

Relaxing Common Belief for Social Networks

Noah Burrell 

University of Michigan, Ann Arbor, MI, USA
burrelln@umich.edu

Grant Schoenebeck 

University of Michigan, Ann Arbor, MI, USA
schoeneb@umich.edu

Abstract

We propose a relaxation of common belief called *factional belief* that is suitable for the analysis of strategic coordination on social networks. We show how this definition can be used to analyze revolt games on general graphs, including by giving an efficient algorithm that characterizes a structural result about the possible equilibria of such games.

This extends prior work on common knowledge and common belief, which has been too restrictive for use in understanding strategic coordination and cooperation in social network settings.

2012 ACM Subject Classification Theory of computation → Social networks; Theory of computation → Network games; Theory of computation → Algorithmic game theory

Keywords and phrases Social networks, network revolt games, common belief

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.45

Related Version A full version of the paper is available at <https://arxiv.org/abs/2009.12936>.

Funding The authors gratefully acknowledge the support of the National Science Foundation under Career Award #1452915.

1 Introduction

Common knowledge and its analogues (e.g. common belief) are fundamental in the analysis of coordination and cooperation in strategic settings. Informally, common knowledge is the phenomenon that, within a population, everyone knows that a proposition is true, everyone knows that everyone knows that it is true, everyone knows that everyone knows that everyone knows that it is true, and so on, ad infinitum. The existence of common knowledge often shows up as an assumption that underlies some kind of strategic coordination. For example, in the standard game theoretic setting, it is typically assumed that the agents playing a game have common knowledge both of the payoffs of the game and of the rationality of each agent. This assumption undergirds the agents' ability to coordinate on equilibria.

A rigorous investigation into the consequences of this type of assumption can be traced back to Robert Aumann [1], who showed that if two rational agents have the same prior, and their posteriors for an event are common knowledge, then their posteriors must be equal (i.e. it is impossible for such agents to “agree to disagree”). Later work considered whether assumptions about common knowledge, and the mathematical consequences of those assumptions, were realistic for trying to explain economic phenomena. In particular, it seemed unrealistic to expect agents to reason about infinite hierarchies of knowledge. Initial attempts to address this critique involved truncating the infinite hierarchy that is required in our informal definition of common knowledge after some large, but finite, number of levels. However, this line of inquiry resulted in the discovery of “common knowledge paradoxes” that arose from examples like Ariel Rubinstein’s Electronic Mail Game [15]. Such examples demonstrated that in situations where common knowledge was required for coordination, if the infinite hierarchy were truncated to be any finite hierarchy, strategic



© Noah Burrell and Grant Schoenebeck;
licensed under Creative Commons License CC-BY
12th Innovations in Theoretical Computer Science Conference (ITCS 2021).
Editor: James R. Lee; Article No. 45; pp. 45:1–45:20



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

agents behaved unexpectedly and unrealistically. That is, they behaved very differently when a proposition was close to common knowledge, in the sense that many – but only finitely many – hierarchies of knowledge were satisfied, then they did when the proposition was actually common knowledge (and therefore infinitely many hierarchies of knowledge were satisfied).

An important breakthrough came from Dov Monderer and Dov Samet in 1989 [8]. Inspired by an earlier version of Rubinstein’s work [15] and the work of Aumann [1], Monderer and Samet proposed an alternative to truncating the infinite hierarchy of knowledge: relaxing the requirement of “knowledge” at each level in the hierarchy. They introduced the notion of *common belief* – an analogous concept to common knowledge that is defined by replacing “everyone knows” with “everyone believes with probability (at least) p ” at each place where it occurs in our informal definition of common knowledge. Under this definition, they showed that common belief approximates common knowledge in precisely the way that truncating the infinite hierarchy did not: When common knowledge of some proposition is relaxed to common belief, agents behave in approximately the same way they did when they had common knowledge of the proposition. For example, common belief can be applied to generalize Aumann’s result that was mentioned above: when rational agents have the same prior and their posteriors about an event are commonly believed, these posteriors must be approximately equal.

In the same paper, they also detailed a key insight that allowed them to address the critique that common knowledge might be too unrealistic to have explanatory power in Economics. They showed that common knowledge has an alternative definition (used implicitly in [1]), that is formally equivalent to the more natural definition, but simpler to reason about. This alternative – inspired by the example of public announcements – defines common knowledge in terms of events that are *evident knowledge*: events for which their occurrence implies knowledge of their occurrence within the entire population. Common belief has a similar formally equivalent, alternative definition in terms of events that are *evident belief*: events for which their occurrence implies belief with probability at least p of their occurrence within the entire population.

Here, it is worth noting that our work most heavily borrows from that of Monderer and Samet and, as such, their paper [8] is highly recommended both as a primer for what follows here and as an introduction to common knowledge and common belief in general.

In a somewhat orthogonal line of research, Michael Suk-Young Chwe studied common knowledge (not common belief) on networks in the context of a revolt game [3, 4]. In his setting, each agent has a threshold that represents the size of the revolt in which she would agree to participate. For example, an agent with a threshold of 3 would need to know that at least 2 other agents would participate in order to herself agree to participate in the revolt. To decide whether this threshold is met, agents learn the thresholds of their neighbors. They use this information (and knowledge of their 2-hop neighborhood in the graph) to reason about which agents have their thresholds satisfied and consequently whether their own threshold is satisfied. Because agents require absolute certainty in this setting, they only consider their neighbors when conducting this reasoning.

The strict requirement that absolute certainty is necessary for an agent to revolt implies that for a group of agents to revolt, it must be common knowledge among them that all of their thresholds are satisfied. Consequently, the problem of finding groups of revolting agents reduces to the problem of finding cliques of a certain size in the network. However, an important innovation of Chwe’s work, in contrast to the settings described above, is that common knowledge in his setting is a local phenomenon occurring within those cliques, not a global phenomenon occurring within the entire population.

Both approaches – that of Monderer and Samet and that of Chwe – are limited in their usefulness in social network settings, because their respective notions are unlikely to apply in many natural graphs that are used to model social networks. For example, an Erdős-Rényi random graph of n vertices with $p = \frac{10}{n}$ will almost certainly be sparse, so population-level phenomenon like common belief will not arise. On the other hand, an Erdős-Rényi random graph of n vertices with $p = \frac{1}{2}$ would be unlikely to contain large cliques, so the phenomenon of local common knowledge would be severely limited despite the fact that each agent, in seeing about half the graph, should have a lot of information about the entire population and might be expected to be able to coordinate with some large fraction of it.

We propose that Monderer’s and Samet’s concept of common belief – itself a relaxation of common knowledge – can be further relaxed to a notion of common belief among a faction (i.e. a minimal-size subset) of the population, while retaining both its mathematical simplicity (in being defined in terms of events that are evident belief) and its economic explanatory power. We refer to this notion as *factional belief*.

Factional belief is a natural application of the ideas of Monderer and Samet to network settings similar to those of Chwe. It retains from Chwe the idea that common knowledge/belief can occur in only a subset of the population and still motivate their behavior. However, it is not prohibitively strict, such that it would be unlikely to occur endogenously in many natural graphs. Factional belief is not necessarily local – agents can and may need to reason about agents outside of their neighborhood. As such, it does not require cliques.

1.1 Our Contributions

- We formally define a notion of factional belief (Section 2), which can be used to analyze revolt games on general graphs (Section 3). Prior notions of common knowledge and common belief were insufficient for this type of analysis.
- We provide an algorithm that characterizes a structural result about the types of equilibria that are possible in instances of the network revolt games described in Section 3 (Section 4) and a natural extension of those games (Section 5).
- We show that, surprisingly, it is sufficient for our algorithm to only have access to the degree sequence of the network; additional details of the network beyond the degree sequence are not relevant.
- In the full version of the paper, we describe algorithms for additional natural extensions of the network revolt game model. Further, we demonstrate the practical utility of our algorithms by applying them to simulated network data to explore how various parameters of networks and of the model relate to the size of revolts that are supported in equilibria of the network revolt game.

1.2 Additional Related Work

The work of Stephen Morris is particularly notable when surveying the literature related to common knowledge and common belief. Morris, often following the work of Monderer and Samet, has done much theoretical work relating to common knowledge and common belief [9, 11, 12, 13]. More recently, he has also collaborated with Benjamin Golub to study higher-order reasoning, reminiscent of the infinite hierarchy of reasoning in the initial definitions of common knowledge and common belief, in network settings [6, 7].

Underlying Morris’ work, above, and our work is our assertion that common knowledge, common belief, and factional belief are useful, not just as mathematical concepts, but for understanding real social and economic phenomena. Here, we briefly outline some relevant

work in applying common knowledge and common belief to explaining such phenomena. One clear direction for future research is to try to similarly apply factional belief as an explanatory tool in these and other settings.

Morris has applied common knowledge/belief to settings such as contagion [10] and global games [14].

Chwe, in 2013, revisited his earlier work and expanded his purview to consider how common knowledge is generated in society [5]. He proposed that the importance of rituals in society can be understood from the perspective that rituals create conditions under which common knowledge can be generated. This theory fits nicely with understanding common knowledge through the lens of evident knowledge events. Another potential direction for future research would be to try to mathematically model this ritualistic generation of common knowledge.

Finally, common knowledge, common belief, and factional belief can be used as tools to help understand the formation and transformation of social norms. In particular, Cristina Bicchieri proposed and meticulously advocated for a definition of social norms of which our definition of factional belief is very reminiscent [2].

2 Defining Factional Belief

For the following definitions, let (Ω, Σ, \Pr) be a probability space, where Ω denotes a set of states, Σ denotes a σ -algebra of events, and \Pr denotes a probability measure on Σ . Let I denote a set of agents.

For each $i \in I$, Π_i is a partition of Ω into measurable sets with positive probability. It is, therefore, a countable partition. For $\omega \in \Omega$, the element of Π_i that contains ω is written as $\Pi_i(\omega)$. Π_i can be interpreted as the information available to agent i . That is, $\Pi_i(\omega)$ is the set of states that are indistinguishable to i when i observes ω . Let $B_i^p(E)$ denote the event that agent i believes in event E with probability at least p . Formally, we write $B_i^p(E) = \{\omega : \Pr[E|\Pi_i(\omega)] \geq p\}$. Lastly, for events E and F , we use the notation $E \subseteq F$ to denote that, whenever E occurs, F occurs.¹

The following examples are helpful to illustrate this notation. When rolling a fair die, with equally probably outcomes in the set $\{1, 2, 3, 4, 5, 6\}$ we have $\{2, 4\} \subseteq \{\text{Outcome is even}\}$. Similarly, when the die has been tossed, but the outcome has not been revealed, we have the event $B_i^{\frac{1}{2}}(\{\text{Outcome is even}\})$ for any agent i .

This notation is borrowed from Monderer and Samet [8], and the rest of the terms and claims in this section are defined and stated, respectively, to be analogous to those from their paper.

► **Definition 1** (Evident (p, μ) -belief). *An event E is an evident (p, μ) -belief if there exists (at least) a μ fraction of agents such that whenever E occurs, those agents assign a probability of at least p to its occurrence. That is:*

$$\exists J \subseteq I \text{ with } |J| \geq \mu|I| \text{ such that for each } j \in J, E \subseteq B_j^p(E).$$

Following, Monderer and Samet, we first define our notion of factional belief in terms of events that are evident (p, μ) -belief. To maintain consistency with their work, we refer to this notion of factional belief as common (p, μ) -belief.

¹ Note that, unlike with knowing, it is possible for an agent to believe something that is not true. In particular, $B_i^p(E)$ need not be a subset of E .

► **Definition 2** (Common (p, μ) -belief). *An event F is common (p, μ) -belief at $\omega \in \Omega$ if there exists an evident (p, μ) -belief event E such that $\omega \in E$ and*

$$\exists J \subseteq I \text{ with } |J| \geq \mu|I| \text{ such that for each } j \in J, E \subseteq B_j^p(F).$$

That is, F is common (p, μ) -belief whenever there is an event (E) that is an evident (p, μ) -belief whose occurrence implies the existence of (at least) a μ fraction of agents that believe with probability at least p in F . Note that any event E that is an evident (p, μ) -belief is trivially also common (p, μ) -belief (with $F = E$).

Now, an important property shared by common knowledge and common belief is the formal equivalence of their definitions in terms of evident events and their intuitive definitions as infinite hierarchies. Common (p, μ) -belief retains this property. In order to state this result formally in Proposition 4, we need to formally define the infinite hierarchy, which is done below in Definition 3.

Informally, each level ($n \geq 1$) in this hierarchy, refers to the event that there exists (at least) a μ fraction of agents who believe with probability (at least) p in the previous level of the hierarchy. The initial level ($n = 0$) is simply the relevant event F . So, written out entirely, the full informal definition would be that an event F is common (p, μ) -belief if there exists a μ fraction of agents who believe F with probability p , there exists a μ fraction of agents who believe with probability p that there exists a μ fraction of agents who believe F with probability p , and so on, ad infinitum.

► **Definition 3.** *For every event F and every $0 \leq p \leq 1$ let*

$$E^{p,\mu}(F) = \bigcap_{n \geq 1} F_\mu^n,$$

where $F_\mu^0 = F$ and F_μ^n is the event “ $\exists J \subseteq I$ with $|J| \geq \mu|I|$ such that $\forall j \in J, B_j^p(F_\mu^{n-1})$ ”.

► **Proposition 4.** *For every event F , every $0 \leq p \leq 1$, and every $0 \leq \mu \leq 1$:*

1. $E^{p,\mu}(F)$ is an evident (p, μ) -belief and $\exists J \subseteq I$ with $|J| \geq \mu|I|$ such that $\forall j \in J, E^{p,\mu}(F) \subseteq B_j^p(F)$.
2. F is common (p, μ) -belief at ω if and only if $\omega \in E^{p,\mu}(F)$.

The proof of this proposition is essentially the same as the proof of the analogous proposition (Proposition 2) in Monderer’s and Samet’s paper [8]. However, the details of the proof are not particularly relevant or instructive with regard to our contributions in this work, so we omit them here and consign them to an appendix of the full version of the paper. The important takeaway from this proposition, is, as noted above, the formal equivalence between the definition of common (p, μ) -belief in terms of evident (p, μ) -belief (Definition 2) and the hierarchical definition (Definition 3). As with the analogous definitions for common knowledge and common belief, the latter definition is more intuitive and perhaps more natural, but the former definition is more mathematically convenient and is what we will reference in what follows. The former definition is also a notion that better corresponds to how agents might be expected to reason about this type of belief in reality, since it is unrealistic to suppose that they consider infinite hierarchies of beliefs.

3 Model

Let $G = (V, E)$ be a graph and I be a set of n agents, such that each vertex $v_i \in V$ corresponds to an agent $i \in I$. We will think of G as representing a social network of strategic agents who are participating in a revolt game. The graph G is common knowledge among the agents.

Nature draws the state of the world $s \in S$ according to a distribution D_S and selects a type $t_i \in T = \{\alpha, \nu\} \cup X$ for each agent i , where $X = \cup_j \{\chi_j\}$ is non-empty (and finite). The types are selected independently at random according to a distribution D_T^s associated with state s . D_S and D_T^s for each $s \in S$ are common knowledge among the agents.

Each agent i will observe the type of each agent $k \in I$ such that $(v_i, v_k) \in E$ (this set of agents constitutes the set of *neighbors* of i). The information resulting from this observation – an agent’s type and the types of all of her neighbors – defines that agent’s *context*. When agent i has the context c , we write it as $c(i) = c$. We use C to denote the set of all contexts that are possible in G .

Ex ante, or, before selection of the state, the assignment of types, and the observation of contexts, agents choose a pure *strategy* $\sigma : C \rightarrow \{R, Y\}$, where $\{R, Y\}$ is the set of *actions*.² A *strategy profile* $(\sigma_1, \sigma_2, \dots, \sigma_n)$ is collection of strategies for each agent. Let $\sigma_{-i} = (\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ denote the strategies of all the agents except for i . Similarly, *ex post*, or, after the selection of types and observation of contexts, an *action profile* (a_1, a_2, \dots, a_n) where $a_i = \sigma_i(c(i))$ is a collection of actions for each agent.

Let $\mathcal{R}(a_1, a_2, \dots, a_n) = \{i \in [n] : a_i = R\}$ denote the set of agents who play the action R (revolt) given their context and strategy. Agent i with type $t_i = t$ receives a payoff according to the function $f_i^t : \{R, Y\}^n \rightarrow [0, 1]$:

$$f_i^\alpha(a_1, a_2, \dots, a_n) = \begin{cases} 1, & \text{if } a_i = R, \\ 0, & \text{otherwise,} \end{cases}$$

$$f_i^\nu(a_1, a_2, \dots, a_n) = \begin{cases} 1, & \text{if } a_i = Y, \\ 0, & \text{otherwise,} \end{cases}$$

and

$$f_i^{\chi_j}(a_1, a_2, \dots, a_n) = \begin{cases} 1 - p_j, & \text{if } |\mathcal{R}(a_1, a_2, \dots, a_n)| \geq \mu_j \cdot n \text{ and } a_i = R, \\ p_j, & \text{if } |\mathcal{R}(a_1, a_2, \dots, a_n)| < \mu_j \cdot n \text{ and } a_i = Y, \\ 0, & \text{otherwise.} \end{cases}$$

where $p_j, \mu_j \in [0, 1]$ for each j and are common knowledge among the agents.

We call σ_i a *best-response* to σ_{-i} when σ_i maximizes agent i ’s *ex ante* expected payoff given σ_{-i} , which, by linearity of expectation, is equivalent to maximizing her expected payoff for each $c \in C$ (using the beliefs she would have about the contexts of the other agents after observing c).

We say that a strategy profile $(\sigma_1, \sigma_2, \dots, \sigma_n)$ is an *equilibrium* when each strategy σ_i is a best-response to σ_{-i} .

Intuitively, the game is defined so that there is a natural strategy for each agent based on type:

- Agents of type α should always revolt.
- Agents of type ν should never revolt.
- Agents of type χ_j should conditionally revolt:
 - Each type χ_j is characterized by a pair of thresholds (p_j, μ_j) that indicate an agent of type χ_j should revolt when she believes that $\Pr[|\mathcal{R}(a_1, a_2, \dots, a_n)| \geq \mu_j] \geq p_j$.

An important thing to note about this game is that when there are agents of type ν , there is not an equilibrium where each agent revolts regardless of her beliefs. Similarly, when there are agents of type α , there is not an equilibrium where each agent chooses not to revolt regardless of her beliefs.

² “Revolt” and “Yield,” respectively.

Even without these, though, there are still potentially multiple equilibria corresponding to revolts of different sizes, and our work does not address equilibrium selection. Consequently, in what follows, we will write that a revolt (of a particular size) is *supported* in equilibrium instead of writing that a revolt will occur. Similarly, we write that agents are *secure enough to revolt* when they sufficiently believe their thresholds are met instead of writing that agents will revolt.

Lastly, in this work we primarily consider the *largest* revolts that are supported in some equilibrium. Such revolts are supported by *symmetric* equilibria, in which each agent adopts the same strategy – namely, the strategy detailed in the bullet points above.

3.1 Motivating Example

To see this model in practice and get a feeling for how agents need to reason about the concepts that we have introduced, we will work through a modest example.

Suppose that G is a grid of n vertices embedded on a torus, so that each the vertex associated with each agent is adjacent to the vertices representing exactly four other agents (and there is no boundary). Thus, a context in this graph consists of the central agent and her type plus the types of each of her neighbors in G . In this example, we will have two types of agents: (1) Agents of type χ , who want to revolt conditionally; they feel secure enough to revolt when the threshold pair $(p = \frac{2}{5}, \mu = \frac{1}{2})$ is satisfied. (2) Agents of type ν , who never want to revolt. We will also have two equally-likely states: an anti-government state A , where agents are of type χ with probability $\frac{4}{5}$ and a pro-government state B where agents are of type ν with probability $\frac{4}{5}$.

We also define the notion of a *candidate agent*. In this setting, for reasons that will become clear in the proof of the following proposition, we refer to agents of type χ with two or more neighbors of type χ as *candidate agents*.

► **Proposition 5.** *In this model, when G is sufficiently large, the event that at least $\frac{1}{2}$ of the agents are candidate agents is an evident $(\frac{2}{5}, \frac{1}{2})$ -belief. When it occurs, it is common $(\frac{2}{5}, \frac{1}{2})$ -belief that supports a revolt of size $\frac{1}{2}$.*

Proof. First, we will compute $\Pr[\text{State} = A|c]$ for each context c and use the computed values to demonstrate that our definition of candidate agents is the correct one for this setting; candidates are likely to feel secure enough to revolt, based on their p -thresholds.

Each agent has 5 independent samples from the probability distribution defined by the state, with which to calculate the probability of each state, given her information. The first step in this is to compute the likelihood of her context given the state, which can be calculated by evaluating the probability mass function of a binomial distribution with the appropriate parameters. Then, the probability for each state given a certain neighborhood can be calculated using Bayes' rule. The results of these calculations are shown in Table 1.

Now, agents of type χ that only have at most one other neighbor of type χ , believe that the state is A with probability at most $\frac{1}{5}$, so they will not feel secure enough to revolt; a revolt of size $\mu = \frac{1}{2}$ is very unlikely to be supported when the state is B , and they do not sufficiently believe that the state is A .

As a result, it is exactly agents of type χ with contexts such that they have two or more neighbors of type χ , which we defined earlier as candidate agents, in which we are interested (cases with $k \geq 3$ in Table 1b). In particular, we are interested in the probability that a majority of agents are candidates.

▷ **Claim.** The probability that at least half the agents are candidates, given that the state is A is at least $\frac{1739}{3125}$.

Proof of Claim. For this, we can try to count all the non-candidate agents and see how likely they are to outnumber the candidates.

Let X count the number of non-candidate agents. Specifically, let $X = \sum_{i \in [n]} X_i$, where X_i is an indicator variable that is 1 if and only if agent i is not a candidate. For each X_i , therefore, the expectation of X_i is equal to the probability that i is not a candidate, which can happen in three ways: (1) i is of type ν , (2) i is of type χ and has no neighbors of type χ , or (3) i is of type χ and has one neighbor of type χ . The probability of (1) is trivially $\frac{1}{5}$. Conditioned on i being of type χ , the respective probabilities of the remaining possibilities are (2) $\frac{1}{625}$ and (3) $\frac{16}{625}$. To obtain unconditional probabilities for (2) and (3), we multiply each conditional probability by $\frac{4}{5}$. Consequently, since the events (1), (2), and (3) are disjoint, we have the following:

$$E[X_i] = \frac{1}{5} + \frac{4}{5} \left(\frac{1}{625} \right) + \frac{4}{5} \left(\frac{16}{625} \right) = \frac{693}{3125}.$$

Therefore, by linearity of expectation, $E[X] = \frac{693n}{3125}$.

Now, we need to upper bound the probability that the number of non-candidate agents is the majority. Ideally for this, we would like to use something tight, like a typical Chernoff-Hoeffding bound. Unfortunately, here, the X_i 's are not independent; the contexts of agents that share neighbors are correlated, which complicates the calculation.

For our more rigorous analysis later, we will insist on tighter probability bounds. But for the sake of simplicity, in this example, Markov's inequality is sufficient:

$$\Pr \left[X \geq \frac{n}{2} \right] \leq \frac{\frac{693n}{3125}}{\frac{n}{2}} = \frac{1386}{3125} \approx 0.44.$$

Consequently, the probability that at least half the agents are candidates, given that the state is A , is at least $(1 - \frac{1386}{3125}) = \frac{1739}{3125}$. \triangleleft

Candidate agents can perform this exact same calculation and, further, they believe that the state is A with probability at least $\frac{4}{5}$. Thus, when the graph is large enough that the context of any individual agent is inconsequential in their reasoning about the total fraction of candidate agents, the probability that they assign to at least half of the agents being candidate agents is at least $\frac{4}{5} \cdot (\frac{1739}{3125}) \geq \frac{2}{5}$.

That is, in sufficiently large graphs, candidate agents believe with probability at least $\frac{2}{5}$ that at least $\frac{1}{2}$ of the agents are candidate agents. As a result, the event that at least $\frac{1}{2}$ of the agents are candidate agents is an evident $(\frac{2}{5}, \frac{1}{2})$ -belief. Consequently, when it occurs, by definition, it is common $(\frac{2}{5}, \frac{1}{2})$ -belief. In this event, a revolt of size $\frac{1}{2}$ is supported, since at least half of the agents have their thresholds satisfied, and consequently feel secure enough to revolt. \blacktriangleleft

4 Applying Factional Belief: A Computational Perspective

The network revolt game model described in Section 3 gives us a useful setting in which to apply our definitions from Section 2 and demonstrate how they can provide insight into strategic coordination. In this section, we explore the interaction between factional belief and strategic coordination from a computational perspective, by trying to answer an elementary question: When can we efficiently determine if strategic coordination (i.e. in our model, a revolt) is supported under given conditions?

■ **Table 1** Results of the probabilistic calculations for the motivating example.

(a) The likelihood, in each state, of having a context with k agents of type χ . (b) The likelihood of state A , given a context with x agents of type χ .

	$\Pr[c \text{State} = A]$	$\Pr[c \text{State} = B]$		$\Pr[\text{State} = A k = x]$
$k = 0$	$\frac{1}{3125}$	$\frac{1024}{3125}$	$x = 0$	$\frac{1}{1025}$
$k = 1$	$\frac{4}{625}$	$\frac{256}{625}$	$x = 1$	$\frac{1}{65}$
$k = 2$	$\frac{32}{625}$	$\frac{128}{625}$	$x = 2$	$\frac{1}{5}$
$k = 3$	$\frac{128}{625}$	$\frac{32}{625}$	$x = 3$	$\frac{4}{5}$
$k = 4$	$\frac{256}{625}$	$\frac{4}{625}$	$x = 4$	$\frac{64}{65}$
$k = 5$	$\frac{1024}{3125}$	$\frac{1}{3125}$	$x = 5$	$\frac{1024}{1025}$

In pursuing an answer to this question, we seek to apply the intuition that we have constructed in our motivating example to a more general setting. Toward that end, a relatively modest generalization of the model used in the example – our Fundamental Case, below – is rich enough to provide an interesting, non-trivial answer to our question and to provide robust intuition that guides us through the various extensions of the model that we discuss in Section 5.

4.1 Fundamental Case: Low-Degree Graphs, Two States, and Three Types

Recall that our model requires a description of possible agent types and possible states (which specify probability distributions over those types). For our initial case, there are three agent types and two possible states.

Our focus will be on agents who want to revolt conditionally – agents of type χ – who have the threshold pair (p, μ) for an arbitrary $0 \leq p \leq 1$ and $0 \leq \mu \leq 1$. In some sense, these are the truly “strategic” agents; they need to coordinate with other agents in order to feel secure enough to revolt. There are also agents who always behave in a prescribed manner, regardless of their contexts or the state: pro-government agents of type ν , who will never revolt, and anti-government agents of type α , who will always revolt.

The possible states are A , an anti-government state, and B , a pro-government state. D_S , the distribution over the states, D_T^A and D_T^B , the distributions over the types in each state, and p and μ , the threshold values for agents of type χ , are commonly known to the agents under the prior P . Intuitively, the labels assigned to the states as being anti- and pro-government correspond to an implication that as the number of agents grows, the size of the largest revolt supported in state A should be larger than in state B . In what follows, we assume that the given labels are assigned correctly, but in practice the correct labels can be determined by running an algorithm that we present later with each possible labeling and comparing the results: When the states are correctly labeled, Algorithm 1 will return X_A and X_B such that $X_A \geq X_B$.

In this more concrete setting, we are able to pose a straightforward question: Given values μ^* and q^* , is a revolt of size (at least) μ^* supported with probability at least q^* ? We refer to this problem as REVOLT.

► **Definition 6** (REVOLT).

Given: (G, P, μ^*, q^*) , where G is a graph with n vertices and $P = (p, \mu, D_T^A, D_T^B, D_S)$ is the common prior.

Question: Is a revolt of size at least μ^* supported with probability at least q^* ?

45:10 Relaxing Common Belief

Note that there is a nuance regarding the timing of the network revolt game model described in Section 3. REVOLT considers the likelihood of revolt of a certain size being supported in a given network *ex ante* – i.e. before the selection of the state and the assignment of types to agents.

Although the question posed by REVOLT is straightforward to state, it is not straightforward to solve efficiently. In fact, we have the following hardness result, which we prove in an appendix of the full version of the paper:

► **Proposition 7.** *REVOLT is NP-hard.*

Consequently, we will instead consider a relaxed version of the problem that we will be able to solve efficiently. In order to define this new problem, we first require p and μ to be strictly between 0 and 1. We will also need two error terms, which will be constants given as input: $\epsilon, \delta > 0$. Lastly, we introduce two additional priors, as follows (recall $P = (p, \mu, D_T^A, D_T^B, D_S)$ is the common prior given as input to both problems):

$$P^- = (p - \delta, \mu - \epsilon, D_T^A, D_T^B, D_S), \quad P^+ = (p + \delta, \mu + \epsilon, D_T^A, D_T^B, D_S).$$

Now, we are ready to define our new problem, PROMISE REVOLT.

► **Definition 8 (PROMISE REVOLT).**

Given: $(G, P, \mu^*, \epsilon, \delta)$, where G is a graph with n vertices, $P = (p, \mu, D_T^A, D_T^B, D_S)$ is the common prior, and $\epsilon, \delta > 0$ are constants.

Output: When exactly one of the following cases is true, output the corresponding symbol:

Ω : A revolt of size $\mu^* - \epsilon$ is supported with probability $q^* \geq 1 - \delta$ under the prior P^- .

A : A revolt of size μ^* is supported with probability $q^* \in [\Pr[\text{State} = A] - \delta, \Pr[\text{State} = A] + \delta]$ under the prior P .

\emptyset : A revolt of size $\mu^* + \epsilon$ is supported with probability $q^* \leq \delta$ under the prior P^+ .

PROMISE REVOLT is named so as to emphasize that this it is a promise problem, in the typical sense. The “promise” is that the given instance is such that exactly one of the cases from the definition of the problem is true. Given that promise, the three cases are mutually exclusive and any solution is required to always output the correct answer. It is exactly this promise to exclude difficult inputs that makes PROMISE REVOLT easier to solve than REVOLT.

Still, in providing an algorithm to solve PROMISE REVOLT, we will see that the problem has a structure that allows us to closely approximate REVOLT by solving PROMISE REVOLT (for large graphs). This structure is discussed in more detail at the end of this subsection (4.1), particularly with respect to Figure 1, which helps to illustrate this intuition.

Lastly, note that we assume all inputs to REVOLT and PROMISE REVOLT are given as rational numbers. This leads us to state our first theorem:

► **Theorem 9.** *Given $\epsilon > 0$, the prior P , and $\mu^* \in [0, 1]$, there exists $\delta(n) \in \frac{1}{\exp(\Omega_{\epsilon, P}(\sqrt[3]{n}))}$ such that for any graph G with n vertices where the largest degree of any vertex is $O_{\epsilon, P}(\sqrt[3]{n})$, Algorithm 3 can be used to solve PROMISE REVOLT($G, P, \mu^*, \epsilon, \delta = \delta(n)$) in polynomial time.*

Before we get to the proof of Theorem 9, we briefly discuss our assumption that each agent has a degree that is $O(\sqrt[3]{n})$. Primarily, this assumption serves to simplify our analysis. For simplicity, it is convenient to make some distinction between “low-degree” agents and “high-degree” agents to highlight the fact that a prototypical high-degree agent would have a lot more information about the state than a prototypical low-degree agent. However, any

■ **Algorithm 1** Finding the expected size of largest revolt supported in each state.

Input: A graph $G = (V, E)$ of n vertices, the prior $P = (p, \mu, D_T^A, D_T^B, D_S)$.
Output: X_A and X_B , the expected size of the largest revolt supported in states A and B , respectively.

Let $e_s(\tau)$ denote the expected fraction of type τ agents in state s .
 Define the set of candidate states $S_C = \{s \in \{A, B\} : e_s(\chi \cup \alpha) \geq \mu\}$.

if $S_C = \emptyset$ **then**
 $X_s = e_s(\alpha) \forall s \in \{A, B\}$.
end

if $S_C = \{A, B\}$ **then**
 $X_s = e_s(\chi \cup \alpha) \forall s \in \{A, B\}$.
end

if $S_C = \{A\}$ **then**
 Let $C(\chi)$ be the set of contexts centered around agents of type χ .
 Define the set of candidate contexts $C_C = \{c \in C(\chi) : \Pr[\text{State} = A|c] \geq p\}$.
 (Below, we slightly abuse notation, treating C_C as if it were a type when writing $e_s(C_C)$.)
 if $e_A(C_C \cup \alpha) \geq \mu$ **then**
 $X_s = e_s(C_C \cup \alpha) \forall s \in \{A, B\}$.
 end
 else
 $X_s = e_s(\alpha) \forall s \in \{A, B\}$.
 end

end
return X_A, X_B .

■ **Algorithm 2** Comparing the expected size of the largest revolt supported in each state.

Input: X_A, X_B (output from Algorithm 1) and μ^* .
Output: Ω , \mathbf{A} , or \emptyset .

if $X_A \geq \mu^*$ and $X_B \geq \mu^*$ **then**
 return Ω
end

if $X_A \geq \mu^*$ and $X_B < \mu^*$ **then**
 return \mathbf{A}
end

if $X_A < \mu^*$ and $X_B < \mu^*$ **then**
 return \emptyset
end

Algorithm 3 Solving PROMISE REVOLT.

Input: $(G, P, \mu^*, \epsilon, \delta)$, where G is a graph of n vertices, $P = (p, \mu, D_T^A, D_T^B, D_S)$ is the common prior, and $\epsilon, \delta > 0$ are constants

Output: Ω , \mathbf{A} , \emptyset , or **Null**.

$X_{A_1}, X_{B_1} \leftarrow$ **Algorithm 1** $(G, P = (p', \mu', D_T^A, D_T^B, D_S))$, where $p' = p + \frac{\delta}{3}$,
 $\mu' = \mu + \frac{\epsilon}{3}$.

$S_1 \leftarrow$ **Algorithm 2** $(X_{A_1}, X_{B_1}, \mu^*)$.

$X_{A_2}, X_{B_2} \leftarrow$ **Algorithm 1** $(G, P = (p', \mu', D_T^A, D_T^B, D_S))$, where $p' = p - \frac{\delta}{3}$,
 $\mu' = \mu - \frac{\epsilon}{3}$.

$S_2 \leftarrow$ **Algorithm 2** $(X_{A_2}, X_{B_2}, \mu^*)$.

if $S_1 = S_2$ **then**
 return S_1

end

else

return **Null**

end

particular choice of cutoff to separate high- and low-degree agents is somewhat arbitrary. Here, we choose $O(\sqrt[3]{n})$, which is mathematically convenient for defining an upper bound on our error function $\delta(n)$ in Theorem 9.

Given this cutoff, we focus first on the case where there are only low-degree agents, which is sufficient to provide the guiding intuition that we will follow in the next section, when we discuss broadening the setting in various ways. One such extension will involve allowing vertices of arbitrary degree.

We now proceed by making several arguments which form the building blocks of the proof of Theorem 9 that follows.

► **Lemma 10.** *Algorithms 1 and 2 terminate in polynomial time with respect to the number of agents.*

Proof. The key for this property of Algorithm 1 is that contexts are identity-agnostic, so the number of contexts is polynomial in n when the number of types is a constant. Therefore, we are able to enumerate all of the possible contexts in polynomial time. For each of these contexts c , we can compute the relevant probabilities – $\Pr[\text{State} = s|c]$ and $\Pr[c|\text{State} = s]$ for both states s – using Bayes' rule. The rest of the steps in the algorithm are linear in the number of contexts, and therefore also computable in polynomial time. Algorithm 2 is trivially computable in constant-time. ◀

Having shown that they are polynomial-time computable, we now demonstrate the correctness of Algorithms 1 and 2 under idealized conditions.

▷ **Claim.** When agents believe with probability 1 that the actual size of the largest supported revolt in each state will exactly equal its expected size, then Algorithm 1 correctly computes the expected size of the largest revolt in each state. When, further, it is true that the actual size of the largest supported revolt in each state will exactly equal its expected size, then Algorithm 2 identifies the set of states in which revolt of size μ^* is supported with no error.

Proof of Claim. Algorithm 1 first defines the set of candidate states – these are the states in which it is possible, but not necessarily the case, that type- χ agents will feel secure enough to revolt, because there are at least μ α - and χ -type agents. If there are no such states, then

only agents of type α will revolt. If both states are candidates, then all α - and χ -type agents will feel secure enough to revolt: Agents of type α always revolt and agents of type- χ have their μ threshold met in both states, by the definition of a candidate state. As a result, their p threshold is also necessarily met, because the probability of the state being either A or B is $1 \geq p$. The most interesting case is when only A is a candidate state. In this case, only type χ agents who p -believe that the state is A will have their p -threshold met. So, similarly to our example from Section 3.1, we call those agents candidate agents (and refer to their contexts as candidate contexts.) If the number of candidate agents and α -type agents, given that the state is A , is at least μ , then all of those agents feel secure enough to revolt. This is true in any state, because even when the state is B , candidate agents, by definition, p -believe that the state is A .

Algorithm 2 simply compares the size of the expected revolt in each state to μ^* to decide in which states, if any, a revolt of expected size at least μ^* is supported. If the actual size of the revolt in any state is exactly equal to its expectation, as we assume, then Algorithm 2 introduces no error. \triangleleft

Now, the assumption that the actual supported revolt exactly equals the expected size of the supported revolt is, of course, too strict. However, we will be able to show that the actual size of the supported revolt concentrates around its expectation, and consequently, the errors in our algorithms decrease quickly as n grows.

Recall that computing the expected size of a supported revolt involves counting the number of some subset of three kinds of agents: α -type agents, χ -type agents, and candidate agents. For α -type and χ -type agents, it is simple to show that the number of such agents concentrates around its expectation because agent types are assigned independently at random.

► **Lemma 11.** *Given $\epsilon > 0$, the probability that, in a given state, the number of agents of type χ and of type α differ from the expected number of agents of type χ and of type α by a multiplicative factor of ϵ is at most $\delta = \delta(n)$ for some $\delta(n) \in \frac{1}{\exp(\Omega_{\epsilon,P}(\sqrt[3]{n}))}$.*

Proof. Because the agent types are drawn independently at random, we can use a standard Chernoff-Hoeffding bound on the expected number of each type of agents. Consequently, the difference between the actual and expected number of each type of agent decays exponentially with n , and therefore we can choose $\delta(n) \in \frac{1}{\exp(\Omega_{\epsilon,P}(\sqrt[3]{n}))}$ so that the difference is trivially smaller. \blacktriangleleft

Counting the number of candidate agents, as foreshadowed in the Motivating Example of Section 3.1, is somewhat more complicated. Here, we need to consider the contexts of agents – not just their type – and agents' contexts are correlated with the contexts of their neighbors and their neighbors' neighbors in G . We will show, though, that as the graph grows, the effect of this correlation is small. In fact, we will still be able to prove an exponentially-decreasing bound on the difference between the actual and expected number of candidate agents.

► **Lemma 12.** *Given $\epsilon > 0$, the probability that, in a given state, the number of candidate agents is less than the expected number of candidate agents by a multiplicative factor of ϵ is at most $\delta = \delta(n)$ for some $\delta(n) \in \frac{1}{\exp(\Omega_{\epsilon,P}(\sqrt[3]{n}))}$.*

Proof. Let X_C be a random variable that denotes the number of candidate agents. We can write $X_C = \sum_{i=1}^n X_{i,C}$ where $X_{i,C}$ is an indicator variable that indicates whether or not agent i is a candidate.

As we have noted above, an agent's status as a candidate is dependent on her context, and as a result, the variables $X_{i,C}$ are not independent. However – given our assumption that the maximum degree of any vertex in G is $O(\sqrt[3]{n})$ – their dependence is constrained enough that we are able to apply a useful exponential bound involving the *fractional chromatic number* $\chi^*(\Gamma)$ of the constraint graph Γ of the random variables $X_{i,C}$.

(See the appendix section of the full version of the paper for additional details).

For our purposes, it is sufficient to use a trivial bound on the fractional chromatic number of any graph. The fractional chromatic number of a graph is at most the *chromatic number* of the graph, which is at most the maximum degree of any vertex plus one. So in Γ , where edges correspond to dependencies between pairs of random variables $(X_{i,C}, X_{j,C})$, the maximum degree of any vertex – and therefore the fractional chromatic number $\chi^*(\Gamma)$ – is $O_{\epsilon,P}(n^{\frac{2}{3}})$. Applying the bound, then, gives

$$\Pr[|X_C - \mathbb{E}[X_C]| \geq \epsilon n] \leq 2 \exp\left(\frac{-2\epsilon^2 n}{\chi^*(\Gamma)}\right) \in \frac{1}{\exp(\Omega_{\epsilon,P}(\sqrt[3]{n}))}. \quad \blacktriangleleft$$

► **Lemma 13.**

1. Suppose Algorithm 1 is run with the given inputs, but with modified $\mu' = \mu + \frac{\epsilon}{3}$ and $p' = p + \frac{\delta}{3}$, and Algorithm 2 is run with the resulting X_A and X_B and the given μ^* . If the result is Ω , then a revolt of size $\mu^* - \epsilon$ is supported with probability at least $1 - \delta$ under the prior P . If the result is \mathbf{A} , then a revolt of size $\mu^* - \epsilon$ is supported with probability at least $\Pr[\text{State} = A] - \delta$ under the prior P .
2. On the other hand, suppose Algorithm 1 is run with the given inputs, with modified $\mu' = \mu - \frac{\epsilon}{3}$ and $p' = p - \frac{\delta}{3}$, and Algorithm 2 is run with the resulting X_A and X_B and the given μ^* . If the result is \emptyset , then a revolt of size $\mu^* + \epsilon$ is supported with probability at most δ under the prior P . If the result is \mathbf{A} , then a revolt of size $\mu^* + \epsilon$ is supported with probability at most $\Pr[\text{State} = A] + \delta$ under the prior P .

Proof. This proof decomposes into two analogous arguments about (1) and (2), which themselves each contain two analogous arguments. We include only the first one in detail here, below, and claim the rest via analogy.

▷ **Claim.** Suppose that, in instance (1), the result is Ω . Then, the probability that a revolt of size $\mu^* - \epsilon$ is supported in both states is at least $1 - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$.

Proof of Claim. There are three cases for the nature of this revolt.

(I) The revolt consists of all α -type agents.

By Lemma 11, we know that, in either state, the probability that the expected number of α -type agents differs from the expected number of α -type agents by a multiplicative factor of ϵ is $\frac{1}{\exp(\Omega(\sqrt[3]{n}))}$. Therefore, the probability that a revolt of size $\mu^* - \epsilon$ is supported in both states is at least $1 - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$.

(II) The revolt consists of all α - and all χ -type agents.

In this case, the presence of χ -type agents slightly complicates the analysis. Actual χ -type agents have slightly easier thresholds to satisfy (p and μ) than the χ -type agents considered by the algorithms (p' and μ'). Our algorithms determined that, in expectation, agents with p' and μ' thresholds would feel secure enough to revolt. Each actual χ -type agent, then, must also feel secure enough to revolt, not just in expectation: By Lemma 11, the probability that the actual number each of α -type and χ -type agents differs from its respective expectation by a multiplicative factor of $\frac{\epsilon}{6}$ is $\frac{1}{\exp(\Omega(\sqrt[3]{n}))}$. Combining these,

then, χ -type agents believe with probability $1 - \frac{1}{\exp(\Omega(\sqrt[3]{n}))} \geq p$ that at least $\mu' - \frac{\epsilon}{3} = \mu$ agents feel secure enough to revolt, and are thus themselves secure enough to revolt. Applying Lemma 11 (from the perspective of Algorithm 2 this time, not the perspective of χ -type agents) we again incur two error terms of $\frac{\epsilon}{6}$ (one each for the χ - and α -type agents). Thus, we can conclude the probability that a revolt of size $\mu^* - \epsilon$ is supported in both states is at least $1 - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$.

(III) The revolt consists of all α -type agents and all candidate agents.

In the final case, we consider candidate agents instead of all agents of type χ , and proceed exactly as we did in the second case, using Lemma 11 to conclude that the number of α -type agents concentrate and Lemma 12 to conclude that the number of candidate agents concentrates. Once again, the result is that the probability that a revolt of size $\mu^* - \epsilon$ is supported in both states is at least $1 - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$. \triangleleft

\triangleright **Claim.** Suppose that, in instance (1), the result is **A**. Then, the probability that a revolt of size $\mu^* - \epsilon$ is supported in both states is at least $\Pr[\text{State} = A] - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$.

Proof of Claim. The proof of this claim is analogous to the previous proof, where events that are assigned probability at least $1 - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$ are instead assigned probability at least $\Pr[\text{State} = A] - \frac{1}{\exp(\Omega(\sqrt[3]{n}))}$. \triangleleft

The proof of the analogous claims for (2) are themselves analogous to the above cases for (1).

Finally, we note here that when χ -type agents decide whether or not their thresholds are satisfied, they are not solely relying on their estimates of the fractions of different kinds of agents, as we describe above. They have additional knowledge, since they see the realized types of the agents in their context. However, for small graphs, $\delta(n)$ can be chosen to account for this. As the graph grows, since each agent has at most $O(\sqrt[3]{n})$ neighbors, the consequences of observing the types of a few adjacent agents is negligible after conditioning on the state. As a result, for large enough n , after the agent reasons about the state, the probabilistic effect of the agent viewing the types in her context is subsumed by the ϵ error term. Furthermore, the appropriate choice of $\delta(n)$ also accounts for the $\frac{1}{\exp(\Omega(\sqrt[3]{n}))}$ terms present in each of the claims stated above, for each n . \blacktriangleleft

Proof of Theorem 9. It follows from Lemma 10 that Algorithm 3 terminates in polynomial time with respect to the inputs G and P . We also note that Algorithm 3 terminates in polynomial time with respect to $\frac{1}{\epsilon}$ and $\frac{1}{\delta}$.

Further, it follows from Lemma 13 that when Algorithm 3 outputs a case, that case is always true. It only remains to show that when exactly one case in the statement of PROMISE REVOLT is true, then Algorithm 3 necessarily outputs that case:

Here, the key is our use of three different potential revolt sizes ($\mu^* - \epsilon$, μ^* , and $\mu^* + \epsilon$) and 3 different priors (P^- , P , and P^+) in defining the three cases of PROMISE REVOLT. By promising that exactly one of those three cases is true, we guarantee that the values of μ and p are sufficiently far – distance at least ϵ for μ and distance at least δ for p – from the crucial decision thresholds in Algorithm 1 (e.g. the value $e_B(\chi \cup \alpha)$, which is used to determine whether or not B is a candidate state)³. This ensures that both calls to Algorithm 1 will return the same values.

³ Whether or not a decision threshold is in the set of *crucial* decision thresholds – the thresholds from which our promise guarantees p and μ are sufficiently far – depends on μ^* . In addition to $e_B(\chi \cup \alpha)$,

We can illustrate this counterfactually:

Let $\mu = e_B(\chi \cup \alpha) + \frac{\epsilon}{2}$ with $e_B(\chi \cup \alpha) > \mu^* > e_B(C_C \cup \alpha)$ and let $e_A(C_C \cup \alpha) > \mu + \frac{\epsilon}{2}$.

Then, A and B would both be candidate states under the prior P^- . As a result, Algorithm 1 (run with prior P^-) would return $X_A = e_A(\chi \cup \alpha)$ and $X_B = e_B(\chi \cup \alpha)$. Note that $X_A > X_B > \mu^*$, so Algorithm 2 would return Ω when run with inputs X_A , X_B , and μ^* . The same analysis holds for P^- with μ and p incremented by $\frac{\epsilon}{3}$ and $\frac{\delta}{3}$, respectively, so applying Lemma 13 implies that the case Ω is true.

On the other hand, only A would be a candidate state under the prior P . Algorithm 1 run with the prior P would return $X_A = e_A(C_C \cup \alpha)$ and $X_B = e_B(C_C \cup \alpha)$. Run with these inputs (and μ^*), Algorithm 2 would return \mathbf{A} . The same analysis holds for P with μ and p incremented *and* decremented by $\frac{\epsilon}{3}$ and $\frac{\delta}{3}$, which by Lemma 13 implies that the case \mathbf{A} is true. We can conclude that these values of μ and μ^* must be excluded by our promise for any input to PROMISE REVOLT for which our assumed constraints hold.

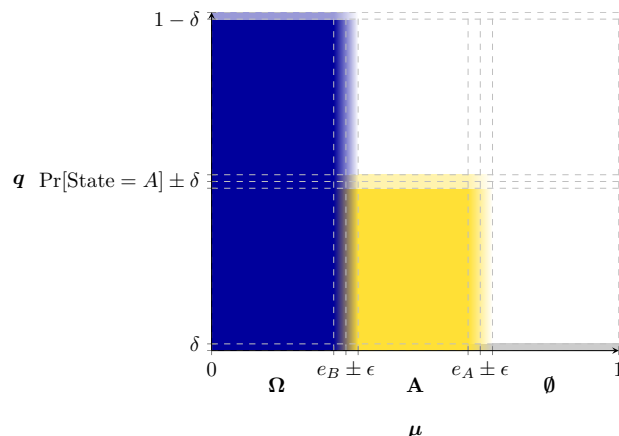
An argument similar to this counterfactual argument suffices to exclude any value of p or μ that is insufficiently far from a crucial decision threshold in Algorithm 1 (along with associated constraints on μ^*) present in an instance of PROMISE REVOLT. ◀

Now we can discuss the information that we gain from solving PROMISE REVOLT. In doing so, it is useful to refer to Figure 1, which contains the following illustration: Given μ^* , we identify all values of q for which a revolt of size μ^* is supported with probability at least q (dark blue and yellow regions), all values of q for which revolt of size μ^* may be supported with probability q (light blue, yellow, and grey regions), and all values of q for which a revolt of size μ^* is supported with probability strictly less than q (white regions). The blurry regions between distinct colors represent the inputs to PROMISE REVOLT for which two of the cases would overlap (and are therefore excluded from the set of inputs by the “promise”).

If we could perfectly solve REVOLT, then we would be able to perfectly define the boundaries – there would be no blurry regions between distinct colors, as in Figure 1. The shape of the resulting figure would characterize the equilibria of the network revolt game in the following sense: The blue column (corresponding to values of μ^* for which only case Ω of PROMISE REVOLT is true) would show the sizes of revolts that are supported in equilibrium with high probability regardless of the state, the next yellow and white column (corresponding to values of μ^* for which only case \mathbf{A} is true) would show the sizes of revolts that are supported in equilibrium with high probability given that the state is A , and the last grey and white column (corresponding to values of μ^* for which case \emptyset is true) would show the sizes of revolts that, with high probability, are not supported in equilibrium (or, equivalently, the sizes of revolts that are supported in equilibrium with low probability).

Although we cannot perfectly solve REVOLT, as shown in Figure 1, solving PROMISE REVOLT with Algorithm 3 allows us to approximate this shape. By choosing ϵ to be very small, we can make the boundary cases only relevant for very small subsets of possible values of μ^* and values of μ in the prior P . And, as we have shown via the proof of Theorem 9, as n grows, $\delta(n)$ quickly becomes very small. Consequently, the range of possibly forbidden values of p in the prior P also quickly becomes small and the probabilities for each class of distinct equilibria described above (when they exist), converge to 1, $\Pr[\text{State} = A]$, and 0, respectively.

it can also include $e_A(C_C \cup \alpha)$ and the value of p below which $e_A(C_C \cup \alpha) \geq \mu$ holds for candidate contexts and above which $e_A(C_C \cup \alpha) < \mu$.



■ **Figure 1** Illustrating what we learn from solving PROMISE REVOLT in an archetypal case. Here, e_B and e_A denote the expected size of the revolt supported in states B and A , respectively.

Lastly, we note that an interesting and surprising corollary to our analysis above is that Algorithms 1, 2, and 3 never use any information about the graph beyond the *degree sequence* of G – an anonymous list that records the degree of each vertex in the graph. That is, our algorithms require a list of the degrees of the agents, but do not require that any degree value be labelled with the identity of any agent, nor do they require any further information about the set of edges in the graph. Rather, the results of the algorithms are valid for any graph that is consistent with the provided degree sequence, because the concentration of the fraction of candidate agents and agents of each type supersedes the additional structure imposed by any concrete edge set consistent with the degree sequence.

We record this fact with the following proposition:

► **Proposition 14.** *To compute PROMISE REVOLT in polynomial time, we only require the degree sequence of the graph G ; the graph itself is not required.*

5 Broadening the Setting

The core intuition from our analysis above actually applies more generally than just to the setting of Section 4.1. In the full version of the paper, we explore settings where the value of interest is the size of the smallest supported revolt in each state (as opposed to the largest) and where there is an arbitrary number of states.

Here, we only discuss the most fundamental of the more general settings – extending our result to general graphs. In doing so, we focus on changes to Algorithm 1; the subsequent changes required to adapt Algorithms 2 and 3 are straightforward.

5.1 General Graphs

Suppose that we impose no restriction on the degree of the vertices in G in the statement of Theorem 9. It turns out that our results still hold. The existence of high-degree χ -type vertices (i.e. those that would not exist when we assume an $O(\sqrt[3]{n})$ bound on the maximum degree of any vertex), complicates the analysis, but only in the case where A is the only

candidate state. When A and B are both candidate states, we only need to calculate the expected number of all χ -type agents, regardless of their degree. When there are no candidate states, χ -type vertices are irrelevant.

When A is the only candidate state, however, there is a key difference: The presence of high-degree agent contexts in the set of candidate contexts would affect our earlier analysis using the concentration bound for the sum of random variables with limited dependence based on the fractional chromatic number of the constraint graph of the random variables (in the proof of Lemma 12), because the contexts of high-degree agents are correlated with many other agents' contexts.

Because of this, though, high-degree agents (with degree at least $c \cdot \sqrt[3]{n}$ for sufficiently large c) have a unique perspective on the graph; their contexts contain a significant amount of information that they can use in determining the state. More concretely, suppose that agent i is a high-degree agent and let $N(i)$ be the set of neighbors of i in G .

Let $X = \sum_{j \in N(i)} X_j$, where X_j is an indicator random variable that is 1 when agent j is of type α or type χ and 0 otherwise.

Now, for any $\epsilon_0 > 0$, applying a standard Chernoff-Hoeffding bound, we have:

$$\Pr [|X - \mathbb{E}[X]| \geq \epsilon_0 (c \cdot \sqrt[3]{n})] \leq 2 \exp\left(\frac{-2(\epsilon_0 c)^2 n^{\frac{2}{3}}}{n^{\frac{1}{3}}}\right) \in \frac{1}{\exp(\Omega_{\epsilon, P}(\sqrt[3]{n}))}. \quad (1)$$

In particular, if we choose ϵ_0 such that $|\mathbb{E}[X|\text{State} = A] - \mathbb{E}[X|\text{State} = B]| > 2\epsilon_0 (c \cdot \sqrt[3]{n})$, then agent i (and by the same argument, any high-degree agent) can correctly determine the state with enough accuracy that their error can be absorbed into the δ error term with an appropriate choice of $\delta(n)$. As a result, when the state is A , all high-degree χ -type agents will behave like candidate agents (recall that A is a candidate state).

However, for computational purposes, we will not lump them in with the low-degree candidate agents, because Lemma 12 can only apply to low-degree candidates. Instead, in the first **if** statement after the condition that A is the only candidate state (the line in Algorithm 1 that reads “**if** $e_A(C_C \cup \alpha) \geq \mu$ **then**”), we calculate $e_A(C_C \cup \alpha \cup H_\chi)$ instead of just $e_A(C_C \cup \alpha)$, where H_χ refers to the set of high-degree χ -type vertices. If this is at least μ , then when we calculate X_A , we include H_χ . However, when we calculate X_B , we only include H_χ if $e_B(C_C \cup \alpha \cup H_\chi) \geq \mu$, since those high-degree agents will not p -believe that the state is A regardless of the state the way that (low-degree) candidate agents will. If $e_B(C_C \cup \alpha \cup H_\chi) < \mu$, we calculate $X_B = e_B(C_C \cup \alpha)$, as described in Algorithm 1.

Next, we must show it is possible for the agents (and the algorithm) to accurately calculate the expected number of high-degree agents of type χ and know that the actual number of such agents sufficiently concentrates around that expected number. Here, we can rely on the fact that the agents and the algorithm know the degree sequence of the graph. There is some subtlety involved: If the number of high-degree agents is small – less than ϵn – then we cannot provide a very useful concentration bound for the number of high-degree agents of type χ . However, we can essentially ignore the high-degree agents in this case and absorb the error we incur by ignoring them into our choice of $\delta(n)$.

On the other hand, if there are at least ϵn high-degree agents, we can again use a standard Chernoff-Hoeffding bound to conclude that the actual number of high-degree χ -type agents will, with high probability, be close to the expected number of such agents. The error here will be smaller than the error from Lemma 12 and so can be absorbed there.

Finally, there is one additional subtlety we must address for high-degree agents. While it is true that, as previously mentioned, their contexts contain a significant amount of information that they can use in determining the state, their contexts actually contain more than just

information about the state – they contain information about the actual realization of types for a significant number of agents. This point is the primary conceptual reason to make the distinction between high-degree and low-degree agents in the first place. Consequently, we need to show that high-degree agents tend to behave as if they only knew the state, when in fact it is possible that the number of agents of a certain type in their context differs greatly from its expectation and as a result they have additional information to use beyond the state of the world. For this, we again appeal to our previous argument that resulted in the bound expressed by the inequality (1), above.

That argument demonstrates that it is highly improbable for the information in a high-degree agent’s context to contradict what their belief would be solely given knowledge of the state. For example, if a high-degree agent uses her context to determine that the state is A , with high probability it will not also be the case that the context that she uses to make that determination has (far) fewer agents of any type than would be expected given that the state is A . Consequently, a high-degree agent will tend to act as if she is just calculating expectations and acting off of them (like a low-degree agent would), even though in actuality she has quite a bit of additional information.

6 Discussion

We proposed that the notion of common belief could be relaxed to a notion of factional belief in order to be more suited to social network settings and gave a natural definition of factional belief, drawing heavily from previous work on common belief. We then applied this definition theoretically and experimentally in a setting inspired by prior work about common knowledge and revolt games on networks to show how this definition moves beyond the limitations of previous work by being applicable in general graphs.

6.1 Open Questions and Future Work

The most clear direction for future work is to continue to apply factional belief in new settings and use it as a tool to understand strategic coordination and cooperation on networks. In particular, as mentioned in the introduction, the work of Stephen Morris provides many examples of applying common belief as a tool for gaining insight into a diverse range of settings. We believe that applying factional belief can yield similar results in social network settings, and that this paper represents an initial step in that process.

Additionally, though, there are some more subtle technical open questions regarding our definition of factional belief that are also worth exploring in future work:

1. Considering our definition of factional belief from the perspective of an infinite hierarchy of reasoning, akin to the initial definition for common knowledge that we provided (3), with regard to the μ fraction of agents at each step in the hierarchy, it is not required by our definition that this be the same μ fraction of agents at each step in the hierarchy. However, it is not clear whether or not this should always be the case. Is it possible to describe an event that is common (p, μ) -belief, but for which the infinite hierarchy refers to a different μ fraction of agents at some step? If so, what are the consequences of this for strategic coordination and cooperation?
2. It is not too difficult to modify the setting described in Section 4 to create a model where the underlying event that supports revolt does not necessarily neatly reduce to a single event that is common (p, μ) -belief. For example, if there are multiple types of conditionally-revolting agents with different p and μ thresholds, the event supporting revolt is more like a μ fraction of agents believe with sufficient probability that their

thresholds are satisfied. This event encompasses common p -beliefs among certain agents of the same type regarding their μ thresholds, but not precisely common (p, μ) -beliefs. Is there a more general definition of factional belief that allows for the existence of different thresholds for different types of agents to still be encompassed in a single event that is a factional belief among all of the agents who feel satisfied enough to revolt?

We believe that answering these questions could yield further insight into factional belief that could inform its application in other settings.

References

- 1 Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. URL: <http://www.jstor.org/stable/2958591>.
- 2 Cristina Bicchieri. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press, 2005. doi:10.1017/CB09780511616037.
- 3 Michael Suk-Young Chwe. Structure and strategy in collective action. *American Journal of Sociology*, 105(1):128–156, 1999. doi:10.1086/210269.
- 4 Michael Suk-Young Chwe. Communication and Coordination in Social Networks. *The Review of Economic Studies*, 67(1):1–16, January 2000. doi:10.1111/1467-937X.00118.
- 5 Michael Suk-Young Chwe. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton University Press, 41 William Street, Princeton, New Jersey 08540, 2013.
- 6 Benjamin Golub and Stephen Morris. Expectations, networks, and conventions, 2017. doi:10.2139/ssrn.2979086.
- 7 Benjamin Golub and Stephen Morris. Higher-order expectations, 2017. doi:10.2139/ssrn.2979089.
- 8 Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, 1989. doi:10.1016/0899-8256(89)90017-1.
- 9 Stephen Morris. Approximate common knowledge revisited. *International Journal of Game Theory*, 28(3):385–408, August 1999. doi:10.1007/s001820050116.
- 10 Stephen Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, January 2000. doi:10.1111/1467-937X.00121.
- 11 Stephen Morris. Coordination, Communication, and Common Knowledge: A Retrospective on the Electronic-mail Game. *Oxford Review of Economic Policy*, 18(4):433–445, December 2002. doi:10.1093/oxrep/18.4.433.
- 12 Stephen Morris. Coordination, timing and common knowledge. *Research in Economics*, 68(4):306–314, 2014. doi:10.1016/j.rie.2014.04.004.
- 13 Stephen Morris and Hyun Song Shin. Approximate common knowledge and co-ordination: Recent lessons from game theory. *Journal of Logic, Language and Information*, 6(2):171–190, April 1997. doi:10.1023/A:1008270519000.
- 14 Stephen Morris, Hyun Song Shin, and Muhamet Yildiz. Common belief foundations of global games. *Journal of Economic Theory*, 163:826–848, 2016. doi:10.1016/j.jet.2016.03.007.
- 15 Ariel Rubinstein. The electronic mail game: Strategic behavior under “almost common knowledge”. *The American Economic Review*, 79(3):385–391, 1989. URL: <http://www.jstor.org/stable/1806851>.