

EXPECTATIONS, NETWORKS, AND CONVENTIONS

BENJAMIN GOLUB AND STEPHEN MORRIS

ABSTRACT. In coordination games of incomplete information, and in speculative over-the-counter financial markets, agents care about what their counterparties, on average, expect *their* counterparties to think, and so forth. *Higher-order average expectations* of a group of agents are formed by starting with their expectations of a random variable, and then repeatedly considering each agent's expectation of the average of others' expectations. The averaging is done according to an arbitrary network describing whom each agent cares about. We offer a unified analysis of these objects and their limits. Special cases include beauty contests under a common prior and the network centrality analysis of network games. Two leading applications illustrate the results. First, we describe a contagion in which asset values are driven to extremes if almost everyone expects his average counterparty to be slightly more optimistic than himself. Second, we study the *tyranny of the uninformed*, in which coordination is determined by the prior beliefs of the agent with the least precise private information, even when there is nearly common certainty of the ex post optimal action to play.

1. INTRODUCTION

Consider a situation in which each agent has strong incentives to match the behavior of others. An outcome that agents coordinate on in such a setting has been called a *convention* in philosophy and economics (see Lewis (1969), Young (1996), and Shin and Williamson (1996)). In deciding how to coordinate, agents will take into account their beliefs about (i) the state of the world, which determines the best action; and (ii) one another's beliefs. Agents may differ from one another, and hence have an incentive to choose differently, for three reasons: first, because they are asymmetrically informed; second, because they interpret the same information differently—that is, they have different priors; and third, because they differ in whom they want to coordinate with. Which conventions emerge in such an environment will depend on

Golub: Department of Economics, Harvard University, Cambridge, U.S.A., bgolub@fas.harvard.edu, website: fas.harvard.edu/~bgolub. Morris: Department of Economics, Princeton University, Princeton, U.S.A., smorris@princeton.edu, website: princeton.edu/~smorris. We are grateful for conversations with Nageeb Ali, Dirk Bergemann, Larry Blume, Ben Brooks, Jason Hartline, Tibor Heumann, Matthew O. Jackson, Bobby Kleinberg, Eric Maskin, Dov Samet, and Omer Tamuz; for comments from Selman Erol, Alireza Tahbaz-Salehi and Muhamet Yildiz, who served as discussants; as well as many questions and comments from seminar and conference participants. We especially thank Ryota Iijima for several inspiring conversations early in this project. Cristian Gradinaru and Georgia Martin provided excellent assistance in preparing the manuscript.

the information asymmetries, the heterogeneous prior beliefs, and the network describing the coordination motives of agents. Our purpose is to characterize this dependence.

We informally describe a simple model of this environment: a coordination game with linear best responses. Nature draws an *external state* θ , and each agent i chooses a real-valued *action* a^i based on some private information. This occurs simultaneously. Agents' payoffs capture two motivations: First, they seek to coordinate with a *basic random variable* $y(\theta)$ —a random variable, common to everyone, that depends only on the external state; second, they seek to take actions that are close to the actions that others take—the various a^j for $j \neq i$. The disutilities they experience are proportional to the squares of the differences between a^i and these various targets; this feature induces best responses linear in an agent's expectations of the various targets. A *network* of weights captures the coordination concerns of the agents—that is, *which* others each agent cares most about coordinating with.

If just the coordination motive were present, with no desire to match the basic random variable, there would be a continuum of equilibria. Indeed, for any action, there would be an equilibrium with everyone choosing that action. The choice of action would be an arbitrary convention. We will mainly be interested in the case where agents put a positive, though small, weight on the accuracy motive—matching the basic random variable—and nearly all the weight on the coordination motive—choosing actions close to others' actions. In this case, there will be a unique equilibrium. Because the weight on others' actions is very high, it can be shown that agents essentially choose a common action. We call this action—the common action played in the limit—the *convention*. If there were common knowledge of the external state θ , the convention would be equal to $y(\theta)$, the value everyone seeks to match. But we are interested in characterizing the convention when there is incomplete information about the state.

The convention will depend on higher-order expectations of the agents. Suppose Ann cares mainly about coordinating with Bob, who cares mainly about coordinating with Charlie. (Recall that the agents all care a little about matching their own expectations of the external variable.) Then Ann's expectation of Bob's expectation of Charlie's expectation of the external variable becomes relevant for Ann's decision. In this scenario, each agent is seeking to coordinate with only one other, but in the general model each seeks to match a (weighted) average of the actions of several others. By an elaboration of the above reasoning about Ann, Bob, and Charlie, *higher-order average expectations* become relevant: Each agent cares about the average of his neighbors' expectations of the average of *their* neighbors' expectations of the external variable, and so on. Thus our analysis of coordination games leads naturally to a study of higher-order

average expectations. We will define the *consensus expectation* to be (essentially) the limit of such higher-order average expectations as the order becomes large. The consensus expectation will equal the convention that obtains in the linear best-response game described above, in the limit as agents' coordination concerns dominate. We will focus on this limit, though many of the techniques we will develop can be extended to study the case where coordination motives are not dominant.

We will report three kinds of substantive results about consensus expectations. Two technical observations are key to establishing these results: (i) agents' higher-order expectations can be represented using a Markov process on the union of agents' signals; (ii) consensus expectations are determined by the stationary distribution of this Markov process. We now present the substantive results, and we discuss the techniques in more detail at the end of the Introduction.

Unifying and Generalizing Network and Asymmetric Information Results. The first results unify and generalize facts known in the literatures on network games and on asymmetric information:

(a) Suppose agents have the same information but may have heterogeneous beliefs about θ —that is, different priors, which are commonly known. Then the consensus expectation is simply a weighted average of agents' heterogeneous prior expectations of the external random variable. The weight on an agent's expectation is his *eigenvector centrality* in the network. This corresponds to the seminal result of Ballester, Calvó-Armengol, and Zenou (2006) on equilibrium actions in certain network games being weighted averages of individuals' ideal points, with someone's weight determined by the extent to which others want to directly and indirectly coordinate with him. The appearance of network centrality here—a statistic of individuals defined from the matrix of coordination weights—is a consequence of the matrix algebra that naturally appears when studying higher-order average expectations.¹

(b) If there is *asymmetric* information, but agents have common prior beliefs, then the consensus expectation is equal to the (common prior) *ex ante* expectation of the external state. Thus consensus expectations are independent of the network structure, and also independent of all features of the information structure except the common prior. This result turns out to be a corollary of the result of Samet (1998a).

¹For recent surveys of economic applications related to network centrality, see Jackson (2008, Section 2.2.4), Acemoglu, Ozdaglar, and Tahbaz-Salehi (2016b), Zenou (2016), and Golub and Sadler (2016).

(c) Embedding both (a) and (b), if agents have both asymmetric information and heterogeneous prior beliefs but a *common prior on signals*,² then the consensus expectation is equal to a weighted average of agents' different ex ante expectations of the basic random variable. Just as in (a), the weight on an agent's expectation is his *eigenvector centrality* in the network.³ This goes beyond existing work on network games with incomplete information due to Calvó-Armengol et al. (2009), de Martí and Zenou (2015), Bergemann et al. (2015b) and Blume et al. (2015): we will discuss these connections in Section 6 when we have introduced the model and key results.

Contagion of Optimism. Our second category of results studies *second-order optimism*. We assume that each agent, given any signal, assesses his average counterparty as more optimistic than himself about the value of the basic random variable, unless he the agent has a first-order expectation that is already very high (close to the highest one induced by any signal). Agents whose expectations are high may be somewhat pessimistic: they may assess the average counterparty as less optimistic than themselves.

We study when arbitrarily slight second-order optimism leads consensus expectations to be very high—near highest possible expectation of y . The proof is via a reduction to a Markov chain inequality. The key subtlety in the analysis is: how much pessimism can be allowed without destroying the contagion of optimism? We give a bound that answers this question, and describe a sense in which this bound is tight (Section 7.4.2). Recent work of Han and Kyle (2016) also discusses a different contagion of optimism in a CARA-normal rational expectations model. We discuss connections with related models in Section 7.4.1.

Tyranny of the Uninformed. Third, we consider a setting where agents start with heterogeneous priors about the external state but share a *common interpretation of signals*.⁴ That is, agents observe signals of the external state. They agree on the probability of any particular signal of a given agent conditional on any external state.⁵ However, their priors over external states may differ, and thus their interim beliefs may not be compatible with a common prior. Given common interpretation of signals, it makes sense to define notions of more and less precisely

²That is, a common prior on how the signal random variables are distributed.

³This decomposition separates, in a suitable sense, the effects of the network and the beliefs. A companion paper, Golub and Morris (2017), gives the necessary and sufficient condition on the information structure for this sort of weighted average decomposition to be possible.

⁴An *interpretation* of a signal random variable is its conditional distribution given the state, in line with the terminology of Kandel and Pearson (1995) and Acemoglu, Chernozhukov, and Yildiz (2016a).

⁵In the environment we study for this application, the signals are conditionally independent given the state, so that signals are correlated only through the state.

informed agents, because the distributions of signals given the external state (and thus the levels of noise in them) are common knowledge.

We show that, in a suitable sense, the consensus expectation approximates the ex ante expectation of the agent whose private information is *least* precise. This is true even if all agents have very precise private signals about the state, as long as the least informed has signals that are less precise than others. The quantitative details of how to define the most and least informed when everyone has different private information are subtle, and rely on a Markov chain connection that we discuss next.

The Interaction Structure and Markov Formalism. The techniques underlying the results discussed above are based on a Markov matrix description of higher-order average expectations. While we defer most of the details until Section 4, when we have more notation, the basic idea is simple. We define a Markov process whose state space is the *union* of all agents' signals. Transition probabilities between any two states combine both the network weights and the subjective probabilities of the agents. In particular, the transition probability from a signal t^i of agent i to a signal t^j of agent j is defined as the product of (i) the network weight that i places on j and (ii) the subjective probability that agent i , given signal t^i , places on t^j . We call the transition matrix of this Markov process the *interaction structure*, and it is our key technical device. This formalism treats beliefs and network weights entirely symmetrically. This symmetric treatment enables the analysis to be reduced to Markov chain results, which provide both a tool and novel insights.⁶ Other work on network games with incomplete information - Calvó-Armengol et al. (2009), de Martí and Zenou (2015), Bergemann et al. (2015b) and Blume et al. (2015) - does not use this general device and must develop more tailored techniques.

The essence of our approach is that the iteration of the Markov matrix associated with the interaction structure enables a brief, explicit description of higher-order average expectations: The n^{th} -order average expectations can be obtained by suitably combining the n -step transition probabilities of the Markov process with the first-order expectations associated to various signals.

To study consensus expectations, we consider the limit as n grows large. Under suitable conditions, in this limit the Markov transition probability to any state—regardless of where the process starts—becomes the stationary probability of that state. This can be used to show that *the stationary distribution of the Markov process determines the consensus expectation*. Indeed,

⁶The symmetric treatment follows Morris (1997). As discussed in detail in Section A.3, our approach echoes Samet (1998a) in using a Markov process to represent incomplete information, although our Markov process is actually a different one in significant ways.

the consensus expectation turns out to be a weighted average of first-order expectations given various signals t^i . The weight on a signal t^i is its weight in the stationary distribution of the Markov process.

Thus, our results on the consensus expectation are proved by studying the stationary distribution of the Markov process and deriving properties of it from more primitive assumptions about the environment. For example, in the analysis of the contagion of optimism, the essential idea is that when second-order optimism holds, probability mass in the Markov process flows on average to signals associated with higher first-order expectations of the basic random variable. It can be shown that, as a consequence, states with high first-order expectations have a larger share of the stationary probability. By our description of the consensus expectation as a weighted average of first-order expectations, with weights given by the stationary probabilities, it follows that the consensus expectation is high.

Other results rely on different reasoning. The most technically involved arguments are the ones associated with the tyranny of the uninformed. These arguments rely on perturbation bounds for Markov chains, which are used to show that the priors of highly informed agents cannot play a substantial role in the stationary distribution that determines the consensus expectation. Overall, our main methodological claim—illustrated by the various applications—is that the structure of higher-order expectations is illuminated by the Markov formalism.

The remainder of the paper is organized as follows. Section 2 presents the environment, defines higher-order average expectations and consensus expectations, and illustrates them with some simple examples. Section 3 motivates higher-order average expectations and consensus expectations by discussing a coordination game and an asset market where they are relevant. Section 4 presents our key technical device, the interaction structure, and the correspondence between higher-order average expectations and statistics of a Markov process. Section 5 relates the interaction structure to the underlying network. Section 6 relates consensus expectations to agents' priors. Together these results unify and extend the known network games and incomplete-information results. Section 7 focuses on higher-order optimism, while Section 8 reports our results on the tyranny of the uninformed. Section A is a discussion of relations to the literature, subtleties, and extensions.

2. MODEL

2.1. The Information Structure.

2.1.1. *States, Signals, and Expectations.* There is a finite set Θ of *states of the world*. There is a finite set N of agents. Associated to each agent $i \in N$ is a finite set T^i of *signals* (i.e., possible signal realizations). and these sets of signals are disjoint across agents. Let $T = \prod_{i \in N} T^i$ be the product of all the signal spaces, with a typical element being a tuple $t = (t^i)_{i \in N}$; let $T^{-i} = \prod_{j \in N \setminus \{i\}} T^j$ be the product of the signal spaces of all the others, viewed from i 's perspective. An agent's signal fully determines all the information he has, including the information he has about others' signals. Let $\Omega = \Theta \times T$ be the set of all *realizations*.

For each i and signal t^i , there is a *belief* $\pi^i(\cdot | t^i) \in \Delta(\Theta \times T^{-i})$ —that is, a probability distribution over $\Theta \times T^{-i}$. This is the *interim* or *conditional* belief that agent i has when he gets signal t^i . We introduce some notation to refer to marginal distributions: $\pi^i(t^j | t^i)$ denotes the probability this belief assigns to agent j 's signal being t^j .⁷ For states $\theta \in \Theta$, the notation $\pi^i(\theta | t^i)$ has an analogous definition. We refer to $\pi = (\pi^i(\cdot | t^i))_{i \in N, t^i \in T^i}$ as the *information structure*.⁸ In situations where only interim beliefs matter, we will use the language of *types*. That is, we will identify each signal with a corresponding (belief) type of the agent. If signal t^i induces a certain belief over $\Theta \times T^{-i}$, we will say that type t^i (of agent i) has that belief. We will call T the *type space*. On the other hand, when we wish to emphasize the ex ante stage and the literal process of drawing signals, we will use the language of signals.

A random variable measurable with respect to i 's information is a function $x^i: T^i \rightarrow \mathbb{R}$, i.e., an element of \mathbb{R}^{T^i} (this set being defined as the set of functions from signals in T^i to real numbers). Given a random variable $z: \Omega \rightarrow \mathbb{R}$, let $E^i z \in \mathbb{R}^{T^i}$ give i 's conditional expectation of z . It is defined by

$$(1) \quad (E^i z)(t^i) = \sum_{(\theta, t^{-i}) \in \Theta \times T^{-i}} \pi^i(\theta, t^{-i} | t^i) z(\theta, t^i, t^{-i}).$$

The summation runs over all (θ, t^{-i}) , and states are weighted using the probabilities assigned by the interim belief $\pi^i(\cdot | t^i)$. We will often abuse notation, as we have done here, by dropping parentheses in referring to elements of Ω in the arguments of beliefs and random variables.

2.1.2. *Priors.* The information structure was defined above in terms of agents' interim beliefs, i.e., their beliefs about external states and others' signals conditional on their own signals. This interim information is enough to define higher-order average expectations and to state our main

⁷That is, $\pi^i(\{(\hat{\theta}, \hat{t}^{-i}) : \hat{\theta} \in \Theta, \hat{t}^j = t^j\})$.

⁸As always, uncertainty about how signals are generated can be built into this description of an information structure. Thus, following Harsanyi (1968), the information structure itself is taken to be common knowledge. For more on this see Aumann (1976, p. 1237) and Brandenburger and Dekel (1993).

results. However, we are interested in the ex ante interpretation of our results: There is a prior stage before agents observe their own signals, and thus where they face uncertainty as to what signals they will observe.

We write $(\mu^i)_{i \in N}$ for agents' ex ante beliefs, with $\mu^i \in \Delta(T^i)$. Combined with conditional beliefs $\pi^i(\cdot | t^i) \in \Delta(\Theta \times T^{-i})$, there is a prior $\mathbf{P}^{\mu^i} \in \Delta(\Omega)$ on the entire space of realizations, assigning to any $(\theta, t) \in \Omega$ a probability⁹

$$(2) \quad \mathbf{P}^{\mu^i}(\theta, t) = \sum_{t^i \in T^i} \mu^i(t^i) \pi^i(\theta, t^{-i} | t^i).$$

If one started from agent i 's prior $\mathbf{P}^{\mu^i} \in \Delta(\Omega)$, one would define conditional beliefs $\pi^i(\cdot | t^i) \in \Delta(\Theta \times T^{-i})$ by updating according to Bayes' rule.

The probability measure \mathbf{P}^{μ^i} gives rise to an ex ante expectation operator,

$$\mathbf{E}^{\mu^i} z = \sum_{\omega \in \Omega} \mathbf{P}^{\mu^i}(\omega) z(\omega) = \sum_{t^i \in T^i} \mu^i(t^i) E^i z.$$

To emphasize when an ex ante perspective is being taken, we adopt the convention that ex ante probabilities, expectations, etc. are in bold.

We will later be interested in what an agent's ex ante beliefs would be if we had fixed his conditional beliefs $\pi^i(\cdot | t^i) \in \Delta(\Theta \times T^{-i})$ but endowed him with alternative prior beliefs. Priors for i other than the true priors μ^i are denoted by $\lambda^i \in \Delta(T^i)$, and we use λ^i in place of μ^i in the notations introduced above.

2.2. The Network. For each pair of agents, i and j , there is a number $\gamma^{ij} \in [0, 1]$, where $\sum_{j \in N} \gamma^{ij} = 1$, with the interpretation that agent i assigns "weight" γ^{ij} to agent j . A matrix Γ , whose rows and columns are indexed by N and whose entries are γ^{ij} , records these weights and is called the *network*. The fact that the weights of any agent add up to 1 corresponds to this matrix being row-stochastic.

The network is to be contrasted with the *information structure* encoded in the interim beliefs $\pi^i(\cdot | t^i)$. One interpretation of the network weight γ^{ij} , which will be used when we discuss coordination games, is that it measures how much agent i cares about the action of j . We define N_i , the *neighborhood* of i , to be the set of j such that $\gamma^{ij} > 0$, and the elements of N_i are i 's *neighbors*. Note that j may be a neighbor of i without i being a neighbor of j .

We now define an important set of statistics arising from the network.

⁹Note that the probability under \mathbf{P}^{μ^i} of any subset of Ω can be written as a sum of probabilities defined in equation (2), and a similar statement holds for the interim probabilities $\pi^i(\cdot | t^i)$.

Definition 1. The *eigenvector centrality weights* of the agents are the entries of the unique row vector $e \in \Delta(N)$ satisfying $e\Gamma = e$ —i.e., for each i ,

$$e^i = \sum_{j \neq i} e^j \gamma^{ji}.$$

Assuming that Γ is irreducible, the Perron–Frobenius Theorem states that the eigenvector centrality weights are well-defined—that there is indeed a unique such vector e . Moreover, the theorem says that all the eigenvector centrality weights are positive.

2.3. Higher-Order Average Expectations. We now define higher-order average expectations. A *basic* random variable is a random variable measurable with respect to the external states, i.e., a function $y : \Theta \rightarrow \mathbb{R}$, or an element of \mathbb{R}^Θ . Consider a random variable $y \in \mathbb{R}^\Theta$ and define¹⁰

$$(3) \quad x^i(1; y) = E^i y$$

for every $i \in N$. This is i 's *first-order expectation*, given i 's own signal, of y .

We can now define the key objects we will focus on: the iterated expectations, or *higher-order average expectations*. For $n \geq 2$, given $(x^i(n))_{i \in N}$, define

$$(4) \quad x^i(n+1; y, \Gamma) = \sum_{j \in N} \gamma^{ij} E^i x^j(n; y, \Gamma).$$

This is i 's subjective expectation of the average of the random variables corresponding to the previous iteration of the process; the average is taken with respect to the network weights.

When we do not wish to emphasize the dependence on y and Γ , or when they are clear from context, we omit these arguments.

Note that equation (4), despite the presence of an iteration, is defined in a static environment: Higher-order average expectations do not correspond to dynamic updating over time, but rather to a hierarchy of beliefs when agents are simultaneously given different information. For this reason, these will figure in the solution of a static game (see Section 3.1.1, and a contrast with dynamics in Section 9).

2.4. Examples.

Example 1. If we have $\gamma^{ij} = 1/|N|$ for all i, j , then every agent is weighting all others equally. Such averages will turn out to be relevant for beauty contests with homogeneous weights: $x^i(n)$

¹⁰Here, abusing notation, we have identified $y \in \mathbb{R}^\Theta$ in the obvious way with a random variable $z \in \mathbb{R}^{\Theta \times T}$, namely, with the random variable z for which $z(\theta, t) = y(\theta)$ for each $(\theta, t) \in \Theta \times T$. Equation (4) relies on a similar understanding.




FIGURE 1. The network of Example 2.

is a random agent's expectation of a random agent's expectation ... of a random agent's expectation of y .

Example 2. Suppose the only nonzero entries of Γ are $\gamma^{i,i+1} = 1$, where indices are interpreted modulo $|N|$, the number of agents. This corresponds to agents being arranged in a cycle, with each paying attention to the one with the next index. Take, for example, $|N| = 3$ (see Figure 1). Then

$$x^1(3) = E^1 E^2 E^3 y.$$

We could continue this process, and then we would essentially look at $(E^3 E^1 E^2)^a E^3 y$, where a is some positive integer (possibly with E^2 or $E^1 E^2$ appended to the front). Our study of higher-order average expectations will allow us to study the limiting properties of this sequence.

2.5. Joint Connectedness: A Maintained Technical Assumption. A key technical assumption—*joint connectedness* of the information structure and network—will be convenient in formulating statements about limits of higher-order average expectations. This assumption will be maintained unless we state otherwise.

Say that a signal t^j (of an agent j) is a neighbor of a signal t^i (of agent i) if agent j is a neighbor of i (i.e., $\gamma^{ij} > 0$) and agent i , when he observes signal t^i , considers signal t^j possible (i.e., $\pi^i(t^j | t^i) > 0$). This defines a binary relation on the set of everybody's signals, $S = \bigcup_{i \in N} T^i$. We say the information structure and network are *jointly connected* if every nonempty, proper

subset $S' \subsetneq S$ contains some signal that is a neighbor of a signal not in S' . We will discuss the content and significance of this assumption below in Section A.1.

2.6. Consensus Expectations: Definition and Existence. An object central to our general theoretical results and the applications will be a kind of limit of higher-order average expectations as we consider many iterations.

Definition 2. For any information structure π , network Γ , and basic random variable y , the *consensus expectation* $c(y; \pi, \Gamma)$ is defined to be any entry of the vector

$$(5) \quad \lim_{\beta \uparrow 1} (1 - \beta) \left(\sum_{n=1}^{\infty} \beta^{n-1} x^i(n; y) \right),$$

for any i , if the limit exists (in the sense of pointwise convergence) and is equal to a constant vector.

The vector in (5) is sometimes called an *Abel average* of the sequence $(x^i(n; y))_{n=1}^{\infty}$ (see, e.g., Kozitsky, Shoikhet, and Zemánek, 2013). Proposition 1 in Section 4 below asserts that the consensus expectation is well-defined under the maintained assumption of joint connectedness.

The consensus expectation is equal to any entry of the simple limit $\lim_{n \rightarrow \infty} x^i(n; y)$ if the latter exists. It also coincides with the Cesàro limit, which is obtained by taking simple averages over many values of n . We will discuss these issues further in Section 4.2.

3. WHY HIGHER-ORDER AVERAGE EXPECTATIONS AND CONSENSUS EXPECTATIONS MATTER

We now discuss two economic problems where higher-order average expectations arise. First, we consider the network game with incomplete information discussed in the Introduction, where equilibrium actions are weighted averages of higher-order average expectations. Second, we describe a stylized asset market with fragmented markets, where asset prices reduce to the solution of the game, and are thus also weighted averages of higher-order average expectations. In each of these two cases, we will (i) show how outcomes are characterized by higher-order average expectations; (ii) motivate the study of consensus expectations—a limit of higher-order average expectations; and (iii) interpret our later results in the context of these applications.

3.1. Coordination. How will a group of agents coordinate their behavior when they have strong incentives to take the same action as others but have different beliefs about what the best action to take is? We consider a class of games with linear best responses where each

agent wants to set her action equal to a weighted average of (i) her expectation of a random variable and (ii) the weighted average of actions taken by others. We show how the equilibrium is determined by higher-order average expectations and then focus on the limit as coordination concerns dominate. There will be a particular single action taken in this limit by all agents after all signals—“the convention.” We first describe the game.

3.1.1. *The Game.* We will consider an incomplete-information game where payoffs depend on the states of the world, Θ . Beliefs and higher-order beliefs about Θ are described by the belief functions introduced in Section 2.1.1.¹¹ The strategic dependencies are encoded in a network Γ . We also assume that $\gamma^{ii} = 0$ for all i .¹²

The game will also depend on y , a basic (i.e., θ -measurable) random variable with support in the interval $[0, M]$.¹³

We will consider the “ β -game” parameterized by $\beta \in [0, 1]$. Each agent i chooses an action $a^i \in [0, M]$, and the best-response action of agent i after observing signal t^i is given by

$$a^i = (1 - \beta)E^i y + \beta \sum_{j \neq i} \gamma^{ij} E^i a^j,$$

where other players’ actions are viewed as random variables that depend on their own signal realizations. The best response can be derived from a quadratic loss function, where the ex post utility of agent i under realized state $\theta \in \Theta$, if action profile $a = (a^i)_{i \in N} \in \mathbb{R}^N$ is played, is

$$u^i(a^i, \theta) = -(1 - \beta)(a^i - y(\theta))^2 - \beta \sum_{j \neq i} \gamma^{ij} (a^i - a^j)^2.$$

A “meetings” interpretation of the weights γ^{ij} is that i has to commit to an action before knowing which agent he will interact with, and i assesses that the probability of interacting with j is γ^{ij} .

3.1.2. *Solution of the Game for Any $\beta < 1$.* To summarize the previous section, the *environment* in which the game is played is described by a tuple consisting of an external random variable, a network, and a coordination weight: (y, Γ, β) . A strategy of agent i in the incomplete-information game, $s^i : T^i \rightarrow \mathbb{R}$, specifies an action for each signal. Write $s^i(t^i)$ for the action

¹¹If we identify types with their (interim) beliefs, we can say that the beliefs are encoded in the type space.

¹²The assumption that the diagonal is 0 is the most natural one for this game. Analogous results hold without this assumption and have game-theoretic interpretations. See Section A.2.

¹³We will focus on the case where agents care about the same basic random variable. But the analysis extends readily to the case where agents care about different random variables, since their heterogeneous expectations of random variables conditional on their signals can be interpreted as agent-specific random variables. See Section A.5 for further discussion.

chosen by agent i upon observing signal t^i . Then agent i 's best response to strategy profile $s = (s^i)_{i \in N}$ is given by

$$(6) \quad \text{BR}^i(s) = (1 - \beta)E^i y + \beta \sum_{j \neq i} \gamma^{ij} E^i s^j.$$

To establish (7), write $R^i(k)$ for the set of i 's pure strategies surviving k rounds of iterated deletion of strictly dominated strategies.¹⁴ The map $\text{BR}(s) : [0, M]^S \rightarrow [0, M]^S$ is a contraction mapping (with Lipschitz constant β). Thus, the sets $R^i(k)$, which are produced by the repeated application of this map to the set $[0, M]$, must converge to a single point satisfying $s = \text{BR}(s)$, which is an equilibrium of our game. A more detailed proof can be found in an Appendix, Section C.1.

This analysis is the asymmetric version of the analysis in Morris and Shin (2002).

3.1.3. Conventions: Equilibrium for $\beta = 1$ and $\beta \uparrow 1$.

Fact 1. *If $\beta < 1$, the β -game in the environment given by (Γ, β, y) has a unique rationalizable strategy profile, and it is given by¹⁵*

$$(7) \quad s_*^i(y, \Gamma, \beta) = (1 - \beta) \left(\sum_{n=1}^{\infty} \beta^{n-1} x^i(n; y, \Gamma) \right).$$

There is a sharp distinction between the game with $\beta < 1$ and the game with $\beta = 1$. In the latter case, there is a continuum of equilibria, one for each $a \in [0, M]$. In these equilibria, agents all choose the same action independent of their signals and thus of the state. To see why, recall that every agent's action must be equal to his (weighted) expectation of others' actions. But now consider the highest action ever played in some equilibrium (i.e., given some signal of some agent). The agent i taking that highest action at some signal t^i must be sure that that highest action is being taken by every other agent j who observes a signal t^j that i considers possible when he observes signal t^i . Now, however, the same logic applies to agent j observing that signal t^j . Continuing in this way, our joint connectedness assumption implies that the

¹⁴One reason to focus on rationalizability is that because we do not have a common prior, there is some inconsistency in using a solution concept (equilibrium) which builds in common prior beliefs about strategic behavior (see Dekel, Fudenberg, and Levine, 2004).

¹⁵This game has a unique equilibrium (even with unbounded action spaces). This follows from the observation that the game is best-response equivalent to a team decision problem, and uniqueness in the team decision problem is shown in Radner (1962). Ui (2009) gives a general statement of this result, expressed in the language of Bayesian potential functions. Since this is a game with strategic complementarities, bounded action spaces imply that the unique equilibrium is the unique action strategy profile surviving iterated deletion of strictly dominated strategies (Milgrom and Roberts (1990)). Our proof of Fact 1 is established by explicitly calculating the iterated elimination of dominated strategies.

highest action must be played by all agents for all signal realizations. This argument and result appear in Shin and Williamson (1996), who label the resulting play—constant across agents and signals—a convention, because each agent is always choosing the same action and is choosing that action because others do.

To summarize: When $\beta < 1$, there is a unique equilibrium, with agents' actions depending on their higher-order expectations of y . When $\beta = 1$, there is a continuum of “conventional” equilibria. What happens as $\beta \uparrow 1$? The play is described by a limit of unique equilibria, which turns out to be well-defined:

$$\lim_{\beta \uparrow 1} (1 - \beta) \left(\sum_{n=1}^{\infty} \beta^{n-1} x^i(n) \right).$$

By an application of the argument of the previous paragraph to the limiting payoffs, under joint connectedness the limit must feature “conventional” play, not depending on one's signal or identity. The existence of the limit and a characterization of the action played in it will be formalized in the next section; the main result is Proposition 1. The limit can be seen as a selection among the continuum of equilibria of the $\beta = 1$ game. It is telling us how conventional play is determined when there is an arbitrarily small amount of dependence of the payoffs on some basic random variable. The basic random variable can be interpreted as a “cue” that orients players' coordination.

Note that the statements made about the game before we started considering the $\beta \uparrow 1$ limit hold for any $\beta \in [0, 1]$. From now on, we will focus on the $\beta \uparrow 1$ limit, motivated by the interpretation of it just given, as a refinement of the coordination game (as well as a parallel motivation we are about to present, based on frequent trade of an asset).¹⁶

Our main results focus on β begin very close to, but not equal to, 1. Some of our results apply to, or have implications for, the more general situation with β much smaller than 1, and we discuss that case when appropriate.

3.1.4. *Conventions with High Coordination Weights: Preview of Main Results.* Our results in Sections 6, 7, and 8 characterize that limit convention in some environments:

1. Under the common prior assumption, the convention is equal to the common ex ante expectation of y . If agents share a common prior on signals, but not necessarily on states,

¹⁶Weinstein and Yildiz (2007a) have argued that, in a fixed linear best response game, very high-order beliefs have only a small impact on rationalizable play; this contrasts with the better-known observation in Weinstein and Yildiz (2007b) that very high-order beliefs can have an arbitrarily high impact in general games. We get sensitivity to high level higher-order beliefs in linear best response games because we are looking at $\beta \uparrow 1$ limit and thus a sequence of different games. Subtleties of such comparisons are also discussed in Morris (2002b).

then the convention is equal to a weighted average of (different) *ex ante* expectations that the agents hold of y , with each agent's expectation weighted by his eigenvector centrality in the interaction network Γ .

2. If all agents always have a small amount of second-order optimism (believing that their average counterparty is a bit more optimistic than they are), the convention will equal the highest possible realization of y —that is, the highest interim expectation ever held by any agent.
3. If there is common interpretation of signals and one agent is sufficiently less informed than all other agents, then the convention will equal the *ex ante* expectation of that least informed agent.

3.2. **Asset Pricing.** Keynes (1936, p. 156) famously likened investment to a “beauty contest” whose outcome depends on higher-order beliefs:

... professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one's judgement, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.”

Keynes is presumably not suggesting that the newspaper competition winner is completely independent of “prettiness” but rather that each competitor has an incentive to try to match the average expectation of prettiness, and then some average expectation of such average expectations, and so on. We will study an asset pricing model where asset prices will correspond to solutions of the coordination game above and thus to the description of investment behavior that Keynes gives.

3.2.1. *Asset Market.* Suppose that there are several populations or classes of traders, indexed by N , and each of these consists of a continuum of infinitesimal *traders*. There is an asset

whose payoff will depend on the realization of a random variable y that is measurable with respect to Θ and that takes values in $[0, M]$. The beliefs and higher-order beliefs of traders in class i about the state space Θ will be given by the belief function π^i defined in the general model; in particular, they share the same belief function. Each trader in class i will also observe the same signal t^i . Thus, they share the same interim beliefs. All traders are risk-neutral and there is no discounting. A single unit of the asset will be traded among all classes of traders. There is a network Γ , which will determine where traders resell their assets in a way we are about to describe.

The trading game works as follows. Time is discrete. At each time t , one trader (say, in class i) enters owning the asset. With probability β , the state is realized and the owner of the asset consumes the realization of the asset (with the interpretation that this corresponds to liquidity needs). He then exits the game. If not (and so with probability $1 - \beta$), a class of traders j is selected randomly (and exogenously). The asset owner believes that class j is selected with probability γ^{ij} . The owner must then sell the asset in a market consisting of all traders of class j who have not yet exited. There is Bertrand competition in market j , with each buyer (i.e., remaining trader in class j) offering a price p and the seller (in class i) deciding to whom to sell the asset. We then enter period $t + 1$ with the chosen buyer in class j holding the asset.

3.2.2. Techniques and Related Models of Asset Pricing. In the special case where the network is uniform, we could have derived the same asset pricing formula in a standard dynamic CARA-normal rational expectations model, with overlapping generations of agents, as studied by Grundy and McNichols (1989) and many others. In each period, the market will shut down with probability $1 - \beta$, and the current old agents will consume a terminal value of the asset. If it does not shut down, the old will sell the asset to the young. In each period, the young will inherit the distribution of signals about the terminal value of the old. The asset price will equal the forward looking risk-adjusted iterated expectation of the value of the asset. If the variance of noise traders in the market increased without bound, there would be no learning in the market and the expected risk-adjusted price would be equal to the iterated average expectation. Allen et al. (2006) show this for a finite truncation of this environment with $\beta = 1$. The dynamic CARA-normal rational expectations model is studied under the common prior assumption. Banerjee and Kremer (2010) and Han and Kyle (2016) have studied the role of heterogeneous prior beliefs in the static version of the model.

This asset market combines features that appear in many other asset pricing models, and we now review some of the connections. Harrison and Kreps (1978) study an asset market where

an asset is re-traded in each period between different risk-neutral agents with heterogeneous prior beliefs. They focus on the minimal price paths, in order to rule out bubbles based purely on everyone's expectation that prices will rise based on calendar time; we achieve a similar effect with our stationarity assumption. We allow asymmetric information but make exogenous the agent to whom another agent must sell. Duffie and Manso (2007) study a random matching model of trade, where traders are matched in pairs at each time period. They focus on information percolation over time with a simple updating rule, while we focus on effects due to higher-order beliefs; our matching technology is also more general. Malamud and Rostek (2016) study markets with an exogenous network structure of access to multiple markets, but endogenize agents' choice of how much to trade in each market.

A key simplification in our model of trading is that each agent is infinitesimal, so any learning about the asset value does not affect anyone's expectations. Steiner and Stewart (2014) obtain the same effect in a model of asymmetric information where agents do not condition on others' information. They give a behavioral interpretation of this restriction via coarse perceptions.¹⁷ Our model and that of Steiner and Stewart (2014) both feature the same dependence of prices only on public information among the agents; the limit where trading becomes frequent is critical to this.

3.2.3. Equilibrium Asset Prices. We will consider symmetric Markov subgame perfect equilibria of the asset trading game described in Section 3.2.1. By "symmetric Markov," we mean that each trader's offer will depend only on the class to which he belongs and the class of the current owner from whom he is buying.

The main result about this asset market is that there is a unique symmetric, Markov, subgame-perfect equilibrium where, whenever the asset is sold in market j , the traders in market j with signal t^j offer a price equal to $s(t^j)$, where $s = s^*(\beta)$ as defined in (7), and owners sell to any trader in class j offering the highest price. In other words, traders always set prices equal to the equilibrium of the linear best-response game of the previous section. To see why, note first that a trader's willingness to pay for the asset does not depend on whom he is buying the asset from. Also, observe that in a symmetric equilibrium, traders must be setting prices equal to their willingness to pay. Thus equilibrium asset prices must satisfy equation (7) i.e., the equilibrium condition from the linear best-response game.

¹⁷One can also give their results an interpretation in terms of heterogeneous beliefs and asymmetric information.

Our analysis does depend on the restriction to symmetric Markov strategies and equilibrium rather than rationalizability as a solution concept.¹⁸ If we did not impose the Markov assumption, there would be “bubble” equilibria, with the asset price growing exponentially. We also used the assumption of equilibrium in our analysis, when we directly assumed that the prices satisfy the equation (7), rather than (as we did in Section 3.1.2) arguing that this condition follows from some weaker solution concept. We used symmetry when we assumed that all members of a given class price the asset the same way whenever they have the opportunity to buy.

3.2.4. Asset Pricing with High-Frequency Trading: Preview of Main Results. Taking the limit $\beta \uparrow 1$ now corresponds to requiring faster and faster trade while holding time preferences fixed. The high-frequency limit price will be the same in every market for every signal. This price turns out to be the consensus expectation.¹⁹ Our main results below have implications for the asset prices which parallel the statements of 3.1.4 applied to the game.

4. THE INTERACTION STRUCTURE

4.1. Interaction Structure. One contribution of this paper is to show that the information structure and the network structure can be seen from a unified perspective—in studying higher-order average expectations and, consequently, for our applications. In particular, we will define an *interaction structure*—a square matrix indexed by the set S comprising the union of everyone’s signals—that simultaneously captures beliefs and the network. This serves two purposes. First, it highlights the symmetry between information and the network. Second, it facilitates relating higher-order average expectations to a Markov matrix and its iteration, which is an important technique for us.²⁰ Indeed, we will use a Markov process representation to establish all the results that follow from results about Markov processes.

Let $S = \bigcup_{i \in N} T^i$ be the union of the (disjoint) sets of signals.²¹ Define $x(n) : S \rightarrow \mathbb{R}$ by $[x(n)](t^i) = [x^i(n)](t^i)$. In words, this one function is a parsimonious way of keeping track of the higher-order average expectations of *all* agents at stage n . A random variable $y : \Theta \rightarrow \mathbb{R}$ that

¹⁸Recall footnote 14 on the comparison of rationalizability and equilibrium in this context.

¹⁹Steiner and Stewart (2014) have used high-frequency limits to show a similar convergence to public random variables.

²⁰Samet (1998a) introduced and used a Markov process as a representation of an information structure. We construct a related, but different, process: Ours simultaneously captures the network and agents’ beliefs and operates on the union of signals instead of realizations. See Golub and Morris (2017) for the exact analogue of Samet’s process.

²¹Recall that this object appeared in the definition of joint connectedness in Section 2.5. It should not be confused with the product set $T = \prod_{i \in N} T^i$, whose elements are signal *profiles*.

depends on the external state is viewed as a vector indexed by Θ , i.e., $y \in \mathbb{R}^\Theta$. The first-order expectation map $y \mapsto x(1)$ can then be viewed as a map $\mathbb{R}^\Theta \rightarrow \mathbb{R}^S$. Using the standard bases for the domain and codomain, we can represent this map via a matrix. Indeed, we can write $x(1) = Fy$, where F is a matrix with rows indexed by T^i and columns indexed by Θ , and whose entries are

$$(8) \quad F(t^i, \theta) = \pi^i(\theta \mid t^i).$$

Even though the rows and columns of this matrix are not ordered, we can define matrix multiplication by stipulating that

$$(Fy)(t^i) = \sum_{\theta \in \Theta} F(t^i, \theta)y(\theta).$$

It is immediate to check that with this definition, $(Fy)(t^i)$ is indeed i 's subjective expectation of y when i receives signal t^i .

Along the same lines, the formula of (4), $x^i(n+1; y) = \sum_{j \in N} \gamma^{jj} E^i x^j(n; y)$, can be described in matrix notation. Equation (4) defines a linear map $\mathbb{R}^S \rightarrow \mathbb{R}^S$ such that $x(n) \mapsto x(n+1)$. Taking the standard basis for \mathbb{R}^S (as both the domain and codomain) we can write $x(n+1) = Bx(n)$, where B is a matrix with rows and columns indexed by S , and entries

$$(9) \quad B(t^i, t^j) = \gamma^{jj} \pi^i(t^j \mid t^i).$$

We call B the *interaction structure*. It captures the weights (arising from both the network and agents' beliefs) that matter for iterating agents' expectations.

Combining the above, we find, for $n \geq 1$, the short formula

$$(10) \quad x(n) = B^{n-1} Fy,$$

which describes the step- n higher-order average expectations. Thus, understanding their behavior boils down to studying powers of the linear operator B . One can check that:

Fact 2. *The interaction structure B is row-stochastic.*

To verify this, note that for each $t^i \in S$ we have

$$\sum_{t^j \in S} B(t^i, t^j) = \sum_{j \in N} \sum_{t^j \in T^j} \gamma^{jj} \pi^i(t^j \mid t^i) = \sum_{j \in N} \gamma^{jj} \sum_{t^j \in T^j} \pi^i(t^j \mid t^i) = 1.$$

The final equality for each t^i follows because the distribution $\pi^i(\cdot \mid t^i)$ is a probability distribution over T^j and Γ is row-stochastic.

We will occasionally emphasize the dependence of the matrices we have defined on $\pi = (\pi^i(\cdot | t^i))_{t^i \in T^i, i \in N}$, and the dependence of B on the network Γ , by writing F_π and $B_{\pi, \Gamma}$, and similarly for derived objects.

The interaction structure B allows us to recover a matrix corresponding to one agent's beliefs about another. For any i and j , if we set γ^{ij} to 1 and all the other entries of Γ to 0, then B restricts naturally to an operator $B^{ij} : \mathbb{R}^{T^j} \rightarrow \mathbb{R}^{T^i}$ sending T^j -measurable random variables to i 's conditional beliefs about them. The entries of the matrix are $B^{ij}(t^i, t^j) = \pi^i(\cdot | t^i)$.

Equation (10) entails a sharp separation between (i) agents' first-order beliefs about Θ , on the one hand, and (ii) their beliefs about each other's signals, as well as the network, on the other. The former beliefs are encoded in F , and the latter are in B .

4.2. The Consensus Expectation via the Interaction Structure. In Section 2.6, we defined the consensus expectation. The formalism we have introduced will allow us to prove Proposition 1, below, on its existence, and in the process also to relate it to properties of the matrix B .

Recalling Definition 2, the consensus expectation is the number in every entry of the following vector:

$$(11) \quad \lim_{\beta \uparrow 1} (1 - \beta) \sum_{n=1}^{\infty} \beta^{n-1} x(n; y)$$

The notation introduced in Section 4.1 above allows us to rewrite this as

$$(12) \quad \lim_{\beta \uparrow 1} (1 - \beta) \left(\sum_{n=0}^{\infty} \beta^n B^n \right) Fy.$$

In this section, we use the formalism we have introduced to explain why this limit exists and why it is a constant vector, as well as to characterize it. The following is our main result on this, which shows that the consensus expectation (recall Definition 2 in Section 2.6) is well-defined.

Proposition 1. *The consensus expectation exists and*

$$(13) \quad c(y; \pi, \Gamma) = \sum_{t^i \in S} p(t^i) E^i[y | t^i],$$

where p is the unique vector in $p \in \Delta(S)$ satisfying $pB = p$. All entries of p are positive, and it is called the vector of agent-type weights.

Thus the consensus expectation of y is a weighted average of the expectations associated with the various signals of each agent, encoded in Fy ; the weight on the expectation of signal

t^i of agent i , or simply type t^i , is given by $p(t^i)$.²² Note that, by definition, p is the stationary distribution of B viewed as a Markov matrix.

A simple but important separation can be read off from the formula of Proposition 1. The vector p , because it is uniquely defined by B (by the Perron–Frobenius Theorem), depends only on the entries of B , which in turn depend only on the network weights γ^{ij} and on agents’ interim marginals on one another’s signals, $\pi^i(t^j | t^i)$. Thus, these features of the model jointly determine the weights $p(t^i)$. Beliefs about Θ enter only through $E^i[y | t^i]$. This reflects the separation noted at the end of Section 4.1. Thus, the interesting effects arising from higher-order beliefs will be characterized by relating p to the information structure; see, for instance, Sections 7 and 8.

For our analysis, we have fixed a y throughout; however, note that if y were arbitrary, Proposition 1 would hold with the *same* p for all y .

To see why Proposition 1 holds, first note that if

$$(14) \quad \lim_{n \rightarrow \infty} B^n Fy$$

exists, then this limit will equal (11). This is because (11) is the weighted mean of terms of the form $B^n Fy$; as $\beta \uparrow 1$, most of the weight is assigned to the terms corresponding to large values of n . To give intuition, here will assume that (14) exists, though our result is more general as shown in the proof of Proposition 1 in Appendix C.2.²³

Joint connectedness will imply that the matrix B is irreducible.²⁴ Thus, by a standard fact about such matrices, every row of B^∞ is p , assuming this limit exists. Writing $\mathbf{1}$ for the function (vector in \mathbb{R}^S) that takes a constant value of 1 on all of S , for any vector $z \in \mathbb{R}^S$ we have

$$(15) \quad \lim_{\beta \uparrow 1} (1 - \beta) \sum_{n=0}^{\infty} \beta^n B^n z = (pz)\mathbf{1},$$

where p is as defined in the statement of Proposition 1. In the analysis of (11), we set $z = Fy$. A variant of the standard Markov chain result shows that (15) holds more generally, even when the limit of the $x(n)$ in (10) does not exist.

²²Recalling Section 2.1.1, we use the terminology of a *type* here for a signal to emphasize the interim perspective: All that matters for higher-order expectations (and hence consensus expectations) are an agent’s interim beliefs (including higher-order beliefs), and agents’ types fully capture these.

²³Sometimes (11) will exist when (14) does not, because for large n , the vector $B^n Fy$ cycles (approximately) among several limit vectors. In this case, (11) takes an average of these vectors. We discuss these issues further in Appendix B.1.

²⁴The meaning of irreducibility in our context is discussed further in Section A.1.

Proposition 1 implies that higher-order average expectations converge in the sense of (11) to a number which is independent of the agent *and* of his signal: the consensus expectation. Thus, in the coordination game, agents' actions in the $\beta \uparrow 1$ limit, where coordination concerns dominate, are equal to a nonrandom consensus.²⁵ Of course, the consensus expectation depends, in general, on all the interim beliefs $(\pi^i)_{i \in N}$ and on the network Γ .

4.3. A Markov Process Interpretation of the Interaction Structure and the Consensus Expectation. The interaction structure B is a row-stochastic or Markov matrix, and corresponds to a Markov process that we construct, with S playing the role of the state space. We can imagine a particle starting at some state $t^i \in S$, and the probability of transitioning to $t^j \in S$ being $\gamma^{ij} \pi^i(t^j | t^i)$.

This process can be useful for understanding the behavior of higher-order average expectations. Fix a signal $t^i \in S$ and consider the Markov process started at this t^i , with its (random) location over time captured by the random variables $W_1 = t^i, W_2, W_3, \dots$. If we define a function $f : S \rightarrow \mathbb{R}$ such that $f(t^i) = (Fy)(t^i)$, then $x^i(n)$, the n^{th} -order average expectation of y , is the expected value of $f(W_n)$. The vector of agent-type weights discussed in Section 4.2 is the stationary distribution of the chain, and the consensus expectation of y is the expected value of $f(W)$ where W is drawn according to the stationary distribution.

The process we have defined provides a physical analogy that is useful for intuition and also suggests proof techniques—see Sections 7 and 8.

5. THE CONSENSUS EXPECTATION AND THE NETWORK

One simple special case of Proposition 1 arises when $|T^i| = 1$ for each i : There is complete information about each agent's signal. In that case, $B = \Gamma$ and so $p = e$, the eigenvector centrality vector of the network Γ . (Recall Definition 1 in Section 2.2.) It follows from (13) that

$$c(y; \pi, \Gamma) = \sum_i e^i E^i y,$$

where, abusing notation, $E^i y$ denotes the interim expectation of y induced by the one signal that agent i ever gets. This relates to network game results of Ballester, Calvó-Armengol, and Zenou (2006), and especially to the limit with high coordination motives studied in Calvó-Armengol, Martí, and Prat (2015), where play is determined by ideal points weighted by eigenvector centralities.

²⁵If there are public events, the consensus is nonrandom once public information is taken into account. See Section A.1.2 for further discussion.

There is a much more general sense in which the eigenvector centralities of the agents figure in the consensus expectation:

Proposition 2. *There are strictly positive priors $(\lambda^i)_{i \in N}$, with $\lambda^i \in \Delta(T^i)$, such that, for all y ,*

$$(16) \quad c(y; \pi, \Gamma) = \sum_i e^i \mathbf{E}^{\lambda^i} y,$$

where the e^i are the eigenvector centralities of the agents.

The expression $\mathbf{E}^{\lambda^i} y$ corresponds to an ex ante expectation of agent i , where the expectation is taken according to a *pseudoprior* λ^i over i 's signals that need not be related to agent i 's actual prior μ^i .

Recalling (13), we can see that this result asserts $e^i \lambda^i(t^i) = p(t^i)$, and indeed its content is that agent i 's agent-type weights sum to his eigenvector centrality, e^i . This is formally stated in the following lemma, which is what we use to prove Proposition 2, and which also relates to the Markov interpretation of consensus expectations in Section 4.3.

Lemma 1. *For each i , the agent-type weights associated with agent i 's types add up to the eigenvector centrality of i :*

$$\sum_{t^i \in T^i} p(t^i) = e^i.$$

Proof. Let $\iota : S \rightarrow N$ map any type t^i to the agent i whose type it is. Check that $V(n) := \iota(W(n))$ is a Markov process on N with transition matrix Γ .²⁶ Now the stationary probabilities of the process W are given by p , and the total stationary probability of the set $T^i \subseteq S$ under W is therefore $\sum_{t^i \in S} p(t^i)$. By the coupling between V and W , this must be equal to the stationary probability of i under V , which is e^i . \square

The proof of Proposition 2 is completed by making the definition $\lambda^i(t^i) = p(t^i)/e^i$, which is legitimate because all the centralities e^i are positive (see comments after Definition 1).

Generally, the pseudopriors λ^i will depend on *both* the information structure π and the network Γ . We will be especially interested in when the λ^i depend only on beliefs. The next section gives some conditions for this, and the issue is discussed more generally in Section A.3.

²⁶The reasoning is as follows: The probability of the event $\{V(n+1) = j\}$ conditional on $\{V(n) = i\}$ is equal to γ^{ij} : For any $t^i \in T^i \subseteq S$, we have

$$\sum_{t^j \in T^j} B(t^i, t^j) = \sum_{t^j \in T^j} \gamma^{ij} \pi^i(t^j | t^i) = \gamma^{ij} \sum_{t^j \in T^j} \pi^i(t^j | t^i) = \gamma^{ij}.$$

5.1. Interpreting the Interaction Structure as a Network. As we mentioned at the end of Section 3.1.1, in the context of the interaction game, the weights γ^{ij} can be interpreted as i 's subjective probabilities of meeting or interacting with various others at the time he has to commit to his action. In light of this interpretation, $B(t^i, t^j) = \gamma^{ij} \pi^i(t^j | t^i)$ can be seen as a subjective probability assessed by agent i , when he has signal t^i , that his partner in the game will have signal t^j : The first factor, γ^{ij} , is i 's probability of meeting j , and $\pi^i(t^j | t^i)$ is the probability, conditional on that meeting, that j has signal t^j . (An agent may be privately informed about his weights or interaction probabilities; see our discussion in Section A.5.2.)

In this sense, the environment can be reduced, from the perspective of each player, purely to incomplete information. Relatedly, we can reduce the analysis purely to networks. To this end, we construct a new environment (whose objects are distinguished by hats) based on the primitives of the original environment. In this environment the new set of *agents*, \widehat{N} , is S , the set of all signals. The network is $\widehat{\Gamma}(t^i, t^j) = \gamma^{ij} \pi^i(t^j | t^i)$; there is complete information about signals (each agent has a singleton type t^i , which is also his agent label); and the first-order beliefs of the new agents replicate those of the corresponding types. Now, the higher-order average expectation vector of this new environment, $\widehat{x}(n; y)$, is the same²⁷ as $x(n, y)$. All statements about higher-order average expectations in the original game of incomplete information can be reinterpreted in this complete-information environment as network quantities. For instance, to get the second-order average expectation of a type t^i , we look at the corresponding agent in the network, and take the average, across all his neighbors, of their neighbors' first-order expectations.

To summarize: We have taken all uncertainty about others' signals, and combined it with the original network weights, to obtain the new network weights $\widehat{\Gamma}$. From this perspective, the game of incomplete information of Section 3.1.1 is reduced to the network game studied by Ballester, Calvó-Armengol, and Zenou (2006). (For a similar reduction, see Morris (1997); see de Martí and Zenou (2015) for computations using a related formalism.) This transformation is essentially the transformation of the game of incomplete information into an agent normal form.

²⁷Under the obvious bijection of indices.

6. UNIFYING AND GENERALIZING NETWORK AND ASYMMETRIC INFORMATION RESULTS

We now study conditions under which the agent-type weights take a particularly simple form. Under these conditions, there are formulas for consensus expectations that decompose nicely into different individuals' prior expectations, weighted by those individuals' centralities.

Recall that agents' priors are given by the profile $(\mu^i)_{i \in N}$ of distributions, with $\mu^i \in \Delta(T^i)$.

Definition 3. There is a *common prior over signals* (CPS) if, for each signal profile $t \in T$ and each $i, j \in N$, we have

$$\mu^i(t^i)\pi^i(t^{-i} | t^i) = \mu^j(t^j)\pi^j(t^{-j} | t^j).$$

CPS does not imply a common prior over the states Θ ; agents may have inconsistent beliefs about θ . A common prior on signals could arise if each agent first observed a signal drawn according to the common prior but interpreted signals differently. However, CPS does imply that there is a common prior over agents' second-order and higher-order beliefs.

Now we can show that under CPS, the distributions λ^i in the representation $c(y; \pi, \Gamma) = \sum_i e^i \mathbf{E}^{\lambda^i} y$ of Proposition 2 are ex ante probability distributions on signals, i.e., $\lambda^i = \mu^i$; the pseudopriors are the *actual* priors. Recall from Section 2.1.2 that bold expectation operators denote ex ante expectations.

Proposition 3. *If there is a common prior over signals, then the consensus expectation is equal to the eigenvector-centrality weighted average of the ex ante expectations of the agents:*

$$c(y; \pi, \Gamma) = \sum_i e^i \mathbf{E}^{\mu^i} y,$$

where μ^i is the prior over i 's signals.

Proposition 3 shows that the consensus expectation is a weighted average of agents' prior expectations, $\mathbf{E}^{\mu^i} y$, weighted by agents' network centralities, e^i . We say in this case that there is a *separability* between the network and the information structure: The network enters only into the centralities, and the information structure determines $\mathbf{E}^{\mu^i} y$. (See Section A.3 for further discussion of this property.) Under complete information about signals but heterogeneous priors about Θ , this yields a reinterpretation of the DeGroot model, as we discuss further in Section 9.

In terms of the generality of the information structure, Proposition 3 goes beyond previous related results that decomposed equilibrium actions into agent-specific quantities weighted by agents' centralities. Results in this category include Calvó-Armengol, Martí, and Prat (2015)

and Bergemann, Heumann, and Morris (2015b) (which rely on Gaussian signals), de Martí and Zenou (2015) (which relies on exchangeable signals), and Blume, Brock, Durlauf, and Jayaraman (2015) (which does not characterize the contribution of higher-order expectation terms). Formal details of each of these models differ in several ways from our model, but in not imposing parametric or symmetry conditions, and in allowing heterogeneous priors about states, our result on the decomposition at the $\beta \uparrow 1$ limit is more general than the others.

If each ex ante expectation is the same, in Proposition 3 it does not matter what the eigenvector centralities are, and we have:

Corollary 1. *If there is a common prior over signals and $\mathbf{E}^{\mu^i} y = \bar{y}$ for all i —that is, agents have a common ex ante expectation of the external random variable—then the consensus expectation is equal to the (common) ex ante expectation.*

$$c(y; \boldsymbol{\pi}, \Gamma) = \bar{y}.$$

Corollary 1 is closely related to Samet (1998a), which shows that if the common prior assumption (over the whole space Ω) holds, then any sequence of expectations (A 's expectation of B 's expectation . . .) of the random variable is equal to the ex ante expectation of the random variable \bar{y} . Since limits of such iterated expectations determine the consensus expectation, it is also equal to \bar{y} . Note, however, that the hypotheses of Corollary 1 are weaker than the full common prior assumption, because they impose no restrictions on the joint distribution of θ and signals.

Proposition 3 is also closely related to Samet (1998a) in the following sense: given a common prior over signals, the highly iterated expectation of any random variable measurable with respect to agent i 's signal is equal to the common prior expectation of that random variable—that is, the expectation of it with respect to the measure μ^i . Our results show that these prior expectations are combined according to agents' network centrality weights. Sections A.3 and A.4.2 elaborate further on these issues, as well as a converse to Samet's result.

7. SECOND-ORDER OPTIMISM

Consider a case in which agents are optimistic about the expectations of those they interact with—that is, they believe that, on average, those others have higher expectations than their own. In this case, consensus expectations are driven to extremes. Sections 3.1.4 and 3.2.4 state the interpretation of this in the game and in the asset market, respectively.




FIGURE 2. The example of Case II, with a counterclockwise network.

7.1. Three Illustrative Cases. To motivate our results on this and to gain intuition, we first consider some extreme cases. These illustrate how the Markov process representation of higher-order expectations and its physical interpretation from Section 4.3 can yield striking results about consensus expectations. Fix a random variable with minimum realization 0 and maximum realization 1. Say that agent i considers j *over-optimistic* if agent i 's expectation of agent j 's expectation is always strictly greater than his own expectation (unless his own expectation is 1, in which case he is sure that agent j 's expectation is 1). Say that agent i thinks that agent j is *over-pessimistic* if agent i 's expectation of agent j 's expectation is always strictly *less* than his own expectation (unless his own expectation is 0, in which case he is sure that agent j 's expectation is 0).

Case I. First, suppose that each agent considers every other agent over-optimistic.²⁸ In this case the consensus expectation must be 1, independent of the network structure.

Case II. Second, suppose that for every agent i , there is an agent he considers over-optimistic and another agent he considers over-pessimistic. Then there is a network structure under which the consensus expectation is 1. We can simply look at the network structure in which each agent puts all weight on agents he thinks are over-optimistic. Symmetrically, there is a network structure in which the consensus expectation is 0. These results do not depend on agents' ex ante expectations—which might take any value between 0 and 1.

Figure 2 illustrates one example of this occurring. There are I agents, indexed by $N = \{1, \dots, I\}$, and their indices are interpreted modulo I . Each agent has many signal realizations,

²⁸This—and other examples we describe here—may involve violating the otherwise maintained joint connectedness assumption (irreducibility of B), but our main result in this section, Proposition 4, does not rely on the joint connectedness of B .

t_k^i , with indices $k \in \{1, \dots, K\}$, with higher- k signals inducing more optimistic first-order beliefs about y . Assume that the most extreme signals lead to expectations 1 and 0. Agent i , when he has signal t_k^i , is certain that agent $i-1$ has signal t_{k+1}^{i-1} , the next more optimistic signal. He is also certain that agent $i+1$ has signal t_{k-1}^{i+1} , the next more pessimistic signal. If k is already extreme (that is, $k = 1$ or K) then we replace $k+1$ (respectively, $k-1$) by k in the above description.

Now the two networks considered are as follows. One has each agent assigning all weight to the agent counterclockwise from him (i.e., to his left, as depicted in Figure 2). The other network has each agent assigning all weight to the one clockwise from him (i.e., to his right). Then in the counterclockwise network, the consensus expectation is 1, and in the clockwise network (not shown), the consensus expectation is 0.

Case III. For our final case, rather than assuming any agent is over-optimistic about any other, assume instead that each agent's expectation of the *average* expectation of others' expectations is greater than his own expectation. As always, averages are taken with respect to network weights, and "greater" is strict except in the case where an agent's expectation is 1. This constitutes a milder form of over-optimism. Note that it is not implied by the assumptions we imposed in either of the above results: While the condition of Case III depends on the network (as in Case II), it allows for the possibility that an agent is never over-optimistic about any other particular agent (recalling that over-optimism is a condition uniform over one's signals). Rather, which agent someone is over-optimistic about may depend on his signal. But again, the consensus expectation is 1.

Markov Process Intuitions. The results in each of the cases above can be established by using the representation of higher-order average expectations via a Markov process, which we presented in Section 4.3.

Let us begin by explaining Case I. If a particle makes transitions over the states S according to the Markov process, then at each step it moves toward strictly more optimistic types of agents, unless it is already at a most optimistic type. Similar arguments can be given for the other cases; see the proof in Section 7.3 for the general argument.

The cases discussed so far involve the unsatisfactory assumption that some types are *certain* that they are the most optimistic. It will often be unreasonable for agents to hold such extreme beliefs, or for the analyst to assume that they do. Thus, we wish to have a result that is more quantitative and more robust. Also, the networks involved in Case II are extreme, not allowing

an agent to put even small amounts of weight on others whom he does not consider over-optimistic (or over-pessimistic). Our general results will relax all these assumptions.

The basic idea behind that generalization is clear: it follows from the arguments above and continuity. But the details are subtle. Indeed, what will be most interesting about the general results we obtain is the nature of the conditions that are involved. How much second-order pessimism can be permitted for the very optimistic agents without losing the contagion of optimism? By relating the situation of second-order optimism to a suitable Markov chain, we are able to give a precise quantitative bound on the permitted degree of “realism” by optimistic types, and this bound turns out to be tight.

7.2. A General Case. We now weaken our assumptions on the most optimistic types, and allow for the possibility that when agents are maximally optimistic they assign only probability $1 - \varepsilon$, for some $\varepsilon > 0$, to any given other being maximally optimistic. But now we assume that when an agent is not maximally optimistic, there is a uniform lower bound, δ , on the degree of over-optimism. With this weakening of our earlier assumptions, the above results remain true with an error of order $\frac{\varepsilon}{\delta}$. Among other things, this allows us to use a network with $\gamma^{ij} > 0$ for all pairs i, j in Case II above.

We state and prove a formal version of our claims in this general case, and then discuss how the claims made about our illustrative cases follow. Critically, in addition to demonstrating the continuity in beliefs we needed, this result gives quantitative bounds on how agents’ interim over-optimism translates into the consensus outcome.

Proposition 4. *Consider an arbitrary information structure π and an arbitrary network Γ (i.e., drop for this result the maintained assumption that B is irreducible). Suppose there exist \bar{f} and $\delta > 0$, $\varepsilon \geq 0$ such that beliefs about neighbors are mildly optimistic in the following sense:*

1. *Every type whose first-order expectation of y is strictly below \bar{f} expects the first-order expectation, averaged across his counterparties, to be at least δ above his own. That is, for every t^i such that $(E^i y)(t^i) < \bar{f}$, we have $\sum_j \gamma^{ij} (E^i E^j y)(t^i) \geq (E^i y)(t^i) + \delta$.*
2. *Every type whose first-order expectation of y is at least \bar{f} expects the first-order expectation, averaged across his counterparties, to be almost as large as his own, with a shortfall of at most ε . That is, for every t^i such that $(E^i y)(t^i) \geq \bar{f}$, we have $\sum_j \gamma^{ij} (E^i E^j y)(t^i) \geq (E^i y)(t^i) - \varepsilon$.*

Then the consensus expectation of y is at least $\frac{\bar{f}}{1+\varepsilon/\delta}$.

The proof of this result, via a suitable Markov chain inequality, is provided in Section 7.3 below.

An important feature of this result is that, fixing the constants δ and ε , its hypotheses do not depend on the finite type space used to represent the environment.

We now return to Cases I–III, with expectations taking values in $[0, 1]$, and describe how to obtain them formally as applications of this result. For Case I, where each agent considers every other one over-optimistic, set $\bar{f} = 1$. Because of finiteness of the type space, there is a δ so that hypothesis (1) of Proposition 4 holds for all types whose first-order expectations of y are strictly below \bar{f} . For this case, we can take $\varepsilon = 0$. Applying Proposition 4, we get that the consensus expectation is 1. For the case of over-pessimism, we simply apply a change of variables from y to $1 - y$ and use the same result to find that the consensus expectation is 0.

For Case II, we constructed two networks. In one network, each agent places all weight on some agent he considers over-optimistic. For this network, the hypotheses of Proposition 4 hold for the same reasons discussed in the previous paragraph, and we conclude that the consensus expectation is 1. In the other network, each agent places all weight on some agent he considers over-pessimistic, and by symmetry the consensus expectation is 0.

Case III is a direct application of the proposition, with $\varepsilon = 0$.

7.3. Markov Chain for Second-Order Optimism. We now analyze the interaction structures corresponding to second-order optimism and discuss how to establish our results using Markov chain arguments.

Consider an arbitrary finite state space S with a Markov kernel, with $B(s, s')$ being the probability of transitioning from state s to s' , and fix a function $f : S \rightarrow \mathbb{R}$.

The purpose of this subsection is to present the following lemma: Assume that for all s such that $f(s)$ is below a certain value \bar{f} , taking one step from s (according to the Markov kernel) to reach a random state W_2 yields a value $f(W_2)$ that is higher by at least δ , in expectation, than $f(s)$. Assume also that if, in contrast, s is chosen such that $f(s)$ exceeds \bar{f} , then the expected value of $f(W_2)$ can decrease relative to $f(s)$ by only a smaller amount, ε . Under these assumptions, we will show that if s is drawn from a stationary distribution of B , the expectation of $f(s)$ is not much below \bar{f} . The lemma we now state makes this quantitative and precise.

We denote by W_1, W_2, \dots the stochastic process induced by the Markov chain. The symbol $\mathbb{P}_{W_1 \sim \nu}$ denotes the probability measure corresponding to this process when W_1 is drawn according to a distribution ν .²⁹ The notation for expectations is analogous.

²⁹When ν is a point measure on s , we write $W_1 = s$ in the subscript as a shorthand.

Lemma 2. *Let B be a Markov chain as described above. Suppose there are real numbers $\delta, \varepsilon > 0$ and \bar{f} such that the following hold:*

1. *For every s such that $f(s) < \bar{f}$, we have $\mathbb{E}_{W_1=s}[f(W_2)] \geq f(s) + \delta$.*
2. *For every s such that $f(s) \geq \bar{f}$, we have $\mathbb{E}_{W_1=s}[f(W_2)] \geq f(s) - \varepsilon$.*

Fix an arbitrary starting state, and let p denote the ergodic distribution over states that is reached starting from that state.³⁰ Then $p(s : f(s) \geq \bar{f}) \geq \frac{1}{1+\varepsilon/\delta}$.

The proof, which appears in Section C, uses the fact that $f(W_2)$ and $f(W_1)$ have the same expectations under the ergodic distribution, and uses the hypotheses of the lemma in this equation to derive the desired inequality. With this result, we can establish all the conclusions about the consensus expectation, as discussed after Proposition 4 above.

7.4. Discussion.

7.4.1. *Related Results.* The result also relates to Harrison and Kreps (1978), who consider the case where risk-neutral agents have heterogeneous beliefs (but symmetric information) and trade and re-trade an asset through time. The asset is always sold to the (endogenously) most optimistic agent at the current history. The price is driven above the highest expectation of the asset’s value held by any agent. Harrison and Kreps (1978) motivate their exercise as a model of “speculation,” and our result has a similar interpretation. In both cases, which agent is most optimistic can vary: in their case, the identity of the most optimistic is determined by the public history of the performance of the asset, whereas for us it is because of asymmetric information.

A closely related paper is that of Izmalkov and Yildiz (2010). They make a primitive assumption similar to our assumption about optimism: Beliefs are all distorted in the same direction. They consider two-agent, two-action coordination games, and show that agents can be induced to take any rationalizable action—including risk-dominated ones—if the degree of optimism is high enough.³¹

Finally, Han and Kyle (2016) report a “contagious optimism” result in a CARA-normal asset pricing model. They study a *static* CARA-normal pricing game in which an agent, in equilibrium, conditions on the information revealed by a counterparty’s trading. While our game is designed to pick up higher-order average expectations, their result depends on different properties of higher-order expectations (certain kinds of hierarchies in which agents wrongly assume

³⁰Note that the chain need not have a unique ergodic distribution, but there is *an* ergodic distribution reached from any initial state.

³¹This observation illustrates a more general point of Weinstein and Yildiz (2007b) that any rationalizable action can be made uniquely rationalizable if a type is perturbed in the product topology.

common knowledge of the mean of an asset value). However, they similarly show that a small amount of optimism can give rise to arbitrarily high asset prices. Another difference is that their result is written for the two-agent case; any extension to many agents would require that the network be uniform, because trade takes place in centralized markets. Because they consider a world with normally distributed uncertainty, there is no upper bound on first-order expectations, and this allows contagious optimism to drive prices up without bound.

7.4.2. *Tightness.* We now construct a chain to show the bound of Lemma 2 is tight. Consider a chain with states $\{0, 1, \dots, m\}$. Let $f(s) = s$ and define:

$$B(s, s') = \begin{cases} (1 + \delta)/2 & \text{if } s' = s + 1 \leq m \\ (1 - \delta)/2 & \text{if } s' = s - 1 \geq 0 \\ \varepsilon/m & \text{if } s = m, s' = 0 \\ 1 - \varepsilon/m & \text{if } s = 0, s' = m \\ 0 & \text{otherwise.} \end{cases}$$

That is, the chain performs a biased random walk when it is in the interval $\{0, 1, \dots, m - 1\}$; if it is at state $s = m$, the “highest” state, it returns with small probability to $s = 0$, the lowest state. While we have described B as a Markov chain, it can be realized as an interaction structure. If the number of states is even, take the network to be two agents, each placing weight 1 on the other, and construct beliefs to make B as described.

It is easy to compute that, up to constant factors, this chain achieves the lower bound of 2. Thus, it really is necessary that ε (pessimism) be bounded as in the formula of the lemma relative to the guaranteed “optimistic drift” δ .

8. TYRANNY OF THE UNINFORMED

In Proposition 3, we gave a sufficient condition (common prior on signals) under which the consensus expectation is the centrality-weighted average of agents’ prior expectations. In this section, we will find conditions on the information structure under which the consensus expectation is (almost) equal to *one* agent’s expectation. That is, rather than influence being shared according to network centrality, it will all be allocated to one agent, in a way that will depend on the information structure. In particular, it will turn out to be the *least informed* agent who accumulates influence.

To motivate these results, we can again consider some extreme cases. First, suppose that one agent is completely ignorant and has no private information, while other agents know the state perfectly. The agents other than the ignorant agent will have degenerate interim beliefs, so nothing about their priors can matter for iterated expectations or the consensus. Thus, if anyone's ex ante beliefs play a role in determining consensus expectations, it must be those of the least informed agent. It turns out that the consensus expectation is simply *equal* to the ignorant agent's prior expectation of y . A simple way to see this is to note that, because the ex ante beliefs of the informed agents don't matter, we may as well take them to be equal to the prior of the ignorant agent; then the conclusion follows by Proposition 1 on the common prior. By continuity, our result continues to hold if the ignorant agent has *almost* no information and the other agents have *almost* perfect information.

Surprisingly, this conclusion remains true when the ignorant agent is only *relatively* ignorant, and when his beliefs are not public as they were in the toy example. The ignorant agent may possess very precise private information about the state. But if others have even more precise (i.e., less noisy) private information, then their priors will still not matter, and only the relatively ignorant agent's priors will determine the consensus expectation.

We now present the statement and proof of the result, and then discuss it and compare it with related results in Section 8.4.

8.1. Common Interpretation of Signals Framework. Fix a complete Γ , i.e., one such that $\gamma^{ij} > 0$ whenever $i \neq j$. We specialize to a framework that we call *common interpretation of signals*, following the terminology of Kandel and Pearson (1995) and Acemoglu, Chernozhukov, and Yildiz (2016a). There is a state $\theta \in \Theta$ that is drawn by nature. Each agent receives conditionally independent signals about it according to a full-support distribution $\eta^i(\cdot | \theta) \in \Delta(T^i)$; these distributions are common knowledge. However, the agents have different full-support priors, $\rho^i \in \Delta(\Theta)$, over the state space. Combined with the conditional distributions encoded in the η^i , these uniquely define a prior distribution over $\Theta \times T$. We denote by \mathbf{E}^{ρ^i} the corresponding prior expectation operator. These primitives also induce in each agent, via Bayes' rule, an interim belief function; for each $t^i \in T^i$, there is a distribution $\pi^i(\cdot | t^i)$ over both the state and over others' signals.

Definition 4. We say that η^i is *at most ε -noisy* if: for every $\theta \in \Theta$, there is exactly one signal t_θ^i satisfying $\eta^i(t_\theta^i) \geq 1 - \varepsilon$, and this t_θ^i also satisfies $\eta^i(t_\theta^i | \theta') \leq \varepsilon$ for all $\theta' \neq \theta$.

This condition requires that for any θ , there is exactly one signal t^i that i receives with very high probability conditional on θ being realized; moreover, no two different θ, θ' can be associated with the same such signal.

Definition 5. We say that η^i is *uniformly at least δ -noisy* if, for every $\theta \in \Theta$ and $t^i \in T^i$, the inequality $\eta^i(t^i | \theta) \geq \delta$ holds.

This condition says that each signal has at least δ probability of being observed under each state, limiting the amount of information that can be inferred from any signal.

8.2. Sufficient Conditions for Tyranny of the Uninformed. Before stating the main proposition, we introduce some quantities that will figure in it. Let $\gamma_{\min} = \min_{i \neq j} \gamma_{ij}$ be the smallest off-diagonal entry of Γ , which is positive by assumption. Let ρ_{\min}^i be the minimal probability assigned to any $\theta \in \Theta$ by the prior $\rho^i \in \Delta(\Theta)$ of agent i . Let $\rho_{\min} = \min_i \rho_{\min}^i$ be the minimum of all of these, across agents. Finally, let $y_{\max} = \max_{\theta \in \Theta} |y(\theta)|$.

Proposition 5. *Suppose that for some $\delta \in (0, 1)$ and $\varepsilon \in (0, 1/2)$,*

1. η^1 is uniformly at least δ -noisy
2. η^i for all $i \neq 1$ is at most ε -noisy.

Then

$$(17) \quad |c(y; B_\pi, F_\pi) - \mathbf{E}^{\rho^1}[y]| \leq \frac{4|\Theta||S|^2}{(\gamma_{\min}\rho_{\min})^2} \cdot \bar{y} \cdot \frac{\varepsilon}{\delta}.$$

This bound is designed for cases where ε is much smaller than δ . It says that if agent 1's information is at least δ -noisy, while all others' information is quite precise (at most ε -noisy), then the difference between the consensus expectation of y and agent 1's expectation of y is small: The upper bound is linear in ε/δ . The constants depend on the sizes of the state space and the signal space S , and on the minimum network and belief weights in the denominator.

We could formulate a version of Proposition 5 without requiring the rather strong assumption of full support of the conditional distributions $\eta^i(\cdot | \theta)$ that is implied by Proposition 5. This is discussed below in Section 8.3.1, once we have a bit more notation.

8.3. Key Steps in the Proof of Proposition 5. We will analyze the consensus expectation in the situation of Proposition 5 by analyzing the interaction structure B and its stationary distribution, p . Indeed, the analysis here is intended as our main illustration of the value of reducing informational questions to questions about the Markov chain corresponding to the interaction structure.

The key insight in proving Proposition 5 is to construct an artificial signal structure $\hat{\eta}$ in which all agents except agent 1 are certain of what θ is. This is done by rounding the signal probabilities $\eta^i(r^i | \theta)$ for $i \neq 1$ to 0 or 1. Along with the priors $(\rho^i)_{i \in N}$ over θ that are part of the setup, this induces an artificial information structure $\hat{\pi} = (\hat{\pi}^i)$. We let $\hat{B} = B_{\hat{\pi}, \Gamma}$.

The proof then proceeds in three steps. First we prove that p , the stationary distribution of B , is well-approximated by that of \hat{B} , which is denoted by \hat{p} . Second we claim that $\hat{\pi}$ can be viewed as having a common prior (corresponding to agent 1's prior beliefs). This is because only agent 1 is uncertain under $\hat{\pi}$ about θ , and so the ex ante beliefs of the others about θ can make no difference; indeed, it can be shown that the other agents' interim beliefs are compatible with agent i 's prior. Thus the consensus expectation of y under \hat{B} is equal to $\mathbf{E}^{\rho^1}[y]$. Finally, we combine these facts to derive the proposition. We carry out these steps below, deferring technical details to Appendix C.4.

The key technique in this argument deserves some extra comment. In the step where we approximate p by \hat{p} , we apply a result of Cho and Meyer (2000) on perturbations of Markov chains. This result, loosely speaking, says the following: As long as the changes in weights in going from \hat{B} to B are small relative to the reciprocal of the *maximum mean first passage time* (MMFPT³²) of \hat{B} , then p is close to \hat{p} . In our application, the change in the interaction structure (corresponding to interim beliefs about θ of the relatively informed agents $i \neq 1$ changing from “slightly uncertain” to “fully certain”) is of order ε , and that is why ε appears in the numerator of the bound in Proposition 5. In the situation of Proposition 5, the MMFPT is of order $1/\delta$, the inverse of the lower bound on the uninformed agent's noise. (That is why δ appears in the denominator in the bound of Proposition 5.) But the technique we have outlined applies more broadly, in any setting where the size of the perturbation to the interaction structure can be bounded relative to the MMFPT. This could be used to weaken the assumptions of Proposition 5, for example to cover cases where noise does not have full support or the network is not complete—see Section 8.3.1 below.

We carry out the details of the proof in Section C.4.

8.3.1. *The Case Where No Player Is at Least δ -Uncertain*. Suppose we did not assume that player 1 is at least δ -uncertain, which entails the strong assumption that there is a lower bound on the conditional probability of seeing any one of his signals, given any possible state. Then

³²The MMFPT in the interaction structure \hat{B} is defined to be the maximum expected time it takes to get from one state to another in the physical process of Section 4.3. It is a measure of the connectedness of \hat{B} as a network; in belief terms, it is a measure of the maximum number of iterations required for there to be contagion of higher-order beliefs between the two “farthest” states in S .

we would define the *uncertainty*, δ , of player 1's information as the minimum of $\eta^i(t^i | \theta)$ over all its nonzero values (as t^i and θ range over all possibilities). Along the same lines, we might wish to relax the assumption that Γ is complete, with every player putting weight on every other. We now discuss how the general principles of our argument would go through and the nature of the subtleties that would arise.

As mentioned in the sketch of the proof above, what really matters in the proof is MMFPTs in \widehat{B} . Assuming \widehat{B} is irreducible, we can still bound these in terms of δ even with the weaker assumptions just discussed. But—as an examination of our bounds on the MMFPT shows—the bounds will involve path lengths in \widehat{B} : the number of steps in \widehat{B} that must be taken to link any two states. Thus, rather than a bound of the form $\max_{z \neq z'} M_{\widehat{B}}(z, z') \leq C\delta^{-1}$, we might have something like $C\delta^{-5}$ in the bound on the right-hand side. The exponent will depend both on the information structure and on Γ . In the end, this will translate into a difference on the right-hand side of (17) in Proposition 5. Indeed, we conjecture that the ratio ε/δ would be replaced by $C\varepsilon/\delta^\kappa$ for a number κ that is increasing in the maximum path length in \widehat{B} . Moreover, this adjustment would be necessary: In the more general setting we are discussing here, it is not possible to write a bound analogous to (17) that depends on ε and our generalized δ only through ε/δ .

While a full exploration of these elaborations is beyond the scope of the present work, our point is to say: (i) the MMFPT technique discussed here does cover less restrictive assumptions on information than we made for the main result; and (ii) the topology of connections among types in the interaction structure \widehat{B} matters in interesting ways for more general results.

8.4. Interpretation and Discussion.

Why Focus on the Least-Informed? The results of this section may seem paradoxical. In models of coordination on a network motivated by organizational questions, a common result is that agents have an incentive to focus on more informed agents, in the sense of paying more attention to them or putting more weight on their signals; see, for example, Calvó-Armengol, Martí, and Prat (2015) and Herskovic and Ramos (2015). Part of the reason for the difference in our result is that asymmetric information gets washed out in our limit of higher-order expectations (recall Proposition 1), rather than being learned or aggregated, and this makes the forces determining influence different.

The Least-Informed Become Effectively More Central. It is also interesting to compare the result on the tyranny of the uninformed with the result of Proposition 3 in Section 6, where we

showed that, under a consistency condition on beliefs, it is an agent's centrality that determines his influence. However, as is seen in our simple benchmark example above, for sufficiently well-informed agents, their priors cannot possibly matter, no matter how central they are in the network; an implication of our result is that asymmetries of this kind cannot be reconciled with common priors over beliefs/signals.

We can get some further intuition for our result by expressing it in the language of our applications. Suppose that agents are making investment decisions, but with strategic complementarities in those decisions. We might say that there is *confidence* in the economy if positive expectations about others' investment are driving agents to invest more. In other words, confidence is founded on common perceptions of what is going on in the economy. Ignorant agents' (prior) views will have a disproportionate role in determining confidence. Similarly, in asset markets with frequent re-trading and random matching, assets will sometimes pass through the hands of ignorant agents. Their views will form a focal point around which market expectations will form.

A Subtlety in the Meaning of "Informed." To interpret and apply our results, it is important to remember that "prior" really means "belief conditional on public information only." (See Section A.1.2, where we note that all our analysis is conditional on public information.) In view of this, we call an agent "uninformed" if the beliefs of that agent are not sensitive to his private information once we have conditioned on public information. This might not correspond to other natural senses of "uninformed," so the distinction is worth keeping in mind.

Least-Informed versus Public. We note in closing that this result is very different from the familiar case of coordinating on something public or "commonly understood" in a beauty contest. The less informed agent's information is *not* public or approximately public. Indeed, in our example, individuals' signals are conditionally independent given the state. A highly informed player's signal provides very good information about the *external state*, but no information better than the prior about the *signals* of the others who are badly informed.

Moreover, in contrast to the standard case of coordinating on a public signal, our result does not hinge on a qualitative matter of determining which information is public (something that, actually, is held constant as we vary the noise rates). It is rather a quantitative matter of how low the noise rate of the relatively informed player must be in order for it to "wash out" of the (public) consensus expectation. As discussed in Section 8.3.1, this can depend in a subtle way on priors and the information structure. In particular, it can happen that the noise of the more

informed player is arbitrarily smaller than the noise of the less informed, and nevertheless the structure of the smaller noise is decisive for the consensus expectation. How small the noise must be in order not to matter depends, in general on the network, priors, and information structure, through quantities that we have described.

9. CONCLUDING DISCUSSION

In Appendix A, we give some detailed discussions and interpretations of important assumptions we have made. Here we briefly summarize some of the key points of those discussions.

Joint Connectedness (Section A.1). The assumption of joint connectedness was a key maintained assumption in our results. In this section, we relate it to beliefs and the network, discussing how it relates to the connectedness of the network and the absence of public events (joint connectedness implies both properties but is not equivalent to their conjunction). We also discuss what can be done without joint connectedness. This comes down to the standard analysis of a Markov matrix where not all states are recurrent.

Heterogeneous Self-Weights (Section A.2). In the linear best-response game, we assumed that all agents put a common weight β on others' actions. If they do not, then our analysis can be adapted to this case by changing the network. In particular, we will show how the linear best-response game with weights $(\beta^1, \dots, \beta^{|N|})$ and network Γ has the same solution as the game with a common coordination weight $\hat{\beta}$ (that depends on $(\beta^1, \dots, \beta^{|N|})$) and an alternative network $\hat{\Gamma}$. The diagonal entries of this network capture the variation in self-weights. This transformation lets us analyze the case of heterogeneous self-weights with our formulas. We give interpretations in terms of both the financial market and the game.

Separability and Connection to Samet (1998a) (Section A.3). In Section 5, we showed that—fixing the information structure and network—there are strictly positive *pseudopriors* $(\lambda_{\pi, \Gamma}^i)_{i \in N}$ such that $c(y; \pi, \Gamma) = \sum_i e^i \mathbf{E}^{\lambda_{\pi, \Gamma}^i} y$. At the same time, we made the observation—which we have now made explicit in the subscripts of λ^i —that those pseudopriors may depend on both the information structure π and the network Γ . We say an information structure π satisfies *separability* if the pseudopriors depend *only* on the information structure. Section 6 that a common prior on signals is sufficient for separability. The environments studied in of Contagion of Optimism and Tyranny of the Uninformed, on the other hand, are not, in general, consistent with separability. In Golub and Morris (2017) we give a necessary and sufficient

condition for separability describing the boundary between these cases exactly, and Section A.3 sketches the essential ideas.

Our results in both this paper and Golub and Morris (2017) relate closely to and build on those of Samet (1998a). The similarity is that, like him, we give conditions under which limiting properties of higher-order expectations depend only on a summary statistic of the information structure (e.g., the pseudoprior). Section A.3 discusses the difference in the results and techniques in detail.

Ex Ante and Interim Interpretation (Section A.4). We take an ex ante perspective in our analysis: At an initial date, agents have prior beliefs—and no information—about a state of the world. They then receive information and update their beliefs. We can interpret the results as answering the question: *How does the consensus expectation change after agents observe their signals?* Our results give conditions under which: (i) the beliefs *do not change* (under common priors over signals); (ii) they change to the most optimistic conceivable beliefs (contagion of optimism); (iii) they change to the beliefs of the least informed (tyranny of the uninformed).

Though we take an ex ante view throughout, consensus expectations can be seen from a purely interim perspective. Indeed, consensus expectations depend only on agents' interim beliefs (across all possible types)—i.e. on the belief functions π . We discuss how certain main results would look if we were to stick to a purely interim interpretation. As in our discussion of separability above, there is a close connection to the characterization of the common prior assumption in purely interim terms given by Samet (1998a). We highlight both how our results can be related to his, and also where a prior perspective makes them distinct. While contagion of optimism has purely interim interpretation, tyranny of the uninformed depends on assumptions about priors and has no simple interim interpretation.

Agent-Specific Random Variables (Section A.5). Our focus throughout the paper has been on agents' higher-order expectations of a random variable of common concern, y . But an equally interesting application considers a case where agents have different preferred actions (which correspond to the different random variables y^i) in the absence of coordination motives, and where one's network neighbors also influence one's choice, with linear best responses assumed (Ballester, Calvó-Armengol, and Zenou, 2006; Calvó-Armengol, Martí, and Prat, 2015; Bergemann, Heumann, and Morris, 2015b). This case can be embedded readily into our formalism. Indeed, we can define our $x^i(n)$ almost identically to capture this case. This embodies an equivalence between different priors over the external states and caring about different random

variables—an equivalence which does *not* extend to higher-order beliefs, as we explain. In discussing this connection, we highlight how our results relate to Calvó-Armengol, Martí, and Prat (2015) and Bergemann, Heumann, and Morris (2015a).

A related point is that there need not be common perceptions or complete information of the network weights γ^{ij} . By allowing these to depend on individuals’ types, we can embed incomplete information about the network into our framework.

Static Higher-Order Expectations, Dynamic Conditional Expectations, Behavioral Learning, and the DeGroot Model. We study higher-order average expectations of a random variable in this paper. These higher-order expectations are interpreted as being computed at a moment of time. We can call them “static higher-order expectations,” as they are properties of the agents’ static beliefs and higher-order beliefs at that moment. All the iteration of computing higher-order expectations occurs “in the agents’ minds” rather than in an interactive dynamic process unfolding over time.

These static higher-order expectations can be contrasted with agents’ “dynamic conditional expectations”: the beliefs formed via a dynamic process of updating expectations *after* observing other agents’ conditional expectations up to that point. In this section, we will use this dichotomy to discuss connections with some important related literatures.

DeGroot (1974) suggested a behavioral model where, at each stage in a process, each of many agents takes a weighted average of the beliefs or estimates of his neighbors. He interpreted this as a heuristic procedure according to which statisticians might average their own estimates or beliefs with the estimates or beliefs of others whose opinions they respect, toward the goal of reaching a reasonable consensus.³³ In the DeGroot model, the vector of agents’ estimates at stage n is $x(n) = \Gamma^n x(0)$, where (as in our model) Γ is an exogenous, fixed stochastic matrix corresponding to the weights agents assign to various others. Under the classical interpretation, the DeGroot model is a dynamic process, where agents start out with different estimates (perhaps based on their private information) and then updating occurs according to a behavioral rule. Economic foundations and implications of this process have been developed by DeMarzo, Vayanos, and Zwiebel (2003), Golub and Jackson (2010), and others.

³³This work grew out of studying aggregation procedures for statistical estimates. Lehrer and Wagner (1981) worked on a related model, seeking normative foundations for agents’ weights in the consensus, based on the problem of aggregating views in a network of peers. Friedkin and Johnsen (1999) studied versions of this model in which each agent persistently weights a fixed opinion, which can be interpreted as a personal ideal point—see Section A.5 for a version of this in our setting. See Golub and Sadler (2016), whose Section 3.5.1 we have partly paraphrased here.

Mathematically, the complete-information special case of our static higher-order expectations model is isomorphic to the classic DeGroot model, in the sense that equation (10) for updating the vector of static higher-order expectations,³⁴ $x(n) = \Gamma^{n-1}Fy$, looks very much like a DeGroot rule of the form $x(n) = \Gamma^n x(0)$. But it has a different interpretation. Our agents start out with different priors, captured by Fy . In the dynamic interpretation, $x(2)$ corresponds to taking the weighted average of neighbors' first-period beliefs. In the static interpretation, $x(2)$ contains agents' expectation of the average first-order expectations of others. In this static interpretation, agents' higher-order expectations are fully Bayesian but based on heterogeneous priors and no asymmetric information, with weights (i.e., the network Γ) which are taken as exogenous.

Indeed, the general incomplete-information version of our model can also be related to the DeGroot model. If we draw a parallel where the types in our model correspond to DeGroot agents, and $x(1)$ is taken to be the profile of initial estimates, then the "DeGroot estimate" of a given *type* at stage n is the n^{th} -order iterated average expectation of that type in our model. In this way, our model can be viewed as an alternative interpretation of DeGroot's formulas.

Despite the formal similarity, substantively, the two interpretations differ very significantly in how they answer a key question in the DeGroot model literature: How does the network Γ affect the ultimate consensus? Recall that in the DeGroot model, the consensus is a weighted average of the agents' initial opinions, with the weight of an agent equal to her eigenvector centrality. (Thus, in DeGroot's model, if high-centrality agents have high first-order expectations, the consensus will also be high.) There is an analogous centrality formula in our setting: Proposition 1. Despite this, in our model, under the common prior assumption, there is *no* interesting dependence of outcomes on Γ , even when the network gives some agents very large network centrality: Higher-order average expectations will always converge to the common prior estimate, independent of the network. It is only when agents have *heterogeneous* priors that the network matters in our model. Thus, whereas in the dynamic learning DeGroot model, the updating implies that centrality always matters, the additional structure present in our model says that it matters (to our outcomes) only in specific circumstances, and not under the common prior assumption.

There is another approach to DeGroot's questions that is different from his own behavioral model and from our interpretation of his equations sketched above. That approach is to study standard Bayesian agents learning dynamically from each other's beliefs, making Bayesian inferences at each stage. In this case we get a very different updating process. Geanakoplos and

³⁴Note that under complete information, $B = \Gamma$.

Polemarchakis (1982) considered this updating process under the common prior assumption. Their finding—in a finite-state model—was that posteriors would converge and there would be common certainty of posteriors in the limit. This model has been generalized in various directions. For example, Parikh and Krasucki (1990) considered the case when one observes posteriors of only some neighbors. Recently, Rosenberg, Solan, and Vieille (2009) and Mueller-Frank (2013) have explored this model further, in the case where agents’ announcements don’t fully reveal their posteriors. This literature provides a fairly rich understanding of dynamically updating conditional expectations with common priors and asymmetric information on a general unweighted graph. Note that it contrasts sharply with our analysis; in the model we have studied in this paper, private information gets “washed out” rather than aggregated as we take n to the infinite limit.

REFERENCES

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016a): “Fragility of Asymptotic Agreement under Bayesian Learning,” *Theoretical Economics*, 11, 187–225.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2016b): “Networks, Shocks, and Systemic Risk,” in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.
- ALLEN, F., S. MORRIS, AND H. S. SHIN (2006): “Beauty Contests and Iterated Expectations in Asset Markets,” *Review of Financial Studies*, 19, 719–752.
- AUMANN, R. (1976): “Agreeing to Disagree,” *Annals of Statistics*, 4, 1236–1239.
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): “Who’s who in Networks. Wanted: the Key Player,” *Econometrica*, 74, 1403–1417.
- BANERJEE, S. AND I. KREMER (2010): “Disagreement and Learning: Dynamic Patterns of Trade,” *The Journal of Finance*, 65, 1269–1302.
- BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2015a): “Information and Volatility,” *Journal of Economic Theory*, forthcoming.
- (2015b): “Networks, Information and Volatility,” Yale University and Princeton University Working paper.
- BLUME, L., W. BROCK, S. DURLAUF, AND R. JAYARAMAN (2015): “Linear Social Interaction Models,” *Journal of Political Economy*, 123, 444–496.
- BRANDENBURGER, A. AND E. DEKEL (1993): “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory*, 59, 189–198.

- CALVÓ-ARMENGOL, A., J. DE MARTÍ, AND A. PRAT (2009): “Endogenous communication in complex organizations,” in *SAET CONFERENCE*, Citeseer.
- CALVÓ-ARMENGOL, A., J. MARTÍ, AND A. PRAT (2015): “Communication and influence,” *Theoretical Economics*, 10, 649–690.
- CHO, G. E. AND C. D. MEYER (2000): “Markov chain sensitivity measured by mean first passage times,” *Linear Algebra and its Applications*, 316, 21–28.
- DE MARTÍ, J. AND Y. ZENOU (2015): “Network games with incomplete information,” *Journal of Mathematical Economics*, 61, 221–240.
- DEGROOT, M. H. (1974): “Reaching a Consensus,” *Journal of the American Statistical Association*, 69, 118–121.
- DEKEL, E., D. FUDENBERG, AND D. K. LEVINE (2004): “Learning to Play Bayesian games,” *Games and Economic Behavior*, 46, 282–303.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *The Quarterly Journal of Economics*, 118, 909–968.
- DUFFIE, D. AND G. MANSO (2007): “Information Percolation in Large Markets,” *The American Economic Review*, 97, 203–209.
- FRIEDKIN, N. E. AND E. C. JOHNSEN (1999): “Social Influence Networks and Opinion Change,” *Advances in Group Processes*, 16, 1–29.
- GEANAKOPOLOS, J. D. AND H. M. POLEMARCHAKIS (1982): “We Can’t Disagree Forever,” *Journal of Economic Theory*, 28, 192—200.
- GOLUB, B. AND M. O. JACKSON (2010): “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–49.
- GOLUB, B. AND S. MORRIS (2017): “Higher-Order Expectations,” Available at SSRN: <http://ssrn.com/abstract=2979089>.
- GOLUB, B. AND E. SADLER (2016): *Learning in social networks*, Oxford University Press, chap. 19, 504–542.
- GRUNDY, B. D. AND M. MCNICHOLS (1989): “Trade and the Revelation of Information through Prices and Direct Disclosure,” *Rev. Financ. Stud.*, 2, 495–526.
- HAN, J. AND A. KYLE (2016): “Speculative Equilibrium with Differences in Higher-Order Beliefs,” .
- HARRISON, J. M. AND D. M. KREPS (1978): “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations,” *The Quarterly Journal of Economics*, 92, 323–336.

- HARSANYI, J. C. (1968): “Games with incomplete information played by ‘Bayesian’ players, Part III. The basic probability distribution of the game,” *Management Science*, 14, 486–502.
- HELLMAN, Z. (2011): “Iterated expectations, compact spaces, and common priors,” *Games and Economic Behavior*, 72, 163–171.
- HERSKOVIC, B. AND J. RAMOS (2015): “Acquiring Information Through Peers,” Mimeo., NYU.
- IZMALKOV, S. AND M. YILDIZ (2010): “Investor Sentiments,” *American Economic Journal*.
- JACKSON, M. O. (2008): *Social and Economic Networks*, Princeton, NJ: Princeton University Press.
- KANDEL, E. AND N. D. PEARSON (1995): “Differential Interpretation of Public Signals and Trade in Speculative Markets,” *Journal of Political Economy*, 103, 831–872.
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*, Macmillan.
- KOZITSKY, Y., D. SHOIKHET, AND J. ZEMÁNEK (2013): “Power convergence of Abel averages,” *Archiv der Mathematik*, 100, 539–549.
- LEHRER, K. AND C. WAGNER (1981): *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*, vol. 21, Springer Science & Business Media.
- LEWIS, D. (1969): *Convention: A Philosophical Study*, Harvard University Press.
- MALAMUD, S. AND M. ROSTEK (2016): “Decentralized Exchange,” Mimeo.
- MEYER, C. D., ed. (2000): *Matrix Analysis and Applied Linear Algebra*, Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
- MILGROM, P. AND J. ROBERTS (1990): “Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities,” *Econometrica*, 58, 1255–1277.
- MORRIS, S. (1994): “Trade with Heterogeneous Prior Beliefs and Asymmetric Information,” *Econometrica*, 62, 1327–1347.
- (1997): “Interaction games: A unified analysis of incomplete information, local interaction and random matching games,” Santa Fe Institute Working Paper.
- (2002a): “Notes on Iterated Expectations,” <http://www.princeton.edu/~smorris/pdfs/hoef.pdf>, Princeton University Working paper.
- (2002b): “Typical Types,” Available at princeton.edu/~smorris/pdfs/typicaltypes.pdf.
- MORRIS, S. AND H. SHIN (2002): “Social Value of Public Information,” *American Economic Review*, 92, 1521–1534.

- MUELLER-FRANK, M. (2013): “A General Framework for Rational Learning in Social Networks,” *Theoretical Economics*, 8, 1–40.
- MYERSON, R. B. (1997): *Game Theory*, Cambridge, Mass.: Harvard University Press.
- NEHRING, K. (2001): “Common priors under incomplete information: a unification,” *Economic Theory*, 18, 535–553.
- PARIKH, R. AND P. KRASUCKI (1990): “Communication, Consensus, and Knowledge,” *Journal of Economic Theory*, 52, 178–89.
- RADNER, R. (1962): “Team Decision Problems,” *The Annals of Mathematical Statistics*, 857–881.
- ROSENBERG, D., E. SOLAN, AND N. VIEILLE (2009): “Informational Externalities and Emergence of Consensus,” *Games and Economic Behavior*, 66, 979–994.
- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge’,” *American Economic Review*, 79, 385–391.
- SAMET, D. (1998a): “Iterated Expectations and Common Priors,” *Games and Economic Behavior*, 24, 131–141.
- (1998b): “Common Priors and Separation of Convex Sets,” *Games and Economic Behavior*, 24, 172–174.
- SHIN, H. AND T. WILLIAMSON (1996): “How Much Common Belief is Necessary for a Convention,” *Games and Economic Behavior*, 13, 252–268.
- STEINER, J. AND C. STEWART (2014): “Price Distortions in High-frequency Markets,” *CEPR Discussion Paper No. DP9817*.
- UI, T. (2009): “Bayesian Potentials and Information Structures: Team Decision Problems Revisited,” *International Journal of Economic Theory*, 5, 271–291.
- WEINSTEIN, J. AND M. YILDIZ (2007a): “Impact of higher-order uncertainty,” *Games and Economic Behavior*, 60, 200–212.
- (2007b): “A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements,” *Econometrica*, 75, 365–400.
- YOUNG, H. P. (1996): “The Economics of Convention,” *The Journal of Economic Perspectives*, 105–122.
- ZENOU, Y. (2016): “Key Players,” in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.

APPENDIX A. DISCUSSION OF ASSUMPTIONS AND VARIANTS OF OUR RESULTS

We now discuss robustness and extensions of our results (Sections A.1 and A.2), as well as their context and broader implications (Sections A.3 through 9). More technical issues are postponed to Appendix B.

A.1. Joint Connectedness. An assumption maintained throughout was a joint connectedness condition (recall Section 2.5), which amounts to the interaction structure B being irreducible. In the present subsection, we no longer treat this condition as a maintained assumption, and examine its content and what can be said without it. Proposition 6 reviews a characterization of the irreducibility condition: It is equivalent to the agent-type vector p being strictly positive. We then relate the condition to properties of the primitives Γ and π . Finally, we discuss results that hold under weakenings of the assumption.

A.1.1. Relations to Beliefs and the Network. Some key properties of the network and beliefs will feature in our characterization of irreducibility. A network Γ is *complete* if $\gamma^{ij} > 0$ for all i and j . Beliefs π have *full support marginals* if $\pi^i(t^j | t^i) > 0$ for all agents i and j , and all signals $t^i \in T^i, t^j \in T^j$. Event $G \subseteq T$ is a *product event* if $G = \prod_{i \in N} G_i$, where $G^i \subseteq T^i$ for each i . Say that a product event $G = \prod_{i \in N} G_i$, is a *public* or *closed* event (under beliefs π) if, for each agent i and each signal $t^i \in G^i$, the following implication holds for any $t^{-i} \in T^{-i}$:

$$\pi^i(t^{-i} | t^i) > 0 \quad \implies \quad (t^i, t^{-i}) \in G.$$

A public or closed event is one that, when it occurs, is common certainty among all the agents: For any observed signal, no probability is assigned any signal outside the event. *Beliefs π are connected* if \emptyset and T are the only public events. This corresponds to the notion of no (nontrivial) common certainty: Every nontrivial product event has a connection (via beliefs placed by some agent) to states outside itself. A subset of agents $J \subseteq N$ is *closed* if $i \in J$ and $\gamma^{ij} > 0$ implies $j \in J$. A network Γ is *connected* if \emptyset and N are the only closed subsets of the agent set N . Recall that a network is a *complete* if $\gamma^{ij} > 0$ for all i and j .

The properties mentioned so far are restrictions on either the beliefs or the network, but not both. The property of joint connectedness from Section 2.5 is a joint restriction on both. The following result relates the two sorts of conditions.

Proposition 6. *The matrix B is irreducible if and only if beliefs and the network are jointly connected. Necessary conditions for this are:*

1. *Beliefs are connected.*




FIGURE 3. An example illustrating that imposing connectedness of the network and of beliefs is not sufficient to ensure joint connectedness, i.e. irreducibility of the interaction structure B .

2. *The network is connected.*

Sufficient conditions for this are:

1. *The network is complete and beliefs are connected.*
2. *The network is connected and beliefs have full support marginals.*

Proof. The “if and only if” part is just a rewriting of the statement that there are no nonempty, proper closed communicating classes in the Markov process corresponding to B . The two sufficient conditions are strengthenings of this property. \square

The following example illustrates that requiring a connected network and connected beliefs separately is not sufficient for irreducibility.

Example. Suppose that there are three agents and each agent observes one of two signals, so that $T^i = \{a^i, b^i\}$. The network is given by a cycle,

$$\Gamma = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix},$$

and the information structure is such that each agent i is sure that agent $i + 1$ has observed the same signal and that agent $i - 1$ has observed a different signal (under mod 3 arithmetic for agent indices). Figure 3 illustrates the network and the information structure. In this example, the network is connected, and the beliefs are connected; nevertheless, B is not irreducible, because beliefs and the network are not jointly connected. The subset $\{a^1, a^2, a^3\}$ places no weight in B on its complement.

The conditions we have discussed have placed restrictions on beliefs directly, rather than on observable consequences. In Appendix B.2, we give a “no trade” behavioral characterization of irreducibility.

A.1.2. Consensus without Irreducibility. All our results have analogues when irreducibility fails. As we saw in Section 4.2, what matters for the limit of $x(n; y)$ as $n \rightarrow \infty$ is the behavior of B^n . This can be characterized quite generally based on the graph described in Section 5. The general result can be found in many textbooks (e.g. Meyer, 2000, Section 8.4), and we summarize it informally. First, consider the set S_A , defined as the set of absorbing states in S according to the transition matrix B . For such t^i , the analysis of $x(n; y)$ can proceed exactly as in Section 4.2, restricting B to the maximal strongly connected component containing t^i .

The simplest case is when S can be partitioned into several strongly connected components (so $S = S_A$) with B having no edges between these components. This occurs, for instance, if there are exactly two public (product) events. Then the analysis can be done on each of these components separately. That is, the analysis can be done *conditional on public information*. More generally, when there are public events, our assertion that the consensus expectation is nonrandom (recall Section 4.2) really means that it is nonrandom conditional on the public event that has occurred (and which, by definition of its being public, is common knowledge).

Now suppose there *are* some nonabsorbing states. For each non-absorbing state $t^i \notin S_A$, the corresponding row of B^∞ is a distribution that allocates mass (in a particular way) across the set S_A of absorbing states.

An important case is relevant to several of our discussions. When S_A consists of exactly one strongly connected component (though it may be a strict subset of S), we can refine our statements above to obtain the following generalization of Proposition 1:

Proposition 7. *If S_A has exactly one strongly connected component, the consensus expectation exists and*

$$(18) \quad c(y; \pi, \Gamma) = \sum_{t^i \in S_A} p(t^i) E^i[y | t^i],$$

where $p \in \Delta(S_A)$, called the vector of agent-type weights, is the stationary distribution of B_{S_A} (B restricted to S_A), i.e. the unique vector in $p \in \Delta(S_A)$ satisfying $pB_{S_A} = p$. Moreover, all entries of p are positive.

This specializes to Proposition 1 in case $S_A = S$. In general, the consensus expectation still exists and is unique, which is what we need for the examples of Section 7, where irreducibility fails to hold.

Our results do not hold if we relax our maintained finiteness assumption: in Appendix B.3 we report an example of Hellman (2011) showing this.

A.2. Heterogeneous Coordination Weights and Their Relation to Self-Weights in the Network. In the linear best-response game, we assumed that all agents put a common weight β on others' actions, and studied the limit $\beta \uparrow 1$. We again maintain the assumption that Γ is irreducible and consider now a more general class of environments, characterized by (Γ, β, y) , where $\beta = (\beta^i)_{i \in N}$ is a profile of agent-specific weights. In the coordination game associated with such an environment, the linear best responses are given by

$$(19) \quad a^i = (1 - \beta^i)E^i y + \beta^i \sum_{j \neq i} \gamma^{ij} E^i a^j.$$

Paralleling our main study, we can ask what happens as $\beta^i \rightarrow 1$ simultaneously across i . As we show in this section, this issue is closely related to “self-weights” γ^{ii} in the network.

First, we consider some simple examples. Suppose $|N| = 2$, with the network

$$\Gamma = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

If $\beta^2 = 1$ and $\beta^1 < 1$, then in the limit $\beta^1 \uparrow 1$, iterating the (modified) best-response equation $a^i = (1 - \beta^i)E^i y + \beta^i \sum_{j \neq i} \gamma^{ij} E^i a^j$ shows that we would have a convention given by³⁵

$$\lim_{n \rightarrow \infty} [E^1 E^2]^n E^1 y.$$

On the other hand, suppose $|N| = 2$, $\beta^1 = 1$, and $\beta^2 < 1$. Now, in the limit $\beta^2 \uparrow 1$, we would symmetrically have a convention given by the (different) simple higher-order expectation

$$\lim_{n \rightarrow \infty} [E^2 E^1]^n E^2 y$$

Thus, with agent-specific self-weights—the corresponding convention will depend on the details of how the limit $(\beta^1, \dots, \beta^{|N|}) \rightarrow (1, \dots, 1)$ is taken—in particular, whose β^i converges faster to 1.

³⁵See Appendix B.1 for discussion of such limits, called *simple* higher-order expectations, and related calculations.

We briefly sketch how our analysis can be adapted to this case by changing the network. In particular, we will show how the linear best-response game with weights $(\beta^1, \dots, \beta^{|M|})$ and network Γ has the same solution as the game with a common coordination weight $\hat{\beta}$ (that depends on $(\beta^1, \dots, \beta^{|M|})$) and an alternative network $\hat{\Gamma}$. The alternative network $\hat{\Gamma}$ may, in general, have nonzero self-weights ($\hat{\gamma}^i > 0$ for some i) even if Γ has zero self-weights ($\gamma^i = 0$ for all i). The transformation applies to any Γ , with or without positive entries on its diagonal.

Note that, in general, consensus expectations were defined allowing the possibility of positive self-weights in Γ , and our analysis of their basic properties (e.g., Propositions 1, 2 and 3) applies in that case as well. Positive self-weights are unnatural in some applications. For instance, in the game with the interpretation that each agent is a single player, one's best response cannot (by definition) depend on one's own action. On the other hand, there are other applications where self-weights have reasonable interpretations. For example, in the game with linear best responses, if we replace each agent with a continuum of identical agents (as we have done in the finance application), it would be natural to think of an agent caring about the average action of individuals like himself (i.e., in the same class). The same holds in the financial trading application, assuming there is a possibility that a player will sell into his own market (whose traders have the same expectation, and thus the same interim beliefs). In those cases, characterizing the average action of each population results in the equilibrium equations we have been studying, but with positive entries permitted on the diagonal of Γ . Formally, we can construct an analogue of the game in Section 3.1.1 and prove, paralleling part of Fact 1, that the game has a unique rationalizable strategy profile. (The proof is by the same contraction argument used to prove Fact 1.)

To see what that strategy profile is, we describe the transformation of any environment to one with an agent-independent common weight on others' actions:

Proposition 8. *Given an environment with a network Γ and a vector $\beta = (\beta^i)_{i \in N}$, define $\hat{\beta} = \max_{i \in N} \beta^i$ and define $\hat{\Gamma}$ by*

$$\hat{\gamma}^i = \frac{\hat{\beta} - \beta^i}{\hat{\beta}(1 - \beta^i)} \text{ and } \hat{\gamma}^j = \gamma^j(1 - \gamma^i).$$

For any y , the environments described by $(\hat{\Gamma}, \hat{\beta}, y)$ and (Γ, β, y) have identical play in their respective unique rationalizable strategy profiles.

A.3. Separability and Connection to Samet (1998a). In Section 5, we showed that—fixing the information structure and network—there are strictly positive *pseudopriors* $(\lambda_{\pi, \Gamma}^i)_{i \in N}$ such

that

$$(20) \quad c(y; \boldsymbol{\pi}, \Gamma) = \sum_i e^i \mathbf{E}^{\lambda_{\boldsymbol{\pi}, \Gamma}^i} y.$$

At the same time, we made the observation—which we have now made explicit in the subscripts of λ^i —that those pseudopriors may depend on both the information structure $\boldsymbol{\pi}$ and the network Γ . We say an information structure $\boldsymbol{\pi}$ satisfies *separability* if the pseudopriors depend *only* on the information structure:

Definition 6. The information structure $\boldsymbol{\pi}$ satisfies *separability* if there exists a profile $(\lambda_{\boldsymbol{\pi}}^i)_{i \in N}$ such that, for every irreducible Γ , we have $\lambda_{\boldsymbol{\pi}, \Gamma}^i = \lambda_{\boldsymbol{\pi}}^i$.

When separability holds, the asymmetric information affects the consensus expectation in an additively separable way, with each agent’s pseudoprior being weighted by his eigenvector centrality. Thus the incomplete information and the network can be analyzed separately. Networks matter only via the network centrality weights, and the information structure $\boldsymbol{\pi}$ affects only the pseudopriors $\boldsymbol{\lambda} = (\lambda_{\boldsymbol{\pi}}^i)_{i \in N}$.

We can illustrate the failure of separability with an example building on the one in Case II of Section 7.1. Suppose we have three agents arranged in a cycle as shown in Figure 2, with each considering his counterclockwise neighbor over-optimistic and his clockwise neighbor over-pessimistic. The network Γ in which all weight goes counterclockwise (i.e., $\gamma^{i, i-1} = 1$ for all i , with indices read modulo 3) gives the maximum consensus expectation. The network—call it Γ' —in which all weight goes clockwise (i.e., $\gamma^{i, i+1} = 1$ for all i , with indices read modulo 3) gives the minimum consensus expectation given the beliefs. Note that all agents are symmetric in each network. Thus, in both networks, by symmetry all agents have the same eigenvector centrality.

If separability held, then the two networks would have the same consensus expectation: We have just said that the centralities are the same across them, and that the information structure also remains the same if we reverse the direction of each link in the network. Since in fact the consensus expectation differs (indeed, differs as much as possible) across the two networks, we have a failure of the separability property.

We have already given one sufficient condition for separability in Section 6: a common prior on signals. Thus the example described above cannot be consistent with a common prior on

signals. In Golub and Morris (2017) we give a *necessary* condition for separability. We now informally report the condition in stages. First, note that for higher-order expectations, and therefore consensus expectations, the only beliefs about others that enter are *marginal* distributions over another’s signal. An agent is never concerned about the correlation in the signals of two or more others. This already suggests that the common prior assumption on signals is more than we need: Recall from Definition 3 that the common prior assumption on signals places strong restrictions on beliefs about *profiles* of signals. In fact, separability is implied by a weaker sufficient condition—one that requires priors about signals to agree only in their marginals on every agent’s signal. Like the existence of a common prior on signals, such a property puts no restrictions on agents’ beliefs about Θ conditional on signals, but it also relaxes substantially the restrictions on beliefs about signals.

In Golub and Morris (2017) we show that an even weaker condition is necessary and sufficient: We call it *higher-order expectation-consistency*. This condition specifies that we can find a “pseudoprior” for each agent with the property that those pseudopriors have the same expectations of all random variables in a certain class. The class consists of all higher-order expectations of random variables that are Θ -measurable. In effect, this necessary and sufficient condition imposes only those restrictions on higher-order beliefs that are relevant to higher-order expectations.

Our results in both this paper and Golub and Morris (2017) relate closely to and build on those of Samet (1998a). Samet showed that—if one fixes a state space and agents’ information on that state space (modeled via a partitional information structure)—then higher-order expectations of all random variables converge. If the common prior assumption holds, they converge to ex ante expectations under the common prior. Our Proposition 3 is a version of this result; critically, however, the reasoning is applied not to the whole state space but to the space of signal profiles. Samet also showed a converse: If all higher-order expectations of any random variable converge to the same number (depending on the random variable) regardless of the order in which they are taken, then the information structure must satisfy the common prior assumption. We do not have a converse in this paper. The characterization of the separability result in Golub and Morris (2017), which we have described above, is tight and thus is the closest analogue to Samet (1998a). There are many conceptual and technical issues that distinguish our notion of separability from the properties that matter in Samet (1998a); these differences are discussed in detail in Golub and Morris (2017).

There is also another important technical and methodological connection to Samet (1998a). We follow Samet (1998a) in representing information structures—as well as a network, which we add to the model—via a Markov process. However, we actually work with a different sort of Markov process than the one in Samet (1998a): Our Markov process operates on the *union* of agents’ types (which we denote by S), whereas Samet’s process applied to our questions operates on *profiles* of agents’ types T .³⁶ There are a number of reasons why the former Markov process (on S) is the appropriate one for our problem. First, it permits a unified or symmetric treatment of networks and asymmetric information, as discussed in Section 5.1. If one adds a network structure to Samet’s Markov formulation, networks and asymmetric information enter in very different ways in the formalism (see Golub and Morris (2017) for a presentation along these lines). Second, and relatedly, our formalism allows us to relate key elements of our analysis to results in the literature on network games. Finally, the Samet (1998a) approach works with matrices whose rows and columns are indexed by $\Omega = \Theta \times \prod_{i \in N} T^i$, which can be much larger than $S = \bigcup_{i \in N} T^i$; thus it can be convenient to have our formalism for doing explicit computations.

A.4. Ex Ante and Interim Interpretation. We take an ex ante perspective in our analysis: At an initial date, agents have prior beliefs—and no information—about a state of the world. This interpretation entails common certainty among the agents of everyone’s prior beliefs and the way agents update their beliefs.³⁷ In this section, we discuss some consequences of our ex ante approach and interim interpretations of our results

A.4.1. Dynamic Interpretation: The Arrival of Information. Under the ex ante perspective, the results of this paper can be given an explicitly dynamic interpretation. Before the arrival of information, there is symmetric information and, therefore, the consensus expectation is equal to the average of agents’ ex ante expectations, weighted by their eigenvector centralities (Section 5). In other words, if the agents had to select actions at that stage, this is what their actions would be equal to. One interpretation of our results is as an answer to the question, *How does the consensus expectation change after agents observe their signals?* We show that common prior over signals is a sufficient condition for *no* change in the consensus expectation (Proposition 3); second-order optimism causes the consensus expectation to increase to the

³⁶Samet works with a partitional formalism; Golub and Morris (2017, Section 6) restates our framework in that formalism.

³⁷Under an interim interpretation, it is without loss of generality to assume common certainty of types’ interim beliefs, i.e. how beliefs are updated: see Aumann (1976, p. 1237) and Brandenburger and Dekel (1993).

highest possible interim belief (Proposition 5); and, under the conditions in the results on the tyranny of the uninformed, the weights on agents' priors change from those induced by the network Γ to a degenerate vector which places all the weight on the least informed.

A.4.2. Interim Interpretation. Though we take an ex ante view throughout, consensus expectations, which emerge from agents' play at the interim stage, cannot depend on agents' ex ante beliefs about their own types. Thus consensus expectations must depend only on agents' interim beliefs (across all possible types). We have emphasized this in our notation, by first expressing the information structure in interim terms (i.e., via the beliefs $\pi^i(\cdot | t^i)$), and only then adding in ex ante beliefs over each agent's signals (the λ^i in Proposition 2).

Let us discuss how certain main results would look if we were to stick to a purely interim interpretation. First, results such as the representation of Proposition 2 would still make sense, but the λ^i would not be interpreted as anyone's beliefs. More substantially, consider Proposition 3. Let us focus on a particularly simple consequence of it: Under the common prior assumption on all of $\Theta \times T$, the consensus expectation of y is the prior expectation of y . To make sense of this in interim terms, we first have to say what the common prior assumption means in interim terms. Samet (1998a) has characterized that assumption as the condition that, for any random variable y , *higher-order expectations converge to the same number, independent of the order in which expectations are taken (as long as each agent appears infinitely often)*; this number can be identified with the common prior expectation of y . Thus an interim statement of the simple consequence of Proposition 3 is: Under the italicized condition, the consensus expectation of y is simply the prior expectation of y . This is natural: We can write the consensus expectation as an average of higher-order expectations, and the irreducibility of Γ ensures that all agents appear infinitely often in each of them. An interim version of Proposition 3 follows from very similar reasoning, with more attention paid to the network, and this is carried out in Golub and Morris (2017).

The consequence of Proposition 3 that we have discussed is similar to Corollary 1 but differs in an important way. Corollary 1 does not depend on there being a common prior on the whole state space (i.e., on signals and beliefs jointly); rather, it requires that the ex ante first-order expectations of y be the same across agents. This assumption does not have an obvious interim interpretation. Thus, the contrast between the "full common prior" result we have discussed in the previous paragraph and the actual result of Corollary 1 helps bring out where an ex ante perspective is important for us.

Our second-order optimism result (Proposition 4) is stated in terms of interim beliefs only (the consensus expectation is equal to the highest possible interim belief), so ex ante beliefs do not play a role in the interpretation. On the other hand, the common interpretation of signals property used in the result on the tyranny of the uninformed (Proposition 5) does not have any natural interim interpretation.³⁸

A.5. Agent-Specific Random Variables. Our focus throughout the paper has been on agents' higher-order expectations of a given random variable, which is the same across all agents. But for many applications of interest, there is a different random variable corresponding to each agent, and then higher-order expectations are taken. For example, a literature on coordination games in networks focuses on the case where agents have different preferred actions (which correspond to the different random variables) in the absence of coordination motives, and where one's network neighbors also influence one's choice, with linear best responses assumed (Ballester, Calvó-Armengol, and Zenou, 2006; Calvó-Armengol, Martí, and Prat, 2015; Bergemann, Heumann, and Morris, 2015b).

This case can be embedded readily into our formalism. Specifically, suppose that instead of being interested in a (common) random variable $y \in \mathbb{R}^\Theta$ measurable with respect to the external state, each agent has a different random variable, $y^i \in \mathbb{R}^\Theta$. Now, in Section 2.3, equation (3) is changed to

$$x^i(1; \mathbf{y}) = E^i y^i.$$

Once $x^i(1; \mathbf{y})$ is set, the higher-order average expectations are defined by the same equation, (4), as before:

$$x^i(n+1; \mathbf{y}) = \sum_{j \in N} \gamma^{ij} E^i x^j(n; \mathbf{y}).$$

Correspondingly, in the matrix notation of Section 4, where the key iteration is $x(n) = B^{n-1} F y$, the vector $F y$ is replaced by a vector $f \in \mathbb{R}^S$, with

$$f(t^i) = E^i [y^i | t^i].$$

The analogue of (12) is

$$\lim_{\beta \uparrow 1} (1 - \beta) \left(\sum_{n=0}^{\infty} \beta^n B^n \right) f,$$

and $B^n f$ has the interpretation that it describes the higher-order average expectations of the agents' first-order expectations of their agent-specific random variables.

³⁸The ex ante properties of Definitions 4 and 5 do imply properties of interim beliefs—see, for example, Lemma 4 in Section 8.3.

One can generalize further and consider a “pure private values” setting: f can be replaced by an *arbitrary* vector $f \in \mathbb{R}^S$, with the interpretation that $f(t^i)$ is the action that agent i would like to take, when he has signal t^i , in the absence of coordination motives—an action he knows. In this case, each agent faces no uncertainty about the random variable of interest to him. Note that the case of different $y^i \in \mathbb{R}^\Theta$ is a special case of this, because in that case $f(t^i)$ is agent i ’s *expectation* of his own y^i given signal t^i .

This brings us closer to Calvó-Armengol, Martí, and Prat (2015) and Bergemann, Heumann, and Morris (2015a). Motivated by a study of endogenous attention allocation, they work with a network version of a setting commonly studied in organizational economics and focus on an analogue of our $\beta \uparrow 1$ limit. They show that it is a weighted average of agents’ heterogeneous ideal points that matters for determining the network consensus. Even though their setting involves normally distributed random variables and linear–quadratic preferences, the core calculations boil down to understanding an analogue of $B^n f$, just as we must in order to study the heterogeneous-values variation of our model we have just presented.

A.5.1. *Equivalence Between Agent-Specific Random Variables and Different Priors over Θ .*

Given any environment with agent-specific random variables and a common prior on signals, we can find another environment with the same prior on signals in which agents all care about the *same* random variable but have heterogeneous beliefs about external states. That is, given any profile $(y^i)_{i \in N}$, we can define a new environment with new beliefs over Θ and a random variable y so that the resulting $f \in \mathbb{R}^S$ mimics that arising from the original environment. Then results such as Proposition 3 can be applied.

This equivalence relies on the common prior on signals assumption: Without a common prior on signals, we could maintain such an equivalence only if the “own random variables” could depend on others’ signals (cf. Myerson, 1997, p. 74). This is related to the essential differences we observed between the model with a common prior over signals (Section 6) and the model without it.

A.5.2. *Type-Dependent Network Weights.* A related extension allows for type-dependence in γ^{ij} . In this case, we take this network weight to depend on the signal of i , and write $\gamma^{ij}(t^i)$. Much of our analysis goes through unchanged: Equation (10) still describes $x(n)$, but now under the definition

$$B(t^i, t^j) = \gamma^{ij}(t^i) \pi^i(t^j | t^i).$$

If we interpret γ^{ij} as i 's probability of meeting or interacting with j , then signal-dependence of these weights corresponds to private information about interactions. The only results that we lose in this generalization are those of Section 6, because there is now no information-independent notion of the network or of centrality. But the limits we study still exist, and much of their structure (e.g., the structure described in Proposition 1, with p the left-hand unit eigenvector of the generalized B) is still present and can be used to study this more general setting.

APPENDIX B. ADDITIONAL DISCUSSION

B.1. Periodicity and Simple Higher-Order Expectations. In defining consensus expectations, or the limit of higher-order average expectations, we considered the *Abel average*

$$(21) \quad \lim_{\beta \uparrow 1} (1 - \beta) \sum_{n=0}^{\infty} \beta^n x(n+1; y) = \lim_{\beta \uparrow 1} (1 - \beta) \left(\sum_{n=0}^{\infty} \beta^n B^n \right) Fy,$$

which is always well defined. It is natural to ask how the higher-order average expectations $x(n; y)$ behave without this averaging, and about the limit

$$(22) \quad \lim_{n \rightarrow \infty} B^n Fy.$$

As long as B is aperiodic,³⁹ the limit (22) exists and is equal to the right-hand side of (21).

Aperiodicity, and the existence of the limit (22), is not relevant for many of the applications reported in the paper. For the linear best-response game and asset pricing, we are explicitly interested in the limit of the weighted sum of higher-order expectations, i.e., (21) above, and not in limits of unweighted higher-order expectations, i.e., (22) above. Nothing about the structure of agent-type weights depends on aperiodicity.

However, periodicity does affect the behavior of the $x(n; y)$ in the limit, and here we discuss how. Suppose that we have a cycle of agents $i_1, i_2, \dots, i_{|N|}, i_1$: that is, that the network Γ has each agent i_k putting weight 1 on agent i_{k+1} . Then the corresponding matrix will not be aperiodic. For example, if there are two agents, $N = \{1, 2\}$, and $\gamma^{12} = \gamma^{21} = 1$, then we have

$$\Gamma = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

³⁹A matrix is said to be aperiodic if, in the associated weighted directed graph, the greatest common divisor of all cycles' lengths is equal to 1. A sufficient condition for this is that the matrix Γ have all positive entries. Even if $\gamma^{ii} = 0$ for all i —a natural special case for some interpretations and applications—and if there are at least 3 agents, $\gamma^{ij} > 0$ for all $j \neq i$ is another sufficient condition for aperiodicity.

and

$$B = \begin{pmatrix} 0 & B^{12} \\ B^{21} & 0 \end{pmatrix}$$

(recall the definition of B^{ij} from Section 4) and B will be periodic and give rise to a two-cycle.

In particular, there will be well-defined limits

$$\lim_{n \rightarrow \infty} [E_2 E_1]^n E_2 y = c_2 \mathbf{1}$$

and

$$\lim_{n \rightarrow \infty} [E_1 E_2]^n E_1(y) = c_1 \mathbf{1}$$

but they will not be equal. In the limit, the vector $x(n)$ will cycle between

$$\begin{pmatrix} c_1 \mathbf{1} \\ c_2 \mathbf{1} \end{pmatrix} \text{ and } \begin{pmatrix} c_2 \mathbf{1} \\ c_1 \mathbf{1} \end{pmatrix}.$$

For more general cycles of agents, we will have limits of the form

$$\lim_{n \rightarrow \infty} [E_{i_1} E_{i_2} \dots E_{i_k}]^n E_{i_1} E_{i_2} \dots E_{i_j} y$$

but they will be different for different values of $j = 1, \dots, k$. We will refer to such expressions as *simple higher-order expectations*. If the network Γ were given by this cycle, then the entries of $B^n F y$ would be the simple higher-order expectations. The general higher-order expectations that we study will end up being complicated weighted sums of such simple higher-order expectations, although we will not in general work with the decompositions.

Without the assumption of finitely many types, behavior more complicated than cycling can arise, and $(1 - \beta) \sum_{n=0}^{\infty} \beta^n x(n+1; y)$ need not converge as $\beta \uparrow 1$. This phenomenon is discussed in Morris (2002b) and Morris (2002a). A related but different lack of convergence plays a role in Han and Kyle (2016): there, because of the lack of finiteness of the type space, arbitrary higher-order expectations can obtain.

B.2. A Behavioral Interpretation of Irreducibility via No Trade. What is the behavioral content of the joint connectedness of beliefs and the network, i.e., the irreducibility of B ?

We report a characterization of the joint connectedness property, and therefore the existence and uniqueness of a distribution of positive agent-type weights. Just as the common prior assumption can be characterized as the non-existence of profitable trades among agents (see

Morris (1994) and Samet (1998b)), the property we are studying here has a no-trade characterization.⁴⁰

Let x^i be a payment rule for agent i , $x^i : T^i \rightarrow \mathbb{R}$, which is measurable with respect to agent i 's signal. A trade consists of a profile of payment rules, $(x^i)_{i \in N}$. The trade generates *strict expected bilateral gains from trade* if

$$x^i(t^i) \leq \sum_{j \neq i} \gamma^{ij} \sum_{t^j \in T^j} \pi^i(t^j | t^i) x^j(t^j)$$

for each agent i and $t^i \in T^i$, with strict inequality for at least one agent i and $t^i \in T^i$. The interpretation is that agent i is committed to making a payment $x^i(t^i)$ as a function of his signal. But he anticipates receiving the payments to which others are committed.

Proposition 9. *There exists a separable trade generating strict expected bilateral gains from trade if and only if beliefs and the network are jointly connected.*

Proof. The existence of a separable trade giving strict expected bilateral gains from trade is equivalent to the requirement that there exists a vector x such that $x > Bx$, where $>$ means a weak inequality on all components and strict inequality on some component. Recall that irreducibility implies the existence of a strictly positive *vector of agent-type weights* p with $pB = p$. Now we have $px > pBx = pBx$, a contradiction. So irreducibility fails. Conversely, suppose that irreducibility fails. Then there exists at least one type $t^i \in S$ that no one assigns positive probability to, so that $\gamma^{ji} \pi^j(t^i | t^j) = 0$ for all t^j . But now if we set $x^i(t^i) < 0$ and $x^j(t^j) = 0$ if $t^j \neq t^i$, then we have a separable trade with strict expected gains. \square

B.3. Non-Existence of Consensus Expectations on Infinite State Spaces. We have maintained the assumption that Θ and all the T^i are finite. In general, without finiteness, there may not be a vector of agent-type weights as defined in Proposition 1. An example offered by Hellman (2011, Section 6) demonstrates this. The example uses a version of the two-player information structure in Rubinstein's (1989) electronic mail game, with T^1 and T^2 both having the cardinality of \mathbb{N} , the natural numbers. If we take a network Γ on two players such that each puts all weight on the other, and construct a suitable infinite analogue of B , Hellman's result implies that there is no invariant measure for B —i.e., no vector $p \in \Delta(S)$ of agent-type weights such that $pB = p$. Therefore, there is no analogue of Proposition 1, which was the foundation for all our results.

⁴⁰See Nehring (2001) for more on the various relations between no-trade conditions, higher-order expectations, and common priors.

We conjecture that if, like Hellman (2011), we require the state space Ω underlying T^1 and T^2 to be compact⁴¹ and the information structure to be everywhere mutually positive in his sense, then we can recover suitable analogues of our results. On the other hand, if the states come from normal distributions and agents receive noisy signals about them, then the relevant type spaces are uncountably infinite and the random variables have unbounded support, but iterated expectations can still be well-behaved. This case is studied in Han and Kyle (2016); see also Morris and Shin (2002).

APPENDIX C. OMITTED PROOFS

C.1. Proof of Fact 1. To establish (7), write $R^i(k)$ for the set of i 's pure strategies surviving k rounds of iterated deletion of strictly dominated strategies. By assumption, $R^i(k) = R^i(0) = [0, M]^{T^i}$. Then using (6),

$$\begin{aligned} R^i(1) &= \{s^i : (1 - \beta)E^i y \leq s^i \leq (1 - \beta)E^i y + \beta M \mathbf{1}\} \\ &= \{s^i : (1 - \beta)x^i(1) \leq s^i \leq (1 - \beta)x^i(1) + \beta M \mathbf{1}\} \end{aligned}$$

For induction, we may assume that for some $k \geq 1$, each $R^i(k)$ for $i \in N$ has the form

$$R^i(k) = \left\{ s^i : (1 - \beta) \left(\sum_{n=1}^k \beta^{n-1} x^i(n) \right) \leq s^i \leq (1 - \beta) \sum_{n=1}^k \beta^{n-1} x^i(n) + \beta^k M \mathbf{1} \right\}.$$

We have already established the base case, $k = 1$. We will argue that then

$$R^i(k+1) = \left\{ s^i : (1 - \beta) \left(\sum_{n=1}^{k+1} \beta^{n-1} x^i(n) \right) \leq s^i \leq (1 - \beta) \left(\sum_{n=1}^{k+1} \beta^{n-1} x^i(n) \right) + \beta^{k+1} M \mathbf{1} \right\}.$$

The reason is that if i conjectures a strategy profile s satisfying

$$(1 - \beta) \left(\sum_{n=1}^k \beta^{n-1} x^i(n) \right) \leq s^j$$

for each $j \neq i$, then since best responses $\text{BR}^i(s)$ are nondecreasing in s , the minimum best response s^i is obtained by applying BR^i to the lower bound

$$(1 - \beta) \left(\sum_{n=1}^k \beta^{n-1} x^i(n) \right),$$

⁴¹Hellman works with a partitional formalism similar to that of Samet (1998a); see Golub and Morris (2017) for a translation of higher-order expectations into this framework.

which yields

$$\begin{aligned} (1-\beta)E^i y + \beta \sum_{j \neq i} \gamma^{ij} E^i (1-\beta) \left(\sum_{n=1}^k \beta^{n-1} x^j(n) \right) \\ = (1-\beta) \left(\sum_{n=1}^{k+1} \beta^{n-1} x^i(n) \right). \end{aligned}$$

The argument for the upper bound is analogous. As $k \rightarrow \infty$, the lower and upper bounds both converge to the $s_*(\beta)$ of (7).

C.2. Existence and Characterization of the Consensus Expectation: Proof of Proposition

1. Recall that p is the unique vector in $p \in \Delta(S)$ satisfying $p = pB$; this vector is uniquely determined and positive by a standard result for irreducible Markov chains. Write

$$(23) \quad x(\beta) = (1-\beta) \sum_{n=0}^{\infty} \beta^n B^n z.$$

We will show that for any $z \in \mathbb{R}^S$, we have

$$(24) \quad \lim_{\beta \uparrow 1} x(\beta) = pz\mathbf{1}.$$

Note that by the Neumann series, which can be used since the spectral radius of βB is $\beta < 1$, we have $\sum_{n=0}^{\infty} (\beta B)^n = (I - \beta B)^{-1}$, where I denotes the identity matrix of appropriate size; in particular, $I - \beta B$ is invertible. So $x(\beta) = (1-\beta)(I - \beta B)^{-1}z$, or, equivalently,

$$(25) \quad (I - \beta B)x(\beta) = (1-\beta)z.$$

The formula (23) says that $x(\beta)$ is an average, because the weights $(1-\beta)\beta^n$ sum to 1, of the vectors $B^n z$. Because B^n is a Markov matrix, no entry of $B^n z$ can exceed the largest value of z in absolute value. So the same is true of $x(\beta)$, and therefore all the $x(\beta)$ lie in a compact set.

Consider a sequence $\beta_k \uparrow 1$. By what we have said, the sequence $(x(\beta_k))_k$ lies inside a compact set. By a standard fact about compact sets, such a sequence converges, and has the limit $pz\mathbf{1}$, if and only if every convergent subsequence of it converges to $pz\mathbf{1}$. So consider a convergent subsequence, $(x(\beta_\kappa))_\kappa$, and let x denote its limit. We will show that $x = pz\mathbf{1}$, which will conclude the proof of (24).

By taking $\beta \uparrow 1$ in (25), we see that x satisfies $x = Bx$, which, given that our matrix B is irreducible, means that $x = a\mathbf{1}$ for some constant a . It remains only to prove that $a = pz$. Premultiplying (25) by p gives $(1-\beta_\kappa)px(\beta_\kappa) = (1-\beta_\kappa)pz$. Canceling $(1-\beta_\kappa)$, we get $px(\beta_\kappa) = pz$. Letting $\kappa \rightarrow \infty$ and recalling that x is defined as the limit of the subsequence

yields $px = pz$. When we plug in $x = a\mathbf{1}$ —the statement that x is a constant vector—we find that $ap\mathbf{1} = pz$. Since p is a probability vector, we have $p\mathbf{1} = \mathbf{1}$, and so we conclude that $a = pz$.

C.3. Proof of Lemma 2. If W_1 is drawn from the ergodic distribution p , the distributions of W_1 and W_2 are the same, and so the expected difference between $f(W_2)$ and $f(W_1)$ is 0:

$$(26) \quad \mathbb{E}_{W_1 \sim p}[f(W_2) - f(W_1)] = 0.$$

On the other hand, using hypotheses (1) and (2) in the second line below, we have

$$\begin{aligned} \mathbb{E}_{W_1 \sim p}[f(W_2) - f(W_1)] &= \sum_{s: f(s) < \bar{f}} p(s) \mathbb{E}_{W_1=s}[f(W_2) - f(s)] + \sum_{s: f(s) \geq \bar{f}} p(s) \mathbb{E}_{W_1=s}[f(W_2) - f(s)] \\ &\geq \delta p(s : f(s) < \bar{f}) - \varepsilon p(s : f(s) \geq \bar{f}). \end{aligned}$$

Combining this result with (26) and using the shorthand $\chi = p(s : f(s) \geq \bar{f})$, we deduce $0 \geq \delta(1 - \chi) - \varepsilon\chi$, from which the lower bound on χ claimed in the proposition follows.

C.4. Proofs of Results on Tyranny of the Uninformed . The key lemma behind our proof of Proposition 5 is:

Lemma 3. *Under the hypotheses of Proposition 5,*

$$\left| \frac{p(s) - \hat{p}(s)}{\hat{p}(s)} \right| \leq \frac{4|\Theta||S|^2}{(\gamma_{\min}\rho_{\min})^2} \cdot \frac{\varepsilon}{\delta}.$$

Proof. The proof relies on Theorem 2.1 of Cho and Meyer (2000), which says that, for any $s \in S$,

$$(27) \quad \left| \frac{p(s) - \hat{p}(s)}{\hat{p}(s)} \right| \leq \frac{1}{2} \left\| B - \hat{B} \right\|_{\infty} \max_{z \neq z'} M_{\hat{B}}(z, z'),$$

where $M_{\hat{B}}(z, z')$ is the mean first passage time⁴² in \hat{B} to z' starting at z ; the norm is the maximum absolute row sum. Two key technical lemmas, stated in Section C.4.1 below, allow us to bound the right-hand side. Using Lemma 4 (summing the upper bounds on absolute differences across any row and taking the maximum over all rows i):

$$\left\| B - \hat{B} \right\|_{\infty} \leq |S| \cdot \frac{4|\Theta||S|\varepsilon}{\min_{i \neq 1} \rho_{\min}^i}.$$

⁴²Consider a Markov chain making transitions according to \hat{B} . The mean first-passage time from z to z' in \hat{B} is denoted by $M_{\hat{B}}(z, z')$ and defined to be the expected number of steps that the chain started at z takes up to its first visit to z' (inclusive).

To finish bounding the right-hand side of (27), it remains to bound $\max_{z \neq z'} M_{\hat{B}}(z, z')$. Lemma 5 does exactly this, giving

$$\max_{z \neq z'} M_{\hat{B}}(z, z') \leq \frac{2}{\delta \rho_{\min}^1 \gamma_{\min}^2}.$$

Recall that γ_{\min} is the minimum off-diagonal entry of Γ —by assumption a positive number. Combining the two inequalities gives the claimed bound. \square

Now we can show how this result implies Proposition 5.

The first step is to show that the consensus expectation under the hatted information structure is equal to the first agent's prior expectation:

$$c(y; B_{\hat{\pi}}, F_{\hat{\pi}}) = \mathbf{E}^{\rho^1}[y].$$

The key to this is to establish that the information structure $(\hat{\pi}^i)_{i \in N}$ is consistent with a common prior over signals. Indeed, we will show that agent 1's prior can be taken to be this common prior. Let $\hat{\mu}^1 \in \Delta(T^1)$ be the prior on T^1 induced by ρ^1 , and let

$$\hat{\mu}^i(t^i) = \sum_{t^1 \in T^1} \hat{\pi}^1(t^i | t^1) \hat{\mu}^1(t^1).$$

For agents $i \neq 1$, the interim beliefs $\hat{\pi}^i(\cdot | t^i)$ are compatible with their respective priors $\hat{\mu}^i$ trivially, because the interim beliefs place probability 0 or 1 on any state, and are compatible with *any* prior—Bayes' rule implies no restrictions. Moreover, with this profile $(\hat{\mu}^i)_{i \in N}$, the information structure $(\hat{\pi}^i)_{i \in N}$ is consistent with a common prior over signals. Now note that the prior over Θ corresponding to any $\hat{\mu}^i$ is ρ^1 . By Proposition 3, the consensus expectation $c(y; B_{\hat{\pi}}, F_{\hat{\pi}})$ is the common prior expectation of y , namely $\mathbf{E}^{\rho^1}[y]$.

The second step is to bound the distance between $c(y; B_{\hat{\pi}}, F_{\hat{\pi}})$, which we have computed, and $c(y; B_{\pi}, F_{\pi})$, which we would like to characterize. It is here that Lemma 3 is relevant:

$$\begin{aligned} |c(y; B_{\pi}, F_{\pi}) - c(y; B_{\hat{\pi}}, F_{\hat{\pi}})| &= \left| \sum_{s \in S} [p(s) - \hat{p}(s)] E^i[y | s] \right| \\ &= \left| \sum_{s \in S} \frac{p(s) - \hat{p}(s)}{\hat{p}(s)} \hat{p}(s) E^i[y | s] \right| && \text{multiply and divide by } \hat{p}(s) \\ &\leq \sum_{s \in S} \left| \frac{p(s) - \hat{p}(s)}{\hat{p}(s)} \right| \hat{p}(s) |E^i[y | s]| && \text{triangle inequality} \\ &\leq \frac{4|\Theta||S|^2}{(\gamma_{\min} \rho_{\min})^2} \cdot \frac{\varepsilon}{\delta} \sum_{s \in S} \hat{p}(s) |E^i[y | s]| && \text{Lemma 3} \end{aligned}$$

$$\leq \frac{4|\Theta||S|^2}{(\gamma_{\min}\rho_{\min})^2} \cdot y_{\max} \cdot \frac{\varepsilon}{\delta}. \quad \text{definition of } y_{\max}$$

This completes the proof of the proposition, except for the technical lemmas, which are the subject of the next section.

C.4.1. Statements of Technical Lemmas. The proof of Lemma 3 used two key bounds. We state both here, and give proofs below.

The first result, which was used to bound $\|B - \widehat{B}\|_{\infty}$, converts hypotheses about the signal structures $(\eta^i)_{i \in N}$ into statements about the agents' interim beliefs (recall that the entries of B are products of network weights from Γ and interim beliefs):

Lemma 4. *For any $t^i, t^j \in S$ with $j \neq i$, we have*

$$|\pi^i(t^j | t^i) - \widehat{\pi}^i(t^j | t^i)| \leq \frac{4|\Theta||S|\varepsilon}{\rho_{\min}^i}.$$

This follows from Bayes' rule, but the exact statement requires a good deal of calculation. The core idea is that $\widehat{\eta}$ is obtained by changing the probabilities in η only slightly. Given full support priors, each $\pi^i(t^j | t^i)$ is continuous in $\eta^i(t^i | \theta)$, so it is natural that the two should be close; our calculation simply gives a quantitative version of this statement.

We also used a bound on mean first-passage times in \widehat{B} :

Lemma 5. *For any two states $z, z' \in S$,*

$$M_{\widehat{B}}(z, z') \leq \frac{2}{\delta \rho_{\min}^1 \gamma_{\min}^2}.$$

The key idea here is that, as a consequence of agent 1 having noisy information, the subjective probability agent 1 puts on any type of any other agent is reasonably high: The lower bound is $\rho_{\min}^1 \delta$, as we establish in the proof. Thus the corresponding weights in \widehat{B} are lower-bounded by $\delta \rho_{\min}^1 \gamma_{\min}$, once we take into account the network part of the weight. The other agents' types have perfect information, so each of them has an edge of weight at least γ_{\min} to a type of agent 1. Thus the Markov chain is well-interconnected by agent 1's types: Starting from any state, one gets to agent 1's types immediately, and then to any other given state in S with substantial probability, so the chain cannot take too long to visit that state (by a standard bound on geometric random variables).

C.4.2. Proof of Lemma 4. The proof relies on the following fact about prior probabilities of signals.

Fact 3. For any $i \neq 1$ and any t^i ,

$$\mu^i(t^i) = \sum_{\theta' \in \Theta} \eta^i(t^i | \theta') \rho^i(\theta') \geq (1 - \varepsilon) \rho_{\min}^i.$$

This bound holds because η^i is assumed to be at most ε -noisy, and so there must be some θ_{t_i} such that $\eta^i(t^i | \theta_{t_i}) \geq 1 - \varepsilon$.

The first step of the proof of Lemma 4 is to write the probabilities in question via sums over states θ . For any t^i, t^j with $j \neq i$, we have

$$\pi^i(t^j | t^i) = \sum_{\theta \in \Theta} \eta^j(t^j | \theta) \pi^i(\theta | t^i)$$

Define $\hat{\pi}^i(t^j | t^i)$ analogously, replacing π^i by $\hat{\pi}^i$ and η^i by $\hat{\eta}^i$. Let

$$H^j(t^j | \theta) = |\eta^j(t^j | \theta) - \hat{\eta}^j(t^j | \theta)|$$

and

$$\Delta^i(\theta | t^i) = |\pi^i(\theta | t^i) - \hat{\pi}^i(\theta | t^i)|.$$

Now note that by the triangle inequality,

$$(28) \quad |\pi^i(t^j | t^i) - \hat{\pi}^i(t^j | t^i)| \leq \sum_{\theta \in \Theta} [\Delta^i(\theta | t^i) + H^j(t^j | \theta) + \Delta^i(\theta | t^i) H^j(t^j | \theta)].$$

Having written the difference we are studying in this way, we will bound it piece by piece. If $j \neq 1$, by definition of “at most ε -noisy,” we have that $|H^j(t^j | \theta)| \leq \varepsilon$. If $j = 1$, then $H^j(t^j | \theta)$ is identically zero. Also, note that $|\Delta^i(\theta | t^i)| \leq 1$. So in all cases, we can bound the last two terms in the brackets by 2ε .

Now, we turn to $\Delta^i(\theta | t^i)$. If $i = 1$, then $\Delta^i(\theta | t^i) = 0$, because 1’s signals are the same in both the original information structure π and the new one $\hat{\pi}$.

So assume $i \neq 1$; we will show that $\Delta^i(\theta | t^i) \leq (|S| - 1) \frac{\varepsilon}{(1 - \varepsilon) \rho_{\min}^i}$, and this will allow us to complete the proof. Let θ_{t_i} be such that $\eta^i(t^i | \theta_{t_i}) \geq 1 - \varepsilon$, which is guaranteed to exist by the definition of “at most ε -noisy.” We will bound $\Delta^i(\theta | t^i)$, considering the cases $\theta \neq \theta_{t_i}$ and $\theta = \theta_{t_i}$ separately. If $\theta \neq \theta_{t_i}$, then by Bayes’ rule,

$$\begin{aligned} \pi^i(\theta | t^i) &= \frac{\eta^i(t^i | \theta) \rho^i(\theta)}{\mu^i(t^i)} \\ &\leq \frac{\eta^i(t^i | \theta) \rho^i(\theta)}{(1 - \varepsilon) \rho_{\min}^i} \end{aligned} \quad \text{by Fact 3}$$

$$\leq \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i} \quad \text{by definition of at least } \varepsilon\text{-nosiy.}$$

Since $\widehat{\pi}^i(\theta | t^i) = 0$, it follows that

$$(29) \quad \Delta^i(\theta | t^i) \leq \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i}$$

By the law of total probability,

$$\pi^i(\theta_{t^i} | t^i) \geq 1 - (|S| - 1) \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i}.$$

Since $\widehat{\pi}^i(\theta_{t^i} | t^i) = 1$, it follows that

$$(30) \quad \Delta^i(\theta_{t^i} | t^i) \leq (|S| - 1) \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i}.$$

This is the looser of the two bounds (29) and (30), so we can say in general that

$$(31) \quad \Delta^i(\theta_{t^i} | t^i) \leq (|S| - 1) \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i}.$$

Putting everything together, it follows that

$$\begin{aligned} |\pi^i(t^j | t^i) - \widehat{\pi}^i(t^j | t^i)| &\leq \sum_{\theta \in \Theta} [\Delta^i(\theta | t^i) + 2\varepsilon] && \text{by (28)} \\ &\leq \sum_{\theta \in \Theta} \left[(|S| - 1) \cdot \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i} + 2\varepsilon \right] \\ &\leq |\Theta| \left((|S| - 1) \frac{\varepsilon}{(1-\varepsilon)\rho_{\min}^i} + 2\varepsilon \right) \\ &\leq |\Theta| \varepsilon \left((|S| - 1) \frac{2}{\rho_{\min}^i} + 2 \right) && \text{using } 1 - \varepsilon \geq \frac{1}{2} \\ &\leq |\Theta| \frac{\varepsilon}{\rho_{\min}^i} (2(|S| - 1) + 2) \end{aligned}$$

The claimed bound follows after noting and that $2|S| \geq 2(|S| - 1) + 2$ because $|S| \geq 2$.

C.4.3. Proof of Lemma 5. The proof requires the following fact:

Fact 4. For any $t^1 \in T^1$ and $t^j \in T^j$ with $j \neq 1$ we have:

$$\widehat{\pi}^1(t^j | t^1) = \frac{\sum_{\theta \in \Theta} \rho^1(\theta) \widehat{\eta}^j(t^j | \theta)}{\sum_{\theta \in \Theta} \rho^1(\theta) \eta^1(t^1 | \theta)} \geq \delta \rho_{\min}^1.$$

To establish this fact, we note that there is some θ_{t^j} such that $\widehat{\eta}^j(t^j | \theta) = 1$, and the denominator is at most 1 since it is the prior probability of the signal t^1 under the information structure associated with $\widehat{\eta}^1$.

Now we prove Lemma 5. Let $(\widehat{W}_n)_n$ be a stochastic process corresponding to the Markov matrix \widehat{B} . Defining the function $\iota : S \rightarrow N$ by $\iota(t^i) = i$, we have a coupling between the chain $(\widehat{W}_n)_n$ and a chain on N , the set of agents, with transition matrix Γ .

Case 1: $z' \notin T^1$. Let us analyze the first passage time to some $z' \notin T^1$. Starting from any $z \in S$, the mean first passage time of the process $(\iota(\widehat{W}_n))_n$ to 1 (the state corresponding to agent 1) is at most $1/\gamma_{\min}$. Then every time the process visits a state in T^1 , it has probability at least $\delta \rho_{\min}^1 \gamma_{\min}$ of visiting z' , by Fact 4. Conditional on not visiting it at this time, we wait on average $1/\gamma_{\min}$ steps for the process to return to a state in T^1 and have another $\delta \rho_{\min}^1 \gamma_{\min}$ chance at visiting z' . Thus, using the formula for the expectation of a geometric random variable, we have

$$M_{\widehat{B}}(z, z') \leq \frac{1}{\delta \rho_{\min}^1 \gamma_{\min}^2} \quad \text{whenever } z' \notin T^1.$$

Case 2: $z' \in T^1$. Let $z' = t^1$. If $z = t^j \notin T^1$, then there is a θ_{t^j} such that $\widehat{\pi}^j(\theta_{t^j} | t^j) = 1$. Then

$$\widehat{\pi}^j(t^1 | t^j) = \eta^1(t^1 | \theta_{t^j}),$$

which is at least δ by the definition of ‘‘at least δ -noisy.’’ Thus, every time the process $(\widehat{W}_n)_n$ visits any state in $S \setminus T^1$, it has probability at least δ of visiting z' . If $z \in T^1$, then the process surely visits the set $S \setminus T^1$ one step later. Thus the process takes at most two steps to be in a position where it has probability δ of visiting z . By the same reasoning discussed above about a geometric random variable, we conclude that

$$M_{\widehat{B}}(z, z') \leq \frac{2}{\delta}.$$

C.5. Proof of Proposition 8. By Fact 1, there is a unique rationalizable strategy profile in environment $(\widehat{\Gamma}, \widehat{\beta}, y)$. In that strategy profile, player i 's action given his signal satisfies

$$a^i = (1 - \widehat{\beta})E^i y + \widehat{\beta} \sum_j \widehat{\gamma}^{ij} E^i a^j.$$

Splitting the $j = i$ term out of the last summation, and then using the definition $\widehat{\gamma}^{ij} = \gamma^{ij} (1 - \gamma^{ii})$, we have

$$a^i = (1 - \widehat{\beta})E^i y + \widehat{\beta} \widehat{\gamma}^{ii} E^i a^i + \widehat{\beta} (1 - \widehat{\gamma}^{ii}) \sum_{j \neq i} \gamma^{ij} E^i a^j.$$

Rearranging and using $E^i a^i = a^i$ gives

$$(1 - \widehat{\beta} \widehat{\gamma}^i) a^i = (1 - \widehat{\beta}) E^i y + \widehat{\beta} (1 - \widehat{\gamma}^i) \sum_{j \neq i} \gamma^{ij} E^i a^j$$

and thus

$$\begin{aligned} a^i &= \frac{1 - \widehat{\beta}}{1 - \widehat{\beta} \widehat{\gamma}^i} E^i y + \frac{\widehat{\beta} (1 - \widehat{\gamma}^i)}{1 - \widehat{\beta} \widehat{\gamma}^i} \sum_{j \neq i} \gamma^{ij} E^i a^j \\ (32) \quad &= (1 - \beta^i) E^i y + \beta^i \sum_{j \neq i} \gamma^{ij} E^i a^j, \end{aligned}$$

where in the last step we have deduced from the formula $\widehat{\gamma}^i = \frac{\widehat{\beta} - \beta^i}{\widehat{\beta}(1 - \beta^i)}$ the fact that $\frac{1 - \widehat{\beta}}{1 - \widehat{\beta} \widehat{\gamma}^i} = 1 - \beta^i$ and $\frac{\widehat{\beta}(1 - \widehat{\gamma}^i)}{1 - \widehat{\beta} \widehat{\gamma}^i} = \beta^i$.

Now (32) is an equilibrium of the coordination game in environment (Γ, β, y) (recall equation (19) defining that game), and so by uniqueness of the rationalizable outcome, the proof is complete.