

Converging on Common Knowledge

Dominik Klein¹ and Rasmus Kræmmer Rendsvig²

¹University of Bamberg and University of Bayreuth

²Center for Information and Bubble Studies, University of Copenhagen

dominik.klein@uni-bayreuth.de, rasmus@hum.ku.dk

Abstract

Common knowledge, as is well known, is not attainable in finite time by unreliable communication, thus hindering perfect coordination. Focusing on the coordinated attack problem modeled using dynamic epistemic logic, this paper discusses unreliable communication protocols from a topological perspective and asks “If the generals may communicate indefinitely, will they then *converge* to a state of common knowledge?” We answer by making precise and showing the following: *common knowledge is attainable if, and only if, we do not care about common knowledge.*

1 Introduction

The concept of common knowledge is ubiquitous in contexts of informationally driven rational agents, whether in philosophy, economics or computer science. Though common knowledge is not a requirement for all rational interaction [Aumann and Brandenburger, 2016], it is a prerequisite for coordination in a plethora of contexts, including social life [Lewis, 1969], game theory [Morris and Shin, 1997; Rubinstein, 1989] and distributed systems [Fagin *et al.*, 1995].

However, as shown by Halpern and Moses [1990], in systems where message passing may fail, attaining common knowledge is impossible. The trouble is exemplified by the **coordinated attack problem**:

Two allied generals desire to overrun a besieged township. Any unilateral attack will certainly fail, but a simultaneous attack is guaranteed victorious. The generals can only communicate by courier, who may face peril on route, and thus fail to deliver its message. Can the generals ensure victory?

Already Akkoyunlu *et al.* [1975] show that even when the courier is always successful, no finite number of messages will suffice to establish an agreement to attack: If $n \geq 1$ is the least number of messages that suffice, loss of the n th message results in non-agreement. Hence the delivery of the n th message must be acknowledge before its sender is certain agreement is reached. For this, an $n + 1$ st message is required, contradicting the assumption. See also [Gray, 1978; Yemini and Cohen, 1979].

Halpern and Moses [1990] show that no correct protocol can ensure victory through a formal representation of the involved higher-order reasoning using interpreted systems.¹ They link the coordinated attack problem and common knowledge by showing that

- Prop. 4: common knowledge to attack is a prerequisite for coordinated attack, and
- Thm. 5: when communication is unreliable, it is impossible to attain common knowledge, and hence
- Cor. 6: the generals will never attack.

The content of Thm. 5 nowadays enjoys slogan-like status in multi-agent dynamics. The result is a hallmark of epistemic logic and the coordinated attack problem has become a stable showcase for formal systems of epistemic dynamics.

Yet, the generals’ epistemic dynamics exhibit a tendency *towards* common knowledge. Since common knowledge is characterizable as having all levels of mutual knowledge, and each successful message delivery yields a raise in the level of mutual knowledge of agreement, we face the following

Question: If the generals may communicate indefinitely, will they then *converge* to a state of common knowledge?

“Yes” may seem a tempting answer: keep adding 1 for long enough and you will reach the first infinity and thus common knowledge. This paper argues that this intuition is both correct and not.

The paper discusses unreliable communication protocols from a topological perspective, focusing on the coordinated attack problem. We cast the analysis in a mathematically expressive framework where *convergent sequences* and *limit points* are natural inhabitants, allowing us to show when and how unreliable communication converges to a state of common knowledge. Jumping ahead, we show that *common knowledge is attainable if, and only if, we do not care about common knowledge.* This statement is made precise by the

¹Common knowledge is also unattainable when messages are delivered reliably, but with unbounded delivery time [Halpern and Moses, 1990], under asynchronous message-passing [Fagin *et al.*, 1995] and in systems with temporal imprecision (*ibid.*). Dimitri [2003], however, identifies variation in the communication structure, that allow to approximate common knowledge.

framework of the paper, where caring about common knowledge means that it is expressible in the formal logical language. Section 2 models the coordinated attack problem using dynamic epistemic logic (DEL). Section 3 introduces a notion of convergence and argues its intuitiveness. Section 4 contains results and Section 5 concludes.

2 Coordinated Attack Problem in DEL

A DEL model typically consists of a sequence of small models, each representing a snapshot of the factual and higher-order information of all agents at a given moment in time. All but the first of these is the result of updating its predecessor with a non-guaranteed communication event. The dynamic model may thus be seen as an evolution

$$x_1 \xrightarrow{e_1} x_2 \xrightarrow{e_2} x_3 \xrightarrow{e_3} x_4 \xrightarrow{e_4} \dots$$

Each snapshot model x_k is self-contained and each e_k is an operation on x_k transforming it into x_{k+1} . Each snapshot model x_k is a multi-agent *epistemic state*—a pointed Kripke model—and each e_k is an *epistemic event*, modeled e.g. using an action model, see e.g. [Baltag and Renne, 2016; van Benthem, 2011; van Ditmarsch *et al.*, 2008]. Details follow.

For generality, we do not define our model directly as an evolution, but as a composition of two maps. The first takes an epistemic state x as input and outputs a message to be sent; the second takes the epistemic state and the selected message as input and outputs an updated epistemic state. The composition of these two functions then maps epistemic states to (updated) epistemic states. A run of the model on an initial state is then obtained by re-applying the map to its own output indefinitely.

To simplify, we will be concerned only with whether the generals attain common knowledge of a given target proposition (working with a single atom is for simplicity only; see Section 5). We thus omit modeling any attacks or lack thereof, implicitly taking such events to be synonymous with attaining common knowledge of the target proposition.

2.1 Languages and Epistemic States

We make use of two languages to describe epistemic states, both defined for two agents, a and b , and a single atom p , representing the target proposition. Our main concern thus is whether a and b can attain common knowledge that p .

Let the language \mathcal{L}_C be given by

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi \mid K_b\varphi \mid C\varphi$$

Let $\mathcal{L} \subseteq \mathcal{L}_C$ be the sub-language of formulas with no occurrences of the common knowledge operator, C .

As semantics for \mathcal{L} and \mathcal{L}_C , we use a special case of pointed Kripke models: an *epistemic state* for \mathcal{L}_C is a pair (M, s) with $M = (S, \sim_a, \sim_b, V(p))$ where S is a non-empty set of states, \sim_i is an equivalence relation on S for $i \in \{a, b\}$ and $V(p) \subseteq S$ is the valuation of the atom p ; finally, $s \in S$ is the actual state. For an epistemic state (M, s) , write M_s . Let \mathcal{G} be (a set representation of) the class of all epistemic states.

With \sim the transitive closure of $\sim_a \cup \sim_b$, the language \mathcal{L}_C is evaluated over epistemic states using standard clauses. We only mention those for the modalities:




Figure 1: The generals' initial epistemic state for coordinated attack, g_1 . The model has two states, s and t , with the double lining indicating s as actual state. In s , general a knows p , but general b does not: $g_1 \models K_a p \wedge \neg K_b p$.

$$M_s \models K_i \varphi \quad \text{iff} \quad \text{for all } t, s \sim_i t \text{ implies } M_t \models \varphi$$

$$M_s \models C\varphi \quad \text{iff} \quad \text{for all } t, s \sim t \text{ implies } M_t \models \varphi$$

The epistemic state of the two generals prior to any communication is illustrated in Figure 1.

2.2 Messages and Acknowledgments

Beyond the first, messages passed between the generals acknowledge receipt of the previous message. The content of the acknowledgments can be encoded either using atomic propositions or epistemic formulas. We choose the latter for simplicity.

Using epistemic formulas, the messages may be given a simple recursive build, suitable irrespective of which general knows the target proposition p . To exemplify, let first general a communicate knowledge of p with the message $K_a p$. Call it m_{a1} , as it is sent by a as first message. If the message is received, this informs b that p . To acknowledge delivery, b returns a second message $m_{b2} := K_b m_{a1} = K_b K_a p$. As b will know $K_a p$ only in the case m_{a1} is delivered, m_{b2} captures acknowledgment of receipt of m_{a1} . In turn, a may acknowledge receipt of m_{b2} by replying with $m_{a3} := K_a m_{b2} = K_a K_b K_a p$, etc. The messages may be given recursively by $m_{a1} := K_a p$, $m_{b1} := K_b p$ and $m_{ak} := K_a m_{bk-1}$, $m_{bk} := K_b m_{ak-1}$. By *the set of messages*, we then refer to $M := \{m_{ik} : i \in \{a, b\}, k \in \mathbb{N}\}$.

2.3 Communication Protocol

We focus on a simple alternating communication protocol akin to that described in the cited literature. The below map takes an epistemic state as input and returns a message. The message returned describes the highest level of higher-order knowledge that p , corresponding to an acknowledgment of the previous message. When both know p or in case of a 'higher-order tie', i.e., when the agents have n levels of mutual knowledge that p without either general knowing this, we take the liberty of letting only general a send a message. A protocol for all epistemic states where either general knows p is captured by the following *message selection map* $\mu : \mathcal{G} \rightarrow M \cup \{\top\}$:

Define a total lexicographic order \preceq on M by $m_{ik} \preceq m_{jk'}$ iff $k < k'$, or $k = k'$ and $i = b, j = a$. For all $x \in \mathcal{G}$ let

$$\mu(x) = \begin{cases} \max_{\preceq} \{m \in M : x \models m\} & \text{if exists} \\ \top & \text{else} \end{cases}$$

No $m \in M$ is redundant: for each, there is an $x \in \mathcal{G}$ such that $\mu(x) = m$. For $x \models K_a p \wedge \neg K_b p$, $\mu(x) = K_a p = m_{a1}$, while for $y \models K_b p \wedge \neg K_a p$, $\mu(y) = K_b p = m_{b1}$. In general, for $x \in \mathcal{G}$ such that $x \models m_{a,k} \wedge \neg m_{b,k}$, the message picked is $\mu(x) = m_{a,k}$. Symmetrically, for y with $y \models m_{b,k} \wedge \neg m_{a,k}$, $\mu(y) = m_{b,k}$.




Figure 2: The unreliable communication ($i: \varphi$). General i successfully sends the message φ to general j , with i uncertain about the success. The double lined success event s represents that the message φ is delivered successfully. In the failure event f , the null information message \top is received by agent j : f represents that the courier did not arrive. The recipient j can tell which event occurred. Assuming that j knows the courier arrives within some fixed time or not at all, this simplistically, but faithfully, represents successful passing of unreliable messages.




Figure 3: The update $g_2 := g_1 \otimes (a: K_a p)$ of the initial epistemic state x_1 (Figure 1) with the first courier-send message. Reflexive loops are omitted. The update g_2 may be divided in two: the state (s, s) as the sole child of the success event s and ‘all the rest’: a full copy of g_1 with all states children of the failure event f . The two parts are connected by linking the two children of the actual state of g_1 with an uncertainty relation of the message’s author: the sender cannot tell whether b received m_{a1} , i.e., the message that $K_a p$, or \top . The update g_2 satisfies $K_b m_{a1}$ and $\neg K_a K_b m_{a1}$.

2.4 Unreliable Message Passing

Beyond deciding on the content of the messages, we must also model the epistemic circumstances of delivery. In particular, the sending agent should consider it possible that the message is not delivered. To model unreliable communication acts, we use a form of semi-private announcements, illustrated in Figure 2.

Define for $i, j \in \{a, b\}, i \neq j$ and $\varphi \in \mathcal{L}_C$, the **unreliable communication** that φ from i to j as the tuple

$$(i: \varphi) = (\{s, f\}, \approx_i, \approx_j, pre(s), pre(f))$$

with a ‘success event’ s and a ‘failure event’ f , relations $\approx_i = \{s, f\}^2, \approx_j = \{(s, s), (f, f)\}$ and preconditions $pre(s) = \varphi, pre(f) = \top$.

To capture the effect of an unreliable communication on an epistemic state, we use **product update**: The **product update** of M_s with $(i: \varphi)$ is the epistemic state $M_s \otimes (i: \varphi)$ defined by

$$M_s \otimes (i: \varphi) := M' s' = (S', \sim'_a, \sim'_b, V', s')$$

with states $S' = \{(t, t) \in S \times \{s, f\} : Mt \models pre(t)\}$, relations $(t, t) \sim'_j (t', t')$ iff $t \sim_j t'$ and $t \approx_j t'$, for $j \in \{a, b\}$, valuation $V'(p) = \{(t, t) \in S' : Mt \models p\}$, and actual state $s' = (s, s)$.

The effect of updating the initial epistemic state x_1 of Figure 1 with the unreliable communication of the first message ($a: m_{a1}$) is depicted and described in Figure 3.

Remark. Semi-private announcements are special cases of **pointed action models**, introduced along with product updates by Baltag *et al.* [1998]; Baltag and Moss [2004]. For introductions, see e.g. [Baltag and Renne, 2016; van Benthem, 2011; van Ditmarsch *et al.*, 2008]. Aucher *et al.* [2012] also

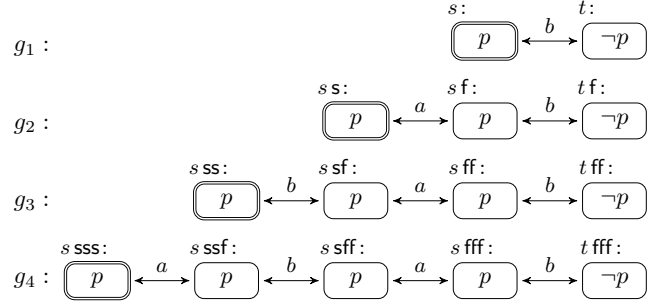


Figure 4: Epistemic states g_1 – g_4 . Reflexive loops are omitted together with commas and parentheses in state names.

use DEL semi-private announcements to model the coordinated attack problem. For alternative approaches, see e.g. [Fagin *et al.*, 1995; Herzig *et al.*, 2015].

2.5 The Dynamic Model

Finally, we define the map $f: \mathcal{G} \rightarrow \mathcal{G}$ representing the generals’ unreliable communication. Given an epistemic state x , the map distinguishes three cases: *i*) the message selected by μ at x ‘belongs’ to a , *ii*) the message ‘belongs’ to b , or *iii*) the message of x is the empty message:

$$f(x) = \begin{cases} x \otimes (a: \mu(x)) & \text{if } \mu(x) = m_{ak} \text{ for some } k \\ x \otimes (b: \mu(x)) & \text{if } \mu(x) = m_{bk} \text{ for some } k \\ x & \text{else} \end{cases}$$

The map f is total on \mathcal{G} , but only represents the communication protocol on the sub-domain where either agent knows p , i.e., the set of epistemic states $[K_a p \vee K_b p]$.

2.6 Running for Common Knowledge

Iteratively applying the map f to the initial epistemic state g_1 produces an infinite, but neatly structured, sequence of epistemic states $\bar{g} := g_1, g_2, \dots$ with $g_{n+1} := f^n(g_1)$. The first elements of \bar{g} are depicted in Figure 4.

The form of the update described in Figure 3 is symptomatic for all latter. The states of g_{n+1} may be divided in two: a single child of the actual state of g_n and event s and a full copy of g_n courtesy of f . The only link between the two components belongs to the general that sends the n th message: it links the two children of g_n ’s actual state.

The sequence $\bar{g} = g_1, g_2, \dots$ adheres to a recursive construction up to isomorphism: $x_1 = (S_1, \sim_{a1}, \sim_{b1}, V_1, s_1)$ with $S_1 = \{s_1, t\}, \sim_{a1} = \{(s, s), (t, t)\}, \sim_{b1} = \{s, t\}^2$, and $V_1(p) = \{s\}$, cf. Figure 1, and $x_k = (S_k, \sim_{ak}, \sim_{bk}, V_k, s_k)$ with

$$\begin{aligned} S_k &= S_{k-1} \cup \{s_k\}, \\ \sim_{ak} &= \begin{cases} \sim_{ak-1} \cup \{s_k, s_{k-1}\}^2 & \text{if } k \text{ is odd,} \\ \sim_{ak-1} \cup \{(s_k, s_k)\} & \text{else} \end{cases} \\ \sim_{bk} &= \begin{cases} \sim_{bk-1} \cup \{s_k, s_{k-1}\}^2 & \text{if } k \text{ is odd,} \\ \sim_{bk-1} \cup \{(s_k, s_k)\} & \text{else} \end{cases} \\ V_k(p) &= V_{k-1}(p) \cup \{s_k\}. \end{aligned}$$

The run \bar{g} mirrors the aforementioned result by Halpern and Moses [1990]: common knowledge is not attained in finite time. For every $n \in \mathbb{N}$, g_n contains a child of the state t , and this child is not in $V_n(p)$. Yet, it is reachable from the actual state of g_n by the transitive closure of $\sim_{an} \cup \sim_{bn}$. Hence

$$\text{for all } n \in \mathbb{N}, g_n \models \neg Cp. \quad (1)$$

2.7 Approaching Common Knowledge

Though the generals at no time attain common knowledge, they are seemingly getting closer: with each trip of the courier, another level of higher-order knowledge that p is created, bringing the generals one step closer to common knowledge. To make things formal, define

$$E^1\varphi := K_a\varphi \wedge K_b\varphi \text{ and } E^{k+1}\varphi := E^1E^k\varphi$$

Then for any epistemic state x , $x \models C\varphi$ iff $x \models E^n\varphi$ for all $n \in \mathbb{N}$, cf. e.g. [Fagin *et al.*, 1995]. For the generals' sequence of epistemic states, $\bar{g} = g_1, g_2, \dots$, it then holds true for all $n \in \mathbb{N}$ that

$$g_{n+1} \models E^n p \wedge \neg E^{n+1} p \quad (2)$$

Hence, as n grows larger, the set of formulas that 'support' common-knowledge-that- p satisfied by g_n grows, while the set of formulas contradicting common-knowledge-that- p satisfied by g_n shrinks. Does this then mean that the generals attain common knowledge as n goes to infinity?

3 Logical Convergence and its Topology

To answer this question, it is natural to turn to the mathematical notions of *convergence* and *limits*. These notions, however, are relative to structure beyond only sequences of models. For an answer to the above question to be illuminating, also this additional structure must be natural.

The nowadays most common type of additional structure comes in the form of a *metric*, or perhaps it's more general cousin, a *topology*. Roughly, a metric is a distance-measure, specifying quantitatively how far any two points are apart. That is, a *metric* on a set X is a function

$$d : X \times X \longrightarrow \mathbb{R}$$

that satisfies the four natural requirements

- i.* positivity: $d(x, y) \geq 0$,
- ii.* identity: $d(x, y) = 0$ iff $x = y$,
- iii.* symmetry: $d(x, y) = d(y, x)$, and
- iv.* the triangular inequality: $d(x, z) \leq d(x, y) + d(y, z)$.

A metric allows to define convergence. Formally, a sequence $\bar{x} = x_1, x_2, \dots \in X$ of points *converges* to a point x in the metric space (X, d) iff for every real number $\varepsilon > 0$, there is a time N after which every point in \bar{x} is within ε distance of x :

$$\forall \varepsilon \in \mathbb{R}_{>0}, \exists N \in \mathbb{N}, \forall n \geq N, d(x_n, x) \leq \varepsilon.$$

The concept of convergence can also be defined in a more abstract setting: a topology is a conceptual abstraction, qualitatively specifying closeness in space. A set imbued with a metric is called a *metric space*, while a set with a topology is

a *topological space*. Topological spaces are generalizations of metric spaces: every metric space gives rise to a topological space, but not *vice versa*. For present purposes, perhaps think of a topological space as obtained from a metric space by throwing away the real values on distances, remembering only which points are closer to some point than others. Formally, a *topology* on a set X is a collection of subsets of X , $\mathcal{T} \subseteq \mathcal{P}(X)$, satisfying three abstract requirements:

1. The full and empty sets are in \mathcal{T} : $X \in \mathcal{T}$ and $\emptyset \in \mathcal{T}$;
2. \mathcal{T} is closed under finite intersections: if $U_i \in \mathcal{T}$ for all $i \in I$ and I is finite, then $\bigcap_{i \in I} U_i \in \mathcal{T}$;
3. \mathcal{T} is closed under arbitrary unions: if $U_j \in \mathcal{T}$ for all $j \in J$, then $\bigcup_{j \in J} U_j \in \mathcal{T}$.

Members of $\mathcal{T} \subseteq \mathcal{P}(X)$ are called *open* sets of X , while the complement of an open set is called *closed*. The notion of convergence, finally, has a topological counterpart: a sequence $\bar{x} = x_1, x_2, \dots \in X$ of points *converges* to a point x in the topological space (X, \mathcal{T}) iff for every $U \in \mathcal{T}$ with $x \in U$, there is a time N after which every point in \bar{x} is in U :

$$\forall U \in \mathcal{T} : x \in U \Rightarrow \exists N \in \mathbb{N}, \forall n \geq N, x_n \in U.$$

To obtain reasonable insights on the convergence of unreliable communication, it does not suffice to fix just any metric or topological space—the space must also capture our intuitions concerning distances between models. Else convergence results are without bite. Consider for example the sequence of reals

$$(1/2^n)_{n \in \mathbb{N}} = 1/2, 1/4, 1/8, 1/16, \dots$$

It is forever decreasing yet always positive, with each incremental point having a smaller difference with 0 than its predecessor: in one intuitive understanding of the reals, the sequence converges to 0. This is also true under the *Euclidean metric* d_e (for all $x, y \in \mathbb{R}$, $d_e(x, y) = |x - y|$), but it is false under the *discrete metric* d_d (for all $x, y \in \mathbb{R}$, if $x \neq y$, $d_e(x, y) = 1$, else 0). This illustrates the importance of choosing a natural notion of distance for results to be of interest: that $1/2, 1/4, 1/8, 1/16, \dots$ converges in (\mathbb{R}, d_e) reflects a strong intuition concerning the nature of the reals, while the conclusion concerning (\mathbb{R}, d_d) seems merely obscure.

Our results pertain to a type of topology that we find highly intuitive. To argue its intuitiveness, we present it through an alternative approach to convergence, namely *sequential convergence*.

3.1 Sequential and Logical Convergence

In the most general formulation, a *sequential convergence* on a set X is a subset \mathbb{L} of $X^{\mathbb{N}} \times X$. When a sequence-element pair (σ, x) belongs to \mathbb{L} , the sequence σ is said to *sequentially converge* to x in \mathbb{L} . Three standard coherence axioms for a sequential convergence are *i*) that a constant sequence converges to its only value, *ii*) that all subsequences of a convergent sequence share its limit, and *iii*) that limits are unique. When \mathbb{L} satisfies suitable coherence axioms, the sequential convergence in turn induces a topology, possibly with standard structural properties—like being Hausdorff. Frič [1997] presents the theory and a worthwhile historical review.

We find the following approach to convergence through sequential convergence highly intuitive: Given that one has chosen a formal language L that represents exactly the features of interest over some class of models, a natural sequential convergence may be defined. The intuition is that a sequence of models x_1, x_2, \dots converges to the model x just in case the models in the sequence steadily become more like the model x . With the features of interest given by formulas of L , this means that the models x_1, x_2, \dots move towards satisfying all the same formulas as x .

The corresponding convergence notion was introduced in [Klein and Rendsvig, 2017] under the label **logical convergence**. A general definition may be given as follows:

Definition 1. Let L be a language binarily interpreted over a set of models X by a satisfaction relation \models . A sequence $x_1, x_2, \dots = \sigma \in X^{\mathbb{N}}$ then **L -converges** to $x \in X$ iff for every L -formula φ for which $x \models \varphi$, there is $N \in \mathbb{N}$ such that $x_n \models \varphi$ for all $n \geq N$.

Logical convergence for the language L induces a sequential convergence: define \mathbb{L}_L to contain exactly the pairs $(\sigma, x) \in X^{\mathbb{N}} \times X$ such that

$$\forall \varphi \in L : x \models \varphi \Rightarrow \exists N \in \mathbb{N} \forall n \geq N, x_n \models \varphi.$$

In the ensuing, specifically the sequential convergences of the languages \mathcal{L} and \mathcal{L}_C will be of interest. We refer simply to \mathcal{L} -convergence and \mathcal{L}_C -convergence.

3.2 Modal Spaces

In applying a logical convergence to a set of epistemic states X , an annoyance may surface: if the set X is not suitably filtered, limits need not be unique. This runs somewhat counter to intuitions: when a sequence σ logically converges to both x and y , this entails that x and y share all features of interest—hence, for all intents and purposes, the models are the same. Yet, x and y need not be set-theoretically identical, making limits non-unique.

There are several ways to restore the uniqueness of limits without interfering with the represented logical dynamics. We choose what we consider the simplest: *quotients by logical equivalence*. In parlance, we follow [Klein and Rendsvig, 2017] in referring to *modal spaces*:

Definition 2. Let L be a language binarily interpreted over a set of models X by a satisfaction relation \models . The **modal space** X_L is the quotient of X under L -equivalence. I.e.,

$$\begin{aligned} X_L &= \{x_L : x \in X\} \text{ for} \\ x_L &= \{y \in X : \forall \varphi \in L, y \models \varphi \text{ iff } x \models \varphi\}. \end{aligned}$$

With X_L given, the **truth set** of $\varphi \in L$ is $[\varphi]_L = \{x_L \in X_L : \forall x \in x_L, x \models \varphi\}$.

As only the languages \mathcal{L} and \mathcal{L}_C are used below, the subscripts of x_L and $[\varphi]_L$ are omitted. This causes no confusion concerning x as $X_{\mathcal{L}} = X_{\mathcal{L}_C}$ for any X , nor for $[\varphi]$ as $[\varphi]_{\mathcal{L}} = [\varphi]_{\mathcal{L}_C}$ for all $\varphi \in \mathcal{L} \subseteq \mathcal{L}_C$ and the subscript is clearly \mathcal{L}_C when φ contains common knowledge operators. Throughout, write $x \models \varphi$ for $x \in [\varphi]$.

As every epistemic state $x \in \mathcal{G}$ is uniquely represented in $\mathcal{G} := \mathcal{G}_{\mathcal{L}} = \mathcal{G}_{\mathcal{L}_C}$, we again refer to the elements $x \in \mathcal{G}$ as epistemic states.

3.3 From Logical Convergence to Topology

The study of convergence and limits is the hallmark of topology. Though sequential convergence—through logical convergence—may be useful in fixing intuitions, as analytical framework it has long been surpassed by general topology.

If one agrees that logical convergence captures a natural notion of convergence and further finds it reasonable to identify models under logical equivalence, then finding a topology which respects one’s intuitions is straightforward. In fact, there is a *unique* topology which respects the intuition behind logical convergence—a unique topology for which logical convergence and topological convergence are equivalent:

Theorem 3 ([Klein and Rendsvig, 2018]). *Let $L \in \{\mathcal{L}, \mathcal{L}_C\}$ and let \mathcal{T} be a topology on X_L . The following are equivalent:*

1. *A sequence x_1, x_2, \dots of points from X_L converges to x in the topological space (X_L, \mathcal{T}) if, and only if, x_1, x_2, \dots logically converges to x in X_L .*
2. *\mathcal{T} is the Stone topology \mathcal{T}_L on X_L i.e. the topology whose open sets are exactly $\{[\varphi] \subseteq X_L : \varphi \in L\}$.*

Both logical convergence and the Stone topology are *language relative*, but for each language, there is a *unique* topology capturing its logical convergence. Importantly, while $X_{\mathcal{L}} = X_{\mathcal{L}_C}$, it is *not* the case that $\mathcal{T}_{\mathcal{L}}$ and $\mathcal{T}_{\mathcal{L}_C}$ are the same topologies. A similar observation about topologies on type spaces is made by Rubinstein [1989] in his seminal work on *almost common knowledge*.²

4 Unreliable Communication in the Limit



Figure 5: The epistemic state $c_p \in \mathcal{G}$ where the generals have common knowledge that p .

In the modal space \mathcal{G} , there is exactly one element that satisfies Cp . This is the point $c_p \in \mathcal{G}$ containing the epistemic state $c_p \in \mathcal{G}$ of Figure 5. Given Theorem 3, whether the generals attain common knowledge in the limit reduces to the question of whether the sequence \bar{g} converges to c_p in the topological space $(\mathcal{G}, \mathcal{T}_{\mathcal{L}_C})$. As it happens, hope for common knowledge in the limit is futile:

Proposition 4. *In $(\mathcal{G}, \mathcal{T}_{\mathcal{L}_C})$, the sequence \bar{g} does not converge.*

Proof. For a contradiction, assume \bar{g} to converges to x in $(\mathcal{G}, \mathcal{T}_{\mathcal{L}_C})$. For every open neighborhood U of x , there must be an $N \in \mathbb{N}$ such that $g_n \in U$ for all $n \geq N$. Both $[Cp]$ and $[\neg Cp]$ are open in $\mathcal{T}_{\mathcal{L}_C}$. They partition \mathcal{G} , so either $x \in [Cp]$ or $x \in [\neg Cp]$. If x is in $[Cp]$, then $[Cp]$ is an open

²Other topological approaches to belief and knowledge relate to the definition of common knowledge (e.g. [Lewis, 1969], see [Vanderschraaf and Sillari, 2014] for an overview) or to belief revision [Baltag *et al.*, 2019]. These use topologies on the sets of worlds within a given model to define semantics for modal operators. In contrast, we assume standard Kripke semantics for modal operators and use topologies defined on a set of models to study convergence.

neighborhood of x : but for no $n \in \mathbb{N}$ does $g_n \models Cp$, cf. Eq. 1, p. 4, so g_n is not in $[Cp]$. The purported limit of \bar{g} can therefore not be in $[Cp]$. If x is in $[\neg Cp]$, then $x \models \neg E^k p$ for some $k \in \mathbb{N}$. I.e., $[\neg E^k p]$ is an open neighborhood of x —but it does not contain g_{k+1} , as $g_{k+1} \models E^k p$ cf. Eq. 2, p. 4. The purported limit of \bar{g} can therefore not be in $[\neg Cp]$ either. Hence \bar{g} does not converge in $(\mathcal{G}, \mathcal{T}_{\mathcal{L}C})$. \square

Proposition 4 may run counter to intuitions: while the generals in finite time attain any level of mutual higher-order knowledge that p , not even in the limit do they attain common knowledge. Topologically, the explanation is simple: the point c_p is isolated in \mathcal{G} and can therefore not be gradually approached. We return to this below.

Emphatically, this non-convergence is not a consequence of the epistemic states *qua* relational structures, but of how these are described by the formal language. Under *exactly the same dynamics* but omitting the common knowledge operator, all higher-order uncertainty that p is successfully eliminated:

Proposition 5. *In $(\mathcal{G}, \mathcal{T}_{\mathcal{L}})$, the sequence \bar{g} converges to c_p .*

Proof. Every φ from \mathcal{L} satisfied by c_p is implied by some formula in $\{E^N p : N \in \mathbb{N}\}$. Thus, for every open neighborhood $U \in \mathcal{T}_{\mathcal{L}}$ of c_p , $[E^N p] \subseteq U$ for some $N \in \mathbb{N}$. Now, for all $n \geq N + 1$, $g_n \models E^N p$, cf. Eq. 2. I.e., $g_n \in [E^N p] \subseteq U$. Hence \bar{g} converges to c_p . \square

Proposition 5 thus shows that omitting common knowledge *syntactically* lets the generals obtain it *semantically*.

Remark. For those who would prefer results concerning dynamics directly on the set of models \mathcal{G} instead of on the modal space \mathcal{G} , we remark that Proposition 4 implies that \bar{g} also diverges in $(\mathcal{G}, \mathcal{T}_{\mathcal{L}C})$. Proposition 5 only implies that if \bar{g} converges in $(\mathcal{G}, \mathcal{T}_{\mathcal{L}C})$, then its limit is some $x \in c_p$.

4.1 Main Results

Propositions 4 and 5 apply to the specific run \bar{g} . Yet, similar conclusions may be drawn more generally for the dynamic model f . Our first main result generalizes the non-convergence conclusion of Proposition 4 by showing that common knowledge, if attainable at all, must be reached in *finite* time:

Proposition 6. *Let $(x_n)_{n \in \mathbb{N}}$ be any convergent sequence from $(\mathcal{G}, \mathcal{T}_{\mathcal{L}C})$. If $\lim_{n \rightarrow \infty} (x_n) \models Cp$, then there exists a $k \in \mathbb{N}$ such that $x_k \models Cp$.*

Proof. By Theorem 3, $(x_n)_{n \in \mathbb{N}}$ converges topologically iff it converges logically. By Definition 1, $\lim_{n \rightarrow \infty} (x_n) \models Cp$ then implies that there is some $k \in \mathbb{N}$ such that $x_k \models Cp$. \square

Our second main result shows that the semantic attainment of common knowledge in Proposition 5 is not an artifact, but a robust consequence of the generals' communication protocol. When either general knows p , the protocol eliminates all uncertainty—semantically, the generals *always* converge to a state of common knowledge. To show this, encode the protocol f on the modal space \mathcal{G} by the map $f : \mathcal{G} \rightarrow \mathcal{G}$ given

by $f(x) = y$ iff $f(x) \in y$, for all $x \in \mathcal{G}$. This map is well-defined as product update preserves bisimulation and hence modal equivalence, cf. [Baltag and Moss, 2004]. We obtain:

Proposition 7. *For all $x \in [K_a p \vee K_b p] \subseteq \mathcal{G}$, $\lim_{n \rightarrow \infty} f^n(x) = c_p$ in the topology $\mathcal{T}_{\mathcal{L}}$.*

The proof uses the following fixed point theorem:

Theorem 8 (Edelstein, 1962). *Every map F on a compact metric space (X, d) for which $d(F(x), F(y)) < d(x, y)$ for all $x, y \in X, x \neq y$, has a unique fixed point in X .*

Corollary 9. *Let F , and (X, d) as in Theorem 8. Then for any $x \in X$, the sequence $x, F(x), F^2(x) \dots$ converges to the unique fixed point of F .*

Call a map satisfying the condition of Theorem 8 **Edelstein contractive**. The property is used as follows: Lemma 10 introduces a metric d_E on \mathcal{G} for which metric space (\mathcal{G}, d_E) is compact and has as metric topology $\mathcal{T}_{\mathcal{L}}$. Lemma 11 shows the second-iterate f^2 of f Edelstein contractive in the subspace $[K_a p \vee K_b p] \subseteq \mathcal{G}$ with fixed point c_p . By Corollary 9, for any x in X , the sequences $x, f^2(x), f^4(x) \dots$ and $f(x), f^3(x), f^5(x) \dots$ converge to c_p . Hence also the sequence $x, f(x), f^2(x) \dots$ does, establishing Proposition 7.

Lemma 10. *The function $d_E : \mathcal{G} \times \mathcal{G} \rightarrow \mathbb{R}$ given by*

$$d_E(x, y) = \begin{cases} 0 & \text{if } x \models E^n p \Leftrightarrow y \models E^n p \text{ for all } n \\ \frac{1}{n+1} & \text{if } n \text{ is the least integer such that} \\ & \text{not } x \models E^n p \Leftrightarrow y \models E^n p \end{cases}$$

is a metric on \mathcal{G} . The metric space (\mathcal{G}, d_E) is compact and has metric topology $\mathcal{T}_{\mathcal{L}}$. The same holds for the subspace $([K_a p \vee K_b p], d_E)$.

Proof. By results of Klein and Rendsvig [2017]. In their terminology, the set $D = \{E^k p\}_{k \in \mathbb{N}}$ is a *finite representative descriptor*, and with the *weight function* $w(E^n p) := \frac{1}{n+1}$ for all n , d_E is obtained as a metric of their family of generalized Hamming distances. Hence d_E has metric topology $\mathcal{T}_{\mathcal{L}}$, which is compact on the modal space \mathcal{G} as S5 is complete w.r.t. \mathcal{G} . The same holds for the subspace $[K_a p \vee K_b p]$ as it is closed in $\mathcal{T}_{\mathcal{L}}$. \square

Lemma 11. *The second-iterate f^2 of f is Edelstein contractive in $([K_a p \vee K_b p], d_E)$ and has fixed point c_p .*

Proof. Let $x, y \in [K_a p \vee K_b p], x \neq y$. Then $d_E(x, y) = \frac{1}{n+1}$ for some n . I.e., both x and y satisfy $E^n p$, but only one satisfies $E^{n+1} p$. Yet, both $f^2(x)$ and $f^2(y)$ satisfy E^{n+1} : Wlog, assume $x \models E^{n+1} p$ and $y \models \neg E^{n+1} p$. Then $f^2(x) \models E^{n+1} p$ as f preserves satisfaction of $E^k p$, for all k . For $f^2(y)$, notice that y must satisfy either 1. $K_a E^n p \wedge \neg K_b E^n p$, 2. $K_b E^n p \wedge \neg K_a E^n p$, or 3. $\neg K_b E^n p \wedge \neg K_a E^n p$. In the first case, $\mu(y) = K_a E^n p$, so $f(y) = y \otimes (a : K_a E^n p)$ satisfies $K_b K_a E^n p$, and hence $K_b E^n p$, so also $E^{n+1} p$. Hence also $f^2(y) \models E^{n+1} p$. The second case is symmetric. In the third case, $\mu(y) = K_a E^{n-1} p$. Hence $f(y)$ satisfies $K_b K_a E^{n-1} p$, so also $K_b E^n p$ —which is $\mu(f(y))$. Then $f^2(y) \models K_a K_b E^n p$, and hence $f^2(y) \models E^{n+1} p$. As both $f^2(x)$ and $f^2(y)$ satisfy E^{n+1} , $d_E(f^2(x), f^2(y)) = \frac{1}{n+2} < d_E(x, y)$. Finally, c_p is a fixed point as $\mu(c_p) = \top$. \square

5 Concluding Remarks

We have modeled and analyzed the coordinated attack problem as a showcase example for asymptotic behavior of long term, asynchronous message passing with uncertain success. The analysis provided employs dynamic epistemic logic in a topological setting. The analysis and our main results concern a highly simplified representation of unreliable communicating, with but two agents and a single atomic proposition. Not much hinges on this choice. Similar techniques work for more complex representations with additional agents and atomic propositions. Generalized results on unreliable communication are obtainable by working with more advanced modal spaces, quotienting out what is not communicated about. As dynamic epistemic logic is a very general framework, the showcased techniques apply also to other instances of long-term informational dynamics.

Our main results imply that the analysis of long-term semantic updates is highly dependent on our choice of language and the corresponding topology. For a language without common knowledge operator, unreliable communication will converge to a state of common knowledge, but the agents are unable to express it. If, however, the language contains a common knowledge operator, then unreliable communication does not converge. Assuming that our choice of language exactly reflects the features we are interested in, this can be put precisely in a slogan-like statement: common knowledge is attainable, if, and only if, we do not care about common knowledge.

Acknowledgements

The Center for Information and Bubble Studies is funded by the Carlsberg Foundation. Both authors were partially supported by the DFG-ANR joint project Collective Attitude Formation [RO 4548/8-1], DK also by the DFG-GAČR joint project From Shared Evidence to Group Attitudes [RO 4548/6-1] and by the NSFC project Logics of Information Flow in Social Networks [17ZDA026].

References

[Akkoyunlu *et al.*, 1975] Eralp Abdurrahim Akkoyunlu, Kattamuri Ekanadham, and R. V. Huber. Some Constraints and Tradeoffs in the Design of Network Communications. In *SOSP '75 Proceedings of the Fifth ACM Symposium on Operating Systems Principles*, pages 67–74, 1975.

[Aucher *et al.*, 2012] Guillaume Aucher, Bastien Maubert, and François Schwarzentruber. Generalized DEL-Sequents. In Luis Fariñas del Cerro, Andreas Herzig, and Jérôme Mengin, editors, *Logics in Artificial Intelligence*, volume 7519 of *Lecture Notes in Computer Science*, pages 54–66. Springer, 2012.

[Aumann and Brandenburger, 2016] Robert J. Aumann and Adam Brandenburger. Epistemic Conditions for Nash Equilibrium. In Horacio Arlo-Costa, Vincent F. Hendricks, and Johan van Benthem, editors, *Readings in Formal Epistemology*, chapter 41. Springer, 2016.

[Baltag and Moss, 2004] Alexandru Baltag and Lawrence S. Moss. Logics for Epistemic Programs. *Synthese*, 139(2):165–224, 2004.

[Baltag and Renne, 2016] Alexandru Baltag and Bryan Renne. Dynamic Epistemic Logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, 2016.

[Baltag *et al.*, 1998] Alexandru Baltag, Lawrence S. Moss, and Sławomir Solecki. The Logic of Public Announcements, Common Knowledge, and Private Suspicions (extended abstract). In *TARK '98: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 43–56. Morgan Kaufmann Publishers, 1998.

[Baltag *et al.*, 2019] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. A Topological Approach to Full Belief. *Journal of Philosophical Logic*, 48(2):205–244, 2019.

[van Benthem, 2011] Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.

[Dimitri, 2003] Nicola Dimitri. Coordination in an Email Game Without “Almost Common Knowledge”. *Journal of Logic, Language and Information*, 12(1):1–11, 2003.

[van Ditmarsch *et al.*, 2008] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, 2008.

[Edelstein, 1962] Michael Edelstein. On Fixed and Periodic Points under Contractive Mappings. *Journal of the London Mathematical Society*, 37:74–79, 1962.

[Fagin *et al.*, 1995] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. The MIT Press, 1995.

[Frič, 1997] Roman Frič. History of Sequential Convergence Spaces. In Charles E. Aull and Robert Lowen, editors, *Handbook of the History of General Topology*, volume 1, pages 343–355. Springer, 1997.

[Gray, 1978] Jim Gray. Notes on Data Base Operating Systems. In Rudolf Bayer, Robert M. Graham, and Gerhard Seegmüller, editors, *Operating Systems*, volume 60 of *Lecture Notes in Computer Science*, pages 393–481. Springer, 1978.

[Halpern and Moses, 1990] Joseph Y. Halpern and Yoram Moses. Knowledge and Common Knowledge in a Distributed Environment. *Journal of the ACM*, 37:549–587, 1990.

[Herzig *et al.*, 2015] Andreas Herzig, Emiliano Lorini, and Faustine Maffre. A Poor Man’s Epistemic Logic Based on Propositional Assignment and Higher-Order Observation. In Wiebe van der Hoek, Wesley H. Holliday, and Wen-fang Wang, editors, *Logic, Rationality, and Interaction*, volume 9394 of *Lecture Notes in Computer Science*, pages 156–168. Springer, 2015.

[Klein and Rendsvig, 2017] Dominik Klein and Rasmus K. Rendsvig. Convergence, Continuity and Recurrence in

Dynamic Epistemic Logic. In Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, editors, *Logic, Rationality, and Interaction*, volume 10455 of *Lecture Notes in Computer Science*, pages 108–122. Springer, 2017.

- [Klein and Rendsvig, 2018] Dominik Klein and Rasmus K. Rendsvig. Metrics for Formal Structures, with an Application to Kripke Models and their Dynamics. In *Logical Dynamics and Dynamical Systems*, pages 177–207. Lund University Press, 2018.
- [Lewis, 1969] David Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.
- [Morris and Shin, 1997] Stephen Morris and Hyun Song Shin. Approximate Common Knowledge and Coordination: Recent Lessons from Game Theory. *Journal of Logic, Language and Information*, 6(2):171–190, 1997.
- [Rubinstein, 1989] Ariel Rubinstein. The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge". *American Economic Review*, 79:385–391, 1989.
- [Vanderschraaf and Sillari, 2014] Peter Vanderschraaf and Giacomo Sillari. Common Knowledge. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, 2014.
- [Yemini and Cohen, 1979] Yechiam Yemini and Danny Cohen. Some Issues in Distributed Processes Communication. In *Proc. of the 1st International Conference on Distributed Computing Systems*, pages 199–203, 1979.