regrets for not adopting the fixed opinion $1/3$ during the whole game play. Selecting $x_1(t) = 1/3$ for all t would incur her total disagreement cost

$$\sum_{\tau=0}^t \frac{1}{2}(1/3)^2 + \frac{1}{2}(1/3 - x_2(t))^2 \simeq 7t/36,$$

which is less than $3t/8$.

The problem with the approach of *Row Dependent Dynamics* is that the opinions of the neighbors that she keeps in her array are *outdated*, i.e. the opinion of a neighbor of agent i has changed since their last meeting. The good news are that as long as this outdatedness is bounded we can still achieve fast convergence to the equilibrium. By bounded outdatedness we mean that there exists a number of rounds B such that all agents have met all their neighbors at least once from $t - B$ to t . The latter is formally stated in Lemma 7, which states that if such a B exists, then the protocol converges exponentially fast to x^* . For convenience, we call such a sequence of B rounds an *epoch*.

Remark 1 Update rule (6), apart from the opinions and the indices of the neighbors that an agent meets, also depends on the exact values of the weights p_{ij} and that is why *Row Dependent dynamics* converge fast. We mention that the lower bound of Sect. 5 still holds even if the agents also use the indices of the neighbors that they meet to update their opinion, since Lemma 5 can be easily modified to cover this case. The latter implies that any update rule that ensures fast convergence would require from each agent i to be *aware* of the i -th row of matrix P .

The idea behind the proof of Lemma 7 is simple: if during each epoch B an agent meets all his neighbors at least once, then certainly at the end of the epoch a full step of the original FJ dynamics will have been computed. This means that the running time will be slower than that of the FJ model by a factor of B .

Lemma 7 *Let $\rho = \min_i a_i$, and $\pi_{ij}(t)$ be the most recent round before round t , that agent i met her neighbor j . If for all $t \geq B$, $t - B \leq \pi_{ij}(t)$ then, for all $t \geq kB$,*

$$\|x(t) - x^*\|_\infty \leq (1 - \rho)^k.$$

Proof To prove our claim we use induction on k . For the induction base $k = 1$,

$$|x_i(t) - x_i^*| = \left| (1 - \alpha_i) \sum_{j \in N_i} p_{ij}(x_j(\pi_{ij}(t)) - x_j^*) \right|$$

Algorithm 2 Row Dependent dynamics

- 1: Initially $x_i(0) = s_i$ for all agent i .
- 2: Each agent i keeps an array M_i of length $|N_i|$, randomly initialized.
- 3: At round $t \geq 0$ each agent i :
 - 4: Meets neighbor with index W_i^t , $\mathbf{P}[W_i^t = j] = p_{ij}$.
 - 5: Suffers cost $(1 - \alpha_i)(x_i(t) - x_{W_i^t}(t))^2 + a_i(x_i(t) - s_i)^2$ and learns $(x_{W_i^t}(t), W_i^t)$.
 - 6: Updates her array M_i and opinion:

$$M_i[W_i^t] \leftarrow x_{W_i^t}(t) \quad x_i(t + 1) = (1 - \alpha_i) \sum_{j \in N_i} p_{ij} M_i[j] + \alpha_i s_i \tag{17}$$

$$\begin{aligned} &\leq (1 - \alpha_i) \sum_{j \in N_i} p_{ij} |x_j(\pi_{ij}(t)) - x_j^*| \\ &\leq (1 - \rho) \end{aligned}$$

Assume that for all $t \geq (k - 1)B$ we have that $\|x(t) - x^*\|_\infty \leq (1 - \rho)^{k-1}$. For $k \geq 2$, we again have that

$$|x_i(t) - x_i^*| \leq (1 - \rho) \sum_{j \in N_i} p_{ij} |x_j(\pi_{ij}(t)) - x_j^*|$$

Since $t - B \leq \pi_{ij}(t)$ and $t \geq kB$ we obtain that $\pi_{ij}(t) \geq (k - 1)B$. As a result, the inductive hypothesis applies, $|x_j(\pi_{ij}(t)) - x_j^*| \leq (1 - \rho)^{k-1}$ and $|x_i(t) - x_i^*| \leq (1 - \rho)^k$. □

In Row Dependent dynamics there does not exist a fixed length window B that satisfies the requirements of Lemma 7. However we can select a length value such that the requirements hold with high probability. To do this observe that agent i should collect the opinions of all of her neighbors, which resembles the process of the *coupons collector problem*. We first state a useful fact concerning this problem, whose proof uses just elementary probability.

Lemma 8 (see e.g. [34]) *Suppose that the collector picks coupons with different probabilities, where n is the number of distinct coupons. Let w be the minimum of these probabilities. If he selects $\ln n/w + c/w$ coupons, then:*

$$\mathbf{P}[\text{collector hasn't seen all coupons}] \leq \frac{1}{e^c}$$

It is now clear that each agent i simply needs to wait to meet the neighbor j with the smallest weight p_{ij} . Therefore, after $\log(1/\delta)/\min_{j \in N_i} p_{ij}$ rounds we have that with probability at least $1 - \delta$ agent i met all her neighbors at least once. Since we want this to be true for all agents, we shall roughly take $B = 1/\min_{i \in V, j \in N_i} p_{ij}$. These calculations become precise in the following lemma.

Lemma 9 Let $\pi_{ij}(t)$ be the most recent round before round t that agent i met agent j and $B = 2 \ln(\frac{nt}{\delta})/w$ where $w = \min_{i \in V} \min_{j \in N_i} p_{ij}$. Then with probability at least $1 - \delta$, for all $\tau \geq B$ and for all $i \in V$ and all $j \in N_i$

$$\tau - B \leq \pi_{ij}(\tau) \leq \tau - 1.$$

Proof Consider an agent i at round $\tau \geq B$ where $B = 2 \ln(\frac{nt}{\delta})/w$ and assume that there exists an agent $j \in N_i$ such that $\pi_{ij}(\tau) < \tau - B$. Agent i can be viewed as a coupon collector that has bought B coupons but has not found the coupon corresponding to agent j . Since $N_i < n$ and $\min_{j \in N_i} p_{ij} \geq w$ by Lemma 8 we have that

$$\mathbf{P}[\text{there exists } j \in N_i \text{ s.t. } \pi_{ij}(\tau) < \tau - B] \leq \frac{\delta}{nt}$$

The proof follows by a union bound for all agents i and all rounds $B \leq \tau \leq t$. \square

Our goal is to prove Theorem 4, showing that the convergence rate of update rule (6) is exponentially fast in expectation (although not as fast as the original FJ dynamics).

Theorem 4 Let $I = (P, s, \alpha)$ be an instance of the opinion formation game of Definition 1 with equilibrium $x^* \in [0, 1]^n$. Then for all rounds $t \geq 6 \ln n/w + 36/w^2 + 9\rho^2/\ln^2 n$,

$$\mathbf{E}[\|x(t) - x^*\|_\infty] \leq \left(\frac{1}{1-\rho} + 1\right) \exp\left(-\rho w \sqrt{t}/(4 \ln(nt))\right),$$

where $x(t) \in [0, 1]^n$ is the opinion vector produced by update rule (6), $\rho = \min_{i \in V} \alpha_i$, $w = \min_{i \in V} \min_{j \in N_i} p_{ij}$.

By direct application of Lemma 7 and Lemma 9, we obtain the following corollary that will be useful in proving Theorem 4.

Corollary 1 Let $x(t)$ be the opinion vector produced by update rule (6) for the instance $I = (P, s, \alpha)$, then with probability at least $1 - \delta$, for all $t \geq 2 \ln(\frac{nt}{\delta})/w$,

$$\|x(t) - x^*\|_\infty \leq \frac{1}{1-\rho} \cdot \exp\left(-\frac{\rho w t}{2 \ln(\frac{nt}{\delta})}\right)$$

where $\rho = \min_{i \in V} \alpha_i$ and $w = \min_{i \in V, j \in N_i} p_{ij}$.

Proof Let $B = 2 \ln(\frac{nt}{\delta})/w$. By Lemma 9 we have that with probability at least $1 - \delta$, for all $i \in V$, $j \in N_i$ and for all $\tau \geq B$,

$$\tau - B \leq \pi_{ij}(\tau)$$

As a result, with probability at least $1 - \delta$ the requirements of Lemma 7 are satisfied. Thus for all $t \geq B$,

$$\|x(t) - x^*\|_\infty \leq (1 - \rho)^{\frac{t}{B}-1} \leq \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho wt}{2 \ln(\frac{nt}{\delta})}\right)$$

□

Corollary 1 states that the convergence happens with high probability. We want to translate this result into one involving the expected error after t iterations of the dynamic. The standard way of doing that is by using the conditional expectations identity. The proof of Theorem 4 is now reduced to choosing a suitable value for the probability δ of the protocol not working. We would like δ to be as small as possible, without blowing up the upper bound on $\|x(t) - x^*\|_\infty$ of Corollary 1.

6.1 The Proof of Theorem 4

Proof Let $u(t) = \|x(t) - x^*\|_\infty$. From Corollary 1 we obtain that for any $\delta \in (0, 1)$, for all rounds $t \geq 2 \ln(\frac{nt}{\delta})/w$,

$$\mathbf{P}\left[u(t) > \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho wt}{2 \ln(\frac{nt}{\delta})}\right)\right] \leq \delta.$$

Since all the parameters of the problem lie in $[0, 1]$, we have $\mathbf{E}[u(t)|u(t) > r] \leq 1$. Now, by conditioning on the event that $u(t) > r$, we get:

$$\begin{aligned} \mathbf{E}[u(t)] &= \mathbf{E}[u(t)|u(t) > r] \mathbf{P}[u(t) > r] + \mathbf{E}[u(t)|u(t) \leq r] \mathbf{P}[u(t) \leq r] \\ &\leq \delta + r, \end{aligned}$$

where $r = \frac{1}{1 - \rho} \exp\left(-\frac{\rho wt}{2 \ln(\frac{nt}{\delta})}\right)$. If we set $\delta = \exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right)$, then:

$$\mathbf{E}[u(t)] \leq \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right) + \exp\left(-\frac{\rho wt}{2 \ln(\frac{nt}{\delta})}\right).$$

We now evaluate r for our choice of probability δ :

$$\begin{aligned} r &= \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho wt}{2 \ln(\frac{nt}{\delta})}\right) = \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho wt}{2 \ln\left(\frac{nt}{\exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right)}\right)}\right) \\ &= \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho wt}{2 \ln(nt) + 2 \frac{\rho w \sqrt{t}}{2 \ln(nt)}}\right) \leq \frac{1}{1 - \rho} \cdot \exp\left(-\frac{\rho wt}{4 \ln(nt) \sqrt{t}}\right) \end{aligned}$$

$$= \frac{1}{1-\rho} \cdot \exp\left(-\frac{\rho w \sqrt{t}}{4 \ln(nt)}\right).$$

Using the previous calculation, we obtain:

$$\begin{aligned} \mathbf{E}[u(t)] &\leq \frac{1}{1-\rho} \cdot \exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right) + \exp\left(-\frac{\rho w \sqrt{t}}{4 \ln(nt)}\right) \\ &\leq \left(\frac{1}{1-\rho} + 1\right) \exp\left(-\frac{\rho w \sqrt{t}}{4 \ln(nt)}\right). \end{aligned}$$

To this end we have established that for $t \geq 2 \ln(\frac{nt}{\delta})/w$ with $\delta = \exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right)$

$$\mathbf{E}[u(t)] \leq \left(\frac{1}{1-\rho} + 1\right) \exp\left(-\frac{\rho w \sqrt{t}}{4 \ln(nt)}\right)$$

The inequality $t \geq 2 \ln(\frac{nt}{\delta})/w$ with $\delta = \exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right)$ can be rewritten as

$$t \geq 2 \frac{\ln n}{w} + 2 \frac{\ln t}{w} + \frac{\rho \sqrt{t}}{\ln(nt)}$$

Notice that

$$\begin{aligned} -\frac{t}{3} &\geq 2 \frac{\ln n}{w} \text{ in case } t \geq 6 \frac{\ln n}{w} \\ -\frac{t}{3 \ln t} &\geq \frac{\sqrt{t}}{3} \geq \frac{2}{w} \text{ in case } t \geq \frac{36}{w^2} \\ -\frac{t}{3} &\geq \frac{\rho \sqrt{t}}{\ln n} \geq \frac{\rho \sqrt{t}}{\ln(nt)} \text{ in case } t \geq \frac{9\rho^2}{\ln^2 n} \end{aligned}$$

As a result, for all $t \geq 6 \frac{\ln n}{w} + \frac{36}{w^2} + \frac{9\rho^2}{\ln^2 n}$ we get that $t \geq 2 \ln(\frac{nt}{\delta})/w$ with $\delta = \exp\left(-\frac{\rho w \sqrt{t}}{2 \ln(nt)}\right)$ and thus

$$\mathbf{E}[u(t)] \leq \left(\frac{1}{1-\rho} + 1\right) \exp\left(-\frac{\rho w \sqrt{t}}{4 \ln(nt)}\right)$$

□

Funding 'Open Access funding provided by the MIT Libraries'

Declarations

Competing Interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alford, J.R., Funk, C.L., Hibbing, J.R., Alford, J.R., Funk, C.L.: Are political orientations genetically transmitted. *Am. Polit. Sci. Rev.* **1**, 153–167 (2005)
2. Abebe, R., Kleinberg, J.M., Parkes, D.C., Sourakakis, C.E.: Opinion dynamics with varying susceptibility to persuasion. CoRR [arXiv:1801.07863](https://arxiv.org/abs/1801.07863) (2018)
3. Bilò, V., Fanelli, A., Moscardelli, L.: Opinion formation games with dynamic social influences. In Cai, Y., Vetta, A. (eds.) *Web and Internet Economics*, pp. 444–458. Springer, Berlin (2016)
4. Bhawalkar, K., Gollapudi, S., Munagala, K.: Coevolutionary opinion formation games. In: *Symposium on Theory of Computing Conference, STOC '13*, Palo Alto, CA, USA, June 1–4, 2013, pp. 41–50 (2013)
5. Bindel, D., Kleinberg, J.M., Oren, S.: How bad is forming your own opinion? In: *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*, Palm Springs, CA, USA, October 22–25, 2011, pp. 57–66 (2011)
6. Bravo, M., Mertikopoulos, P.: On the robustness of learning in games with stochastically perturbed payoff observations. *Games Econ. Behav.* **103**, 41–66 (2017)
7. Bertsekas, D.P., Tsitsiklis, J.N.: *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific (1997)
8. Cesa-Bianchi, N., Lugosi, G.: Potential-based algorithms in on-line prediction and game theory. *Mach. Learn.* **51**(3), 239–261 (2003)
9. Cheung, Y.K., Cole, R.: A unified approach to analyzing asynchronous coordinate descent and tatonnement. CoRR, [arXiv:1612.09171](https://arxiv.org/abs/1612.09171) (2016)
10. Chen, P.-A., Chen, Y.-L., Chi-Jen, L.: Bounds on the price of anarchy for a more general class of directed graphs in opinion formation games. *Oper. Res. Lett.* **44**(6), 808–811 (2016)
11. Cohen, J., Héliou, A., Mertikopoulos, P.: Hedging under uncertainty: regret minimization meets exponentially fast convergence. In: *Proceedings of Algorithmic Game Theory—10th International Symposium, SAGT 2017*, L'Aquila, Italy, September 12–14, 2017, pp. 252–263 (2017)
12. Chierichetti, F., Kleinberg, J.M., Oren, S.: On discrete preferences and coordination. In: *ACM Conference on Electronic Commerce, EC '13*, Philadelphia, PA, USA, June 16–20, 2013, pp. 233–250 (2013)
13. DeGroot, M.H.: Reaching a consensus. *J. Am. Stat. Assoc.* **69**, 118–121 (1974)
14. Deffuant, G., Neau, D., Amblard, F., Weisbuch, G.: Mixing beliefs among interacting agents. *Adv. Complex Syst.* **3**(1–4), 87–98 (2000)
15. Epitropou, M., Fotakis, D., Hofer, M., Skoulakis, S.: Opinion formation games with aggregation and negative influence. In: *Algorithmic Game Theory—10th International Symposium, SAGT 2017*, L'Aquila, Italy, September 12–14, 2017, *Proceedings*, pp. 173–185 (2017)
16. Even-Dar, E., Mansour, Y., Nadav, U.: On the convergence of regret minimization dynamics in concave games. In: *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, Bethesda, MD, USA, May 31–June 2, 2009, pp. 523–532 (2009)
17. Ferraioli, D., Goldberg, P.W., Ventre, C.: Decentralized dynamics for finite opinion games. *Theor. Comput. Sci.* **648**, 96–115 (2016)
18. Friedkin, N.E., Johnsen, E.C.: Social influence and opinions. *J. Math. Sociol.* **15**(3–4), 193–206 (1990)
19. Fotakis, D., Palyvos-Giannas, D., Skoulakis, S.: Opinion dynamics with local interactions. In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI 2016*, New York, NY, USA, 9–15 July 2016, pp. 279–285 (2016)
20. Freund, Y., Schapire, R.E.: Adaptive game playing using multiplicative weights. *Games Econ. Behav.* **29**(1), 79–103 (1999)

21. Foster, D.P., Vohra, R.V.: Calibrated learning and correlated equilibrium. *Games Econom. Behav.* **21**(1), 40–55 (1997)
22. Gourdon, X., Sebah, P.: The euler constant: γ . *Young* **1**, 2n (2004)
23. Ghaderi, J., Srikant, R.: Opinion dynamics in social networks with stubborn agents: equilibrium and convergence rate. *Automatica* **50**(12), 3209–3215 (2014)
24. Gionis, A., Terzi, E., Tsaparas, P.: Opinion maximization in social networks. In: Proceedings of the 13th SIAM International Conference on Data Mining, May 2–4, 2013. Austin, Texas, USA, pp. 387–395 (2013)
25. Hazan, E., Agarwal, A., Kale, S.: Logarithmic regret algorithms for online convex optimization. *Mach. Learn.* **69**(2), 169–192 (2007)
26. Hazan, E.: Introduction to online convex optimization. *Found. Trends Optim.* **2**(3–4), 157–325 (2016)
27. Héliou, A., Cohen, J., Mertikopoulos, P.: Learning with bandit feedback in potential games. In: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4–9 December 2017, Long Beach, CA, USA, pp. 6372–6381 (2017)
28. Hegselmann, R., Krause, U.: Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Soc. Soc. Simul.* **5**, 1 (2002)
29. Jackson, M.O.: *Social and Economic Networks*. Princeton University Press, Princeton (2008)
30. Kleinberg, R., Piliouras, G., Tardos, E.: Multiplicative updates outperform generic no-regret learning in congestion games: Extended abstract. In: Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC '09, pp. 533–542. ACM, New York (2009)
31. Kleinberg, R., Piliouras, G., Tardos, É.: Load balancing without regret in the bulletin board model. *Distrib. Comput.* **24**(1), 21–29 (2011)
32. Krackhardt, D.: A plunge into networks. *Science* **326**(5949), 47–48 (2009)
33. Musco, C., Musco, C., Tsourakakis, C.E.: Minimizing polarization and disagreement in social networks. CoRR, [arXiv:1712.09948](https://arxiv.org/abs/1712.09948) (2017)
34. Motwani, R., Raghavan, P.: *Randomized Algorithms*. Cambridge University Press, Cambridge (1995)
35. Mertikopoulos, P., Staudigl, M.: Convergence to nash equilibrium in continuous games with noisy first-order feedback. In: 56th IEEE Annual Conference on Decision and Control, CDC 2017, Melbourne, Australia, December 12–15, 2017, pp. 5609–5614 (2017)
36. Sergiu, H., Andreu, M.-C.: A simple adaptive procedure leading to correlated equilibrium. *Econometrica* **68**(5), 1127–1150 (2000)
37. Syrgkanis, V., Agarwal, A., Luo, H., Schapire, R.E.: Fast convergence of regularized learning in games. In: NIPS, pp. 2989–2997 (2015)
38. Tsybakov, A.B.: *Introduction to Nonparametric Estimation*, 1st ed. Springer, New York (2008)
39. Wald, A.: Contributions to the theory of statistical estimation and testing hypotheses. *Ann. Math. Stat.* **10**(4), 299–326 (1939)
40. Yildiz, M.E., Ozdaglar, A.E., Acemoglu, D., Saberi, A., Scaglione, A.: Binary opinion dynamics with stubborn agents. *ACM Trans. Econ. Comput.* **1**(4), 19:1–19:30 (2013)
41. Yu, B., Assouad, F., Cam, L.: In: *Festschrift for Lucien Le Cam*, pp. 423–435. Springer, New York (1997)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.