

Topological Semantics for Common Inductive Knowledge

Siddharth Namachivayam

April 2026

Department of Philosophy

Carnegie Mellon University

Pittsburgh, PA 15213

Thesis Committee:

Prof. Adam Bjorndahl

Prof. Kevin Kelly

Submitted in fulfillment of the requirements for the degree:

Master of Science in Logic, Computation, and Methodology

Contents

1	Introduction	3
2	Pre-Theory	4
3	Topology & Learning	9
4	Semantics	17
5	Syntax	21
6	Soundness	23
7	Inductive Coordinated Attack	27
8	Future Directions & Conclusion	30
9	Bibliography	31

Topological Semantics for Common Inductive Knowledge

Siddharth Namachivayam

Department of Philosophy
Carnegie Mellon University*
Pittsburgh, Pennsylvania
snamachi@andrew.cmu.edu

Consider a community of scientists whose labs are each capable of conducting a different set of experiments. The scientists want to work together to confirm a new hypothesis, but to ensure blindness, their labs generally prohibit the scientists from communicating with each other. Further, each scientist can only make so many retractions to their lab before having to cease inquiry and suspend judgement forever. How might the scientists coordinate whether to affirm or suspend judgement on this hypothesis in light of their private experiments so that their labs are guaranteed to converge to the same conclusion and that this conclusion will not be a false positive? Call this problem ‘inductive coordinated attack.’ In this paper, we develop a logic for solving inductive coordinated attack by determining when and how a hypothesis can become what we call ‘common inductive knowledge.’ We begin by precisifying Lewis’ account of common knowledge in *Convention* which describes the generation of higher-order expectations between agents as hinging upon agents’ inductive standards and a shared witness. Our language has a rather rich syntax in order to capture equally rich notions central to Lewis’ account; for instance, we speak of an agent ‘having inductive reason to believe’ a proposition and one proposition ‘indicating’ to an agent that another proposition holds. This syntax affords a novel topological semantics which, following Kelly 1996’s approach in *The Logic of Reliable Inquiry*, takes as primitives agents’ information bases. In particular, we endow each agent with a ‘switching tolerance’ meant to represent their personal inductive standards for learning. After establishing soundness of our proof system with respect to this semantics, we conclude by showing how our logic can be used to solve inductive coordinated attack.

1 Introduction

In *Convention*, David Lewis presents 3 conditions for a witness W to generate common knowledge of a proposition P .¹ These are:

1. Everyone has reason to believe W .
2. W indicates to everyone that P .
3. W indicates to everyone that 1.

where W indicates to an agent i that P iff:

if i had reason to believe W , i would thereby have reason to believe P .

*The author would like to thank Adam Bjorndahl and Kevin Kelly for advising this project. Preliminary progress was presented in 2025 at the Indian Statistical Institute’s conference on Recent Trends in Logic and Game Theory as well as the Prolog Conference Series’ 12th Workshop on Combining Logic and Probability. This work also received financial support in the summer of 2024 from Carnegie Mellon University’s Institute for Complex Social Dynamics.

¹Strictly speaking, Lewis defines what it means for W to be a ‘basis’ for common knowledge of P . We will instead call W a ‘witness,’ reserving the term ‘basis’ to be used in its topological sense.

For Lewis, what W indicates to i (and whether W generates common knowledge of P) crucially depends on i 's inductive standards. After making appropriate 'ancillary' assumptions about the nature of these inductive standards, Lewis uses his 3 conditions to derive that:

- (1) Everyone has reason to believe P .
- (2) Everyone has reason to believe (1).
- (3) Everyone has reason to believe (2).
-
- (n) Everyone has reason to believe ($n - 1$).

However, Lewis is rather opaque about what his 'ancillary' assumptions exactly are. Understandably, subsequent formal treatments of common knowledge largely tended to overlook this critical inductive dimension in their accounts. Additionally, most formal treatments have also neglected the role of witnesses in generating common knowledge. For instance, at a pre-theoretic level, it is widespread in the game theory literature to speak of common knowledge/belief of P and mean no more than the obtaining of some infinite collection of propositions like:

- (1)' Everyone knows/believes P .
- (2)' Everyone knows/believes (1)'.
- (3)' Everyone knows/believes (2)'.
-
- (n)' Everyone knows/believes ($n - 1$)'.

As a result, the original Lewisian picture of a witness generating common knowledge has largely been swept under the rug. Cubitt & Sugden 2003 are a notable exception; they attempt to painstakingly reconstruct Lewis' vision and provide an account of what his 'ancillary' assumptions about agents' inductive standards might be. While such an exercise is not without its merits, the primary purpose of this paper is not to argue about what Lewis did or did not think. Instead, we will aim to provide a logic for capturing how common inductive knowledge can be generated by a witness that dispenses with the need for Lewis' 'ancillary' assumptions as much as possible while still maintaining the spirit of his original picture.

2 Pre-Theory

Sharpening Lewis' view, our primary thesis will be to clearly distinguish the notion of 'having *inductive* reason to believe W ' from the notion of 'having W as *deductive* reason to believe P .' At a high level, imagine agents receive possibly distinct streams of evidence and have different 'switching tolerances' for how many times they are willing to change their mind while limit deciding a proposition. In this setting, when we say an agent i has reason to believe a proposition, we intend to convey that there is some (possibly transitory) state of affairs where i is justified in believing this proposition. Namely, we say that agent i has *inductive* reason to believe W iff i can 'feasibly' and 'efficiently' limit decide W after affirming it in light of some true evidence E , where 'feasibly' means doing so within agent i 's switching tolerance while 'efficiently' means doing so with as few mind changes as methodologically possible. This way of modeling inductive justification can be seen as an alternative to Bayesianism, where instead of conditioning a credence function, an agent's primary guiding epistemic norm is to follow Ockham's razor (Kelly 2004). On this view, methods which make unnecessary mind changes are needlessly complex. Hence, an agent should chose to believe W in light of true evidence E only if doing so is part of a 'simple' method which minimizes the number times they will have to change their mind.

Next, we say that i has W as *deductive* reason to believe P iff:

- i has inductive reason to believe W .
- If E is some true evidence which gives i inductive reason to believe W then $W \wedge E$ entails P .

We take the latter requirement above to constitute the meaning of the phrase ' W indicates to i that P ' since it can be glossed as saying 'if i had inductive reason to believe W , i would thereby have W as deductive reason to believe P .'² According to these definitions then, we have as primitive rule of inference:

$$\frac{i \text{ has } A \text{ as deductive reason to believe } B}{(i \text{ has inductive reason to believe } A) \wedge (A \text{ indicates to } i \text{ that } B)}$$

Now observe for any proposition W and true evidence E that $W \wedge E$ entails W . Hence, W always indicates to i that W . Similarly if P is valid then for any proposition W and true evidence E we have that $W \wedge E$ entails P , i.e. W indicates to i that P . Thus, we have as additional primitive rules:

$$\frac{}{A \text{ indicates to } i \text{ that } A} \quad \frac{B}{A \text{ indicates to } i \text{ that } B}$$

Next suppose W indicates to i that P and W also indicates to i that $P \rightarrow Q$. If E is some true evidence which gives i inductive reason to believe W then $W \wedge E$ entails both P and $P \rightarrow Q$, i.e. $W \wedge E$ entails Q . Hence, we also have that W indicates to i that Q . This gives us the primitive rule:

$$\frac{[(A \text{ indicates to } i \text{ that } B) \wedge (A \text{ indicates to } i \text{ that } B \rightarrow C)]}{(A \text{ indicates to } i \text{ that } C)}$$

The above rule is quite powerful. For instance, along with our first three primitive rules, it can be used to show all the following derived rules:

$$\frac{(i \text{ has } A \text{ as deductive reason to believe } B) \wedge [i \text{ has } A \text{ as deductive reason to believe } (B \rightarrow C)]}{i \text{ has } A \text{ as deductive reason to believe } C}$$

$$\frac{[(A \text{ indicates to } i \text{ that } B) \wedge (A \text{ indicates to } i \text{ that } C)]}{(A \text{ indicates to } i \text{ that } B \wedge C)}$$

$$\frac{(i \text{ has } A \text{ as deductive reason to believe } B) \wedge (i \text{ has } A \text{ as deductive reason to believe } C)}{i \text{ has } A \text{ as deductive reason to believe } (B \wedge C)}$$

$$\frac{A \rightarrow B}{A \text{ indicates to } i \text{ that } B}$$

$$\frac{B \rightarrow C}{(A \text{ indicates to } i \text{ that } B) \rightarrow (A \text{ indicates to } i \text{ that } C)}$$

$$\frac{B \rightarrow C}{(i \text{ has } A \text{ as deductive reason to believe } B) \rightarrow (i \text{ has } A \text{ as deductive reason to believe } C)}$$

Additionally, we will maintain throughout that our agents are *valid reasoners*, i.e. if W holds and i has W as deductive reason to believe P then P also holds.

²We stress that indication is *not* mere entailment. Nevertheless, it may be helpful for readers to think of indication as just being mere entailment on a first pass of this section.

$$[A \wedge (i \text{ has } A \text{ as deductive reason to believe } B)] \rightarrow B$$
$$[A \wedge (i \text{ has } A \text{ as deductive reason to believe } B)] \rightarrow (i \text{ has some true reason to believe } B)$$

$$(i \text{ has some true reason to believe } A) \rightarrow A$$

1. Everyone has inductive reason to believe W .
2. W indicates to everyone that P .
3. W indicates to everyone that 1 and 2.
-
- n . W indicates to everyone that 1 and $(n-1)$.

$$(A \text{ generates common inductive knowledge of } B) \rightarrow (i \text{ has } A \text{ as deductive reason to believe } B)$$

³As an example, suppose Alice has agreed to deposit assets of a particular value in a designated account as collateral to Bob and her account initially starts with sufficient collateral. At every future time step, Bob can see the assets in Alice's account and determine whether their total value continues to constitute a sufficient amount of collateral. Unfortunately, the proposition $P = \text{'Alice's account eventually always has sufficient collateral'}$ is not limit decidable for Bob. However, for any finite natural k , the proposition $W_k = \text{'Alice's account dips below having sufficient collateral at most } k \text{ times but eventually always has sufficient collateral'}$ is limit decidable for Bob. In particular, if Bob has a switching tolerance of at least $2k + 1$, then Bob can feasibly and efficiently limit decide W_k after affirming it in light of initially seeing that Alice's account has sufficient collateral, i.e. Bob has inductive reason to believe W_k . Suppose in reality that Alice's account dips below having sufficient collateral exactly 1 time before bouncing back to eventually always having sufficient collateral and Bob has a switching tolerance of 5. Then W_2 is true and so Bob can attain stable justified true belief in the proposition 'Alice's account dips below having sufficient collateral at most 2 times but eventually always has sufficient collateral.' Additionally, since W_2 entails P , Bob has W_2 as deductive reason to believe P and can therefore also attain stable justified true belief in P . In this sense, Bob can inductively know 'Alice's account eventually always has sufficient collateral' even though he cannot limit decide this fact.

Now, note 1 and 3 entail ‘Everyone has W as deductive reason to believe 1 and 2.’ Since 1 and 2 entail ‘Everyone has W as deductive reason to believe P ,’ 1 and 3 also entail ‘Everyone has W as deductive reason to believe ‘everyone has W as deductive reason to believe P .’ Similarly, starting with 1 and 4, one can recursively derive ‘Everyone has W as deductive reason to believe ‘everyone has W as deductive reason to believe ‘everyone has W as deductive reason to believe P .’” Hence, **L1** entails what we will call **L2**:

- (1) Everyone has W as deductive reason to believe P .
- (2) Everyone has W as deductive reason to believe (1).
- (3) Everyone has W as deductive reason to believe (2).
-
- (n) Everyone has W as deductive reason to believe ($n - 1$).

Next, note (2) entails ‘ W indicates to everyone that (1).’ Since (1) entails 1 and 2, (2) also entails ‘ W indicates to everyone that 1 and 2,’ i.e. (2) entails 1 and 3. Similarly, starting with (3), one can recursively derive ‘ W indicates to everyone that 1 and 3,’ i.e. (3) entails 1 and 4. Thus, **L1** and **L2** are logically equivalent and so either can be taken to constitute what it means for W to generate common inductive knowledge of P . While **L1** evidently bears more structural similarity to Lewis’ account of how common inductive knowledge is generated, we will work frequently with **L2** for purposes of clarity. For instance, **L2** makes it obvious that if W generates common inductive knowledge of P then for each (n) $\in$ **L2** everyone has W as deductive reason to believe (n), and so given our prior pre-theoretic commitments, everyone has W as deductive reason to believe ‘ W generates common inductive knowledge of P .’ This motivates the primitive rule of inference:

$$\frac{A \text{ generates common inductive knowledge of } B}{i \text{ has } A \text{ as deductive reason to believe } (A \text{ generates common inductive knowledge of } B)}$$

In fact, for any proposition Q , it can be seen that if Q entails ‘everyone has W as deductive reason to believe P ’ and Q entails ‘everyone has W as deductive reason to believe Q ’ then Q entails every proposition in **L2** holds. To demonstrate, using the first entailment and the consequent of the second entailment, we have Q entails ‘everyone has W as deductive reason to believe ‘everyone has W as deductive reason to believe P .’ Similarly, using this new entailment with the consequent of the second entailment, one can recursively derive Q entails ‘everyone has W as deductive reason to believe ‘everyone has W as deductive reason to believe ‘everyone has W as deductive reason to believe P .’” Thus, we stipulate as yet another primitive rule:

$$\frac{C \rightarrow (\text{Everyone has } A \text{ as deductive reason to believe } C) \quad C \rightarrow (\text{Everyone has } A \text{ as deductive reason to believe } B)}{C \rightarrow (A \text{ generates common inductive knowledge of } B)}$$

How then ought we define common inductive knowledge? Lewis suggests we say P can become common inductive knowledge iff :

some state of affairs holds such that it generates common inductive knowledge of P .⁴

While we agree with Lewis that this condition is sufficient for common inductive knowledge, we will insist that it is not necessary.

⁴Again, strictly speaking, Lewis actually says that P is common knowledge iff some state of affairs holds such that it is a basis for common knowledge of P .

To motivate why, suppose that W holds and generates common inductive knowledge of P as Lewis suggests. Then, using **L2**, we immediately have **L3**:

(1)' Everyone has some true reason to believe P .

(2)' Everyone has some true reason to believe (1)'.

(3)' Everyone has some true reason to believe (2)'.

... ..

(n)' Everyone has some true reason to believe ($n-1$)'.

Recall Lewis seeks to show his 3 conditions for a witness to generate common inductive knowledge of P entail exactly this sort of infinite list. But **L3** is not equivalent to Lewis' definition of common inductive knowledge for each agent i could have a different true reason for believing each proposition in **L3**. On our view, Lewis' definition is unnecessarily stringent by his own lights and **L3** is closer to a desirable definition of common inductive knowledge. However, **L3** is not without its own shortcomings. For if every proposition in **L3** holds, although we can conclude for each (n)' $\in$ **L3** everyone has some true reason to believe (n)', we cannot further conclude everyone has some true reason to believe every (n)' $\in$ **L3**. After all, agents may not have a principled procedure for creating a new 'combined' true reason to believe every (n)' $\in$ **L3**. Thus, to guarantee our notion of common inductive knowledge is a 'fixed point,' our tertiary thesis will be that what it means for P to potentially become common inductive knowledge is constituted by satisfying the primitive rules of inference:

$$\frac{(A \text{ can become common inductive knowledge}) \rightarrow A}{\frac{A \text{ can become common inductive knowledge} \rightarrow}{i \text{ has some true reason to believe } (A \text{ can become common inductive knowledge})} \frac{A \rightarrow (\text{Everyone has some true reason to believe } A) \quad A \rightarrow B}{A \rightarrow (B \text{ can become common inductive knowledge})}}$$

In particular, since 'everyone has some true reason to believe' can be alternatively glossed as 'everyone can inductively know,' this means ' P can become common inductive knowledge' is the weakest proposition entailing P such that if it is true then everyone can inductively know it.

As a check, we now briefly argue why Lewis' definition of common inductive knowledge entails ours. So, suppose P can become Lewisian common inductive knowledge. Then some 'witness' W holds so that W generates common inductive knowledge of P . By our earlier inference rules, this means everyone has W as deductive reason to believe P and also that everyone has W as deductive reason to believe ' W generates common inductive knowledge of P .' Since W holds and everyone is a valid reasoner, it follows that P holds and so ' P can become Lewisian common inductive knowledge' entails P . Additionally, since W entails W , it follows that everyone has W as deductive reason to believe ' W holds and generates common inductive knowledge of P .' Using the fact W holds then, this means ' P can become Lewisian common inductive knowledge' entails 'everyone has some true reason to believe ' P can become Lewisian common inductive knowledge.' Therefore, by our last primitive rule, we can conclude ' P can become Lewisian common inductive knowledge' entails ' P can become common inductive knowledge,' i.e.:

$$\frac{[W \wedge (W \text{ generates common inductive knowledge of } P)]}{\rightarrow} P \text{ can become common inductive knowledge}$$

In the coming section, we provide the reader with an overview of concepts from topological learning theory which will feature prominently in our semantics. We begin by reviewing standard results from the field. This review is intended to fix notation and recall basic facts. Readers who are already familiar with the topological difference hierarchy and its epistemic interpretation may wish to skip ahead. Such readers can turn directly to page 12 where we begin introducing several new pieces of terminology. These new terms and the theorems we prove about them will be essential for understanding the rest of the paper.

3 Topology & Learning

Let $\Omega \neq \emptyset$ be the set of all possible worlds. We say that a *piece of evidence* E for agent i is a non-empty subset of possible worlds in Ω which eliminates all worlds outside E as possibilities. Let $\mathcal{E}_i$ consist of all pieces of evidence agent i might ever learn. Further, given a world $w \in \Omega$, let $\mathcal{E}_i(w)$ denote those elements of $\mathcal{E}_i$ which contain w . Intuitively, $\mathcal{E}_i(w)$ represents exactly those pieces of evidence which agent i eventually learns in world w . We will assume that $\mathcal{E}_i$ forms an *information basis* for i , that is:

1. $\forall w \in \Omega, \mathcal{E}_i(w) \neq \emptyset$.
2. $\forall w \in \Omega$ and $\forall E_1, E_2 \in \mathcal{E}_i(w)$ there exists $E_3 \in \mathcal{E}_i(w)$ such that $E_3 \subseteq E_1 \cap E_2$.

The first requirement above simply says that agent i learns some piece of evidence in every world. In our view, this is a very weak assumption. For even if agent i never rules out any possible world in Ω , then Ω itself must be a piece of evidence for i . The second requirement is more substantive and corresponds to the assumption that evidence is cumulative. Namely, this requirement says that if agent i learns both E_1 and E_2 as evidence then they will also learn some piece of evidence E_3 at least as strong as $E_1 \cap E_2$. We now note that our definition of an *information basis* is identical to that of a *topological basis*. Hence, the collection $\mathcal{T}_i$ generated by taking arbitrary unions of pieces of evidence in $\mathcal{E}_i$ forms a topology over Ω .

In particular, it turns out the limit decidability of a set for agent i will be closely related to whether the set can be expressed as the nested difference of a certain number of descending open sets in the topology $\mathcal{T}_i$.⁵ However, to illustrate why, we will need to introduce quite a bit of terminology. For starters, say that:

- A *decision method* for i is a map $m_i : \mathcal{E}_i \rightarrow \{\text{Yes}, \text{No}\}$.
- A *t-switching sequence starting with 'Yes'* for m_i is a finite downward sequence $E_0 \supseteq \dots \supseteq E_t$ of pieces of evidence in $\mathcal{E}_i$ such that $m_i(E_{2k}) = \text{Yes}$ and $m_i(E_{2k+1}) = \text{No}$.
- A *t-switching sequence starting with 'No'* for m_i is a finite downward sequence $E_0 \supseteq \dots \supseteq E_t$ of pieces of evidence in $\mathcal{E}_i$ such that $m_i(E_{2k}) = \text{No}$ and $m_i(E_{2k+1}) = \text{Yes}$.
- A decision method for i *has at most n switches after saying 'Yes'* if it is a decision method which contains no t -switching sequences starting with 'Yes' for any $t > n$.
- A decision method for i *has at most n switches after saying 'No'* if it is a decision method which contains no t -switching sequences starting with 'No' for any $t > n$.
- A decision method for i *has a bounded number of switches* if it is a decision method which either has at most n switches after saying 'Yes' or has at most n switches after saying 'No' for some n .

⁵We define the nested difference of a finite sequence of sets $S_0, \dots, S_n$ to be given by $S_0 \setminus (S_1 \setminus (\dots S_n \dots))$.

Lemma 3.0: Given a decision method m_i for agent i which has a bounded number of switches, there exists a function $\sigma_{m_i} : \Omega \rightarrow \{\text{Yes}, \text{No}\}$ which $\forall w \in \Omega$ sets:

$$\sigma_{m_i}(w) = \begin{cases} \text{Yes} & \exists E \in \mathcal{E}_i(w) \forall E' \in \mathcal{E}_i(w), E' \subseteq E \rightarrow m_i(E') = \text{Yes} \\ \text{No} & \exists E \in \mathcal{E}_i(w) \forall E' \in \mathcal{E}_i(w), E' \subseteq E \rightarrow m_i(E') = \text{No} \end{cases}$$

Proof

First we show σ_{m_i} is a well-defined partial function. So suppose towards a contradiction it were not. Then for some $w^* \in \Omega$ both piecewise conditions above are satisfied. Then there $\exists E_1, E_2 \in \mathcal{E}_i(w^*)$ such that $\forall E' \in \mathcal{E}_i(w^*), E' \subseteq E_1 \rightarrow m_i(E') = \text{Yes}$ and $E' \subseteq E_2 \rightarrow m_i(E') = \text{No}$. Further, since $\mathcal{E}_i$ is an information basis, $\exists E_3 \in \mathcal{E}_i(w^*)$ such that $E_3 \subseteq E_1 \cap E_2$. But this implies that both $m_i(E_3) = \text{Yes}$ and $m_i(E_3) = \text{No}$, giving us our desired contradiction.

Now we show σ_{m_i} is total. So suppose towards a contradiction it were not. Then for some $w^* \in \Omega$ neither piecewise condition above is satisfied. Further, since $\mathcal{E}_i$ is an information basis, we know $\exists E_0 \in \mathcal{E}_i(w^*)$. WLOG suppose $m_i(E_0) = \text{Yes}$. Because the first piecewise is not satisfied, $\exists E_1 \in \mathcal{E}_i(w^*)$ so that $E_1 \subseteq E_0$ and $m_i(E_1) = \text{No}$. However, because the second piecewise condition is not satisfied, $\exists E_2 \in \mathcal{E}_i(w^*)$ so that $E_2 \subseteq E_1$ and $m_i(E_2) = \text{Yes}$. In this way we can create an infinite downward sequence $E_0 \supseteq E_1 \supseteq E_2 \dots$ of pieces of evidence in $\mathcal{E}_i$ such that $m_i(E_{2k}) = \text{Yes}$ and $m_i(E_{2k+1}) = \text{No}$. This contradicts the fact that m_i is a decision method which has a bounded number of switches, as desired. $\square$

Conceptually speaking, given a decision method m_i for agent i which has a bounded number of switches, $\sigma_{m_i}(w)$ returns the output that m_i converges to in world w . Since $\sigma_{m_i}^{-1}(\text{Yes}) = W$ iff $\sigma_{m_i}^{-1}(\text{No}) = \Omega \setminus W$:

- Agent i can limit decide W in n switches after saying ‘Yes’ just in case there is a decision method m_i for i such that m_i has at most n switches after saying ‘Yes’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$.
- Agent i can limit decide W in n switches after saying ‘No’ just in case there is a decision method m_i for i such that m_i has at most n switches after saying ‘No’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$.

Theorem 3.0: Agent i can limit decide W in n switches after saying ‘Yes’ if and only if W can be expressed as the nested difference of $n + 1$ descending open sets in $\mathcal{T}_i$.

Proof

First we prove the ‘only if’ direction. So assume agent i can limit decide W in n switches after saying ‘Yes.’ Then there is a decision method m_i for i such that m_i has at most n switches after saying ‘Yes’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Let $I_0 = \{E \in \mathcal{E}_i : m_i(E) = \text{Yes}\}$. Next, recursively define I_k by setting:

$$I_k = \begin{cases} \{E' \in \mathcal{E}_i : \exists E \in I_{k-1}, E' \subseteq E \text{ and } m_i(E') = \text{No}\} & k \text{ is odd} \\ \{E' \in \mathcal{E}_i : \exists E \in I_{k-1}, E' \subseteq E \text{ and } m_i(E') = \text{Yes}\} & k \text{ is even} \end{cases}$$

Note that $I_{n+1} = \emptyset$ since if $\exists E_{n+1} \in I_{n+1}$ then we could find a $(n + 1)$ -switching sequence starting with ‘Yes’ for m_i of the form $E_0 \supseteq \dots \supseteq E_{n+1}$ where each $E_k \in I_k$. Now let $O_k = \bigcup_{E \in I_k} E$. Each O_k is open as it is a union of basis elements. In particular, $\{O_k\}_{k \in \mathbb{N}}$ forms a descending sequence of open sets with $O_{n+1} = \emptyset$ and so it suffices show $O_0 \setminus (O_1 \setminus (\dots O_n \dots)) = W$.

Suppose $w \in O_0 \setminus (O_1 \setminus (\dots O_n \dots))$. Let k^* be the largest k for which $w \in O_k$. Since O_{k^*} is not empty, we know $k^* < n + 1$. Additionally, since $w \in O_0 \setminus (O_1 \setminus (\dots O_n \dots))$, we know k^* is even. Hence, $\exists E^* \in I_{k^*}$ such that $w \in E^*$ and $m_i(E^*) = \text{Yes}$. Assume for sake of contradiction that $\exists E' \in \mathcal{E}_i(w)$ such that $E' \subseteq E^*$ and $m_i(E') = \text{No}$. Then $E' \in I_{k^*+1}$ and so $w \in E' \subseteq O_{k^*+1}$, contradicting maximality of k^* . Thus, E^* is an element of $\mathcal{E}_i(w)$ such that $\forall E' \in \mathcal{E}_i(w)$ we have $E' \subseteq E^* \rightarrow m_i(E') = \text{Yes}$, i.e. $w \in \sigma_{m_i}^{-1}(\text{Yes}) = W$.

Now suppose $w \in W = \sigma_{m_i}^{-1}(\text{Yes})$. Then $\exists E_0 \in \mathcal{E}_i(w)$ such that $m_i(E_0) = \text{Yes}$. In particular, $E_0 \in I_0$ and so $w \in E_0 \subseteq O_0$. Let k^* be the largest k for which $w \in O_k$. Since O_{k^*} is not empty we know $k^* < n+1$. Suppose towards a contradiction that k^* is odd. Then, $\exists E^* \in I_{k^*}$ such that $w \in E^*$ and $m_i(E^*) = \text{No}$. Further E^* is an element of $\mathcal{E}_i(w)$ such that $\forall E' \in \mathcal{E}_i(w)$, if $E' \subseteq E^*$ then $m_i(E') = \text{No}$ (otherwise we would have $E' \in I_{k^*+1}$ and $w \in E' \subseteq O_{k^*+1}$, contradicting maximality of k^*). But this means $w \in \sigma_{m_i}^{-1}(\text{No})$. Contradiction. Thus k^* is even and so we can conclude $w \in O_0 \setminus (O_1 \setminus (\dots O_n) \dots)$.

Next we prove the ‘if’ direction. So suppose W can be expressed as the nested difference of $n+1$ descending open sets $O_0 \supseteq \dots \supseteq O_n$ in $\mathcal{T}_i$. Given $E \in \mathcal{E}_i$, define k_E to be the largest k for which $E \subseteq O_k$ (if there is no such k set $k_E = -1$). Let m_i be the decision method for i defined by $\forall E \in \mathcal{E}_i$ setting:

$$m_i(E) = \begin{cases} \text{No} & k_E \text{ is odd} \\ \text{Yes} & k_E \text{ is even} \end{cases}$$

Assume for sake of contradiction there is a $(n+1)$ -switching sequence starting with ‘Yes’ for m_i of the form $E_0 \supseteq \dots \supseteq E_{n+1}$. Since k_{E_0} is at least 0 and $k_{E_0} < k_{E_1} < \dots < k_{E_{n+1}}$, we know $k_{E_{n+1}} \geq n+1$. However, we also know $k_{E_{n+1}} \leq n$ by definition. Contradiction. Hence, m_i has at most n switches after saying ‘Yes’ and so it suffices to show $O_0 \setminus (O_1 \setminus (\dots O_n) \dots) = \sigma_{m_i}^{-1}(\text{Yes})$.

Suppose $w \in O_0 \setminus (O_1 \setminus (\dots O_n) \dots)$. Let k^* be the largest k for which $w \in O_k$. Then $\exists E^* \in \mathcal{E}_i(w)$ such that $E^* \subseteq O_{k^*}$. Further, E^* is an element of $\mathcal{E}_i(w)$ such that $\forall E' \in \mathcal{E}_i(w)$, if $E' \subseteq E^*$ then $k_{E'} = k^*$ (otherwise we would have $k_{E'} > k^*$ and $w \in E' \subseteq O_{k_{E'}}$, contradicting maximality of k^*). Since $w \in O_0 \setminus (O_1 \setminus (\dots O_n) \dots)$, we know k^* is even. Thus $\forall E' \in \mathcal{E}_i(w)$, if $E' \subseteq E^*$ then $m_i(E') = \text{Yes}$, i.e. $w \in \sigma_{m_i}^{-1}(\text{Yes})$.

Now suppose $w \in \sigma_{m_i}^{-1}(\text{Yes})$. Then $\exists E_0 \in \mathcal{E}_i(w)$ such that $m_i(E_0) = \text{Yes}$. In particular, $w \in E_0 \subseteq O_0$. Let k^* be the largest k for which $w \in O_k$. Suppose for contradiction that k^* is odd. We know $\exists E^* \in \mathcal{E}_i(w)$ such that $E^* \subseteq O_{k^*}$. Further, E^* is an element of $\mathcal{E}_i(w)$ such that $\forall E' \in \mathcal{E}_i(w)$, if $E' \subseteq E^*$ then $k_{E'} = k^*$ (otherwise we would have $k_{E'} > k^*$ and $w \in E' \subseteq O_{k_{E'}}$, contradicting maximality of k^*). But since k^* is odd this means $\forall E' \in \mathcal{E}_i(w)$, if $E' \subseteq E^*$ then $m_i(E') = \text{No}$, i.e. $w \in \sigma_{m_i}^{-1}(\text{No})$. Contradiction. Thus k^* is even and so we can conclude $w \in O_0 \setminus (O_1 \setminus (\dots O_n) \dots)$. $\square$

Lemma 3.1: Agent i can limit decide W in n switches after saying ‘No’ if and only if agent i can limit decide $\Omega \setminus W$ in n switches after saying ‘Yes.’

Proof

We only show one direction as the other direction follows from a parallel argument. So assume agent i can limit decide W in n switches after saying ‘No.’ Then there is a decision method m_i for i such that m_i has at most n switches after saying ‘No’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Consider the decision method m'_i obtained by $\forall E \in \mathcal{E}_i$ setting:

$$m'_i(E) = \begin{cases} \text{Yes} & m_i(E) = \text{No} \\ \text{No} & m_i(E) = \text{Yes} \end{cases}$$


Note that m'_i has at most n switches after saying ‘Yes.’ Further $\sigma_{m'_i}^{-1}(\text{Yes}) = \sigma_{m_i}^{-1}(\text{No}) = \Omega \setminus W$. Thus agent i can limit decide $\Omega \setminus W$ in n switches after saying ‘Yes.’ $\square$

Corollary 3.0: Agent i can limit decide W in n switches after saying ‘No’ if and only if W can be expressed as the complement of a nested difference of $n+1$ descending open sets in $\mathcal{T}_i$.

Proof

Immediate by the Theorem 3.0 and Lemma 3.1. $\square$

Going forward, we will say a set is n -open in $\mathcal{T}_i$ iff it is the nested difference of n descending open sets in $\mathcal{T}_i$. Additionally, say a set is n -closed in $\mathcal{T}_i$ iff it is the complement of a set which is n -open in $\mathcal{T}_i$. Finally, say a set is n -clopen in $\mathcal{T}_i$ iff it is both n -open and n -closed in $\mathcal{T}_i$. These topological complexity classes are then ordered by inclusion according to the below *topological difference hierarchy*:⁶



We have already seen an epistemic interpretation for when $n > 0$ of what it means to be n -open, n -closed, and n -clopen in $\mathcal{T}_i$. But what does it mean for a set to be 0-open or 0-closed in $\mathcal{T}_i$?⁷ One interpretation is as follows: say an information basis $\mathcal{E}_i$ for i has a *starting point* if $\Omega \in \mathcal{E}_i$, i.e. agent i receives Ω as a piece of evidence before learning anything stronger. If $\mathcal{E}_i$ has a starting point, a decision method m_i for i must make a choice of whether $m_i(\Omega) = \text{No}$ or $m_i(\Omega) = \text{Yes}$. Intuitively then, a set W is 0-open in $\mathcal{T}_i$ if and only if agent i can limit decide W in 0 switches *starting from* ‘No,’ i.e. $W = \emptyset$. Similarly, a set W is 0-closed in $\mathcal{T}_i$ if and only if agent i can limit decide W in 0 switches *starting from* ‘Yes,’ i.e. $W = \Omega$.

We now temporarily assume agent i ’s information basis $\mathcal{E}_i$ has a starting point and formalize what it means for agent i to be able to limit decide W in n switches starting from ‘Yes’/‘No’ by saying:

- A decision method m_i for i *starts from* ‘Yes’ (‘No’) if $m_i(\Omega) = \text{Yes}$ (No).
- A decision method for i *has at most n switches starting from* ‘Yes’ (‘No’) if it is a decision method which starts from ‘Yes’ (‘No’) and has at most n switches after saying ‘Yes’ (‘No’).
- A decision method for i *has at most n switches* if it is a decision method which has at most n switches starting from ‘Yes’ or at most n switches starting from ‘No.’
- Agent i *can limit decide W in n switches starting from* ‘Yes’ (‘No’) if there is a decision method m_i for i such that m_i has at most n switches starting from ‘Yes’ (‘No’) and $\sigma_{m_i}^{-1}(\text{Yes}) = W$.
- Agent i *can limit decide W in n switches* if agent i can limit decide W in n switches starting from ‘Yes’ or in n switches starting from ‘No.’

⁶For historical context, Putnam’s 1965 paper was the first to try and characterize the complexity of what he called ‘trial and error predicates,’ albeit in a computability theoretic context. Kevin Kelly later recast the problem in general topological terms...

⁷No set is n -open, n -closed, or n -clopen for $n < 0$.

The above notions are also closely related to the topological difference hierarchy. To see precisely how, given a set W , let $n_{i,W}$ be the smallest n for which agent i can limit decide W in n switches (if there is no such n , let $n_{i,W} = \infty$). Additionally, assume agent i has finite *switching tolerance* $n_i \geq 0$ so that i adopts only those decision methods which have at most n_i switches. Because i can ‘feasibly’ limit decide W iff $n_{i,W} \leq n_i$ and does so ‘efficiently’ iff their method has at most $n_{i,W}$ switches:

- Agent i has reason to limit decide W starting from ‘Yes’ (‘No’) just in case $n_{i,W} \leq n_i$ and agent i can limit decide W in $n_{i,W}$ switches starting from ‘Yes’ (‘No’).

Theorem 3.1: Agent i has reason to limit decide W starting from ‘Yes’ if and only if $\exists k \leq n_i$ such that W is k -closed in $\mathcal{T}_i$ but not $(k-1)$ -open.

Proof

First we show the ‘if’ direction.

Suppose $\exists k \leq n_i$ so that W is k -closed in $\mathcal{T}_i$ but not $(k-1)$ -open.

If $k = 0$ we know W is 0-closed, and hence, i can limit decide W in 0 switches starting from ‘Yes,’ i.e. $n_{i,W} = 0$. Since $0 = n_{i,W} \leq n_i$, we can conclude agent i has reason to limit decide W starting from ‘Yes.’ If $k = 1$ we know W is not 0-open, and hence, i cannot limit decide W in 0 switches starting from ‘No.’ However, because W is 1-closed, i can limit decide W in 0 switches after saying ‘No.’ Thus, there is a method m_i for i such that m_i has at most 0 switches after saying ‘No’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Further, we know $m_i(\Omega) \neq \text{‘No’}$ as otherwise i could limit decide W in 0 switches starting from ‘No.’ Hence $m_i(\Omega) = \text{‘Yes’}$ and so i can limit decide W in 1 switch starting from ‘Yes.’ Therefore, either $n_{i,W} = 0$ or $n_{i,W} = 1$. Further $1 = k \leq n_i$ and so $n_{i,W} \leq n_i$. If $n_{i,W} = 1$ then we can conclude agent i has reason to limit decide W starting from ‘Yes.’ Similarly, if $n_{i,W} = 0$ then i can limit decide W in 0 switches starting from ‘Yes’ (as i can’t limit decide W in 0 switches starting from ‘No’). So, in this case as well, we can conclude agent i has reason to limit decide W starting from ‘Yes.’

If $k \geq 2$ we know W is not $(k-1)$ -open, and hence, i cannot limit decide W in $k-2$ switches after saying ‘Yes.’ However because W is k -closed, i can limit decide W in $k-1$ switches after saying ‘No.’ Thus, there is a method m_i for i such that m_i has at most $k-1$ switches after saying ‘No’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Further, for any such m_i , we know $m_i(\Omega) \neq \text{‘No’}$ as otherwise i could limit decide W in $k-2$ switches after saying ‘Yes.’ Hence $m_i(\Omega) = \text{Yes}$ and so i can limit decide W in k switches starting from ‘Yes’ but cannot in $(k-1)$ switches starting from ‘No.’ Therefore, $n_{i,W} \leq k \leq n_i$. If $n_{i,W} = k$ then we can conclude agent i has reason to limit decide W starting from ‘Yes.’ Similarly, if $n_{i,W} < k$ then i can limit decide W in $n_{i,W}$ switches starting from ‘Yes’ (as i can’t limit decide W in $n_{i,W}$ switches starting from ‘No’). So in these cases as well, we can conclude i has reason to limit decide W starting from ‘Yes.’

Now we show the ‘only if’ direction.

Suppose agent i has reason to limit decide W starting from ‘Yes.’ Then $n_{i,W} \leq n_i$. Further, i can limit decide W in $n_{i,W}$ switches starting from ‘Yes’ and, hence, after saying ‘Yes’. Thus, W is $n_{i,W}$ -closed in $\mathcal{T}_i$ and so it suffices to show W is not $(n_{i,W}-1)$ -open in $\mathcal{T}_i$.

If $n_{i,W} = 0$ then, since W is never (-1) -open, we are done.

If $n_{i,W} = 1$ then, by definition, i cannot limit decide W in 0 switches. In particular, i cannot limit decide W in 0 switches starting from ‘No.’ Hence $W \neq \emptyset$ and so W is not 0-open as desired.

If $n_{i,W} \geq 2$ then, by definition, i cannot limit decide W in $(n_{i,W}-1)$ switches. Suppose towards a contradiction W is $(n_{i,W}-1)$ -open. Then i can limit decide W in $n_{i,W}-2$ switches after saying ‘Yes.’ Thus, there is a method m_i for i such that m_i has at most $n_{i,W}-2$ switches after saying ‘Yes’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Further, we know $m_i(\Omega) \neq \text{Yes}$ as otherwise i could limit decide W in $(n_{i,W}-2)$ switches starting from ‘Yes.’ Hence $m_i(\Omega) = \text{No}$. But this means agent i can limit decide W in $(n_{i,W}-1)$ switches starting from ‘No.’ Contradiction. Hence W is not $(n_{i,W}-1)$ -open, as desired. $\square$

Lemma 3.2: Agent i has reason to limit decide W starting from ‘No’ if and only if agent i has reason to limit decide $\Omega \setminus W$ starting from ‘Yes.’

Proof

We only show one direction as the other direction follows from a parallel argument. So assume agent i has reason to limit decide W starting from ‘No.’ Then $n_{i,W} \leq n_i$ and there is a decision method m_i for i such that m_i has at most $n_{i,W}$ switches starting from ‘No’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Consider the decision m'_i obtained by $\forall E \in \mathcal{E}_i$ setting:

$$m'_i(E) = \begin{cases} \text{Yes} & m_i(E) = \text{No} \\ \text{No} & m_i(E) = \text{Yes} \end{cases}$$

Note that m'_i has at most $n_{i,W}$ switches starting from ‘Yes.’ Further $\sigma_{m'_i}^{-1}(\text{Yes}) = \sigma_{m_i}^{-1}(\text{No}) = \Omega \setminus W$. Hence $n_{i,\Omega \setminus W} \leq n_{i,W} \leq n_i$. Suppose towards a contradiction $n_{i,\Omega \setminus W} < n_{i,W}$. Then we can construct a decision method m''_i for i such that m''_i has at most $n_{i,\Omega \setminus W}$ switches and $\sigma_{m''_i}^{-1}(\text{Yes}) = \Omega \setminus W$. Further, we can construct a decision method m'''_i from m''_i (analogously to how we constructed m'_i from m_i) such that m'''_i has at most $n_{i,\Omega \setminus W}$ switches and $\sigma_{m'''_i}^{-1}(\text{Yes}) = W$. Contradiction. Hence $n_{i,\Omega \setminus W} = n_{i,W}$ and we can conclude agent i has reason to limit decide $\Omega \setminus W$ starting from ‘Yes.’ $\square$

Corollary 3.1: Agent i has reason to limit decide W starting from ‘No’ if and only if $\exists k \leq n_i$ such that W is k -open in $\mathcal{T}_i$ but not $(k-1)$ -closed.

Proof

Immediate by Theorem 3.1 and Lemma 3.2. $\square$

We now no longer assume $\mathcal{E}_i$ has a starting point. Instead, fix $E \in \mathcal{E}_i$ and let $\mathcal{E}_i|E = \{E' \in \mathcal{E}_i : E' \subseteq E\}$. If, in some state of affairs, E comprises the set of all worlds i considers possible then $\mathcal{E}_i|E$ consists of all those pieces of evidence which refine E that i might learn going forward. In particular, $\mathcal{E}_i|E$ forms a basis for the subspace topology $\mathcal{T}_i|E$ over E where $E \in \mathcal{E}_i|E$. Therefore, we can consistently relativize every bit of terminology we have introduced thus far to E by replacing Ω with E as the set of all possible worlds and replacing $\mathcal{E}_i$ with $\mathcal{E}_i|E$ so that i ’s information basis has a starting point. In this context, a decision method m_i for i is a map $m_i : \mathcal{E}_i|E \rightarrow \{\text{Yes}, \text{No}\}$. Furthermore, we can show that agent i has reason to limit decide W starting from ‘Yes’ iff $\exists k \leq n_i$ such that W is k -closed in $\mathcal{T}_i|E$ but not $(k-1)$ -open. Similarly, we can show that i has reason to limit decide W starting from ‘No’ iff $\exists k \leq n_i$ such that W is k -open in $\mathcal{T}_i|E$ but not $(k-1)$ -closed. Implicit in both these statements however is that W is a subset of E . So, given an arbitrary subset W of Ω , one must take its intersection with E for such statements to even make sense (after all $W \cap E$ consists of those worlds where W is true and i learns evidence E). As a consequence then, in our original context, we say:

- Agent i has reason to limit decide W starting from ‘Yes’ in light of evidence E if $\exists k \leq n_i$ such that $W \cap E$ is k -closed in $\mathcal{T}_i|E$ but not $(k-1)$ -open.
- Agent i has reason to limit decide W starting from ‘No’ in light of evidence E if $\exists k \leq n_i$ such that $W \cap E$ is k -open in $\mathcal{T}_i|E$ but not $(k-1)$ -closed.

For concision, we abbreviate the above phrases in natural language by writing:

- E gives i inductive reason to believe W .
- E gives i inductive reason to disbelieve W .

and denote them by the predicates $\text{Yes}_i(W|E)$ and $\text{No}_i(W|E)$ respectively.

Lemma 3.3: Suppose $(X_j)_{j \in J}$ is a collection of $(n+1)$ -open sets in $\mathcal{T}_i$. Furthermore, suppose that each X_j is the nested difference of a descending sequence of open sets $O_0^j \supseteq \dots \supseteq O_n^j$. If for each $m \in J$ we have $X_m = O_0^m \cap \bigcup_{j \in J} X_j$ then $\bigcup_{j \in J} X_j$ is the nested difference of the descending sequence of open sets $\bigcup_{j \in J} O_0^j \supseteq \dots \supseteq \bigcup_{j \in J} O_n^j$ and so $\bigcup_{j \in J} X_j$ is $(n+1)$ -open in $\mathcal{T}_i$.⁸

Proof

We proceed by induction on n . When $n = 0$, each $X_j = O_0^j$ and so we can immediately establish the base case as $\bigcup_{j \in J} X_j = \bigcup_{j \in J} O_0^j$. For the inductive step, assume the claim holds when $n = k - 1$. Take $(X_j)_{j \in J}$ to be a collection of $(k+1)$ -open sets in $\mathcal{T}_i$ so that each X_j is the nested difference of a descending sequence of open sets $O_0^j \supseteq \dots \supseteq O_k^j$. Suppose for each $m \in J$ we have $X_m = O_0^m \cap \bigcup_{j \in J} X_j$. Next, for each $j \in J$, let X'_j be the nested difference of the descending sequence of open sets $O_1^j \supseteq \dots \supseteq O_k^j$. We claim now that:

$$\bigcup_{j \in J} X_j = \left(\bigcup_{j \in J} O_0^j \right) \setminus \left(\bigcup_{j \in J} X'_j \right).$$

Suppose $w \in \left(\bigcup_{j \in J} O_0^j \right) \setminus \left(\bigcup_{j \in J} X'_j \right)$ then there is some $s \in J$ such that $w \in O_0^s$. Since we are guaranteed $w \notin X'_s$ this means $w \in X_s$ and so $w \in \bigcup_{j \in J} X_j$. Now suppose $w \in \bigcup_{j \in J} X_j$. Then there is some $s \in J$ such that $w \in X_s$. In particular, we know $w \in O_0^s$. Entertain for contradiction there exists some $m \in J$ so that $w \in X'_m$. Since $X_m = O_0^m \cap \bigcup_{j \in J} X_j$ we know $O_0^m \cap X_s \subseteq X_m$. Since $w \in X'_m$ we have $w \notin X_m$ and hence $w \notin O_0^m \cap X_s$. However, since $X'_m \subseteq O_0^m$, this is only possible if $w \notin X_s$. Contradiction. Therefore $w \notin \bigcup_{j \in J} X'_j$ and so $w \in \left(\bigcup_{j \in J} O_0^j \right) \setminus \left(\bigcup_{j \in J} X'_j \right)$. Next we claim for each $m \in J$ that:

$$X'_m = O_0^m \cap \left(\bigcup_{j \in J} X_j \right)^c.$$

Suppose $w \in O_0^m \cap \left(\bigcup_{j \in J} X_j \right)^c$. Then $w \in O_0^m$ and $w \notin \bigcup_{j \in J} X_j$. Since $X_m = O_0^m \cap \bigcup_{j \in J} X_j$, we know $w \notin X_m$. Further, because $w \in O_0^m$, this is only possible if $w \in X'_m$. Now suppose $w \in X'_m$. Then $w \in O_0^m$ and $w \notin X_m$. Since $X_m = O_0^m \cap \bigcup_{j \in J} X_j$, this is only possible if $w \notin \bigcup_{j \in J} X_j$. Hence, $w \in O_0^m \cap \left(\bigcup_{j \in J} X_j \right)^c$. Now, using the fact $X'_m \subseteq O_0^m$ and $X'_m \subseteq O_0^m \subseteq \bigcup_{j \in J} O_0^j$, we have for each $m \in J$ that:

$$X'_m = \left(O_0^m \cap \left(\bigcup_{j \in J} X_j \right)^c \right) \cap \left(O_1^m \cap \bigcup_{j \in J} O_0^j \right).$$

Since $O_1^m \subseteq O_0^m$ and $\left(\bigcup_{j \in J} X_j \right)^c \cap \left(\bigcup_{j \in J} O_0^j \right) = \bigcup_{j \in J} \left(O_0^j \cap \left(\bigcup_{j' \in J} X_{j'} \right)^c \right) = \bigcup_{j \in J} X'_j$, we can rearrange the above to obtain for each $m \in J$ that $X'_m = O_1^m \cap \bigcup_{j \in J} X'_j$. By hypothesis then, we have $\bigcup_{j \in J} X'_j$ is the nested difference of the descending sequence of open sets $\bigcup_{j \in J} O_1^j \supseteq \dots \supseteq \bigcup_{j \in J} O_k^j$. Finally, since $\bigcup_{j \in J} X_j = \left(\bigcup_{j \in J} O_0^j \right) \setminus \left(\bigcup_{j \in J} X'_j \right)$, we can conclude $\bigcup_{j \in J} X_j$ is the nested difference of the descending sequence of open sets $\bigcup_{j \in J} O_0^j \supseteq \dots \supseteq \bigcup_{j \in J} O_k^j$. $\square$

⁸In general, when $n > 0$, the union of two $(n+1)$ -open sets in $\mathcal{T}_i$ may not itself be $(n+1)$ -open. For instance, suppose we are working in Cantor space. Take A to be the set of sequences which have at least 1 and at most 5 non-0 entries. Take B to be the set of sequences which have at least 7 and at most 13 non-0 entries. One can verify that while A and B are both 2-open, $A \cup B$ is not 2-open (although, it is 4-open). Similarly, when $n > 1$, the intersection of two $(n+1)$ -open sets in $\mathcal{T}_i$ may not itself be $(n+1)$ -open. Again, suppose we are working in Cantor space. Take C to be the set of sequences which have at least 3 and at most 9 non-0 entries. Take D to be the set of sequences which have at least 11 and at most 15 non-0 entries. Defining A and B as before, one can verify that while $A \cup B$ and $C \cup D$ are both 4-open, $(A \cup B) \cap (C \cup D)$ is not 4-open (although, it is 6-open).

Lemma 3.4: If W is $(n+1)$ -open in $\mathcal{T}_i$ then $\forall w \in W$ there $\exists E \in \mathcal{E}_i(w)$ such that $W \cap E$ is n -closed in the subspace topology $\mathcal{T}_i|E$.

Proof

Since W is $(n+1)$ -open there exists a descending sequence of open sets $O_0 \supseteq \dots \supseteq O_n$ in $\mathcal{T}_i$ such that $W = O_0 \setminus (O_1 \setminus (\dots O_n \dots))$. Suppose $w \in W$. Then $w \in O_0$ and so there exists $E^* \in \mathcal{E}_i(w)$ such that $E^* \subseteq O_0$. Additionally:

$$W \cap E^* = (O_0 \cap E^*) \setminus ((O_1 \cap E^*) \setminus (\dots (O_n \cap E^*) \dots)) = E^* \setminus ((O_1 \cap E^*) \setminus (\dots (O_n \cap E^*) \dots))$$

, i.e. $W \cap E^*$ is n -closed in $\mathcal{T}_i|E^*$. $\square$

Theorem 3.2: W is (n_i+1) -open in $\mathcal{T}_i$ iff $\forall w \in W$ there $\exists E \in \mathcal{E}_i(w)$ such that $\text{Yes}_i(W|E)$ holds.

Proof

First we show the ‘if’ direction. Suppose $\forall w \in W$ there $\exists E \in \mathcal{E}_i(w)$ such that $\text{Yes}_i(W|E)$ holds. In particular, for each $w \in W$, take E_w to be an element of $\mathcal{E}_i(w)$ for which $\text{Yes}_i(W|E_w)$ holds so that $W = \bigcup_{w \in W} W \cap E_w$. Note for each $w \in W$ there is a decision method $m_{i,w} : \mathcal{E}_i|E_w \rightarrow \{\text{Yes}, \text{No}\}$ such that $m_{i,w}$ has at most n_i switches starting from ‘Yes’ and $\sigma_{m_{i,w}}^{-1}(\text{Yes}) = W \cap E_w$. Therefore, $W \cap E_w$ is the nested difference of n_i+1 descending open sets $O_0^w \supseteq \dots \supseteq O_{n_i}^w$ in $\mathcal{T}_i|E_w$ with $O_0^w = E_w$ (as per the construction in Theorem 3.0). Thus, for each $w' \in W$, we have $W \cap E_{w'} = O_0^{w'} \cap W = O_0^{w'} \cap \bigcup_{w \in W} (W \cap E_w)$. Applying Lemma 3.3 then, we can conclude $W = \bigcup_{w \in W} W \cap E_w$ is (n_i+1) -open in $\mathcal{T}_i$.

Now we show the ‘only if’ direction. Suppose W is (n_i+1) -open and $w \in W$. By Lemma 3.4, we can define $k_w \leq n_i$ to be the least k for which $\exists E \in \mathcal{E}_i(w)$ so that $W \cap E$ is k -closed in $\mathcal{T}_i|E$ and let $E_w \in \mathcal{E}_i(w)$ be an information state for which $W \cap E_w$ is k_w -closed in $\mathcal{T}_i|E_w$.

If $k_w = 0$ then $W \cap E_w$ is 0-closed in $\mathcal{T}_i|E_w$ but not (-1) -open and so $\text{Yes}_i(W|E_w)$ holds.

If $k_w = 1$ then $W \cap E_w$ is 1-closed in $\mathcal{T}_i|E_w$ but not 0-open (as $w \in W \cap E_w$) and so $\text{Yes}_i(W|E_w)$ holds.

If $k_w \geq 2$ then $W \cap E_w$ is k_w -closed in $\mathcal{T}_i|E_w$. Suppose for contradiction that $W \cap E_w$ is (k_w-1) -open in $\mathcal{T}_i|E_w$. Then there is a decision method $m_i : \mathcal{E}_i|E_w \rightarrow \{\text{Yes}, \text{No}\}$ such that m_i has at most k_w-2 switches after saying ‘Yes’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W \cap E_w$. Because $w \in W \cap E_w$, there $\exists E' \in (\mathcal{E}_i|E_w)(w)$ for which $m_i(E') = \text{Yes}$. Consider the decision method $m'_i : \mathcal{E}_i|E' \rightarrow \{\text{Yes}, \text{No}\}$ which $\forall E \in \mathcal{E}_i|E'$ sets $m'_i(E) = m_i(E)$. Clearly, m'_i has at most k_w-2 switches starting from ‘Yes.’ Consider arbitrary $w' \in E'$. If $\sigma_{m_i}(w') = \text{Yes}$ then $\exists E'' \in (\mathcal{E}_i|E_w)(w')$ so that $\forall E \in (\mathcal{E}_i|E'')(w')$ we have $m_i(E) = \text{Yes}$. In particular, if we take $E''' \in (\mathcal{E}_i|E')(w')$ to be a subset of $E' \cap E''$ containing w' , then this means $\forall E \in (\mathcal{E}_i|E''')(w')$ we have $m'_i(E) = m_i(E) = \text{Yes}$. Hence, $\sigma_{m'_i}(w') = \text{Yes}$. Symmetrically, if $\sigma_{m_i}(w') = \text{No}$ then $\sigma_{m'_i}(w') = \text{No}$. Thus, it follows $\forall w' \in E'$ that $\sigma_{m'_i}(w') = \sigma_{m_i}(w')$. Namely, $\sigma_{m'_i}^{-1}(\text{Yes}) = W \cap E'$ and so $W \cap E'$ is the nested difference of k_w-1 descending open sets $O_0 \supseteq \dots \supseteq O_{k_w-2}$ in $\mathcal{T}_i|E'$ with $O_0 = E'$ (as per the construction in Theorem 3.0), i.e. $W \cap E'$ is (k_w-2) -closed in $\mathcal{T}_i|E'$. However, this contradicts minimality of k_w . Therefore, $W \cap E_w$ is k_w -closed in $\mathcal{T}_i|E_w$ but not (k_w-1) -open and so $\text{Yes}_i(W|E_w)$ holds. $\square$

Corollary 3.2: W is (n_i+1) -closed in $\mathcal{T}_i$ iff $\forall w \in \Omega \setminus W$ there $\exists E \in \mathcal{E}_i(w)$ such that $\text{No}_i(W|E)$ holds.

Proof

Immediate by Theorem 3.2 and the fact $\text{No}_i(W|E)$ holds iff $\text{Yes}_i(\Omega \setminus W|E)$ holds. $\square$

In the next section, we will semantically interpret each of the pre-theoretic phrases we introduced earlier. For now, this concludes our overview of topological learning theory. We refer the curious reader to Kelly 1996 for a more comprehensive treatment of the subject as well as its relationship to computability.

4 Semantics

Define a frame $\mathcal{F}$ to be a pair $(\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ where $\Omega \neq \emptyset$ and for each agent $i \in N$ we have that $(\mathcal{E}_i, n_i)$ consists of i 's information basis $\mathcal{E}_i$ over Ω together with i 's switching tolerance n_i .

Next, given a frame $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$, for each $i \in N$ define the operator $\mathbf{R}_i : 2^\Omega \rightarrow 2^\Omega$ so that $\forall W \subseteq \Omega$:

$$\mathbf{R}_i(W) = \{w \in \Omega : \exists E \in \mathcal{E}_i(w), \text{Yes}_i(W|E)\}$$

, i.e. $\mathbf{R}_i(W)$ corresponds to those worlds where i has inductive reason to believe W .⁹

Additionally, mirroring our discussion in the pre-theory section, given $W \subseteq \Omega$ define the operators $\mathbf{I}_{i@W} : 2^\Omega \rightarrow 2^\Omega$ and $\mathbf{B}_{i@W} : 2^\Omega \rightarrow 2^\Omega$ by $\forall P \subseteq \Omega$ setting:

$$\mathbf{I}_{i@W}(P) = \{w \in \Omega : \forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P\} \text{ and } \mathbf{B}_{i@W}(P) = \mathbf{R}_i(W) \cap \mathbf{I}_{i@W}(P)$$

, i.e. $\mathbf{I}_{i@W}(P)$ and $\mathbf{B}_{i@W}(P)$ correspond to the worlds where W indicates to i that P and i has W as deductive reason to believe P respectively.

Now, define $\mathbf{S}_i : 2^\Omega \rightarrow 2^\Omega$ so that $\forall P \subseteq \Omega$:

$$\mathbf{S}_i(P) = \{w \in \Omega : \exists W \subseteq \Omega, w \in W \cap \mathbf{B}_{i@W}(P)\}$$

, i.e. $\mathbf{S}_i(P)$ corresponds to those worlds where i has some true reason to believe P /can inductively know P .

Further, $\forall k \in \mathbb{N}^+$, let $\mathbf{E}_W^k(P)$ denote the result of applying the map $X \mapsto \bigcap_{i \in N} \mathbf{B}_{i@W}(X)$ iteratively k times on P . Then, given $W \subseteq \Omega$, define $\mathbf{G}_W : 2^\Omega \rightarrow 2^\Omega$ by $\forall P \subseteq \Omega$ setting:

$$\mathbf{G}_W(P) = \bigcap_{k \in \mathbb{N}^+} \mathbf{E}_W^k(P)$$

, i.e. $\mathbf{G}_W(P)$ corresponds to those worlds where W generates common inductive knowledge of P .

Finally, recall from the pre-theory section ‘ P can become common inductive knowledge’ is supposed to be the weakest proposition entailing P such that if it holds then everyone can inductively know it. However, this notion is well-defined relative to our frame $\mathcal{F}$ if and only if we can guarantee the existence of a greatest postfix point for the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$. We prove this is the case below.

Lemma 4.0: $\mathbf{S}_i$ is monotone.

Proof

Suppose $P_1 \subseteq P_2$ and $w \in \mathbf{S}_i(P_1)$. Then for some $W \subseteq \Omega$ there exists $E^* \in \mathcal{E}_i(w)$ such that $\text{Yes}_i(W|E^*)$. Further, $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P_1$. In particular, $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P_2$ as $P_1 \subseteq P_2$. Hence, $w \in \mathbf{S}_i(P_2)$. $\square$

Theorem 4.0: The map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$ has a greatest (post)fixed point.

Proof

Since each $\mathbf{S}_i$ is monotone, the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$ is monotone and so Tarski's theorem guarantees the existence of a greatest fixed point which is equal to its greatest postfix point. $\square$

⁹Note we can use $\mathbf{R}_i$ to now state Theorem 3.2 more concisely. Namely, we have that W is $(n_i + 1)$ -open in $\mathcal{F}_i$ iff it is a postfix point of the map $X \mapsto \mathbf{R}_i(X)$.

But what can be said of this fixed point beyond its existence? In order to provide a more ‘constructive’ characterization of what it means for P to potentially become common inductive knowledge, we now study how each operator $\mathbf{S}_i$ relates to the topology $\mathcal{T}_i$ generated by $\mathcal{E}_i$.

Lemma 4.1: Let int_i be the interior operator with respect to the topology $\mathcal{T}_i$. If $n_i = 0$ then $\forall P \subseteq \Omega$ we have that $\mathbf{S}_i(P) = \text{int}_i(P)$.

Proof

Firstly, note $\text{Yes}_i(W|E)$ holds for some piece of evidence $E \in \mathcal{E}_i$ iff $W \cap E$ is 0-closed in $\mathcal{T}_i|E$, i.e. $E \subseteq W$. Now suppose $w \in \text{int}_i(P)$. Then $\exists W \in \mathcal{E}_i(w)$ such that $W \subseteq P$. Thus we immediately have that $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P$. Further, since $w \in W$ and $\text{Yes}_i(W|W)$ holds, we can conclude $\exists E \in \mathcal{E}_i(w)$ such that $\text{Yes}_i(W|E)$ holds and therefore $w \in \mathbf{S}_i(P)$. Next suppose $w \in \mathbf{S}_i(P)$. Then $\exists W \subseteq \Omega$ containing w for which $\exists E^* \in \mathcal{E}_i(w), E^* \subseteq W$ and $\forall E \in \mathcal{E}_i(w), E \subseteq W \rightarrow W \cap E \subseteq P$. In particular, this means $W \cap E^* \subseteq P$ and so $E^* \subseteq P$, i.e. $w \in \text{int}_i(P)$. $\square$

This leads us to ask whether, in general, $\mathbf{S}_i$ acts like an ‘interior’ operator with respect to some topology over Ω . Towards that end, the Kuratowski interior axioms say a map $\mathbf{i} : 2^\Omega \rightarrow 2^\Omega$ can be viewed as the interior operator with respect to some topology $\mathcal{T}_i$ over Ω iff:

1. $\mathbf{i}$ preserves the total space, i.e. we have $\mathbf{i}(\Omega) = \Omega$
2. $\mathbf{i}$ is intensive, i.e. $\forall P \subseteq \Omega$ we have $\mathbf{i}(P) \subseteq P$
3. $\mathbf{i}$ is idempotent, i.e. $\forall P \subseteq \Omega$ we have $\mathbf{i}(P) = \mathbf{i}(\mathbf{i}(P))$
4. $\mathbf{i}$ preserves binary intersections, i.e. $\forall P_1, P_2 \subseteq \Omega$ we have $\mathbf{i}(P_1) \cap \mathbf{i}(P_2) = \mathbf{i}(P_1 \cap P_2)$

Indeed, it turns out that $\mathbf{S}_i$ satisfies all four of Kuratowski’s axioms.

Theorem 4.1: $\mathbf{S}_i$ preserves the total space.

Proof

It suffices to show $\Omega \subseteq \mathbf{S}_i(\Omega)$. Suppose $w \in \Omega$. Since $\mathcal{E}_i(w)$ is non-empty, $\exists E^* \in \mathcal{E}_i(w)$. In particular, $\Omega \cap E^* = E^*$ is 0-closed in $\mathcal{T}_i|E^*$ and so $\text{Yes}_i(\Omega|E^*)$ holds. Further $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(\Omega|E) \rightarrow \Omega \cap E \subseteq \Omega$. Since $\Omega \subseteq \Omega$, we can conclude $w \in \mathbf{S}_i(\Omega)$. $\square$

Theorem 4.2: $\mathbf{S}_i$ is intensive.¹⁰

Proof

Suppose $w \in \mathbf{S}_i(P)$. Then for some $W \subseteq \Omega$ containing w , we have that $\exists E^* \in \mathcal{E}_i(w)$ such that $\text{Yes}_i(W|E^*)$ holds. Further, $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P$. Thus $W \cap E^* \subseteq P$. Since $w \in E^*$, this means we can conclude $w \in P$. $\square$

Theorem 4.3: $\mathbf{S}_i$ is idempotent.

Proof

Because $\mathbf{S}_i$ is intensive we know $\mathbf{S}_i(\mathbf{S}_i(P)) \subseteq \mathbf{S}_i(P)$ and so it suffices to show $\mathbf{S}_i(P) \subseteq \mathbf{S}_i(\mathbf{S}_i(P))$. Suppose $w \in \mathbf{S}_i(P)$. Then there exists $W \subseteq \Omega$ containing w such that $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P$. Take arbitrary $E' \in \mathcal{E}_i(w)$ for which $\text{Yes}_i(W|E')$ holds. Then, $\exists k \leq n_i$ so that $W \cap E'$ is k_i -closed in $\mathcal{T}_i|E'$ but not $(k_i - 1)$ -open. Hence, $\text{Yes}_i(W \cap E'|E')$ also holds. Take $w' \in W \cap E'$. Because $E' \in \mathcal{E}_i(w')$, we know

¹⁰Putting together Theorems 4.1 and 4.2, we can transfinitely iterate the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$ starting with Ω and obtain its greatest fixed point is the transfinite limit of the statements ‘Everyone has some true reason to believe P ,’ ‘Everyone has some true reason to believe ‘Everyone has some true reason to believe P .’

$\exists E \in \mathcal{E}_i(w')$ for which $\text{Yes}_i(W \cap E' | E)$ holds. Further, $\forall E \in \mathcal{E}_i(w'), \text{Yes}_i(W \cap E' | E) \rightarrow (W \cap E') \cap E \subseteq P$ as $W \cap E' \subseteq P$. Hence, $w' \in \mathbf{S}_i(P)$ and so $W \cap E' \subseteq \mathbf{S}_i(P)$. Since E' was arbitrary, this means that $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W | E) \rightarrow W \cap E \subseteq \mathbf{S}_i(P)$. Also, because $w \in \mathbf{S}_i(P)$, there exists $E^* \in \mathcal{E}_i(w)$ for which $\text{Yes}_i(W | E^*)$ holds. Therefore, we can conclude $w \in \mathbf{S}_i(\mathbf{S}_i(P))$. $\square$

By monotonicity of $\mathbf{S}_i$, we already have $\forall P_1, P_2 \subseteq \Omega$ that $\mathbf{S}_i(P_1 \cap P_2) \subseteq \mathbf{S}_i(P_1)$ and $\mathbf{S}_i(P_1 \cap P_2) \subseteq \mathbf{S}_i(P_2)$, i.e. $\mathbf{S}_i(P_1 \cap P_2) \subseteq \mathbf{S}_i(P_1) \cap \mathbf{S}_i(P_2)$. Hence, for $\mathbf{S}_i$ to preserve intersections, it suffices to show $\forall P_1, P_2 \subseteq \Omega$ that $\mathbf{S}_i(P_1) \cap \mathbf{S}_i(P_2) \subseteq \mathbf{S}_i(P_1 \cap P_2)$. Intuitively then, whenever $\mathbf{S}_i(P_1)$ and $\mathbf{S}_i(P_2)$ hold, we want a principled procedure which picks out some true reason i has to believe P_1 as well as some true reason i has to believe P_2 and combines them to produce some true reason i has for believing $P_1 \cap P_2$. We now prove a few lemmas that will help construct exactly this procedure.

Lemma 4.2: $\mathbf{S}_i(P)$ is the union of all $(n_i + 1)$ -open sets in $\mathcal{T}_i$ which entail P .

Proof

Let $P(n_i)$ be the collection of all $(n_i + 1)$ -open sets in $\mathcal{T}_i$ which entail P . Suppose $w \in \bigcup_{W \in P(n_i)} W$. Then $\exists W \subseteq \Omega$ containing w such that W is $(n_i + 1)$ -open in $\mathcal{T}_i$ and $W \subseteq P$. Further, as $W \subseteq P$, we immediately have $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W | E) \rightarrow W \cap E \subseteq P$. In particular, by Theorem 3.2, $\exists E^* \in \mathcal{E}_i(w)$ for which $\text{Yes}_i(W | E)$ holds. Hence, $w \in \mathbf{S}_i(P)$. Now suppose $w \in \mathbf{S}_i(P)$. Then for some $W' \subseteq \Omega$ containing w there exists $E^* \in \mathcal{E}_i(w)$ such that $\text{Yes}_i(W' | E^*)$ holds and $W' \cap E^* \subseteq P$. In particular, $W' \cap E^*$ is n_i -closed in $\mathcal{T}_i | E^*$ and so it is $(n_i + 1)$ -open in $\mathcal{T}_i$. Finally, because $w \in W' \cap E^*$, we have that $W' \cap E^* \in P(n_i)$ and so $w \in \bigcup_{W \in P(n_i)} W$. Thus, $\mathbf{S}_i(P) = \bigcup_{W \in P(n_i)} W$ is the union of all $(n_i + 1)$ -open sets in $\mathcal{T}_i$ entailing P . $\square$

Lemma 4.3: $\forall n > 0, P$ is a union of $(n + 1)$ -open sets in $\mathcal{T}_i$ iff P is a union of 2-open sets in $\mathcal{T}_i$.

Proof

Since the ‘if’ direction is trivial, we focus on the ‘only if’ direction. Suppose P is a union of $(n + 1)$ -open sets in $\mathcal{T}_i$, i.e. there exists a collection $\mathcal{C}$ of $(n + 1)$ -open sets in $\mathcal{T}_i$ such that $P = \bigcup_{W \in \mathcal{C}} W$. Fix $W \in \mathcal{C}$. Since W is $(n + 1)$ -open in $\mathcal{T}_i$ there exists a descending sequence of open sets $W_0 \supseteq \dots \supseteq W_n$ in $\mathcal{T}_i$ such that $W = W_0 \setminus (W_1 \setminus (\dots W_n \dots))$. For ease of argument, let $W_{n+1} = \emptyset$. Note $w \in W$ iff there is an even $k \neq n + 1$ for which $w \in W_k \setminus W_{k+1}$, i.e. $W = \bigcup_{k \leq n: k \text{ is even}} W_k \setminus W_{k+1}$. Since each $W_k \setminus W_{k+1}$ is 2-open in $\mathcal{T}_i$, it follows that P is a union of 2-open sets in $\mathcal{T}_i$. $\square$

Lemma 4.4: If $n_i > 0$ then $\mathbf{S}_i(P)$ is the union of all 2-open sets in $\mathcal{T}_i$ which entail P .

Proof

Let $P(2)$ be the collection of all 2-open sets in $\mathcal{T}_i$ which entail P . Suppose $w \in \bigcup_{W \in P(2)} W$. Then $\exists W \subseteq \Omega$ containing w such that W is 2-open in $\mathcal{T}_i$ and $W \subseteq P$. Note W is certainly $(n_i + 1)$ -open in $\mathcal{T}_i$ as it is 2-open. Hence, by Lemma 4.2, $w \in \mathbf{S}_i(P)$. Now suppose $w \in \mathbf{S}_i(P)$. Then, by Lemma 4.2, w belongs to a set W' which is $(n_i + 1)$ -open in $\mathcal{T}_i$ and entails P . Additionally, by Lemma 4.3, w belongs to a set W which is 2-open in $\mathcal{T}_i$ and entails W' . Finally, because $W \subseteq P$, we can conclude $w \in \bigcup_{W \in P(2)} W$. Thus, $\mathbf{S}_i(P) = \bigcup_{W \in P(2)} W$ is the union of all 2-open sets in $\mathcal{T}_i$ entailing P . $\square$

The above implies that whenever $n_i > 0$ if $w \in \mathbf{S}_i(P)$ then w is contained in a 2-open set W of $\mathcal{T}_i$ which entails P . Thus, if i has some true reason to believe P then i also has some true 2-open set as reason to believe P . We will think of the principled procedure we are constructing as picking out such 2-open sets as reasons in case $n_i > 0$. In the next lemma, we show that if one ‘combines’ two 2-open reasons by taking their intersection then the resulting reason is also 2-open in $\mathcal{T}_i$.

Lemma 4.5: If W_1 and W_2 are 2-open in $\mathcal{T}_i$ then so is $W_1 \cap W_2$.

Proof

Since W_1 and W_2 are 2-open in $\mathcal{T}_i$, there exist descending sequences of open sets $O_0 \supseteq O_1$ and $O'_0 \supseteq O'_1$ such that $W_1 = O_0 \setminus O_1$ and $W_2 = O'_0 \setminus O'_1$. In particular:

$$W_1 \cap W_2 = (O_0 \cap O'_0) \setminus (O_1 \cup O'_1) = (O_0 \cap O'_0) \setminus ((O_0 \cap O'_0) \cap (O_1 \cup O'_1))$$

Since $(O_0 \cap O'_0) \supseteq ((O_0 \cap O'_0) \cap (O_1 \cup O'_1))$, we can conclude $W_1 \cap W_2$ is 2-open in $\mathcal{T}_i$. $\square$

Conveniently then, using Lemmas 4.4 and 4.5, we can show...

Theorem 4.4: $\mathbf{S}_i$ preserves binary intersections.

Proof

By monotonicity, it suffices to show $\mathbf{S}_i(P_1) \cap \mathbf{S}_i(P_2) \subseteq \mathbf{S}_i(P_1 \cap P_2)$. When $n_i = 0$ we know $\mathbf{S}_i$ is just int_i and so we are done. Hence, suppose $n_i > 0$ and $w \in \mathbf{S}_i(P_1) \cap \mathbf{S}_i(P_2)$. By Lemma 4.4, for some $W_1, W_2 \subseteq \Omega$ containing w which are 2-open in $\mathcal{T}_i$, we have that $W_1 \subseteq P_1$ and $W_2 \subseteq P_2$. By Lemma 4.5, $W_1 \cap W_2$ is 2-open in $\mathcal{T}_i$. Thus w is contained in a 2-open set of $\mathcal{T}_i$ which entails $P_1 \cap P_2$, i.e. $w \in \mathbf{S}_i(P_1 \cap P_2)$. $\square$

As promised then, by Theorems 4.1-4.4, $\mathbf{S}_i$ satisfies the Kuratowski interior axioms. If agent i is a deductive learner with $n_i = 0$ then the topology $\mathcal{T}_{\mathbf{S}_i}$ over Ω generated by $\mathbf{S}_i$ is just $\mathcal{T}_i$. Alternatively, if agent i is a truly inductive learner with $n_i > 0$, Lemma 4.4 and Lemma 4.5 tell us $\mathcal{T}_{\mathbf{S}_i}$ can be generated by taking the 2-open sets of $\mathcal{T}_i$ as basis elements. In particular, so long as $n_i > 0$, the topology $\mathcal{T}_{\mathbf{S}_i}$ is invariant of n_i . Having detailed the properties of $\mathbf{S}_i$, we now return to the question of how we can more constructively characterize the worlds where P can become common inductive knowledge.

Theorem 4.5: Let int_N be the interior operator with respect to the intersection topology $\bigcap_{i \in N} \mathcal{T}_{\mathbf{S}_i}$. The greatest fixed point of the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$ is $\text{int}_N(P)$.

Proof

Suppose X is a fixed point of the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$, i.e. $X = P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$. Clearly X is a subset of P . Additionally for each $i \in N$, we have $X \subseteq \mathbf{S}_i(X)$ and so $X = \mathbf{S}_i(X)$ by intensivity. But this means X is open in each $\mathcal{T}_{\mathbf{S}_i}$, i.e. X is open in $\bigcap_{i \in N} \mathcal{T}_{\mathbf{S}_i}$. Hence, it follows that $X \subseteq \text{int}_N(P)$. Next, since $\text{int}_N(P)$ is a subset of P and open in each $\mathcal{T}_{\mathbf{S}_i}$, we know $\text{int}_N(P) = P \cap \bigcap_{i \in N} \mathbf{S}_i(\text{int}_N(P))$. As a result, $\text{int}_N(P)$ is the greatest fixed point of the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$. $\square$

Therefore, define the operator $\mathbf{C} : 2^\Omega \rightarrow 2^\Omega$ so that $\forall P \subseteq \Omega$:


$$\mathbf{C}(P) = \text{int}_N(P)$$

, i.e. $\mathbf{C}(P)$ corresponds to those worlds where P can become common inductive knowledge. Note if each $n_i > 0$ then $\bigcap_{i \in N} \mathcal{T}_{\mathbf{S}_i}$ is invariant of the n_i and so $\mathbf{C}(P)$ is as well. Thus, if agents are truly inductive learners, then the set of worlds where P can become common inductive knowledge is invariant of agents' switching tolerances. We view this property of our semantics as a feature and not a bug. After all, a good notion of common inductive knowledge should be fairly robust to changes in agents' inductive standards.

On the other hand, recall from the pre-theory section that Lewis says P can become common inductive knowledge if and only if:

some state of affairs holds such that it generates common inductive knowledge of P .

Lewis' account lacks the robustness of our definition as, for any witnessing state of affairs which holds, whether that witness generates common inductive knowledge substantially varies with agents' switching tolerances. To illustrate, suppose $\Omega = \{w_1, w_2, w_3\}$ and take $P = \{w_1, w_3\}$. Further, suppose we have $|N| = 2$ agents named Alice and Bob (who we abbreviate by writing a and b respectively). Finally, suppose Alice and Bob each have information bases over Ω given by the below diagrams:



If $n_a = n_b = 1$ it is easy to see the only witness W that Bob has as deductive reason to believe P in any world is $W = \{w_1, w_3\}$. One can also compute straightforwardly that $\mathbf{B}_{a@W}(P) = \{w_3\}$. However, this means $\mathbf{B}_{b@W}(\mathbf{B}_{a@W}(P)) = \emptyset$ and so $\mathbf{G}_W(P) = \emptyset$. Hence, under Lewis' definition, P cannot become common inductive knowledge in any world. Alternatively if $n_a = 2$, one can compute that $\mathbf{B}_{a@W}(P) = \mathbf{B}_{b@W}(P) = \Omega$ and so $\mathbf{G}_W(P) = \Omega$. Moreover, on Lewis' account, we suddenly have P can become common inductive knowledge in every world P holds as $W \cap \mathbf{G}_W(P) = P$. In fact, one can generalize the above example so Lewis' definition shifts from saying P cannot become common inductive knowledge in any world to saying P can become common inductive knowledge in every world P holds at arbitrarily high switching tolerances for Alice. By contrast, for all values of $n_a, n_b > 0$, our definition of common inductive knowledge robustly maintains that $\mathbf{C}(P) = P$. This can be seen visually using the topological characterization of $\mathbf{C}$ in Theorem 4.5 almost immediately. For a more epistemic argument, note $W_1 = \{w_1\} = \Omega \setminus \{w_2, w_3\}$ and $W_3 = \{w_3\}$ are 2-open sets in $\mathcal{T}_a$ which entail P . Hence, Alice has W_1 as true reason to believe P at w_1 and W_3 as true reason to believe P at w_3 , i.e. $\mathbf{S}_a(P) = P$. Similarly, $W_1 \cup W_3 = \Omega \setminus \{w_2\}$ is a 2-open set in $\mathcal{T}_b$ which entails P . Therefore, Bob has $W_1 \cup W_3$ as true reason to believe P at both w_1 and w_3 , i.e. $\mathbf{S}_b(P) = P$. Thus, P is the weakest proposition entailing P such that if it is true then Alice and Bob can inductively know it, i.e. $\mathbf{C}(P) = P$.

In the sequel, we present the formal syntax for our logical language as well as a proof system with respect to which it is sound. Readers eager to see the relevance of our logic to distributed consensus should feel free to skip over to section 7.

5 Syntax

Let PROP be a countable set of primitive propositions and N be a finite set of agents. Our logical language $\mathcal{L}$ will be defined by the Backus-Naur form:

$$\varphi ::= p \mid \top \mid \perp \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \rightarrow \varphi_2 \mid \varphi_1 \leftrightarrow \varphi_2 \mid \mathbb{R}_i \varphi \mid \mathbb{I}_{i@ \varphi_1} \varphi_2 \mid \mathbb{B}_{i@ \varphi_1} \varphi_2 \mid \mathbb{S}_i \varphi \mid \mathbb{G}_{\varphi_1} \varphi_2 \mid \mathbb{C} \varphi$$

where $p \in \text{PROP}$ and $i \in N$.

The table on the following page summarizes how each of the modal connectives above correspond to each of the pre-theoretic phrases/semantic operators defined in the prior sections.

Notation	Gloss
$\mathbb{R}_i\varphi$	i has inductive reason to believe φ
$\mathbb{I}_{i@ \varphi_1}\varphi_2$	φ_1 indicates to i that φ_2
$\mathbb{B}_{i@ \varphi_1}\varphi_2$	i has φ_1 as deductive reason to believe φ_2
$\mathbb{S}_i\varphi$	i has some true reason to believe φ /can inductively know φ
$\mathbb{G}_{\varphi_1}\varphi_2$	φ_1 generates common inductive knowledge of φ_2
$\mathbb{C}\varphi$	φ can become common inductive knowledge

Table 1: Gloss of epistemic modalities

More formally, a model $\mathcal{M}$ for $\mathcal{L}$ is a pair $(\mathcal{F}, v)$ where $\mathcal{F}$ is a frame $(\Omega, (\mathcal{C}_i, n_i)_{i \in N})$ and $v : \text{PROP} \rightarrow 2^\Omega$ is a valuation function over primitive propositions. Then $\forall w \in \Omega$, given a primitive proposition $p \in \text{PROP}$ we write $(\mathcal{M}, w) \models p$ iff $w \in v(p)$, i.e. $v(p)$ yields an interpretation of those worlds where p holds in $\mathcal{M}$. Next, we extend v so that its domain consists of all formulas φ in $\mathcal{L}$ by interpreting classical connectives in their standard manner and recursively writing:

$$\begin{aligned}
(\mathcal{M}, w) &\models \mathbb{R}_i\varphi \text{ iff } w \in \mathbf{R}_i(v(\varphi)) \\
(\mathcal{M}, w) &\models \mathbb{I}_{i@ \varphi_1}\varphi_2 \text{ iff } w \in \mathbf{I}_{i@ v(\varphi_1)}(v(\varphi_2)) \\
(\mathcal{M}, w) &\models \mathbb{B}_{i@ \varphi_1}\varphi_2 \text{ iff } w \in \mathbf{B}_{i@ v(\varphi_1)}(v(\varphi_2)) \\
(\mathcal{M}, w) &\models \mathbb{S}_i\varphi \text{ iff } w \in \mathbf{S}_i(v(\varphi)) \\
(\mathcal{M}, w) &\models \mathbb{G}_{\varphi_1}\varphi_2 \text{ iff } w \in \mathbf{G}_{v(\varphi_1)}(v(\varphi_2)) \\
(\mathcal{M}, w) &\models \mathbb{C}\varphi \text{ iff } w \in \mathbf{C}(v(\varphi))
\end{aligned}$$

, while $\forall \varphi \in \mathcal{L}$ setting $v(\varphi) = \{w \in \Omega : (\mathcal{M}, w) \models \varphi\}$.

Our proof system takes not only all propositional tautologies to be axioms, but also all instances of the following primitive schemas:

$$\mathbf{Ax}_R : \mathbb{R}_i\varphi \leftrightarrow \mathbb{R}_i\mathbb{R}_i\varphi$$

$$\mathbf{Ax}_{I1} : \mathbb{I}_{i@ \varphi}\varphi$$

$$\mathbf{Ax}_{I2} : \mathbb{I}_{i@ \varphi_1}\mathbb{I}_{i@ \varphi_1}\varphi_2 \rightarrow \mathbb{I}_{i@ \varphi_1}\varphi_2$$

$$\mathbf{Ax}_{I3} : \mathbb{I}_{i@ (\mathbb{R}_i\varphi_1)}\varphi_2 \rightarrow \mathbb{I}_{i@ \varphi_1}\varphi_2$$

$$\mathbf{Ax}_{B1} : \mathbb{B}_{i@ \varphi_1}\varphi_2 \leftrightarrow (\mathbb{R}_i\varphi_1 \wedge \mathbb{I}_{i@ \varphi_1}\varphi_2)$$

$$\mathbf{Ax}_{S1} : (\varphi_1 \wedge \mathbb{B}_{i@ \varphi_1}\varphi_2) \rightarrow \mathbb{S}_i\varphi_2$$

$$\mathbf{Ax}_{S2} : \mathbb{S}_i\varphi \rightarrow \varphi$$

$$\mathbf{Ax}_{G1} : \mathbb{G}_{\varphi_1}\varphi_2 \rightarrow \mathbb{B}_{i@ \varphi_1}\varphi_2$$

$$\mathbf{Ax}_{C1} : \mathbb{C}\varphi \rightarrow \varphi$$

$$\mathbf{Ax}_{I4} : \mathbb{I}_{i@ \varphi_1}\varphi_2 \rightarrow \mathbb{I}_{i@ \varphi_1}\mathbb{R}_i(\varphi_1 \wedge \varphi_2)$$

$$\mathbf{Ax}_{I5} : (\mathbb{I}_{i@ \varphi_1}\varphi_2 \wedge \mathbb{I}_{i@ (\varphi_1 \wedge \varphi_2)}\varphi_3) \rightarrow \mathbb{I}_{i@ \varphi_1}\varphi_3$$

$$\mathbf{Ax}_{I6} : (\mathbb{I}_{i@ \varphi_1}\varphi_2 \wedge \mathbb{I}_{i@ \varphi_1}(\varphi_2 \rightarrow \varphi_3)) \rightarrow \mathbb{I}_{i@ \varphi_1}\varphi_3$$

$$\mathbf{Ax}_{B2} : \mathbb{B}_{i@ \varphi_1}\varphi_2 \rightarrow \mathbb{R}_i(\varphi_1 \wedge \varphi_2)$$

$$\mathbf{Ax}_{S3} : \mathbb{S}_i\varphi \leftrightarrow \mathbb{S}_i\mathbb{S}_i\varphi$$

$$\mathbf{Ax}_{S4} : (\mathbb{S}_i\varphi_1 \wedge \mathbb{S}_i\varphi_2) \leftrightarrow \mathbb{S}_i(\varphi_1 \wedge \varphi_2)$$

$$\mathbf{Ax}_{G2} : \mathbb{G}_{\varphi_1}\varphi_2 \rightarrow \mathbb{B}_{i@ \varphi_1}\mathbb{G}_{\varphi_1}\varphi_2$$

$$\mathbf{Ax}_{C2} : \mathbb{C}\varphi \rightarrow \mathbb{S}_i\mathbb{C}\varphi$$

We also take as primitive rules of inference:

$$\begin{aligned}
\mathbf{R}_R : \frac{\varphi}{\mathbb{R}_i\varphi} \quad \mathbf{R}_I : \frac{\varphi_2}{\mathbb{I}_{i@ \varphi_1}\varphi_2} \\
\mathbf{R}_G : \frac{\varphi_3 \rightarrow \bigwedge_{i \in N} \mathbb{B}_{i@ \varphi_1}\varphi_3 \quad \varphi_3 \rightarrow \bigwedge_{i \in N} \mathbb{B}_{i@ \varphi_1}\varphi_2}{\varphi_3 \rightarrow \mathbb{G}_{\varphi_1}\varphi_2} \quad \mathbf{R}_C : \frac{\varphi_1 \rightarrow \bigwedge_{i \in N} \mathbb{S}_i\varphi_1 \quad \varphi_1 \rightarrow \varphi_2}{\varphi_1 \rightarrow \mathbb{C}\varphi_2}
\end{aligned}$$

In the next section, we show soundness of this proof system with respect to our semantics. We leave it as an open conjecture whether the proof system is also complete.

6 Soundness

We begin by showing soundness of each axiom schema.

AX_R : $\mathbb{R}_i \varphi \leftrightarrow \mathbb{R}_i \mathbb{R}_i \varphi$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W be any subset of Ω . It suffices to show $\mathbf{R}_i(W) = \mathbf{R}_i(\mathbf{R}_i(W))$. Note that $\mathbf{R}_i(W) = \{w \in \Omega : \exists E \in \mathcal{E}_i(w), \text{Yes}_i(W|E)\}$ is the union of all information states E in $\mathcal{E}_i$ for which $\text{Yes}_i(W|E)$ holds. Hence, $\mathbf{R}_i(W)$ is 1-open in $\mathcal{T}_i$ and therefore also $(n_i + 1)$ -open. By Theorem 3.2, this means $\mathbf{R}_i(W) \subseteq \mathbf{R}_i(\mathbf{R}_i(W)) = \{w \in \Omega : \exists E \in \mathcal{E}_i(w), \text{Yes}_i(\mathbf{R}_i(W)|E)\}$. We claim also that $\mathbf{R}_i(\mathbf{R}_i(W)) \subseteq \mathbf{R}_i(W)$. So suppose $w \in \mathbf{R}_i(\mathbf{R}_i(W))$. Then $\exists E^* \in \mathcal{E}_i(w)$ so that $\text{Yes}_i(\mathbf{R}_i(W)|E^*)$ holds, i.e. $\exists k \leq n_i$ so that $\mathbf{R}_i(W) \cap E^*$ is k -closed but not $(k - 1)$ -open in $\mathcal{T}_i|E^*$. Since $\mathbf{R}_i(W)$ is 1-open in $\mathcal{T}_i$, we know $\mathbf{R}_i(W) \cap E^*$ is also 1-open in $\mathcal{T}_i|E^*$. Therefore, it must be that $\mathbf{R}_i(W) \cap E^*$ is 0-closed in $\mathcal{T}_i|E^*$, i.e. $E^* \subseteq \mathbf{R}_i(W)$. Thus, since $w \in E^*$, we conclude $w \in \mathbf{R}_i(W)$. $\square$

AX_{I1} : $\mathbb{I}_{i@} \varphi$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W be any subset of Ω . It suffices to show $\forall w \in \Omega$ that $w \in \mathbf{I}_{i@}(W)$. Take arbitrary $w \in \Omega$. Note $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq W$ as we always know $W \cap E \subseteq W$. By definition then, $w \in \mathbf{I}_W(W)$. $\square$

AX_{I2} : $\mathbb{I}_{i@} \varphi_1 \mathbb{I}_{i@} \varphi_2 \rightarrow \mathbb{I}_{i@} \varphi_1 \varphi_2$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{I}_{i@}(P) \subseteq \mathbf{I}_{i@}(\mathbf{I}_{i@}(P))$. Take arbitrary $w \in \mathbf{I}_{i@}(\mathbf{I}_{i@}(P))$ and arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W|E^*)$ holds then $W \cap E^* \subseteq \mathbf{I}_{i@}(P)$. Then $\forall w' \in W \cap E^*$, since $w' \in \mathbf{I}_{i@}(P)$ and $E^* \in \mathcal{E}_i(w')$, it must be that $w' \in P$ and so $W \cap E^* \subseteq P$. Since E^* was arbitrary, $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P$. By definition then, $w \in \mathbf{I}_{i@}(P)$. $\square$

AX_{I3} : $\mathbb{I}_{i@}(\mathbb{R}_i \varphi_1) \varphi_2 \rightarrow \mathbb{I}_{i@} \varphi_1 \varphi_2$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{I}_{i@}(\mathbf{R}_i(W))(P) \subseteq \mathbf{I}_{i@}(P)$. Take arbitrary $w \in \mathbf{I}_{i@}(\mathbf{R}_i(W))(P)$ and arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W|E^*)$ holds then we know $E^* \subseteq \mathbf{R}_i(W)$, i.e. $\text{Yes}_i(\mathbf{R}_i(W)|E^*)$ holds. Further, this means $\mathbf{R}_i(W) \cap E^* \subseteq P$. As $E^* \subseteq \mathbf{R}_i(W)$, it follows $E^* \subseteq P$ and so $W \cap E^* \subseteq P$. Since E^* was arbitrary, $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P$. By definition then, $w \in \mathbf{I}_{i@}(P)$. $\square$

AX_{I4} : $\mathbb{I}_{i@} \varphi_1 \varphi_2 \rightarrow \mathbb{I}_{i@} \mathbb{R}_i(\varphi_1 \wedge \varphi_2)$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{I}_{i@}(P) \subseteq \mathbf{I}_{i@} \mathbf{R}_i(W \cap P)$. Take arbitrary $w \in \mathbf{I}_{i@}(P)$ and arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W|E^*)$ holds then note $W \cap E^* \subseteq P$ and so $W \cap E^* = (W \cap P) \cap E^*$. Hence if $\text{Yes}_i(W|E^*)$ holds then

$\text{Yes}_i(W \cap P | E^*)$ also holds. Take arbitrary $w' \in W \cap E^*$. Since $E^* \in \mathcal{E}_i(w')$ we know that $w' \in \mathbf{R}_i(W \cap E^*)$. Further, since w' was arbitrary, this means $W \cap E^* \subseteq \mathbf{R}_i(W \cap P)$. Finally, since E^* was arbitrary, we have that $\forall E \in \mathcal{E}_i(w)$, $\text{Yes}_i(W | E) \rightarrow W \cap E \subseteq \mathbf{R}_i(W \cap P)$. By definition then, $w \in \mathbf{I}_{i@W}(\mathbf{R}_i(W \cap P))$. $\square$

$$\mathbf{Ax}_{I5} : (\mathbb{I}_{i@ \varphi_1} \varphi_2 \wedge \mathbb{I}_{i@(\varphi_1 \wedge \varphi_2)} \varphi_3) \rightarrow \mathbb{I}_{i@ \varphi_1} \varphi_3$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P, Q be any subsets of Ω . It suffices to show $\mathbf{I}_{i@W}(P) \cap \mathbf{I}_{i@(W \cap P)}(Q) \subseteq \mathbf{I}_{i@W}(Q)$. Take arbitrary $w \in \mathbf{I}_{i@W}(P) \cap \mathbf{I}_{i@(W \cap P)}(Q)$ and arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W | E^*)$ holds then $W \cap E^* \subseteq P$ and so $W \cap E^* = (W \cap P) \cap E^*$. Hence, if $\text{Yes}_i(W | E^*)$ holds then $\text{Yes}_i(W \cap P | E^*)$ also holds and so $W \cap E^* = (W \cap P) \cap E^* \subseteq Q$. Since E^* was arbitrary, this means $\forall E \in \mathcal{E}_i(w)$, $\text{Yes}_i(W | E) \rightarrow W \cap E \subseteq Q$. By definition then, $w \in \mathbf{I}_{i@W}(Q)$. $\square$

$$\mathbf{Ax}_{I6} : (\mathbb{I}_{i@ \varphi_1} \varphi_2 \wedge \mathbb{I}_{i@ \varphi_1} (\varphi_2 \rightarrow \varphi_3)) \rightarrow \mathbb{I}_{i@ \varphi_1} \varphi_3$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P, Q be any subsets of Ω . It suffices to show $\mathbf{I}_{i@W}(P) \cap \mathbf{I}_{i@W}(P \rightarrow Q) \subseteq \mathbf{I}_{i@W}(Q)$. Take arbitrary $w \in \mathbf{I}_{i@W}(P) \cap \mathbf{I}_{i@W}(P \rightarrow Q)$ and arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W | E^*)$ holds then $W \cap E^* \subseteq P$ and $W \cap E^* \subseteq (\Omega \setminus P) \cup Q$. Hence, $W \cap E^* \subseteq P \cap ((\Omega \setminus P) \cup Q) = Q$. Since E^* was arbitrary, this means $\forall E \in \mathcal{E}_i(w)$, $\text{Yes}_i(W | E) \rightarrow W \cap E \subseteq Q$. By definition then, $w \in \mathbf{I}_{i@W}(Q)$. $\square$

$$\mathbf{Ax}_{B1} : \mathbb{B}_{i@ \varphi_1} \varphi_2 \leftrightarrow (\mathbb{R}_i \varphi_1 \wedge \mathbb{I}_{i@ \varphi_1} \varphi_2)$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{B}_{i@W}(P) = \mathbf{R}_i(W) \cap \mathbf{I}_{i@W}(P)$, which holds by definition. $\square$

$$\mathbf{Ax}_{B2} : \mathbb{B}_{i@ \varphi_1} \varphi_2 \rightarrow \mathbb{R}_i(\varphi_1 \wedge \varphi_2)$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{B}_{i@W}(P) \subseteq \mathbf{R}_i(W \cap P)$. Take arbitrary $w \in \mathbf{B}_{i@W}(P)$. Then $w \in \mathbf{R}_i(W)$ and so $\exists E^* \in \mathcal{E}_i(w)$ for which $\text{Yes}_i(W | E^*)$ holds. Since $w \in \mathbf{I}_{i@W}(P)$, this means $W \cap E^* \subseteq P$ and thus $W \cap E^* = (W \cap P) \cap E^*$. Hence, $\text{Yes}_i(W \cap P | E^*)$ also holds and so $w \in \mathbf{R}_i(W \cap P)$. $\square$

$$\mathbf{Ax}_{S1} : (\varphi_1 \wedge \mathbb{B}_{i@ \varphi_1} \varphi_2) \rightarrow \mathbb{S}_i \varphi_2$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $W \cap \mathbf{B}_{i@W}(P) \subseteq \mathbf{S}_i(P)$, which holds by definition. $\square$

$$\mathbf{Ax}_{S2} : \mathbb{S}_i \varphi \rightarrow \varphi$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and P be any subset of Ω . It suffices to show $\mathbf{S}_i(P) \subseteq P$. We proved exactly this in Theorem 4.2. $\square$

$$\mathbf{Ax}_{S3} : \mathbb{S}_i \varphi \leftrightarrow \mathbb{S}_i \mathbb{S}_i \varphi$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and P be any subset of Ω . It suffices to show $\mathbf{S}_i(P) = \mathbf{S}_i(\mathbf{S}_i(P))$. We proved exactly this in Theorem 4.3. $\square$

$$\mathbf{Ax}_{S4} : (\mathbb{S}_i \varphi_1 \wedge \mathbb{S}_i \varphi_2) \leftrightarrow \mathbb{S}_i(\varphi_1 \wedge \varphi_2)$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and P_1, P_2 be any subsets of Ω . It suffices to show $\mathbb{S}_i(P_1) \cap \mathbb{S}_i(P_2) = \mathbb{S}_i(P_1 \cap P_2)$. We proved exactly this in Theorem 4.4. $\square$

$$\mathbf{Ax}_{G1} : \mathbb{G}_{\varphi_1} \varphi_2 \rightarrow \mathbb{B}_{i@ \varphi_1} \varphi_2$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W}(P)$, which holds by definition. $\square$

$$\mathbf{Ax}_{G2} : \mathbb{G}_{\varphi_1} \varphi_2 \rightarrow \mathbb{B}_{i@ \varphi_1} \mathbb{G}_{\varphi_1} \varphi_2$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and W, P be any subsets of Ω . It suffices to show $\mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W}(\mathbf{G}_W(P))$. Take arbitrary $w \in \mathbf{G}_W(P) = \bigcap_{k \in \mathbb{N}^+} \mathbf{E}_W^k(P)$. Since we know $w \in \mathbf{E}_W^1(P) = \bigcap_{j \in N} \mathbf{B}_{j@W}(P)$ it follows $w \in \mathbf{B}_{i@W}(P) \subseteq \mathbf{R}_i(W)$. Additionally, since $\forall k \in \mathbb{N}^+$ we know $w \in \mathbf{E}_W^{k+1}(P) = \bigcap_{j \in N} \mathbf{B}_{j@W}(\mathbf{E}_W^k(P))$ it follows $w \in \mathbf{B}_{i@W}(\mathbf{E}_W^k(P)) \subseteq \mathbf{I}_{i@W}(\mathbf{E}_W^k(P))$. Take arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W|E^*)$ holds then $\forall k \in \mathbb{N}^+$ we know $W \cap E^* \subseteq \mathbf{E}_W^k(P)$, i.e. we have $W \cap E^* \subseteq \mathbf{G}_W(P)$. Hence, as E^* was arbitrary, $w \in \mathbf{I}_{i@W}(\mathbf{G}_W(P))$ and so $w \in \mathbf{B}_{i@W}(\mathbf{G}_W(P))$. $\square$

$$\mathbf{Ax}_{C1} : \mathbb{C} \varphi \rightarrow \varphi$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and P be any subset of Ω . It suffices to show $\mathbf{C}(P) \subseteq P$. Recall $\mathbf{C}(P)$ is a fixed point of the map $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$. Thus, $\mathbf{C}(P) = P \cap \bigcap_{i \in N} \mathbf{S}_i(\mathbf{C}(P)) \subseteq P$. $\square$

$$\mathbf{Ax}_{C2} : \mathbb{C} \varphi \rightarrow \mathbb{S}_i \mathbb{C} \varphi$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame, i be an agent in N , and P be any subset of Ω . It suffices to show $\mathbf{C}(P) \subseteq \mathbf{S}_i(\mathbf{C}(P))$. Recall $\mathbf{C}(P)$ is a fixed point of the map $X \mapsto P \cap \bigcap_{j \in N} \mathbf{S}_j(X)$. Thus, $\mathbf{C}(P) = P \cap \bigcap_{j \in N} \mathbf{S}_j(\mathbf{C}(P)) \subseteq \mathbf{S}_i(\mathbf{C}(P))$ as desired. $\square$

We now show soundness of each inference rule.

$$\mathbf{R}_R : \frac{\varphi}{\mathbb{R}_i \varphi}$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and i be an agent in N . It suffices to show $\forall w \in \Omega$ that $w \in \mathbf{R}_i(\Omega)$. Take arbitrary $w \in \Omega$. Then there exists some $E^* \in \mathcal{E}_i(w)$. Further, since $\Omega \cap E^* = E^*$ is 0-closed in $\mathcal{F}_i|E^*$, we know $\text{Yes}_i(W|E^*)$ holds. Therefore $w \in \mathbf{R}_i(\Omega)$. $\square$

$$\mathbf{R}_I : \frac{\varphi_2}{\mathbb{I}_{i@ \varphi_1} \varphi_2}$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and i be an agent in N and W be any subset of Ω . It suffices to show $\forall w \in \Omega$ that $w \in \mathbf{I}_{i@W}(\Omega)$. Note $\forall E \in \mathcal{E}_i(w)$, $\text{Yes}_i(W|E) \rightarrow W \cap E \subseteq \Omega$ as we always know $W \cap E \subseteq \Omega$. By definition then, $w \in \mathbf{I}_W(\Omega)$. $\square$

To simplify our argument as to why $\mathbf{R}_G$ is sound, we next prove two key lemmas...

Lemma 6.0: Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and W, P be any subsets of Ω . Then the map $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ is monotone.

Proof

Suppose $X_1 \subseteq X_2$. Take arbitrary $w \in \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X_1))$. Since $\forall i \in N$ we know $w \in \mathbf{B}_{i@W}(P)$ it follows $\forall i \in N$ that $w \in \mathbf{R}_i(W)$. Take arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W|E^*)$ holds then $W \cap E^* \subseteq X_1$ and hence $W \cap E^* \subseteq X_2$. Hence, as E^* was arbitrary, we have $\forall i \in N$ that $w \in \mathbf{I}_{i@W}(X_2)$. Thus, $\forall i \in N$ it follows $w \in \mathbf{B}_{i@W}(X_2)$. In particular, $w \in \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X_2))$. Therefore, the map $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ is monotone. $\square$

Lemma 6.1: Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and W, P be any subsets of Ω . Then the map $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ commutes with intersection.

Proof

Let $(X_j)_{j \in J}$ be an arbitrary collection of subsets of Ω . Note first that $\bigcap_{j \in J} \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X_j)) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \bigcap_{j \in J} \mathbf{B}_{i@W}(X_j)) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \bigcap_{j \in J} \mathbf{I}_{i@W}(X_j))$. Now observe $w \in \bigcap_{j \in J} \mathbf{I}_{i@W}(X_j)$ iff $\forall j \in J$ we have $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq X_j$. Therefore, $w \in \bigcap_{j \in J} \mathbf{I}_{i@W}(X_j)$ just in case it holds that $\forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq \bigcap_{j \in J} X_j$, i.e. iff $w \in \mathbf{I}_{i@W}(\bigcap_{j \in J} X_j)$. Hence, we have $\bigcap_{j \in J} \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X_j)) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{I}_{i@W}(\bigcap_{j \in J} X_j)) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(\bigcap_{j \in J} X_j))$. Thus, the map $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ commutes with intersection. $\square$

$$\mathbf{R}_G : \frac{\varphi_3 \rightarrow \bigwedge_{i \in N} \mathbb{B}_{i@P_1} \varphi_3 \quad \varphi_3 \rightarrow \bigwedge_{i \in N} \mathbb{B}_{i@P_1} \varphi_2}{\varphi_3 \rightarrow \mathbb{G}_{P_1} \varphi_2}$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and W, P, Q be any subsets of Ω . It suffices to show if $Q \subseteq \bigcap_{i \in N} \mathbf{B}_{i@W}(Q)$ and $Q \subseteq \bigcap_{i \in N} \mathbf{B}_{i@W}(P)$ then $Q \subseteq \mathbf{G}_W(P)$. By Lemma 6.0, Tarski's implies the map $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ has a greatest (post)fixed point; by Lemma 6.1, if we let $\mathbf{e}_W^k(P)$ be the result of applying the map $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ iteratively k times on Ω , Kleene's theorem implies this greatest (post)fixed point is given by $\bigcap_{k \in \mathbb{N}^+} \mathbf{e}_W^k(P)$. Observe, by monotonicity, that $\Omega \supseteq \mathbf{e}_W^1(P) \supseteq \mathbf{e}_W^2(P) \dots$ forms a descending sequence of subsets of Ω . We will now show $\forall k \in \mathbb{N}^+$ that $\mathbf{e}_W^k(P) = \mathbf{E}_W^k(P)$ and so $\mathbf{G}_W(P) = \bigcap_{k \in \mathbb{N}^+} \mathbf{e}_W^k(P)$. We proceed by induction on k . When $k = 1$, we know $\mathbf{e}_W^1(P) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(P)) = \mathbf{E}_W^1(P)$ establishing the base case. For the inductive step, assume $\forall k \leq t$ we have $\mathbf{e}_W^k(P) = \mathbf{E}_W^k(P)$. Note $\mathbf{e}_W^{t+1}(P) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(\mathbf{e}_W^t(P))) = \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(\mathbf{E}_W^t(P))) = \mathbf{E}_W^1(P) \cap \mathbf{E}_W^{t+1}(P)$. Take arbitrary $w \in \mathbf{E}_W^{t+1}(P) = \bigcap_{i \in N} \mathbf{B}_{i@W}(\mathbf{E}_W^t(P))$. Then $\forall i \in N$ we have $w \in \mathbf{R}_i(W)$ and $w \in \mathbf{I}_{i@W}(\mathbf{E}_W^t(P))$. Take arbitrary $E^* \in \mathcal{E}_i(w)$. If $\text{Yes}_i(W|E^*)$ holds then $W \cap E^* \subseteq \mathbf{E}_W^t(P)$. By hypothesis, we know $\mathbf{E}_W^t(P) \subseteq \dots \subseteq \mathbf{E}_W^1(P)$. Hence, if $\text{Yes}_i(W|E^*)$ holds then $W \cap E^* \subseteq \mathbf{E}_W^1(P)$. Namely $\forall w' \in W \cap E^*$, since $w' \in \mathbf{E}_W^1(P) \subseteq \mathbf{I}_{i@W}(P)$ and $E^* \in \mathcal{E}_i(w')$, it must be that $w' \in P$ and so $W \cap E^* \subseteq P$. Hence, $w \in \mathbf{I}_{i@W}(P)$ and therefore $w \in \mathbf{E}_W^1(P) = \bigcap_{i \in N} \mathbf{B}_{i@W}(P)$. In particular, as $\mathbf{E}_W^{t+1}(P) \subseteq \mathbf{E}_W^1(P)$, this means $\mathbf{e}_W^{t+1}(P) = \mathbf{E}_W^{t+1}(P)$ thereby completing the inductive step. Since $Q \subseteq \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(Q))$, it is a postfixed point of $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$. As the greatest (post)fixed point of $X \mapsto \bigcap_{i \in N} (\mathbf{B}_{i@W}(P) \cap \mathbf{B}_{i@W}(X))$ is $\mathbf{G}_W(P)$, it follows $Q \subseteq \mathbf{G}_W(P)$. $\square$

$$\mathbf{R}_C : \frac{\varphi_1 \rightarrow \bigwedge_{i \in N} \mathbb{S}_i \varphi_1 \quad \varphi_1 \rightarrow \varphi_2}{\varphi_1 \rightarrow \mathbb{C} \varphi_2}$$

Proof

Let $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$ be an arbitrary frame and W, P be any subsets of Ω . It suffices to show if $W \subseteq \bigcap_{i \in N} \mathbf{S}_i(W)$ and $W \subseteq P$ then $W \subseteq \mathbf{C}(P)$. Since $W \subseteq P \cap \bigcap_{i \in N} \mathbf{S}_i(W)$, by Theorem 4.2, we know $W = P \cap \bigcap_{i \in N} \mathbf{S}_i(W)$, i.e. it is a fixed point of $X \mapsto P \cap \bigcap_{i \in N} \mathbf{S}_i(X)$. By definition then, $W \subseteq \mathbf{C}(P)$. $\square$

7 Inductive Coordinated Attack

Consider the following inductive variant of the coordinated attack problem:¹¹

Each agent $i \in N$ is deciding whether to attest P is true or defer judgement. Let n_i be the maximum number of mind changes agent i is willing to tolerate after first attesting P is true.

- If P is true, each agent i ideally wants everyone to converge on attesting P is true.
- However, if P is false or some agent i does not converge on attesting P is true, then anyone who converges on attesting P is true is perpetually shamed.

What non-trivial attestation protocols are available to the agents, if any, so that they can guarantee none of them will be perpetually shamed?

More formally, suppose we have a frame $\mathcal{F} = (\Omega, (\mathcal{E}_i, n_i)_{i \in N})$. An *attestation strategy* for i is a map $s_i : \mathcal{E}_i \rightarrow \{\text{Yes}, ?\}$ such that the decision method $m_i : \mathcal{E}_i \rightarrow \{\text{Yes}, \text{No}\}$ which $\forall E \in \mathcal{E}_i$ sets $m_i(E) = \text{Yes}$ iff $s_i(E) = \text{Yes}$ has at most n_i switches after saying ‘Yes.’ Additionally, as with decision methods, define $\sigma_{s_i} : \Omega \rightarrow \{\text{Yes}, ?\}$ to be the function which for each $w \in \Omega$ returns $\sigma_{s_i}(w)$ as the output s_i converges to in world w . Finally, an *attestation protocol* for N is a tuple $(s_i)_{i \in N}$ of attestation strategies for each $i \in N$.

We say an attestation protocol $(s_i)_{i \in N}$ *solves coordinated attack for P* iff it satisfies:

1. **Validity-** $\forall i \in N, \sigma_{s_i}^{-1}(\text{Yes}) \subseteq P$.
2. **Agreement-** $\forall i, j \in N, \sigma_{s_i}^{-1}(\text{Yes}) = \sigma_{s_j}^{-1}(\text{Yes})$.
3. **Nontriviality-** $\bigcap_{i \in N} \sigma_{s_i}^{-1}(\text{Yes}) \neq \emptyset$.

Validity ensures that agents will never falsely converge on attesting P is true. Agreement ensures that agents will all converge on the same output in every world. Validity and agreement together ensure that no one is perpetually shamed. A trivial way to satisfy validity and agreement is to take the attestation protocol where each agent always defers judgement on P . To rule out such pathological protocols, the nontriviality condition further ensures there is at least one world where everyone converges on attesting P is true. In particular, given an attestation protocol $(s_i)_{i \in N}$ which solves consensus for P , we will say that $\bigcap_{i \in N} \sigma_{s_i}^{-1}(\text{Yes})$ is the *success set* of $(s_i)_{i \in N}$.

Theorem 7.0: There exists an attestation protocol $(s_i)_{i \in N}$ which solves coordinated attack for P with success set W if and only if W is a non-empty $(n_i + 1)$ -open subset of P in each $\mathcal{T}_i$.

Proof

First we show the ‘if’ direction. So suppose W is a non-empty $(n_i + 1)$ -open subset of P in each $\mathcal{T}_i$. By Theorem 3.0, $\forall i \in N$ there is a decision method m_i has at most n_i switches after saying ‘Yes’ and $\sigma_{m_i}^{-1}(\text{Yes}) = W$. Consider the attestation protocol $(s_i)_{i \in N}$ which $\forall i \in N$ sets $s_i(E) = \text{Yes}$ iff $m_i(E) = \text{Yes}$. Clearly, $\forall i \in N$ we have $\sigma_{s_i}^{-1}(\text{Yes}) = \sigma_{m_i}^{-1}(\text{Yes}) = W \subseteq P$. Hence, because W is non-empty, it immediately follows that $(s_i)_{i \in N}$ solves coordinated attack for P with success set W .

Now we show the ‘only if’ direction. Suppose there exists an attestation protocol $(s_i)_{i \in N}$ which solves coordinated attack for P with success set W . Agreement guarantees $\forall i \in N$ that $\sigma_{s_i}^{-1}(\text{Yes}) = W$. Non-triviality and validity further guarantee W is a non-empty subset of P . Additionally, $\forall i \in N$ the decision method $m_i : \mathcal{E}_i \rightarrow \{\text{Yes}, \text{No}\}$ which sets $m_i(E) = \text{Yes}$ iff $s_i(E) = \text{Yes}$ has at most n_i switches after saying ‘Yes.’ Since $\sigma_{m_i}^{-1}(\text{Yes}) = \sigma_{s_i}^{-1}(\text{Yes}) = W$, this means W is $(n_i + 1)$ -open in each $\mathcal{T}_i$. $\square$

¹¹Halpern & Moses 1990 were the first to analyze coordinated attack using a logic for common knowledge. To our awareness, this variant (despite its clear applications for designing distributed consensus protocols) has not received any such treatment.

Theorem 7.0 implies that the possible success sets of an attestation protocol solving coordinated attack for P are precisely those non-empty subsets of P which are $(n_i + 1)$ -open in each $\mathcal{T}_i$. Additionally, given any such subset W of P , Theorem 3.0 allows us to explicitly construct at least one attestation protocol solving coordinated attack for P with success set W . In effect then, if one regards two attestation protocols solving coordinated attack for P as equally successful whenever they have the same success sets, the problem of solving coordinated attack for P can be reduced to the task of simply finding those non-empty subsets of P which are $(n_i + 1)$ -open in each $\mathcal{T}_i$.

Theorem 7.1: W is a $(n_i + 1)$ -open subset of P in each $\mathcal{T}_i$ if and only if W is a fixed point of the map $X \mapsto X \cap \mathbf{G}_X(P)$. Further, $\forall W \subseteq \Omega$, $W \cap \mathbf{G}_W(P)$ is itself such a fixed point.

Proof

Suppose W is a $(n_i + 1)$ -open subset of P in each $\mathcal{T}_i$. By our proof of $\mathbf{Ax}_{I1}$, we know $\forall i \in N$ that $\mathbf{I}_{i@W}(W) = \Omega$. Moreover, by Theorem 3.2, we have $W \subseteq \mathbf{R}_i(W) = \{w \in \Omega : \exists E \in \mathcal{E}_i(w), \text{Yes}_i(W|E)\}$. Hence, $W \subseteq \mathbf{R}_i(W) \cap \mathbf{I}_{i@W}(W) = \mathbf{B}_{i@W}(W)$ and so $W \subseteq \bigcap_{i \in N} \mathbf{B}_{i@W}(W)$. Next, $\forall i \in N$ we have that $\mathbf{I}_{i@W}(P) = \{w \in \Omega : \forall E \in \mathcal{E}_i(w), \text{Yes}_i(W|E) \rightarrow W \cap E \subseteq P\} = \Omega$ as $W \subseteq P$. In particular, it follows $W \subseteq \mathbf{R}_i(W) \cap \mathbf{I}_{i@W}(P) = \mathbf{B}_{i@W}(P)$ and so $W \subseteq \bigcap_{i \in N} \mathbf{B}_{i@W}(P)$. Therefore, using our proof of $\mathbf{R}_G$, this means $W \subseteq \mathbf{G}_W(P)$, i.e. $W = W \cap \mathbf{G}_W(P)$.

Now suppose $W = W \cap \mathbf{G}_W(P)$. By our proof of $\mathbf{Ax}_{G1}$, we know $\forall i \in N$ that $\mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W}(P)$. Hence, $W \subseteq W \cap \mathbf{B}_{i@W}(P)$. By our proof of $\mathbf{Ax}_{S1}$, we have $W \cap \mathbf{B}_{i@W}(P) \subseteq \mathbf{S}_i(P)$. Further, by Theorem 4.2, $\mathbf{S}_i(P) \subseteq P$. Therefore, $W \subseteq P$. Also by definition, we know that $\mathbf{B}_{i@W}(P) \subseteq \mathbf{R}_i(W)$. Hence, $W \subseteq \mathbf{R}_i(W) = \{w \in \Omega : \exists E \in \mathcal{E}_i(w), \text{Yes}_i(W|E)\}$ and so W is $(n_i + 1)$ -open in $\mathcal{T}_i$ by Theorem 3.2. Thus, W is a $(n_i + 1)$ -open subset of P in each $\mathcal{T}_i$.

Finally consider arbitrary $W \subseteq \Omega$. We claim $W \cap \mathbf{G}_W(P) = (W \cap \mathbf{G}_W(P)) \cap \mathbf{G}_{W \cap \mathbf{G}_W(P)}(P)$. By our proof of $\mathbf{Ax}_{G2}$, we know $\forall i \in N$ that $\mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W}(\mathbf{G}_W(P))$. Also, we know $\mathbf{B}_{i@W}(\mathbf{G}_W(P)) \subseteq \mathbf{R}_i(W \cap \mathbf{G}_W(P))$ by our proof of $\mathbf{Ax}_{B2}$. Thus, $W \cap \mathbf{G}_W(P) \subseteq \mathbf{R}_i(W \cap \mathbf{G}_W(P))$. Next, by our proof of $\mathbf{Ax}_{I1}$, we know $\forall i \in N$ that $\mathbf{I}_{i@W \cap \mathbf{G}_W(P)}(W \cap \mathbf{G}_W(P)) = \Omega$. Since $\mathbf{R}_i(W \cap \mathbf{G}_W(P)) \cap \mathbf{I}_{i@W \cap \mathbf{G}_W(P)}(W \cap \mathbf{G}_W(P)) = \mathbf{B}_{i@W \cap \mathbf{G}_W(P)}(W \cap \mathbf{G}_W(P))$, we have $W \cap \mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W \cap \mathbf{G}_W(P)}(W \cap \mathbf{G}_W(P))$. In particular, this means $W \cap \mathbf{G}_W(P) \subseteq \bigcap_{i \in N} \mathbf{B}_{i@W \cap \mathbf{G}_W(P)}(W \cap \mathbf{G}_W(P))$. By our proof of $\mathbf{Ax}_{B1}$, we know $\mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W}(P)$ and so $W \cap \mathbf{G}_W(P) \subseteq W \cap \mathbf{B}_{i@W}(P)$. By our proof of $\mathbf{Ax}_{S1}$, $W \cap \mathbf{B}_{i@W}(P) \subseteq \mathbf{S}_i(P)$. Further, by Theorem 4.2, we have $\mathbf{S}_i(P) \subseteq P$. Thus, $W \cap \mathbf{G}_W(P) \subseteq P$. Using this fact then, $\forall i \in N$ it follows that we have $\mathbf{I}_{i@W \cap \mathbf{G}_W(P)}(P) = \{w \in \Omega : \forall E \in \mathcal{E}_i(w), \text{Yes}_i(W \cap \mathbf{G}_W(P)|E) \rightarrow (W \cap \mathbf{G}_W(P)) \cap E \subseteq P\} = \Omega$. Since $\mathbf{R}_i(W \cap \mathbf{G}_W(P)) \cap \mathbf{I}_{i@W \cap \mathbf{G}_W(P)}(P) = \mathbf{B}_{i@W \cap \mathbf{G}_W(P)}(P)$, we know $W \cap \mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W \cap \mathbf{G}_W(P)}(P)$. In particular, this means $W \cap \mathbf{G}_W(P) \subseteq \bigcap_{i \in N} \mathbf{B}_{i@W \cap \mathbf{G}_W(P)}(P)$. Therefore, appealing to our proof of $\mathbf{R}_G$, we have $W \cap \mathbf{G}_W(P) \subseteq \mathbf{G}_{W \cap \mathbf{G}_W(P)}(P)$ and so $W \cap \mathbf{G}_W(P) = (W \cap \mathbf{G}_W(P)) \cap \mathbf{G}_{W \cap \mathbf{G}_W(P)}(P)$. $\square$

Hence, we can use the map $X \mapsto X \cap \mathbf{G}_X(P)$ to pick out any subset of P which is $(n_i + 1)$ -open in each $\mathcal{T}_i$ and thereby construct (modulo success sets) all those attestation protocols which solve coordinated attack for P . Additionally, if one defines the operator $\mathbf{L} : 2^\Omega \rightarrow 2^\Omega$ by $\forall P \subseteq \Omega$ setting:

$$\mathbf{L}(P) = \{w \in \Omega : \exists W \subseteq \Omega, w \in W \cap \mathbf{G}_W(P)\}$$

then $\mathbf{L}(P)$ corresponds to those worlds which belong to the success set of some attestation protocol that solves coordinated attack for P . Note these are also exactly those worlds where some state of affairs holds such that it generates common inductive knowledge of P ! Thus, $\mathbf{L}(P)$ further corresponds to those worlds where P can become Lewisian common inductive knowledge. As we argued in the pre-theory section, Lewis' definition of common inductive knowledge entails ours. In what follows, we provide a simplified proof of this fact before using it to analyze inductive coordinated attack.

Lemma 7.0: $\mathbf{L}(P) \subseteq \mathbf{C}(P)$.

Proof

Suppose $w \in \mathbf{L}(P)$. Then $\exists W \subseteq \Omega$ for which $w \in W \cap \mathbf{G}_W(P)$. By Theorem 7.1, $\forall i \in N$ we have $W \cap \mathbf{G}_W(P)$ is $(n_i + 1)$ -open in each $\mathcal{T}_i$. Hence, by Lemma 4.2, $W \cap \mathbf{G}_W(P) \subseteq \mathbf{S}_i(W \cap \mathbf{G}_W(P))$. In particular, $W \cap \mathbf{G}_W(P) \subseteq \bigcap_{i \in N} \mathbf{S}_i(W \cap \mathbf{G}_W(P))$. While proving soundness of $\mathbf{Ax}_{G1}$, we showed that $\forall i \in N$ we have $\mathbf{G}_W(P) \subseteq \mathbf{B}_{i@W}(P)$. Hence, $W \cap \mathbf{G}_W(P) \subseteq W \cap \mathbf{B}_{i@W}(P)$. By our proof of $\mathbf{Ax}_{S1}$, we know $W \cap \mathbf{B}_{i@W}(P) \subseteq \mathbf{S}_i(P)$. Further, by Theorem 4.2, $\mathbf{S}_i(P) \subseteq P$. Thus, $W \cap \mathbf{G}_W(P) \subseteq P$. Appealing to our proof of $\mathbf{R}_C$, this means $W \cap \mathbf{G}_W(P) \subseteq \mathbf{C}(P)$ and therefore $w \in \mathbf{C}(P)$ as desired. $\square$

By Lemma 7.0, since the success set of any attestation protocol which solves coordinated attack for P is a subset of $\mathbf{L}(P)$, we have that $\mathbf{C}(P)$ is also an upper bound for any such success set. This upper bound may be unattainable as illustrated by our example on page 21 where $\mathbf{L}(P) \subset \mathbf{C}(P)$. Nevertheless, if Ω is finite and each n_i is sufficiently high, then whenever $\mathbf{C}(P)$ is non-empty we can show that there is an attestation protocol solving coordinated attack for P whose success set attains this upper bound. Hence, asymptotically speaking, an attestation protocol solving coordinated attack for P is ‘welfare-maximizing’ iff its success set is $\mathbf{C}(P)$ as the worlds outside $\mathbf{C}(P)$ do not belong to the success set of any such attestation protocol. We now prove an important lemma necessary to establish this result...

Lemma 7.1: Suppose Ω is finite. For each i there exists an n so that $\mathbf{C}(P)$ is $(n + 1)$ -open in $\mathcal{T}_i$. In particular, we can define $n_{i,P}^*$ to be the least such n for every agent $i \in N$.

Proof

Since Ω is finite $\mathcal{T}_i$ only has finitely many 2-open sets. By Theorem 4.5 then, $\mathbf{C}(P)$ is a finite union of 2-open sets in $\mathcal{T}_i$. Let $W_0 = O_0 \setminus O'_0, \dots, W_t = O_t \setminus O'_t$ be a finite sequence of 2-open sets where each $O_k, O'_k \in \mathcal{T}_i$ and $O_k \supseteq O'_k$ so that $\mathbf{C}(P) = \bigcup_{k \leq t} W_k$. Consider the decision method $m_i : \mathcal{E}_i \rightarrow \{\text{Yes}, \text{No}\}$ which $\forall E \in \mathcal{E}_i$ sets:

$$m_i(E) = \begin{cases} \text{Yes} & \exists k \leq t \text{ such that } E \subseteq O_k \text{ and } W_k \cap E \neq \emptyset \\ \text{No} & \text{else} \end{cases}$$

We will first show m_i has at most $2t + 1$ switches after saying ‘Yes.’ Towards that end, suppose for sake of contradiction there exists a $(2t + 2)$ -switching sequence $E_0 \supseteq E'_0 \supseteq \dots E_t \supseteq E'_t \supseteq E_{t+1}$ for m_i starting from ‘Yes.’ Then $\forall s \leq t + 1$, because $m_i(E_s) = \text{Yes}$, there $\exists k_s \leq t$ such that $E_s \subseteq O_{k_s}$ and $W_{k_s} \cap E_s \neq \emptyset$. Further, since $m_i(E'_s) = \text{No}$ and $E'_s \subseteq E_s \subseteq O_{k_s}$, we know $W_{k_s} \cap E'_s = \emptyset$. Now note $\forall s \leq t$ if $s' > s$ then $E'_s \supseteq E_{s'}$ and so $W_{k_s} \cap E_{s'} = \emptyset$ as well, i.e. $k_{s'} \neq k_s$. Hence, every element of the sequence $k_0, \dots, k_{t+1}$ is distinct. Contradiction. Thus m_i has at most $2t + 1$ switches after saying ‘Yes.’

Now we show $\sigma_{m_i}^{-1}(\text{Yes}) = \mathbf{C}(P)$. Suppose $w \in \mathbf{C}(P)$. Then $\exists k_w \leq t$ so that $w \in W_{k_w} = O_{k_w} \setminus O'_{k_w}$. Hence $\exists E_w \in \mathcal{E}_i(w)$ such that $E_w \subseteq O_{k_w}$. Furthermore $\forall E' \in \mathcal{E}_i(w)$, we have $W_{k_w} \cap E' \neq \emptyset$ as $w \in E'$ and $w \in W_{k_w}$. Thus, it follows $\forall E' \in \mathcal{E}_i(w)$ if $E' \subseteq E_w$ then $m_i(E') = \text{Yes}$, i.e. $w \in \sigma_{m_i}^{-1}(\text{Yes})$. On the other hand, suppose that $w \notin \mathbf{C}(P)$. Take arbitrary $k \leq t$. We have that either $w \notin O_k$ or $w \in O'_k$. In case $w \notin O_k$ then $\forall E' \in \mathcal{E}_i(w)$ we know $E' \not\subseteq O_k$. Since $\mathcal{E}_i(w)$ is non-empty, we can pick any $E_k^* \in \mathcal{E}_i(w)$ so that $\forall E' \in \mathcal{E}_i(w)$ if $E' \subseteq E_k^*$ then $E' \not\subseteq O_k$. Alternatively, in case $w \in O'_k$, we can pick $E_k^* \in \mathcal{E}_i(w)$ so that $\forall E' \in \mathcal{E}_i(w)$ if $E' \subseteq E_k^*$ then $E' \subseteq O'_k$ and so $W_k \cap E' = (O_k \setminus O'_k) \cap E' = \emptyset$. Let $E^* \in \mathcal{E}_i(w)$ be an information state such that $E^* \subseteq \bigcap_{k \leq t} E_k^*$. By construction, $\forall E' \in \mathcal{E}_i(w)$ if $E' \subseteq E^*$ then $\forall k \leq t$ either $E' \not\subseteq O_k$ or $W_k \cap E' = \emptyset$ as $E' \subseteq E_k^*$, i.e. $m_i(E') = \text{No}$. Hence, $\sigma_{m_i}(w) = \text{No}$ and so $w \notin \sigma_{m_i}^{-1}(\text{Yes})$.

Thus, taking $n = 2t + 1$ and applying Theorem 3.0, we can conclude $\mathbf{C}(P)$ is $(n + 1)$ -open in $\mathcal{T}_i$. $\square$

Theorem 7.2: Suppose Ω is finite and each $n_i \geq n_{i,P}^*$. If $\mathbf{C}(P)$ is non-empty then there is an attestation protocol $(s_i)_{i \in N}$ solving coordinated attack for P whose success set is $\mathbf{C}(P)$. Further, $(s_i)_{i \in N}$ is ‘welfare-maximizing’ in the sense that $\mathbf{L}(P) = \mathbf{C}(P)$.¹²

Proof

By Lemma 7.1, $\mathbf{C}(P)$ is $(n_{i,P}^* + 1)$ -open in each $\mathcal{T}_i$. Hence, $\mathbf{C}(P)$ is $(n_i + 1)$ -open in each $\mathcal{T}_i$. Furthermore, since $\mathbf{C}(P)$ is non-empty by assumption and is always a subset of P , it follows by Theorem 7.0 that there is an attestation protocol $(s_i)_{i \in N}$ solving coordinated attack for P whose success set is $\mathbf{C}(P)$. Additionally, since $\mathbf{C}(P)$ is the success set of an attestation protocol which solves coordinated attack for P , we have that $\mathbf{C}(P) \subseteq \mathbf{L}(P)$. Thus, applying Lemma 7.0, we obtain $\mathbf{L}(P) = \mathbf{C}(P)$. $\square$

8 Future Directions & Conclusion

In this paper, we gave a topological semantics for how common inductive knowledge can be generated by a shared witness. A natural extension would be to study how common inductive knowledge can be generated by a set of separate ‘local’ witnesses for each agent. In the deductive case, Gonczarowski & Moses 2024 have recently introduced a notion of asynchronous common knowledge which does essentially this; we believe our setup will be useful for generalizing their notion to an inductive context. Future work also ought to analyze how our logic changes when agents have unbounded switching tolerances. While we omitted answering this question in the present paper for sake of exposition, we expect our results will by and large carry over when dealing with arbitrary ordinals. Yet another extension of our logic worth pursuing is to model learners who receive noisy data. Genin & Kelly 2017 examine how one might do topological learning theory in statistical settings and thus provide a potential path forward.

While we demonstrated that the proof system presented in this paper is sound with respect to the proposed semantics, a corresponding completeness result has not yet been obtained. We conjecture completeness holds, although it is possible that additional axioms or rules of inference will be necessary to fully capture our semantics. Regardless, any completeness result will have to confront a central technical difficulty stemming from the presence of a non-normal modality in our language—namely, that of ‘having inductive reason to believe.’ This operator fails to satisfy the distribution axiom and places standard completeness techniques for normal modal logics out of reach. A natural first step is therefore to restrict attention to a normal fragment of our language. One might then aim to establish completeness for that fragment by translating our semantics into an appropriate class of Kripke frames.

Lastly, we showed that our semantics lends itself to finding ‘attestation protocols’ which solve an inductive variant of the well-known coordinated attack problem. The solutions to this problem ensure agents are coordinated in whether they converge on correctly attesting a proposition P is true or whether they converge on deferring judgement. We hope future work finds such attestation protocols useful for designing novel and practical distributed consensus algorithms. For instance, imagine an ‘attestation aggregator’ will be listening to each agent’s output but is worried that some agents may be dishonest or faulty. If less than half of agents fall into this category then the aggregator can tell agents to follow an attestation protocol which solves coordinated attack for P and, via majority vote, guarantee that it converges to attesting P is true if and only if all honest agents converge to correctly attesting P is true. To conclude, the logic being developed here provides a fruitful foundation for reasoning about consensus in inductive settings and we are excited to see further progress along the directions outlined above.

¹²This also means Lewis’ definition of common inductive knowledge coincides with ours when each n_i is sufficiently high.

9 Bibliography

References

- [1] P. Cousot & R. Cousot (1979): *Constructive Versions of Tarski's Fixed Point Theorems*. 82, Pacific Journal of Mathematics, doi:<https://doi.org/10.2140/pjm.1979.82.43>.
- [2] R. Cubitt & R. Sugden (2003): *Common Knowledge, Salience, and Convention: A Reconstruction of David Lewis' Game Theory*. 19, Economics and Philosophy, doi:<https://doi.org/10.1017/S0266267103001123>.
- [3] K. Genin & K. Kelly (2017): *The Topology of Statistical Inquiry*. Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017), doi:<https://doi.org/10.4204/EPTCS.251.17>.
- [4] Y. Gonczarowski & Y. Moses (2024): *Common Knowledge, Regained*. Proceedings of the 25th ACM Conference on Economics and Computation (EC 2024), doi:<https://doi.org/10.1145/3670865.3673442>.
- [5] J. Halpern & Y. Moses (1990): *Knowledge and Common Knowledge in a Distributed Environment*. 37, Journal of the ACM, doi:<https://doi.org/10.1145/79147.79161>.
- [6] K. Kelly (1996): *The Logic of Reliable Inquiry*. Oxford University Press.
- [7] K. Kelly (2004): *Justification as Truth-Finding Efficiency: How Ockham's Razor Works*. 14, Minds and Machines, doi:<https://doi.org/10.1023/B:MIND.0000045993.31233.63>.
- [8] D. Lewis (1969): *Convention: A Philosophical Study*. Harvard University Press.
- [9] H. Putnam (1965): *Trial and Error Predicates and the Solution to a Problem of Mostowski*. 30, Journal of Symbolic Logic, doi:<https://doi.org/10.2307/2270581>.
- [10] A. Tarski (1955): *A Lattice-Theoretical Fixpoint Theorem and Its Applications*. 5, Pacific Journal of Mathematics, doi:<https://doi.org/10.2140/pjm.1955.5.285>.