Higher-Order Expectations*

Benjamin Golub[†] Stephen Morris[‡]

August 31, 2017

Abstract

We study *higher-order expectations* paralleling the Harsanyi (1968) approach to higher-order beliefs—taking a basic set of random variables as given, and building up higher-order expectations from them. We report three main results. First, we generalize Samet's (1998a) characterization of the common prior assumption in terms of higher-order expectations, resolving an apparent paradox raised by his result. Second, we characterize when the limits of higher-order expectations can be expressed in terms of agents' heterogeneous priors, generalizing Samet's expression of limit higher-order expectations via the common prior. Third, we study higher-order *average* expectations—objects that arise in network games. We characterize when and how the network structure and agents' beliefs enter in a separable way.

1 Introduction

Samet (1998a) considered sequences of higher-order expectations such as: Ann's expectation of a random variable; Bob's expectation of Ann's expectation; Charlie's expectation of Bob's expectation of Ann's; and so on. Samet showed that any such sequence converges to a constant random variable whose value is common certainty—in other words, limits of higher-order expectations are public. He also showed that *order-independence* is equivalent to the *consistency* of beliefs. *Order-independence* holds if, for all random variables, the limits of higher-order expectations are independent of the order in which expectations are taken. Agents' interim belief functions are *consistent* if there is a single prior distribution that is simultaneously compatible with each agent's updating. Consistency is the interim content of the common prior assumption. Thus,

*We thank Dov Samet for many spirited exchanges. We also benefited from the input of Nageeb Ali, Dirk Bergemann, Ben Brooks, Jason Hartline, Tibor Heumann, Omer Tamuz, and Muhamet Yildiz. We especially thank Ryota Iijima for several inspiring conversations early in this project.

[†]Department of Economics, Harvard University, Cambridge, U.S.A., bgolub@fas.harvard.edu.

[‡]Department of Economics, Princeton University, Princeton, U.S.A., smorris@princeton.edu.

Samet's (1998a) result on order-independence characterizes a formulation of the common prior assumption in terms of the properties of higher-order expectations.[1] Viewed from another perspective, it provides a simple way of calculating limits of higher-order expectations under suitable assumptions: just find the common prior, and take expectations with respect to that.

In Samet's analysis, there is a space Ω modeling all uncertainty, and all the conditions he studies (in particular, order-independence) are imposed on all random variables defined on this space—that is, on all the random variables the modeler can express. In contrast, we focus on a set of random variables constructed as in the Harsanyi (1968) construction of higher-order beliefs. That is, we fix a set of *parameters* or *external states*; then we define the *basic* random variables as the ones measurable with respect to these external states. Then we inductively build up an *expectations space* consisting of agents' expectations and higher-order expectations of those basic random variables. The expectations space is the natural one to study when the properties of interest are higher-order expectations of a given set of random variables, as is the case in the study of some network games and asset markets. We study analogues of Samet's conditions on this space.

As we will explain, this seemingly small change from Samet's (1998a) model yields a structure that contrasts sharply with his. We report three main results concerning higher-order expectations in this framework. The first result studies the content of order-independence when it is imposed on the expectations space rather than on all random variables. Because in a special case the expectations space is equal to the space of all random variables, our result generalizes Samet's. We will argue that our approach resolves an apparent paradox raised by Samet's characterization—namely, that a property of only higher-order expectations (order-independence) can completely characterize a property of beliefs (the common prior assumption), even though higher-order expectations are a coarse measure of agents' beliefs about any given random variable. Our second result describes the *values* of higher-order expectations and shows when these values can be expressed in terms of heterogeneous priors associated with different agents; this result generalizes Samet's expression of limit higher-order expectations via the common prior. The third result focuses on objects closely related to higher-order expectations—higher-order average expectations—that arise in network games and certain asset markets. Under suitable assumptions, we decompose the values of higher-order average expectations in a way that *separates* the effects of the network structure and agents' beliefs; moreover, we characterize exactly when such a decomposition is possible.

[1] We follow Harsanyi (1968) in using consistency rather than the common prior assumption, both to highlight its interim interpretation and to provide a benchmark for notions of consistency that we will introduce later, which do not have natural interpretations from the ex ante perspective. We also differ from Samet (1998a) in using the terminology of "higher-order" rather than "iterated" expectations; we use the former to highlight the analogy with higher-order beliefs, which will play an important role in this paper.

The Expectations Space and Expectation-Consistency We approach higher-order expectations as Harsanyi (1968) approached higher-order beliefs. First, we fix a set of *external states*.[2] Then we define all the random variables measurable with respect to those states, and inductively construct expectations and higher-order expectations of those random variables. These variables together make up the *expectations space*. As we will see, this contrasts with Samet's (1998a) approach of imposing conditions on higher-order expectations of *all* random variables measurable with respect to the ambient state space Ω—the set of all states the modeler can conceive of.

In contrast to Samet's focus on consistency defined in terms of beliefs, our first result uses the property of *expectation-consisten*cy. This property holds if there is a profile of distributions, each a prior for one agent,[3] which yield the same expectations for all random variables in the expectations space. Our first result shows that there is order-independence with respect to the expectations space if and only if beliefs are expectation-consistent. We call any of these measures an *expectation-consistent prior* for the environment.

The property of expectation-consistency is a significant weakening of consistency—i.e., the requirement that all distributions in the profile of priors mentioned earlier actually be the *same* distribution. This can be seen in two ways. First, because only expectations matter, expectation-consistency depends only on each agent's *marginal* belief about each other agent's types, but not on his beliefs about profiles of others' types. Second, it will turn out that expectation-consistency is automatically satisfied when there are only two agents, whereas a consistent prior is, of course, not guaranteed to exist even in the two-agent case.

We thus show an equivalence between order-independence (when the basic space of random variables is fixed) and a certain property of beliefs that is weaker than the common prior assumption. That this property should be weaker is intuitive, because an agent's expectation of a given random variable is a coarser measure of his beliefs than the agent's full beliefs themselves; thus, we would not expect properties of higher-order expectations (order-independence) to characterize properties of higher-order beliefs (consistency). The result of Samet (1998a) can seem paradoxical in that it is inconsistent with this intuition. We resolve the tension as follows: Samet (1998a) quantifies over arbitrary random variables in his definition of order-independence. The random variables being considered can then depend in an arbitrary way on all agents' beliefs and higher-order beliefs, and thus the condition of order-independence imposed on this large set of random variables is very strong. We impose the analogous condition on the smaller expectations space, and that characterizes a less demanding notion of consistency.

We are interested in this result for three reasons. First, as we have just stated, it resolves the apparent paradox in the striking result of Samet (1998a). Second, it identifies

[2] These were called *parameters* in the work of Mertens and Zamir (1985). Our terminology highlights the interpretation that they reflect outcomes separate from the agents' perceptions.

[3] That is, a prior that is compatible with that agent's belief updating.

the relevant space of higher-order expectations, in analogy with the space of higher-order beliefs described in Mertens and Zamir (1985). Third, it is a foundation for our second set of results on the values of higher-order expectations even when no expectation-consistent prior exists. We discuss these next.

Characterizing Limits of Higher-Order Expectations when Expectation-Consistency Does Not Hold Our second set of results studies the *values* of higher-order expectations, with and without expectation-consistency. Samet (1998a) shows that under consistency, the limit of higher-order expectations of a random variable is the common prior expectation of that random variable.[4] An analogous result holds when we replace consistency with expectation-consistency. More importantly, we show that limits of higher-order expectations can be characterized under conditions substantially more general than expectation-consistency.

To describe this characterization, consider the *higher-order expectations space* consisting of agents' (first-order) expectations, and all higher-order expectations of them. This space is a subset of the expectations space discussed earlier, because it excludes random variables of the "zeroth order"—ones that are measurable with respect to the external space and do not involve taking any expectations. We say that beliefs are *higher-order expectation-consistent* if there is a profile of agents' priors that agree in their expectations of every element of the higher-order expectations space. Higher-order expectation-consistency is a significant weakening of expectation-consistency, because it imposes no consistency on agents' beliefs about the external state space, and in this sense allows arbitrary heterogeneous priors; but it does impose consistency on beliefs about others' expectations.

We characterize this property by relating it to *higher-order-independence*. Order independence says that a limiting higher-order expectation, such as $\cdots E_{\text{Charlie}} E_{\text{Bob}} E_{\text{Ann}} f$, of a random variable f does not depend on the sequence at all. Higher-order independence says that the limiting higher-order expectation can depend only on the identity of the rightmost agent in the sequence of expectations—in this case, Ann.

We show that higher-order expectation-consistency is equivalent to higher-order-independence. Moreover, we can describe the limiting higher-order expectations in a simple way. They are simply prior expectations of the relevant random variable, evaluated according to a certain prior distribution associated with the rightmost agent; no aspect of the order other than the identity of the rightmost agent matters. The prior distribution in question is one that corresponds to that agent's information and thus is a prior distribution over his types (i.e., partition elements). We call this the *public prior* on that agent. An interpretation is that, under higher-order expectation-consistency, there is a consistent public view about each agent's first-order beliefs. Very high-order expectations depend exclusively on this public prior about the agent who is the rightmost in the sequence.

[4] In interim terms, it is the expectation under the consistent prior, which is uniquely determined under Samet's maintained assumptions.

In the case of expectation-consistency, all the public priors are the same, and taking expectations with respect to them yields expectations corresponding to the uniquely identified expectation-consistent prior.

Higher-Order Average Expectations in Networks and the Separability of Networks and Beliefs Our third set of results considers a relative of higher-order expectations that comes up in network games and other applications: higher-order *average* expectations. Fix a random variable on the external state space. *Higher-order average expectations* of a group of agents are formed by starting with their expectations of a random variable, and then considering each agent's expectation of a weighted average of others' expectations. The averaging, for each agent, is done according to arbitrary weights on others. This profile of weights is interpreted as a *network*.

In Golub and Morris (2017), we have shown that such higher-order average expectations characterize behavior in certain coordination games. In particular, suppose agents place some weight on coordinating with the actions of others and some weight on coordinating with the realized value of a basic random variable. The actions and the random variable enter linearly in each agent's best response. As coordination motives dominate, an agent's equilibrium play converges to the limit of his n^{th}-order average expectation of the basic random variable as n tends to infinity. The network with respect to which the higher-order average expectations are defined reflects the agents' coordination motives: the weight agent i places on j in the network corresponds to how much i wants to coordinate with j. Under suitable conditions, play is deterministic and equal across players, and we call the common action played the *consensus expectation*.

Motivated by this application, we are interested in the interaction of two factors in determining consensus expectations: (i) agents' beliefs and (ii) the network structure. It turns out that the condition of higher-order consistency that figured in our second set of results implies that these two ingredients interact in a very simple way: The consensus expectation is the weighted average of various agents' public prior expectations, with a given agent's contribution to the average weighted by his *eigenvector centrality* in the network. In this case we say that consensus expectations are *separable*. Our main result in this section is that separability is actually *equivalent* to higher-order expectation-consistency. The main formula we derive in this section echoes formulas known in the study of deterministic network games (Ballester et al., 2006; Calvó-Armengol et al., 2015) but in a much more general setting. In this sense we identify a class of environments in which those formulas generalize nicely.

Our first and second sets of results, before we come to higher-order average expectations, are mathematically straightforward extensions of Samet (1998a). It is their conceptual content that we want to stress. The third set of results, about higher-order average expectations, is more subtle; deriving them requires a number of steps involving a connection between higher-order average expectations and a Markov process. The probabilistic approach turns out to be useful, in that a key lemma uses the idea of a

Markov chain reversal. These subtleties are discussed in more detail in Section 6.2.

We describe the model in the next section. In Section 3, we describe the expectations space and higher-order expectations space and explain how these two spaces give rise to a coarser notion of redundancy. In Section 4, we collect our results about order-independence and compare them with the results of Samet (1998a). In Section 5, we describe higher-order versions of our order-independence and limit characterization results. In Section 6, we describe the relation between backward-looking and forward-looking higher-order expectations and use them to give our characterizations of consensus expectations and separability.

2 Model

2.1 States, Types, Priors, and Expectations

Let $I = \{1, 2, \ldots, n\}$, with $n \geq 2$, be a set of *agents*. Let Θ be a finite set of *external states*. We are interested in the agents' beliefs about Θ, their expectations of random variables defined on Θ, and higher-order beliefs and expectations. The model has a finite set of *states* Ω, whose elements $\omega \in \Omega$ fully resolve all uncertainty. We denote by $\theta(\omega)$ the external state at $\omega \in \Omega$. The mapping θ generates a partition of Ω,[5] and we will refer to this partition as $\boldsymbol{\Theta}$ (note the different typesetting of this $\boldsymbol{\Theta}$). For each $i \in I$, there is a partition $\mathscr{P}_i$ of Ω representing agent i's information, and $P_i(\omega)$ is the element of $\mathscr{P}_i$ containing ω. Now, let[6] $t_i : \Omega \to \Delta(\Omega)$ be a function such that, for each $\omega \in \Omega$:

1. $t_i(\omega)$ is a probability distribution with support contained in $P_i(\omega)$, i.e., such that $(t_i(\omega))(P_i(\omega)) = 1$;

2. for any $P_i \in \mathscr{P}_i$, the value $t_i(\omega) \in \Delta(\Omega)$ is the same at all $\omega \in P_i$.

We call t_i the *type function* or *belief function* of agent i.

This is the model Samet (1998a) uses; the only (formal) difference from his formulation is that we introduce external states and their associated partition $\boldsymbol{\Theta}$ into the description of his model. Just as Samet does, we maintain the assumptions that $t_i(\omega)(\omega) > 0$ for every $\omega \in \Omega$ (a *common support* assumption)[7] and that the meet of agents' partitions is the trivial partition consisting of Ω and the empty set (this amounts to an assumption of *no nontrivial common certainty*).[8]

A *random variable* is a function $f : \Omega \to \mathbb{R}$, also written $f \in \mathbb{R}^\Omega$. We say f is measurable with respect to a family of sets $\mathscr{P} \subseteq 2^\Omega$ if, for any $r \in \mathbb{R}$, the preimage $f^{-1}(r)$ is in $\mathscr{P}$.

[5] That is, it partitions Ω into sets which are mapped by θ to the same element of Θ.

[6] $\Delta(\Omega)$ denotes the set of probability measures, or probability distributions, over Ω.

[7] This assumption is with some loss of generality, but analogous results can be proved in its absence.

[8] This assumption is without loss of generality: if it failed, our results would be true when conditioned on the realized element of the meet of agents' partitions.

We denote by $m\mathscr{P}$ the vector space of all random variables measurable with respect to $\mathscr{P}$. In the reverse direction, given a set of random variables V, define

$$\pi V = \left\{ f^{-1}(r) \colon f \in V, r \in \mathbb{R} \right\},$$

which is the coarsest partition with respect to which that random variable is measurable.

For any probability measure $p \in \Delta(\Omega)$ and any random variable $f : \Omega \to \mathbb{R}$, the expectation of f under p is $pf = \sum_{\omega\in\Omega} p(\omega) f(\omega)$. To define the subjective expectation $E_i f$ for any random variable $f : \Omega \to \mathbb{R}$, we let $(E_i f)(\omega) := t_i(\omega) f$. This is the expectation of f taken with respect to the belief that i holds when ω occurs. It is a random variable because it depends on which element of $\mathscr{P}_i$ the realized state ω falls in.

2.2 Priors

A *prior* for agent i is a convex combination of his posterior beliefs, or "types" (Samet, 1998b). That is, $p_i \in \Omega(\Delta)$ is a prior for i if there is a tuple of nonnegative real numbers $(\alpha(P_i))_{P_i \in \mathscr{P}_i}$ such that

$$p_i = \sum_{P_i \in \mathscr{P}_i} \alpha(P_i) t_i(P_i) \quad \text{and} \quad \sum_{P_i \in \mathscr{P}_i} \alpha(P_i) = 1. \tag{1}$$

A "prior" refers to a hypothetical ex ante belief with the property that the agent's interim beliefs could be derived by Bayes-updating of the prior. However, throughout the paper we take an interim approach to beliefs and expectations; the ex ante stage never plays an explicit role.

2.3 Higher-Order Expectations

A *sequence*[9] of agents is a map $\sigma : \{1, 2, \ldots\} \to I$. Following Samet's terminology, an I-*sequence* is defined to be a sequence in which each agent in I appears infinitely often. Let $\mathscr{I}$ be the set of all I-sequences.

For any sequence σ, we will define the *higher-order (backward) expectation* operator $S(k;\sigma) : \mathbb{R}^\Omega \to \mathbb{R}^\Omega$ by

$$S(k;\sigma) = E_{\sigma(k)} \cdots E_{\sigma(2)} E_{\sigma(1)}.$$

"Backward" in the definition refers to the fact that the entries of the sequence appear in reverse order when read from left to right. We define $S(0;\sigma)$ to be the identity operator for any sequence σ.

Note that $S(k;\sigma)$ is a composition of linear operators, and hence is linear. A fact observed and used extensively by Samet (1998a) is that expectations and their products are not just linear but correspond to Markov kernels (see also Gaifman (1986)). When expressed in terms of the standard basis for $\mathbb{R}^\Omega$, they are Markov matrices—nonnegative

[9] Sequences are understood to be infinite unless otherwise noted.

matrices where every row sums to 1. We will identify the E_i and products of them with the corresponding Markov matrices whenever convenient.

3 Expectations and Higher-Order Expectation Spaces

Before analyzing properties of expectations and higher-order expectations, we devote this section to identifying what we mean by expectations and higher-order expectations. First, fixing a probability space, we will define spaces of random variables corresponding to expectations and higher-order expectations. This exercise is necessary for us to identify the properties of these objects that will be important in our analysis. Second, we will construct hierarchies of expectations, paralleling the standard exercise for hierarchies of beliefs, independent of the arbitrary choices of representation.

3.1 Expectation and Higher-Order Expectation Spaces of Random Variables

This subsection defines spaces of random variables corresponding to expectations and higher-order expectations.[10]

We first let

$$V_i^0 = m\Theta$$

be the space of the *basic* random variables; *basic* is thus defined, throughout, to mean Θ-measurable. For $k \geq 1$, we now inductively define the k^{th}*-order expectations* of agent i:

$$V_i^k = \left\{ E_i f : f \in V_j^{k-1} \text{ for some } j \in I \right\}.$$

For $k \geq 2$, this is is the set of random variables which are i's expectations about various agents' $(k-1)^{\text{th}}$-order expectations. Define

$$V^\infty = \bigcup_{\substack{i \in I \\ k \geq 0}} V_i^k$$

to be the union of all the sets inductively defined so far. We will call this set the *expectation space*, noting that it also includes $V_i^0 = m\Theta$, the basic random variables that involve taking *no* expectation. Observe that

$$V^\infty = \bigcup_{\substack{\sigma \in \mathscr{I}, k \geq 0 \\ f \in m\Theta}} S(k;\sigma) f$$

We can also inductively define a space corresponding to expectations and higher-

[10] Recall the presentation of expectations as random variables from Section 2.1.

order expectations of the variables in V^∞:

$$\widetilde{V}^\infty = \left\{E_i f : i \in I, f \in V^\infty\right\}$$

We will call this space the *higher-order expectation* space. Observe that

$$\widetilde{V}^\infty = \bigcup_{\substack{\sigma\in\mathscr{I},k\geq 1\\ f\in m\Theta}} S(k;\sigma) f. \tag{2}$$

From this it is clear that $\widetilde{V}^\infty \subsetneq V^\infty$.

3.2 Expectation Hierarchies

The spaces of random variables defined in the previous section offer the most concise way of representing higher-order expectations for the purpose of stating our results in this paper. But, for several reasons, it is useful to also have an alternative representation of expectations and higher-order expectations. First, the definition of the previous section references the space Ω, but this space is being used as an ambient setting for purposes of an analyst's representation and does not have meaning in its own right. Second, we would like to make a careful analogy between (i) beliefs and higher-order beliefs and their classical representation as belief hierarchies and (ii) expectations and higher-order expectations and their representation as expectation hierarchies. Finally, we would like to address the issue of redundancy—i.e., the idea that Ω may make distinctions not relevant for players' expectations and higher-order expectations. This issue is addressed by collapsing states in Ω with identical higher-order expectation hierarchies.

3.2.1 Belief Hierarchies

We will briefly review standard material on belief hierarchies in order to compare them with expectation hierarchies. To describe belief hierarchies, we must first describe a space containing them. We write T_i^k for a set containing all possible k^{th}-level belief types of agent i. These sets can be defined inductively as follows: Let $T_i^1 = \Delta(\Theta)$ and, for all $k \geq 1$, let

$$T_i^{k+1} = \Delta\left(\left(\prod_{j\neq i} T_j^k\right) \times \Theta\right)$$

Thus a $(k+1)^{\text{th}}$-level belief type is a joint distribution over the k^{th}-level belief types of the other $n-1$ agents and the external states Θ. A belief hierarchy is now a sequence

$$\left(t_i^1, t_i^2, \ldots\right) \in H_i^{\text{belief}} := T_i^1 \times T_i^2 \times \cdots$$

For each agent i, there is natural mapping, $\tau_i : \Omega \to H_i^{\text{belief}}$, from our ambient space Ω into i's belief hierarchies. We will spell out this mapping, again using an inductive construction. Let $\tau_i^1 : \Omega \to T_i^1$ be given by

$$\tau_i^1[\omega]\left(\widehat{\theta}\right) = (t_i(\omega))\left(\{\omega' \mid \theta(\omega') = \widehat{\theta}\}\right).$$

That is, $\tau_i^1[\omega]\left(\widehat{\theta}\right)$ gives agent i's first-order belief at ω. For $k \geq 1$, let $\tau_i^{k+1} : \Omega \to T_i^{k+1}$ be defined by

$$\tau_i^{k+1}[\omega]\left(\widehat{t}_{-i}^k, \widehat{\theta}\right) = (t_i(\omega))\left(\left\{\omega' \mid \tau_j^k[\omega'] = \widehat{t}_j^k \text{ for all } j \neq i \text{ and } \theta(\omega') = \widehat{\theta}\right\}\right). \tag{3}$$

This specifies an agent's $(k+1)^{\text{th}}$-level beliefs by giving the probability he places on the occurrence of a profile $\widehat{t}_{-i}^k$ combined with a state $\widehat{\theta}$.

Finally, let agent i's *belief hierarchy* be

$$\tau_i[\omega] = \left(\tau_i^k[\omega]\right)_{k=1}^{\infty}.$$

Thus, for each state of the world ω, we have given an explicit description of agent i's higher-order beliefs.

We say that there is *no redundancy* if for any pair ω and ω' , there is an agent i with different hierarchies at those states:

$$\tau_i[\omega] \neq \tau_i[\omega'].$$

If this condition fails for some states ω and ω', we will say that those states are belief-redundant.

An important set of results about belief hierarchies concerns the existence and properties of a *universal space* of belief hierarchies with the properties that (i) all type spaces can be mapped into the universal space of belief hierarchies in a natural way;[11] and (ii) there is a homeomorphism between an agent's belief hierachies and the set of probability distributions over external states and others' hierarchies (Mertens and Zamir, 1985). We omit a formal statement of this result.

3.2.2 Expectation Hierarchies

We will construct expectation hierarchies analogously. This will coarsen the set of belief hierarchies of the previous section, because distinct higher-order beliefs can yield the same higher-order expectations. The description of expectation hierarchies will be simpler than the description of belief hierarchies. It does not need to be inductive: The space of expectations of order $n+1$ does not need to be defined in terms of expectations of order n, but can be described explicitly. However, at the end of this section, we will briefly discuss how we could define the same objects in a way that more closely parallels the belief hierarchy case.

[11] In particular, a mapping such that belief-redundant states map to the same universal type.

We will represent an agent's higher-order expectations of $\boldsymbol{\Theta}$-measurable random variables[12] via probability distributions on Θ, noting that many distributions will correspond to the same expectations. An agent's first-order expectation type is represented by a first-order belief type—that is, an element of $\Delta(\Theta)$. Agent i's second-order expectation type will specify i's expected first-order type of each other agent—that is, for each $j \neq i$, a distribution over $\Delta(\Theta)$ corresponding to i's expectation of j's expectation of $\boldsymbol{\Theta}$-measurable random variables. Thus a second-order expectation type can be represented by an element of $(\Delta(\Theta))^{n-1}$. An agent's third-order expectation type will specify that agent's expectation of the second-order type of each other agent j, which we have just described, and so on.

To represent these objects in general, write $\mathscr{S}(k)$ for the set of *finite* sequences of elements of I which are of length k and let $\mathscr{S}(k;i)$ be the set of all sequences $s \in \mathscr{S}(k)$ with $s(k) = i$.[13] Now, the set of k^{th}-level expectation types of agent i is defined as

$$X_i^k = (\Delta(\Theta))^{\mathscr{S}(k,i)}. \tag{4}$$

An element of X_i^k (a k^{th}-level expectation type) can be written as

$$x_i^k = \left(x_i^k(s)\right)_{s \in \mathscr{S}(k,i)},$$

where $x_i^k(s) \in \Delta(\Theta)$. (Note that k is fixed and the indexing in the tuple is over *sequences*.) As we will state formally below, for a given sequence s the distribution $x_i^k(s)$ corresponds to the measure associated with the higher-order expectation

$$E_{s(k)} \cdots E_{s(2)} E_{s(1)},$$

which, as in Section 2.3, is a backward expectation. An *expectation hierarchy* is then an element of

$$\left(x_i^1, x_i^2, \dots\right) \in H_i^{\text{expectation}} := X_i^1 \times X_i^2 \times \cdots$$

For each agent i, there natural mapping, $\xi_i : \Omega \to H_i^{\text{expectation}}$, from our ambient space Ω into i's expectation hierarchies,. An explicit construction follows. For any sequence $s \in \mathscr{S}(k;i)$, we will define $\xi_i^{k,s} : \Omega \to X_i^k$, and for every $\omega \in \Omega$ and $s \in \mathscr{S}(k;i)$, we will define $x_i^k(s)[\omega] \in \Delta(\Theta)$. In order to do this, we must specify the value the measure $\mu = x_i^k(s)[\omega]$ assigns to every $F \in \boldsymbol{\Theta}$; we do this as follows:

$$\mu(F) = (E_{s(k)} E_{s(k-1)} \cdots E_{s(2)} E_{s(1)} \mathbf{1}_F)[\omega]. \tag{5}$$

[12] Recall from Section 3.1 that *basic* is defined to mean $\boldsymbol{\Theta}$-measurable.

[13] Recall our introduction of infinite sequences in Section 2.1; we will follow the convention that the letter s is used for a typical finite sequence and σ is used for a typical infinite sequence.

Finally, putting into one (infinite) tuple these values across all lengths k, we get

$$\xi_i[\omega] = \left(x_i^k[\omega]\right)_{k=1}^{\infty}$$

Thus, for each state of the world ω, we have given an explicit description of agent i's higher-order expectation hierarchy. It can be seen that, by construction, ξ_i gives rise to a coarser partition than τ_i.

We now have a corresponding notion of "redundancy": We will say that there is *no expectation redundancy* if, for all ω and ω' , there is an agent i with

$$\xi_i[\omega] \neq \xi_i[\omega']$$

If this condition fails for some states ω and ω', we will say that those states are *expectation-redundant*.

Expectation types at any level k were defined directly using sequences of length k, but could also have been done inductively to more closely parallel our definition of belief types. For a given level k, an agent's expectation hierarchy corresponds to that agent's first-order expectation of other agents' expectation hierarchies of level $k-1$, where we collapse everything to ultimately yield a measure over $\Delta(\Theta)$. Proceeding this way gives an expectations version of equation (3) and a closer analogy to our construction of belief hierarchies. In the case of expectations, it is possible to roll together the inductive construction into (4), as we have done; there is no good analogue of this for beliefs.

We leave the construction of a universal expectation space as an open question.

4 Limits of Higher-Order Expectations and Order-Independence

4.1 Limits of Higher-Order Expectations and Samet's Result

We first define the relevant notion of convergence of higher-order expectations. Recall that σ refers to an arbitrary infinite sequence.

Condition 1 (Convergence to a Deterministic Limit)**.** Convergence to a deterministic limit occurs with respect to sequence σ if

i) $S(k;\sigma)$ converges[14] to some limit operator $S(\infty;\sigma)$;

ii) for any $f \in \mathbb{R}^{\Omega}$, the limit $S(\infty;\sigma)f$ is nonrandom (i.e., constant).

[14] A sequence of operators $(S_k)_{k=1}^{\infty}$ converges to S if it converges in the operator norm: if $\sup_{f:\|f\|_\infty \leq 1} \left\|(S_k - S)f\right\|_\infty \to 0$. Here $\|\cdot\|_\infty$ denotes the supremum norm.

Proposition 1 in Samet (1998a) established that there is convergence to a deterministic limit for every I-sequence σ. Given convergence to a deterministic limit, there is a linear map $\ell(\sigma) : \mathbb{R}^{\Omega} \to \mathbb{R}$ such that $\ell(\sigma) = S(\infty;\sigma)$. Because $\ell(\sigma)$ is a limit of expectation operators, we can identify $\ell(\sigma)$ with a limit measure on Ω. The limit measure will, in general, depend on the I-sequence σ.

Samet (1998a) then studies when $\ell(\sigma)$ is independent of the sequence σ, defining order-independence with respect to all random variables.

Condition 2 (Full Order-Independence)**.** For any random variable $f \in m\Omega$, the limit expectation $l(\sigma)f$ along any I-sequence σ does not depend on σ.

The following definition is the interim version of the common prior assumption.

Condition 3 (Belief-Consistency)**.** There exists a $p \in \Delta(\Omega)$ such that p is a prior for each agent.

The main result in Samet (1998a) is then:

Proposition 1. *Full order-independence and belief-consistency are equivalent.*

This result can be deduced from our results below. As discussed in the introduction, a paradoxical feature of this result is that full order-independence concerns expectations and higher-order expectations only, while belief-consistency imposes a condition on all beliefs and higher-order beliefs. Given a random variable, expectations of it are strictly less informative about beliefs than the beliefs themselves; the expectation of any random variable f is determined by all beliefs about the events $\{\omega : f(\omega) = a\}$ but not vice versa. Thus, it cannot be only by reference to higher-order expectations of a random variable f that consistency of beliefs *about* f is established.

The resolution of the paradox is that by allowing *any* random variable on Ω as the initial variable whose higher-order expectations are taken, Samet is already allowing this random variable to depend on beliefs and higher-order beliefs in an arbitrary way. The condition of order-independence is imposed with any such variable as the initial one. In the approach pursued in this paper, the key contrast is this: We impose conditions such as order-independence not on all of $m\Omega$ but only on the (typically strict) subset V^{∞}—that is, the set of higher-order expectations of basic random variables. The next subsection studies the generalization of Samet's result to this setting.

4.2 Order-Independence and Expectation-Consistency

Relative to the condition of full order-independence, the following condition is imposed on a smaller class of random variables ($m\boldsymbol{\Theta}$ rather than $m\Omega$).

Condition 4 (Order-Independence)**.** For any $\boldsymbol{\Theta}$-measurable random variable $f \in m\boldsymbol{\Theta}$, the limit expectation $l(\sigma)f$ along any I-sequence σ does not depend on σ.

We will characterize this in terms of a more primitive condition on beliefs. It is a weakening of belief-consistency that requires beliefs to be consistent only in the sense of giving the same expectations of random variables in V^∞. A fuller discussion of the contrast with belief-consistency is in Section 4.3.

Condition 5 (Expectation-Consistency)**.** There exists a tuple $(p_i)_{i\in I}$ of priors ($p_i \in \Delta(\Omega)$ is a prior for agent i), such that, for each i and each $f \in V^\infty$, we have $p_i f = p_j f$ for all i and j. In this case, we call $(p_i)_{i\in I}$ the *expectation-consistent priors.*

The following generalization of Samet (1998a) will be derived later from Proposition 3:

Proposition 2. *Order-independence and expectation-consistency are equivalent.*

This result is tight in the following sense: In the absence of expectation-consistency, the condition of order-independence cannot hold in general: $\ell(\sigma)$ may unavoidably depend on $\sigma(1)$. For example, suppose agent $\sigma(1)$'s conditional probability of some event $F \subseteq \Theta$ is always in $[\alpha, \beta]$, no matter what the state—i.e., $\left[E_{\sigma(1)}\mathbf{1}_F\right](\omega) \in [\underline{\alpha}, \overline{\alpha}]$ for all $\omega \in \Omega$. Then all higher-order expectations of $\mathbf{1}_F$ along the sequence also take values in $[\underline{\alpha}, \overline{\alpha}]$, since they are expectations of agent $\sigma(1)$'s first-order belief. It follows that the limit measure $\ell(\sigma)$ must put a probability on Fthat lies in $[\underline{\alpha}, \overline{\alpha}]$. Since agent $\sigma(1)$ and the interval $[\underline{\alpha}, \overline{\alpha}]$ were arbitrary, it follows that order-independence cannot hold. Our main result, in the next section, exactly characterizes when the dependence on the initial agent just described is the *only* dependence there is—that is, when $\ell(\sigma)$ depends *only* on $\sigma(1)$. This reduces to Proposition 2 of Samet (1998a) in the case where the partition $\boldsymbol{\Theta}$ is the finest one on Ω.

4.3 Comparison of Expectation-Consistency with Consistency

Expectation-consistency is a significant weakening of belief-consistency. We can highlight this by decomposing the weakening into three parts.

First, belief-consistency requires that players' priors agree on *all* states, even about the relative probabilities of states that are belief-redundant in the sense that they are indistinguishable based on agents' beliefs and higher-order beliefs about any random variable defined on Ω. (Recall Section 3.2.1.)

Second, even if we defined Ω to rule out belief redundancies (by requiring that the coarsest common refinement of agents' partitions consist of singletons), it could still be that the agents' priors agree on the partitions corresponding to all beliefs and higher-order beliefs about $\boldsymbol{\Theta}$ but do not agree on every subset of Ω. This would correspond to the agents disagreeing about some random variables "orthogonal" to θ.

Finally, even if this situation is also ruled out—by assuming that all uncertainty in Ω corresponds to beliefs and higher-order beliefs about $\boldsymbol{\Theta}$[15]—expectation-consistency

[15] That is, the mapping $\tau_i : \Omega \to H_i^{\text{belief}}$ of Section 3.2.1 is one-to-one.

is still weaker than belief-consistency because expectations and higher-order expectations of Θ-measurable random variables do not fully determine an agent's beliefs. For example, consider two situations: (i) agent 1 believes that agents 2 and 3 always have identical types and (ii) agent 1 believes that agents 2 and 3 have types that are conditionally independent given the realization of Θ. Holding marginals fixed, higher-order expectations cannot detect this difference.

5 Higher-Order-Independence and Higher-Order Expectation-Consistency

Section 4 studied the situation in which agents' priors are consistent in terms of their beliefs (or expectations) about Θ. In the remainder of the paper, we focus on the possibility of some disagreement. Our goal is to characterize the values of higher-order expectations, especially when they can be represented via *heterogeneous* priors associated with various agents. A key condition we will study is higher-order-independence:

Condition 6 (Higher-Order-Independence)**.** For any Θ-measurable random variable $f \in m\Theta$, the limit expectation $l(\sigma)f$ depends on σ only through its initial index $\sigma(1)$: for any two infinite sequences σ and σ', if $\sigma(1) = \sigma'(1)$, then $\ell(\sigma)f = \ell(\sigma')f$.

A more accurate name would be *higher-order order-independence*, since the condition requires that the limit expectation depend on the sequence only through the identity of agent $\sigma(1)$, and that it be otherwise independent of the sequence in which higher-order expectations are taken.

We will again characterize higher-order-independence in terms of a more primitive condition on beliefs. This condition imposes consistency on higher-order expectations about random variables measurable with respect to some agent's types:

Condition 7 (Higher-Order Expectation-Consistency)**.** There exists a tuple $(p_i)_{i\in I}$ of priors ($p_i \in \Delta(\Omega)$ is a prior for agent i) such that, for each i and each $f \in \widetilde{V}^\infty$, we have $p_i f = p_j f$ for all i and j. If this last condition holds, we call $(p_i)_{i\in I}$ *higher-order expectation-consistent priors.*

Higher-order expectation-consistency requires that there be priors for all of the agents that assign the same expectation to all random variables measurable with respect to any agent's higher-order expectations (which include first-order expectations). Note that the *same* family of priors must satisfy this requirement simultaneously for all the pairs of agents. Under our maintained assumptions, there is a unique profile of higher-order expectation-consistent priors when they exist.[16] As with expectation-consistency, it is an agent's beliefs over each other agent's types, and not the correlation between them, that matters. But this property is a considerable weakening of

[16] This follows from Proposition 3.

expectation-consistency. First, it puts no restrictions on first-order expectations. Second, as we will be able to argue when we present Corollary 1 below, it always holds when there are two agents (while, of course, expectation-consistency may fail even with two agents.)

Now we have:

Proposition 3. *1. Higher-order-independence and higher-order expectation-consistency are equivalent.*

2. When either condition holds, the expectation of any $\boldsymbol{\Theta}$-measurable random variable under $\ell(\sigma)$ is equal to its expectation under $p_{\sigma(1)}$ in any profile of higher-order expectation-consistent priors.

Proof. Suppose higher-order expectation-consistency holds, with the tuple $(p_i)_{i\in I}$ of agent priors. Fix some I-sequence σ. We will show that $\ell(\sigma)f$ is equal to $p_{\sigma(1)}f$ for each $\boldsymbol{\Theta}$-measurable random variable f—that is, higher-order-independence holds.

By the definition of a prior, for any $\boldsymbol{\Theta}$-measurable random variable f we have $p_{\sigma(1)}f = p_{\sigma(1)}E_{\sigma(1)}f$. Because $E_{\sigma(1)}f \in \widetilde{V}^{\infty}$, higher-order expectation-consistency implies that the right-hand side is also equal to $p_{\sigma(2)}E_{\sigma(1)}f$. Combining these facts, we have

$$p_{\sigma(1)}f = p_{\sigma(2)}E_{\sigma(1)}f. \tag{6}$$

Applying the same conclusion with the random variable $E_{\sigma(1)}f$ in place of f, agent $\sigma(2)$ in place of $\sigma(1)$, and agent $\sigma(3)$ in place of $\sigma(2)$, we see that $p_{\sigma(2)}E_{\sigma(1)}f = p_{\sigma(3)}E_{\sigma(2)}E_{\sigma(1)}f$. (We used here that $E_{\sigma(2)}E_{\sigma(1)}f$ is in $\widetilde{V}^{\infty}$.) Combining that result with (6) shows that

$$p_{\sigma(1)}f = p_{\sigma(3)}E_{\sigma(2)}E_{\sigma(1)}f.$$

We can continue in this way inductively to conclude that $p_{\sigma(1)}f = p_{\sigma(k)}S(k;\sigma)f$ for every k, and consequently, since $S(\infty;\sigma)f$ is a constant $\ell(\sigma)f$, we have $p_{\sigma(1)}f = \ell(\sigma)f$. Since f was arbitrary, this establishes the claim.

Conversely, suppose higher-order-independence holds. Our goal is to construct a family of priors $(p_i)_{i\in I}$ so that for any $g \in \widetilde{V}^{\infty}$ and any i and j, we have $p_i g = p_j g$.

Define $p_i = \ell(\sigma)E_i$, where σ is an arbitrary I-sequence. We prove two claims. First, that for random variables $g \in \widetilde{V}^{\infty}$, the expectation $p_i g$ does not depend on the choice of σ. Second, that beyond being well-defined, $p_i g$ does not depend on the choice of i.

By equation (2) in Section 3.1, any $g \in \widetilde{V}^{\infty}$ can be expressed as

$$g = E_{s(k)} \cdots E_{s(2)}E_{s(1)}f, \tag{7}$$

where $f \in m\boldsymbol{\Theta}$ and $s(\cdot)$ is some *finite* sequence. Thus it suffices to show that $p_i g (= \ell(\sigma)E_i g)$ does not depend on the choice of σ or i. By (7) we can rewrite this quantity as

$$p_i g = \ell(\sigma)E_i E_{s(k)} \cdots E_{s(2)}E_{s(1)}f.$$

By higher-order-independence, this does not depend on σ, showing that p_i is well-defined. But the same property of higher-order-independence also shows that the right-hand side cannot depend on any index in the sequence of higher-order expectations except $s(1)$, and in particular not on i. Therefore, $p_i g = p_j g$ for any i and j. □

Now, we show that this result readily implies Proposition 2: First, suppose that expectation-consistency holds. Then higher-order expectation-consistency holds with $p_i = p$ for some single p, and so Proposition 3 yields that higher-order independence holds—and moreover, that for any σ and $f \in m\Theta$, we have $\ell(\sigma)f = p_{\sigma(1)}f = pf$. Thus we in fact have full order-independence. Conversely, suppose order-independence holds. Then higher-order independence holds. Proposition 3 then yields that higher-order expectation-consistency holds with some profile of priors $(p_i)_{i\in N}$ for which $\ell(\sigma)f = p_{\sigma(1)}f$ holds whenever $f \in m\Theta$. Now, because full order-indepenence holds, we know that the left-hand side of this equation cannot depend on $\sigma(1)$. In other words, there is some p that satisfies $\ell(\sigma)f = pf$ whenever $f \in m\Theta$. That verifies expectation-consistency.

5.1 Interpretation and Discussion

Because expectation-consistency is related to belief-consistency and the common prior assumption, the notion of higher-order expectation-consistency resembles a common prior *on individuals' beliefs* about $\boldsymbol{\Theta}$—a condition we might call *Consistency of Higher-Order Beliefs*. For example, if beliefs arise from signals, then expectation-consistency is implied by a common prior about the signals agents receive. We pursue this interpretation in detail in Golub and Morris (2017).

However, higher-order expectation-consistency is weaker than consistency of higher-order beliefs. A corollary of Proposition 3 illustrates this:

Corollary 1. *If there are two agents, then higher-order expectation-consistency always holds.*

The reason is as follows: For any random variable $f \in m\boldsymbol{\Theta}$, the only two limit higher-order expectations to consider are

$$\cdots E_1 E_2 E_1 E_2 E_1 f \qquad \text{and}$$
$$\cdots E_2 E_1 E_2 E_1 E_2 f$$

and so Samet's (1998a) result that both limits exist gives higher-order-independence. This proves the corollary.

In contrast, even with two agents it is possible not to have consistency of higher-order beliefs. The essential reason is that $\widetilde{V}^{\infty}$ contains only events measurable with respect to each individual's information, while consistency of higher-order beliefs would require consistency on product events as well. For example, consider a case with two agents where one is sure that there is common certainty of $f = 1$ and the other is sure that there is common certainty of $f = 2$ (both random variables being degenerate or constant). There is clearly not consistency of higher-order beliefs, but on the other hand higher-order expectation-consistency does hold.

6 Higher-Order Average Expectations and Separability

In this section, we apply the characterizations derived earlier to analyze *higher-order average expectations*, which come up in a variety of applications, including network games.

An average expectation of a random variable is defined to be a weighted average of agents' expectations. One can then take average expectations of individuals' average expectations. Indeed, we can proceed in this way on and consider the *limit* of such higher-order average expectations. As we will discuss below, these objects arise in applications, and they turn out to have nice connections to the higher-order expectations we have been discussing.

In defining higher-order average expectations, we must specify what weights are used in taking averages. For this, we fix a nonnegative row-stochastic matrix Γ whose rows and columns are indexed by I; assume it is irreducible and aperiodic. Let $\boldsymbol{\Gamma}$ be the class of all such matrices.[17] In taking averages from i's perspective, the weights placed on others come from row i of Γ, so that $\Gamma(i,j)$ is the weight that i places on j.

In this section, we define and characterize higher-order average expectations. We then show how to establish their properties—especially an important one called *separability*—using results from the previous section.

6.1 Definitions and Main Result

Define the operators $A_i : m\Omega \to m\Omega$ by

$$\begin{aligned} A_i(1;\Gamma) &= E_i; \\ A(k+1;\Gamma) &= \sum_{j\in I} \Gamma(i,j) A_j(k;\Gamma), \qquad k = 1,2,\ldots. \end{aligned}$$

The operator $A_i(1;\Gamma)$ is agent i's expectation operator; $A_i(2;\Gamma)$ is agent i's expectation of the average of others' expectations; $A_i(3;\Gamma)$ is agent i's expectation of the average expectation of the average expectation, and so on. In each case, in taking averages from i's perspective, the weights placed on others come from row i of matrix Γ.

Our motivation for studying higher-order average expectations comes from network games with linear best responses and incomplete information (Bergemann et al. (2015),

[17] Note that $\boldsymbol{\Gamma}$ is italicized. The irreducibility and acyclicity assumptions do involve loss of generality, but analogous results hold with these assumptions relaxed. We will discuss below where we use each of these.

Blume et al. (2015), and Golub and Morris (2017)). In these games, predictions—e.g., Nash equilibria or rationalizable outcomes—depend on higher-order average expectations. More precisely, fix a random variable f. Suppose agents choose actions a_i in a compact interval containing the maximum and minimum values of f. The matrix Γ that we introduced at the start of Section 6 has the interpretation that $\Gamma(i,j)$ is the weight that a given player puts on others' actions; here we assume $\Gamma(i,i)=0$. For a constant $\beta \in (0,1)$, the best response of agent i is

$$\mathrm{BR}_i(a_{-i}) = (1-\beta)E_i f + \beta \sum_{j\neq i} \Gamma(i,j) E_i a_j.$$

According to this best-response function, the agent wants to take an action equal to a convex combination of his subjective expectation of the random variable f and the average of his expectations of others' actions. As discussed in **?**,[18] agent i's action in any rationalizable action profile is given by

$$\sum_{k=1}^{\infty} \beta^{k-1} A_i(k;\Gamma) f. \tag{8}$$

The limit $\beta \uparrow 1$ is a particular focus of Golub and Morris (2017). It is clear that behavior in this high-coordination limit is determined by $A_i(k;\Gamma)f$ for large k, which motivates our study of this limit.

One basic question about the limit is whether it exists. In order to study it, we first define an adjusted version of convergence to a deterministic limit.

Condition 8 (Convergence to a Deterministic Limit for Higher-Order Average Expectations)**.** This condition occurs with respect to network Γ if the following hold:

i) For each i, the sequence $A_i(k;\Gamma)$ converges, as $k\to\infty$, to some limit operator $c(\Gamma)$, independent of i;

ii) for any $f \in \mathbb{R}^{\Omega}$, the limit $c(\Gamma)f$ is nonrandom (i.e., constant).

We will show that this convergence always holds under our maintained assumptions.[19]

Suppose there is convergence to a deterministic limit. Because $c(\Gamma)$ is a limit of expectation operators, we will, as before, identify it with a limit measure $c(\Gamma)$ on Ω. We will refer to $c(\Gamma)$, in either sense, as the *consensus expectation*. The analogue of order-independence is now:

[18] See Ui (2009) for a more general version of this fact.

[19] This requires our maintained assumption of acyclicity as well as irreducibility. To illustrate the need for irreducibility, take $n=2$ and $\Gamma(1,2)=\Gamma(2,1)=1$. Then $A_1(0;\Gamma)=E_1$, $A_1(1;\Gamma)=E_1E_2$, $A_1(3;\Gamma)=E_1E_2E_1$,and so on. This sequence will converge to a two-cycle. Without acyclicity in such cases, we would still have convergence of the Abel average $\lim_{\beta\uparrow 1}\sum_{k=1}^{\infty}\beta^{k-1}A_j(k,\Gamma)f$, which is what is relevant for the game (see equation (8)).

Condition 9 (Network-Independence)**.** If both Γ and Γ' are irreducible and aperiodic, then $c(\Gamma) = c(\Gamma')$.

Recall our study of the condition of higher-order expectation-consistency, which was equivalent to higher-order-jndependence. Under that condition, the limit of higher-order expectations (the consensus expectation) along a sequence was determined by a public prior on the type of the agent appearing rightmost in the sequence. In the network context, there is no longer a single sequence, so there is no direct analogue of this condition. It turns out that an appropriate version of the condition is *separability*. This condition states that the consensus expectation is a linear combination of network-independent priors (one for each agent) with weights on these priors determined by the network.

Condition 10 (Separability)**.** There exist a tuple $(p_i)_{i \in I}$ of priors ($p_i \in \Delta(\Omega)$ is a prior for agent i), and real-valued functions $e_i(\Gamma)$ such that for all networks Γ

$$c(\Gamma) = \sum_i e_i(\Gamma) p_i \tag{9}$$

Separability is an important property in the context of our network game. If separability holds, equilibrium play in the $\beta \uparrow 1$ limit is determined by the network and the beliefs in a simple way: The numbers $e_i(\Gamma)$, which are weights depending on the network, determine the relative importance of the measures p_i.

Separability plays a central role in Golub and Morris (2017). If it holds, our results there extend those in the literatures on network games and on incomplete information, as we explain there. If separability fails, qualitatively new properties emerge, because beliefs and the network become "entangled" in a more complex way than separability permits. It is therefore important to understand the dichotomy between separability and non-separability. Here we can give an exact characterization of separability.[20]

Before stating our main result on this, we define eigenvector centrality.

Definition. Given a matrix $\Gamma \in \boldsymbol{\Gamma}$, the eigenvector centrality vector $e(\Gamma)$ is defined to be the unique vector $e \in \Delta(I)$ satisfying $e\Gamma = e$.

Since $\Gamma \in \boldsymbol{\Gamma}$, uniqueness of e in the definition follows from the Perron–Frobenius Theorem.

Now we can state the following proposition, which summarizes network analogues of Propositions 1, 2, and 3.

Proposition 4.

1. *Higher-order average expectations satisfy convergence to a deterministic limit.*

2. *Beliefs satisfy expectation-consistency if and only if network-independence holds.*

[20] Our formalism in Golub and Morris (2017) is not well-suited to this, because (among other things) it does not handle redundancy well. For discussion of redundancy, see Sections 3.2.1 and 3.2.2.

3. *Beliefs satisfy higher-order expectation-consistency if and only if separability holds. Moreover, in this case:*

   3.i. *the $(p_i)_{i \in N}$ in the definition of separability are higher-order expectation-consistent priors;*

   3.ii. *$e(\Gamma)$ is the eigenvector centrality vector of Γ.*

The rest of this section is devoted to proving this result.

6.2 Key Ideas Behind Proposition 4

We start with a high-level overview of the key ideas involved in Proposition 4. The first step is to reinterpet higher-order average expectations via an artificial random process. It turns out that higher-order average expectations are the same as *expected* higher-order expectations along a sequence that is drawn *randomly* in a certain way, according to a Markov process reflecting players' network weights. In particular, we select a starting agent, say Ann, then draw the next agent in proportion to Ann's coordination weights, and continue in this way to construct a random sequence. The expected operator we get (averaging across random sequences) turns out to be the higher-order average expectation.

The second step is to understand limiting higher-order expectations along such random sequences when the sequences are long. One key tool is Proposition 3, in which we showed that, under higher-order expectation-consistency, the limiting higher-order expectation along a sequence depends exclusively on the public prior corresponding to the rightmost agent in the sequence of expectations, called $\sigma(1)$ in that result.[21]

The third step is to characterize the distribution of the final agent in the random sequence. Assuming that the matrix of network weights (and thus the Markov process on agents) is irreducible, the Markov process has an ergodic distribution. Moreover, the entry corresponding to an agent in this ergodic distribution is (just by definition) his eigenvector centrality in the network of weights. Thus, the probability that a long sequence ends in a given agent is that agent's ergodic weight, or eigenvector centrality. Putting these facts together makes higher-order average expectations converge to a centrality-weighted sum of public prior expectations, no matter which agent appears leftmost in the sequence. While this description covers the essence of the final step, there are some technical subtleties in carrying it out, which we defer to Section 6.4.1.

6.3 A Markovian Formalism for Higher-Order Average Expectations

We now formalize the outline, starting with a Markovian formalism for higher-order average expectations.

[21] This involves some restrictions on the sequences, which will be satisfied when they are drawn in the Markovian way we describe.

We have claimed that the $A_i(k;\Gamma)$ can be related to a Markov process on agents. To make this precise, it will be useful to define an object $(A_i(k;\Gamma))_{i\in I}$ that keeps track of all the k^{th}-order average expectations simultaneously, and to define a Markov iteration on it. To carry this out, let the *higher-order average expectation profile* $A(k,\Gamma)$ be defined as follows: For any $f \in \mathbb{R}^{\Omega}$, define $A(k,\Gamma)f \in \mathbb{R}^{I\times\Omega}$ by stipulating that the (i,ω) entry of $A(k;\Gamma)f$ is $(A_i(k;\Gamma)f)(\omega)$. This defines a linear operator $A(k;\Gamma) : \mathbb{R}^{I\times\Omega} \to \mathbb{R}^{I\times\Omega}$. Next, define a Markov kernel B_Γ on $I \times \Omega$ by

$$B_\Gamma((i,\omega),(i',\omega')) = \Gamma(i,i')\cdot(t_i(\omega))(\omega').$$

The transition probability from state (i,ω) to state (i',ω') is defined to be the product of (i) the weight individual i places on i' and (ii) the subjective probability that type $t_i(\omega)$ of i places on state ω'.[22]

It can be checked from the definition of $A(k;\Gamma)$ that

$$A(k+1;\Gamma) = B_\Gamma A(k;\Gamma),$$

so that

$$A(k;\Gamma) = B_\Gamma^{k-1} A(1;\Gamma), \tag{10}$$

where the superscript $k-1$ corresponds to the $(k-1)$-fold application of B_Γ, i.e., the corresponding matrix power.

6.3.1 Convergence

Using this formalism, we can state and prove our result on the convergence of higher-order average expectations (under our maintained assumptions).

Lemma 1. *Convergence to a deterministic limit holds.*

Proof. The irreducibility of B_Γ can be deduced from the assumed irreducibility of Γ and the maintained assumption of no nontrivial common certainty. Similarly, aperiodicity of B_Γ can be shown from the aperiodicity of Γ along with the maintained assumption of common support. Then it follows from the standard Markov chain convergence result that $A(k;\Gamma)f$ converges to a constant vector as $k \to \infty$, which (by definition) yields convergence to a deterministic limit for higher-order average expectations. □

6.3.2 Form of the Limit

Standard results on the Markov iteration summarized in (10) can deliver more information about higher-order average expectations. Under our maintained assumptions,

[22] It is a consequence of this that the transition probability from any (i,ω) where $\omega \in P_i \in \mathscr{P}_i$ to $\{i'\} \times \mathscr{P}_{i'}$ where $P_{i'} \in \mathscr{P}_{i'}$ is equal to $\Gamma(i,i')$ multiplied by the subjective probability that i (given that he has type P_i) places on type $P_{i'}$ of i'. Here we have identified types of agents with elements of the corresponding partitions. Golub and Morris (2017) has a notation that is more suited to expressing the subjective probabilities that various types place on one other.

there is a uniquely determined stationary distribution $\mu(\Gamma) \in \Delta(I \times \Omega)$, with no zero entries, such that, if viewed as a row vector in $\mathbb{R}^{I\times\Omega}$, it satisfies the stationarity equation

$$\mu(\Gamma)B_\Gamma = \mu(\Gamma). \tag{11}$$

For any i, let $\mu_i(\Gamma)$ refer to the row vector in $\mathbb{R}^\Omega$ consisting of entries (i,ω) of $\mu(\Gamma)$ as ω ranges over Ω. Then, recalling (10) and our definition of the higher-order average expectation profile $A(0,\Gamma)$, we may rewrite $c(\Gamma)f$ as follows:

$$c(\Gamma)f = B_\Gamma^\infty A(1,\Gamma)f = \sum_i \mu_i(\Gamma)E_i f. \tag{12}$$

Proposition 4 can be established based on this expression, as we now outline.

The following lemma allows us to analyze the entries of $\mu_i(\Gamma)$.

Lemma 2. *If Γ is irreducible, then for each i we have*

$$\sum_{\omega\in\Omega} \mu_{(i,\omega)}(\Gamma) = e_i(\Gamma),$$

where $e(\Gamma) \in \Delta(I)$ is the (strictly positive) stationary distribution of Γ when Γ is viewed as a Markov transition matrix on the agents.

That is, the entries of $\mu_i(\Gamma)$ add up to the number $e_i(\Gamma)$, the eigenvector centrality of i. One proof that can be adapted to the present setting appears as Lemma 1 in Golub and Morris (2017). We give a self-contained proof here.

Proof. Define $\widehat{\mu}(\Gamma) \in \Delta(I)$ by $\widehat{\mu}_i = \sum_{\omega\in\Omega} \mu_{(i,\omega)}(\Gamma)$. We will show that $\widehat{\mu}(\Gamma) = e(\Gamma)$.

Define a matrix M whose rows are indexed by $I \times \Omega$ and whose colums are indexed by I. Let

$$M((i,\omega), j) = \mathbf{1}_{i=j}.$$

Post-multiply both sides of the stationarity condition $\mu(\Gamma)B_\Gamma = \mu(\Gamma)$ by M to get

$$\mu(\Gamma)(\mathbf{1}_{|\Omega|} \otimes \Gamma) = \widehat{\mu}(\Gamma),$$

where $\mathbf{1}_{|\Omega|}$ is the vector of all ones with length $|\Omega|$, and $\mathbf{1}_{|\Omega|} \otimes \Gamma$ denotes the Kronecker product of the column vector $\mathbf{1}_{|\Omega|}$ with the matrix Γ. This is equivalent to

$$\widehat{\mu}(\Gamma)\Gamma = \widehat{\mu}(\Gamma).$$

Since for an irreducible Markov matrix Γ, there is just one probability distribution that is a left-hand Perron eigenvector of Γ, it follows that $\widehat{\mu}(\Gamma) = e(\Gamma)$. □

We can now define $\overline{\mu}_i(\Gamma) \in \Delta(\Omega)$ to be $\mu_i(\Gamma)$ divided by $e_i(\Gamma)$; because Γ is irreducible, all the entries $e_i(\Gamma)$ of the left-hand unit eigenvector Γ are positive. Then

$$\mu_{(i,\omega)}(\Gamma) = e_i(\Gamma)\overline{\mu}_{(i,\omega)}(\Gamma), \tag{13}$$

and we may rewrite (12) as

$$c(\Gamma)f = \sum_i e_i(\Gamma)\overline{\mu}_i(\Gamma)E_i f. \tag{14}$$

The content of statement (3) of Proposition 4 is that beliefs satisfy higher-order expectation-consistency if and only if $\overline{\mu}_i(\Gamma)$ is an expectation-consistent prior for i (independent of Γ).

In the following subsection, we introduce some useful methods based on higher-order expectations for studying (12), and then use them, along with the formulas established here, to deduce statements (2) and (3) of Proposition 4.

6.4 Connection between Higher-Order Expectations and Higher-Order Average Expectations

In Section 2.3, we studied *higher-order backward expectations* of the form

$$S(k;\sigma) = E_{\sigma(k)} \cdots E_{\sigma(1)}.$$

It turns out that higher-order average expectations relate more closely to *higher-order forward expectations*. To study these, define

$$T(\sigma;k) := E_{\sigma(1)}E_{\sigma(2)} \cdots E_{\sigma(k)}.$$

Observe that

$$A_i(1;\Gamma) = E_i; \qquad A_i(k+1;\Gamma) = \mathbb{E}_{j\sim\Gamma(i,\cdot)}[A_j(k;\Gamma)]. \tag{15}$$

In words, the first-order average expectation is just E_i. The $(k+1)^{\text{th}}$-order average expectation is the expected k^{th}-order expectation of agent j, where j is drawn according to the probability distribution in row i of Γ.

Now, fix $i \in I$ and let σ_i denote a *random* sequence of agents starting at i and generated by a Markov process on I with transition matrix Γ. Now, starting with the observation that $T(\sigma_i;1) = A_i(1;\Gamma)$, we can show, by repeatedly applying (15), that

$$A_i(k;\Gamma) = \mathbb{E}[T(\sigma_i;k)].$$

Here the expectation $\mathbb{E}$ is with respect to the distribution over sequences σ induced by the Markov process just described. Thus, a higher-order average expectation is the expected value of a higher-order forward expectation, where the sequence of agents is generated according to a Markov process.

Using this way of writing the higher-order average expectations, we will establish the key lemma behind Proposition 4:

Lemma 3. *If beliefs satisfy higher-order expectation-consistency, then separability holds. Moreover, in this case, the $p_i (i \in I)$ in the definition of separability are higher-order* *expectation-consistent priors, and $e_i(\Gamma)$ is the i^{th} entry of the stationary distribution of Γ viewed as a Markov chain.*

6.4.1 The Idea behind the Proof of the Lemma

The lemma corresponds to the final step in our outline in Section 6.2. Recall that the final step just requires only that we understand the expected value of higher-order forward expectations as we generate a random forward sequence of agents according to the Markov process Γ. However, there are technical subtleties involved in doing this.

If we think of that expected value as a weighted average of higher-order expectations (each for some deterministic sequence), there are more and more such sequences to average over as the order of the higher-order average expectation increases. The leftmost agent, $\sigma(1)$ in $E_{\sigma(1)}E_{\sigma(2)}\cdots E_{\sigma(k)}$, is always the same, but the sequences become very different further along. It is not obvious how to understand this profusion of higher-order expectations along different sequences using our main tool so far, a result on the limit of higher-order *backward* expectations along a single infinite sequence.

The technical problem can be solved by using the idea of Markov chain reversal. That is, rather than thinking of a random sequence—e.g., $\sigma(1),\sigma(2),\ldots,\sigma(k)$—as being constructed starting with $\sigma(1)$, we instead think of it as starting with $\sigma(k)$ and being built up backwards. It is possible to do this in a way that keeps the distribution of sequences the same—that is the essence of Markov chain reversal. The advantage of the backward approach is that sequences can now be naturally grouped into sets that have the same (long) segment at the end (closest to the random variable). That allows us to apply our prior results on the limits of backward expectations.

6.4.2 Completion of the Proof of Lemma 3

Proof of Lemma 3. Suppose higher-order expectation-consistency holds. Recall that

$$A_i(k;\Gamma) = \mathbb{E}[T(\sigma_i;k)].$$

By the law of iterated expectations, breaking the outcome up according to the value of the final index $\sigma(k)$, we have

$$\mathbb{E}[T(\sigma_i;k)] = \sum_{j\in I}\mathbb{P}(\sigma_i(k)=j)\mathbb{E}[T(\sigma_i;k)\mid \sigma_i(k)=j].$$

Hence

$$\lim_{k\to\infty}\mathbb{E}[T(\sigma_i;k)] = \lim_{k\to\infty}\sum_{j\in I}\mathbb{P}(\sigma_i(k)=j)\mathbb{E}[T(\sigma;k)\mid \sigma_i(k)=j].$$

Now, since all operators involved are uniformly bounded (in total variation norm), and $\mathbb{P}(\sigma_i(k)=j)$ converges to $e_i(\Gamma)$ as $k\to\infty$ by the basic properties of a Markov chain, we may use the Dominated Convergence Theorem to write

$$\lim_{k\to\infty} \mathbb{E}[T(\sigma_i;k)] = \sum_{j\in I} e_j(\Gamma) \lim_{k\to\infty} \mathbb{E}[T(\sigma_i;k) \mid \sigma_i(k) = j]. \tag{16}$$

Recall that σ_i is a random sequence conditioned to start at i. Define $\widetilde{\sigma}_j$ to be a random sequence that starts at j and makes transitions according to $\widetilde{\Gamma}$, the time-reversal[23] of Γ. The transition matrix of this time-reversed process is

$$\widetilde{\Gamma}(i,j) = \frac{e_j(\Gamma)\Gamma(j,i)}{e_i(\Gamma)}.$$

Let $\rho(\sigma)$ be the reversal of a sequence. The probability distribution over sequences $\rho(\sigma_i)$ conditional on the sequence σ_i ending at j is the same as the distribution over sequences $\widetilde{\sigma}_j$ conditioned to end at i. This is by the basic properties of Markov chain reversal. Then we may write

$$\begin{aligned}\lim_{k\to\infty} \mathbb{E}[T(\sigma_i;k)] &= \sum_{j\in I} e_j(\Gamma) \lim_{k\to\infty} \mathbb{E}[T(\sigma_i;k) \mid \sigma_i(k) = j] \qquad \text{by (16)}\\ &= \sum_{j\in I} e_j(\Gamma) \lim_{k\to\infty} \mathbb{E}[S(\widetilde{\sigma}_j;k) \mid \widetilde{\sigma}_j(k) = i].\end{aligned}$$

A Markov process on I with an irreducible and aperiodic transition matrix generates an I-sequence almost surely. Thus by higher-order expectation-consistency and Proposition 3, we deduce that almost surely $S(\widetilde{\sigma}_j, k)$ converges to the expectation-consistent prior p_j as $k \to \infty$, irrespective of what $\widetilde{\sigma}_j(k)$ is. □

6.5 Completion of the Proof of Proposition 4

The final major step in the proof of Proposition 4 is the converse of Lemma 3:

Lemma 4. *If separability holds, then beliefs satisfy higher-order expectation-consistency.*

Proof. Assume that separability holds with a tuple $(p_i)_{i\in I}$ of priors. By using the separability condition with a suitable Γ, we will show that for each i and j,

$$p_i E_j = p_i, \tag{17}$$

which implies higher-order expectation-consistency.

Recall that separability states that for any $\Gamma \in \boldsymbol{\Gamma}$,

$$c(\Gamma) = \sum_{i'\in I} e_{i'}(\Gamma) p_{i'}.$$

[23] See (1.31) in Levin et al. (2009), and the book more generally, for basic properties of time reversals.

Plugging the above into the equation $c(\Gamma)B_\Gamma = c(\Gamma)$, we have

$$\sum_{i'} e_{i'}(\Gamma) p_{i'} B_\Gamma = \sum_{i} e_{i'}(\Gamma) p_{i'}. \tag{18}$$

Fix i and j, as well as $\alpha \in (0,1)$. For small $\delta, \varepsilon > 0$, consider $\Gamma_{\delta,\varepsilon}$ defined as follows (we describe the essence of the construction after the definition).

1. $\Gamma_{\delta,\varepsilon}(i', j') = \varepsilon$ for $i' \in \{i, j\}$, $j' \notin \{i, j\}$;
2. $\Gamma_{\delta,\varepsilon}(i', j') = 1/n$ for all $i' \notin \{i, j\}$ any $j' \in I$.
3. $\Gamma_{\delta,\varepsilon}(i, i) = \delta\alpha$;
4. $\Gamma_{\delta,\varepsilon}(j, j) = \delta(1-\alpha)$.
5. $\Gamma_{\delta,\varepsilon}(i, j)$ and $\Gamma_{\delta,\varepsilon}(j, i)$ are defined so that the corresponding rows sum to 1.

The idea is this: There are two distinguished states, i and j. The constant ε will be smaller than δ. (1) says that the probability of transitioning from $\{i, j\}$ to a state outside that set is "small," of order ε. (2) says that the transition probabilities from all states outside the distinguished pair $\{i, j\}$ are uniform. (3) says that the probability of a self-transition at i is $\delta\alpha$, and the probability of a self-transition at j is $\delta(1-\alpha)$. The purpose of this is to ensure that the ratio of these holding probabilities is $\alpha/(1-\alpha)$.

Now consider

$$\overline{e} = \lim_{\delta\downarrow 0}\lim_{\varepsilon\downarrow 0} e(\Gamma_{\delta,\varepsilon}).$$

It can be shown that $\overline{e}_i(\Gamma) = \alpha$, while $\overline{e}_j(\Gamma) = 1-\alpha$, and all other entries of $\overline{e}$ are zero.[24] Now,

$$\lim_{\delta\downarrow 0}\lim_{\varepsilon\downarrow 0} p_i B_{\Gamma_{\delta,\varepsilon}} = p_j \quad \text{and} \quad \lim_{\delta\downarrow 0}\lim_{\varepsilon\downarrow 0} p_j B_{\Gamma_{\delta,\varepsilon}} = p_i.$$

Thus, taking $\lim_{\delta\downarrow 0}\lim_{\varepsilon\downarrow 0}$ on both sides of (18), we get

$$\alpha p_i E_j + (1-\alpha) p_j E_i = \alpha p_i + (1-\alpha) p_j.$$

Since this must hold for all α, we conclude that $p_i E_j = p_i$, as desired. □

What we have shown now implies all the claims in Proposition 4. The only thing we have not yet discussed is how to establish statement (2) of that result. If expectation-consistency holds with $p_i = p$ for all i, and then by statement (3) we have network-independence. Conversely, if network-independence holds, then $c(\Gamma) = c$ for some fixed c, and then separability holds (with $p_i = c$ for each i), and so (again by statement (3) of the proposition) expectation-consistency holds.

[24] For any $\delta \geq 0$, the chain $\Gamma_{\delta,0}$ has one recurrent class (namely, $\{i, j\}$) and the stationary distribution is uniquely defined and continuous around such a chain (Schweitzer (1968)). The stationary distribution of $\Gamma_{\delta,0}$ depends only on the restriction of this matrix to i, j, and it can easily be computed to be $\overline{e}$.

References

BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): "Who's who in Networks. Wanted: the Key Player," *Econometrica*, 74, 1403–1417. 4

BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2015): "Networks, Information and Volatility," Yale University and Princeton University Working Paper. 6.1

BLUME, L., W. BROCK, S. DURLAUF, AND R. JAYARAMAN (2015): "Linear Social Interaction Models," *Journal of Political Economy*, 123, 444–496. 6.1

CALVÓ-ARMENGOL, A., J. MARTÍ, AND A. PRAT (2015): "Communication and Influence," *Theoretical Economics*, 10, 649–690. 4

GAIFMAN, H. (1986): "A Theory of Higher Order Probabilities," in *Theoretical Aspects of Reasoning About Knowledge: Proceedings of the 1986 Conference*, ed. by J. Y. Halpern, 275–292: Morgan Kaufman. 9

GOLUB, B. AND S. MORRIS (2017): "Expectations, Networks and Conventions," Available at SSRN: https://ssrn.com/abstract=2979086. 4, 5.1, 6.1, 18, 19, 20, 22, 6.3.2

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. (document), 1, 1

LEVIN, D. A., Y. PERES, AND E. L. WILMER (2009): *Markov Chains and Mixing Times*, Providence, RI: American Mathematical Society, with a chapter on coupling from the past by James G. Propp and David B. Wilson. 23

MERTENS, J. AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29. 2, 3, 11

SAMET, D. (1998a): "Iterated Expectations and Common Priors," *Games and Economic Behavior*, 24, 131–141. (document), 1, 1, 2, 3, 3, 4, 6, 9, 4.1, 3, 5, 4.2, 5.1

——— (1998b): "Common Priors and Separation of Convex Sets," *Games and Economic Behavior*, 24, 172–174. 2.2

SCHWEITZER, P. J. (1968): "Perturbation Theory and Finite Markov Chains," *Journal of Applied Probability*, 5, 401. 24

UI, T. (2009): "Bayesian Potentials and Information Structures: Team Decision Problems Revisited," *International Journal of Economic Theory*, 5, 271–291. 18