Chapter 1

Information Partitions and Knowledge

Section 1.1 introduces the partitional model of information and three definitions of common knowledge. Sections 1.2 presents Aumann (1976)'s result that agents cannot agree to disagree. Section 1.3 presents Rubinstein (1989)'s email game. Section 1.4 defines common p-belief.

We assume a finite state space in Sections 1.1-1.4 to ease exposition, and discuss in Section 1.5 how these results extend more generally.

1.1 Common Knowledge

An unknown state ω takes values in the finite set Ω. Agents $i \in \mathcal{I}$ share a *common prior* that the state ω is distributed according $P \in \Delta(\Omega)$. Each agent i's *information partition* Π_i is a partition of Ω, with the property that for any realization of the state ω, agent i is informed that the state belongs to $\Pi_i(\omega)$.

Assumption 1. *Every partition element has strictly positive probability under the prior; that is, $P(\Pi_i(\omega)) > 0$ for every agent $i \in \mathcal{I}$ and state $\omega \in \Omega$.*

DEFINITION 1.1 (Knowledge). *The set of states at which agent i* knows *the event $A \subseteq \Omega$ to be true is*

$$K_i(A) = \{\omega : \Pi_i(\omega) \subseteq A\}.$$

No agent can think that an event is true if it is not; that is, $K_i(A) \subseteq A$ for every agent i and event A.

DEFINITION 1.2 (Mutual Knowledge). *The set of states at which the event $A \subseteq \Omega$ is* mutual knowledge *is*

$$K(A) = \bigcap_{i \in \mathcal{I}} \{\omega : \Pi_i(\omega) \subseteq A\}$$

i.e., all agents know A to be true.

EXAMPLE 1.1. Suppose the set of states is $\Omega = \{1, 2, 3, 4, 5, 6\}$, and there are two agents with information partitions $\Pi_1 = \{\{1, 2, 3\}, \{4, 5\}, \{6\}\}$ and $\Pi_2 = \{\{1, 2\}, \{3, 4\}, \{5\}, \{6\}\}$. Let $A = \{3, 4, 5, 6\}$. Then, the set of states at which agent 1 knows A to be true is $K_1(A) = \{4, 5, 6\}$, the set of states at which agent 2 knows A to be true is $K_2(A) = \{3, 4, 5, 6\}$, and the set of states at which both agents know A to be true is $K(A) = \{4, 5, 6\}$.

The knowledge operators K_i and K can be applied to events that themselves represent knowledge or mutual knowledge of a state, thus building up higher-order knowledge (agent 1 knows that agent 2 knows that. . .).

EXERCISE 1.1 (U). *Suppose there are two agents indexed to $i = 1, 2$.*

  (a) *Prove that $K_1(K_2(A)), K_2(K_1(A)) \subseteq K(A)$ for every event $A \subseteq \Omega$.*

  (b) *Provide an example in which $K(A) \nsubseteq K_1(K_2(A))$, demonstrating that even if both players know an event to be true, either can fail to know that the other knows it.*

EXERCISE 1.2 (U). *Prove that $\neg K_i(\neg K_i(A)) = K_i(A)$ for every event $A \subseteq \Omega$ (where $\neg A$ denotes the complement of A.)*

An implicit assumption is made that all agents know the state space Ω and the information partitions $(\Pi_i)_{i \in \mathcal{I}}$. This assumption is less strong than it might initially seem, since we can always redefine states and expand the state space to accommodate uncertainty about other players' partitions, as in the following example.

EXAMPLE 1.2. Let $\Omega = \{1, 2\}$, $\mathcal{I} = \{1, 2\}$, and $\Pi_1 = \Pi_2 = \{\{1\}, \{2\}\}$. Suppose we want to model the situation where agent 1 has uncertainty over whether agent 2's information is the complete partition Π_2 or the trivial partition $\Pi_2' = \{\{1, 2\}\}$. One way to do this is to expand the state space: Define $\widetilde{\Omega} = \Omega \times \{c, t\} = \{\{1, c\}, \{1, t\}, \{2, c\}, \{2, t\}\}$ and revise the agents' information partitions to be

$$\widetilde{\Pi}_1 = \{\{(1, c), (1, t)\}, \{(2, c), (2, t)\}\}$$
$$\widetilde{\Pi}_2 = \{\{(1, c)\}, \{(2, c)\}, \{(1, t), (2, t)\}\}$$

Then, for example, at state $(1, c)$ both agents know $\omega = 1$ to be true, but agent 1 does not know whether agent 2 knows it.

The event A is *common knowledge* at state ω if both agents know it to be true, know the other to know it to be true, ad infinitum. There are at least three equivalent ways to define this.

The First Definition.   The most direct approach is to recursively define higher-order levels of knowledge.

Dᴇꜰɪɴɪᴛɪᴏɴ 1.3 (Common Knowledge, Definition 1). *For any event $A \subseteq \Omega$, define $\mathscr{A}^1 := \bigcap_{i \in \mathcal{I}} K_i(A)$ to be the set of states at which every agent knows A, and recursively define*

$$\mathscr{A}^k := \bigcap_{i \in \mathcal{I}} K_i(\mathscr{A}^{k-1})$$

for each $k \geq 2$. (For example, $\mathscr{A}^2$ is the set of states at which every agent knows that every agent knows A.) The set of states at which A is common knowledge *is $\mathscr{A}^\infty := \bigcap_{n \geq 1} \mathscr{A}^n$.*

Exᴇʀᴄɪsᴇ 1.3 (G). *Consider the informational environment of Example 1.1. Find the smallest value of k with the property that $\mathscr{A}^{k'} = \mathscr{A}^{k'+1}$ for all $k' \geq k$.*

The Second Definition. Alternatively, we can define common knowledge using the meet of the players' information partitions. If two partitions Π and Π' satisfy

$$\Pi'(\omega) \subseteq \Pi(\omega) \quad \forall \omega \in \Omega$$

then we say that Π is a *coarsening* of Π' (corresponding to weakly less information at every state), and Π' is a *refinement* of Π (corresponding to weakly more information at every state). If Π' a coarsening of both partitions Π_1 and Π_2, then it is a *common coarsening* of Π_1, Π_2.

Dᴇꜰɪɴɪᴛɪᴏɴ 1.4. *Let $\Pi_1 \wedge \Pi_2$ denote the finest common coarsening of Π_1, Π_2, i.e., the common coarsening of these partitions that is moreover a refinement of every other common coarsening of Π_1 and Π_2.*

Dᴇꜰɪɴɪᴛɪᴏɴ 1.5. *For any sequence of information partitions $(\Pi_1, \ldots, \Pi_{|\mathcal{I}|})$, let $\mathscr{P}_2 = \Pi_1 \wedge \Pi_2$, and for each $k > 2$, recursively define $\mathscr{P}_k = \mathscr{P}_{k-1} \wedge \Pi_k$. The* meet *of $(\Pi_1, \ldots, \Pi_{|\mathcal{I}|})$ is $\bigwedge_{i \in \mathcal{I}} \Pi_i \equiv \mathscr{P}_{|\mathcal{I}|}$.*

Exᴇʀᴄɪsᴇ 1.4 (G). *Prove that for any sequence of information partitions $(\Pi_1, \ldots, \Pi_{|\mathcal{I}|})$, the meet $\mathscr{P}^n$ is invariant to permutations of players indices.*

Exᴀᴍᴘʟᴇ 1.3. Consider Example 1.1. Stack the two information partitions on top of one another, and suppose an ant is placed on one of the states in an agent's partition (see Figure 1.1).
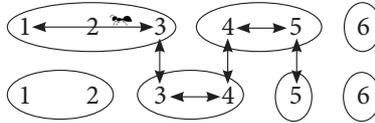

Figure 1.1: $\Pi_1 \wedge \Pi_2(\omega)$ includes all states that an ant seeded at ω can reach.

The ant's movements obey two laws: The ant can move from side to side within an information partition element, and it can jump across the players' information partitions along the same state. The ant's full range of motion when seeded at any state ω then recovers the member of the meet that includes that state. So in this example, we have $\Pi_1 \wedge \Pi_2 = \{\{1, 2, 3, 4, 5\}, \{6\}\}$.

EXERCISE 1.5 (G). *Formalize the statements in the example above by proving that two points x' and x'' belong to the same element of $\bigwedge_{i \in \mathcal{I}} \Pi_i$ if and only if there is a sequence $(x_0, x_1, x_2, \ldots, x_n, x_{n+1})$, with $x_0 = x'$ and $x_{n+1} = x''$, such that for every $0 \le m \le n$, x_m and x_{m+1} belongs to the same element of Π_i for some $i \in \mathcal{I}$.*

DEFINITION 1.6 (Common Knowledge, Definition 2). *An event $A \subseteq \Omega$ is common knowledge at state $\omega \in \Omega$ if $\bigwedge_{i \in \mathcal{I}} \Pi_i(\omega) \subseteq A$.*

REMARK 1.1. It is immediate that the set Ω is common knowledge at every $\omega \in \Omega$.

The Third Definition. Our final definition of common knowledge starts from the definition of an *evident* event which, upon its occurrence, is known to all agents.

DEFINITION 1.7 (Evident Events). *The event $A \subseteq \Omega$ is* evident *(or* public*) if $A \subseteq K(A)$.*

DEFINITION 1.8 (Common Knowledge, Definition 3). *The event $A \subseteq \Omega$ is common knowledge at ω if and only if there is an evident event E such that $\omega \in E$ and $E \subseteq K(A)$.*

EXERCISE 1.6 (G). *Let $\mathcal{I} = \{1, 2\}$. Prove that an event $E \subseteq \Omega$ is evident if and only if it is a union of elements of the meet $\Pi_1 \wedge \Pi_2$.*

These three definitions of common knowledge are equivalent (see for example Monderer and Samet (1989)).

1.2 Agreeing to Disagree

Often we are interested not only in agents' knowledge (which depends only on the agents' information partitions) but also in agents' posterior beliefs (which depend additionally on the prior P). At any state ω and for any event $A \subseteq \Omega$, agent i's posterior probability of event A is pinned down by Bayes' rule (see Section 2.2):

$$P(A \mid \Pi_i(\omega)) = \frac{P(A \cap \Pi_i(\omega))}{P(\Pi_i(\omega))}$$

Our assumption that every partition element has strictly positive prior probability ensures that this expression is well-defined.

One event of interest is the one in which a player's posterior belief takes on a particular value. Fixing an event A and a number $p \in [0, 1]$, define $A_p = \{\omega : P(A \mid \Pi_i(\omega)) = p\}$ to be the set of states at which player i assigns posterior probability p to the event A being true. If player 1 announces that he assigns probability p to A, then all other agents know that the state must belong to A_p.

EXAMPLE 1.4. Consider the informational environment of Example 1.1 with a uniform prior on Ω, and define the event $A = \{2,3\}$. Agent 2 has four partition elements, $\{1,2\}$, $\{3,4\}$, $\{5\}$, $\{6\}$, and assigns to A a posterior probability of $1/2$, $1/2$, 0, and 0 (respectively) on these partition elements. So the set of states $A_{1/2}$ at which agent 2 assigns probability $1/2$ to event A being true, is $A_{1/2} = \{1,2,3,4\}$.

The following theorem shows that whenever players' posterior beliefs about an event are common knowledge (e.g., because players have publicly announced these beliefs), then these posterior beliefs must be identical. So disagreement cannot be sustained whenever players' beliefs are commonly known.

Theorem 1.1 (Aumann (1976)). *Suppose $\mathcal{I} = \{1,2\}$. Fix any state $\omega \in \Omega$ and event $A \subseteq \Omega$. If it is common knowledge at ω that agent 1 assigns (posterior) probability q_1 to event A, while agent 2 assigns (posterior) probability q_2 to the same event, then $q_1 = q_2$.*

The result is stated for two agents, but the proof below directly extends for an arbitrary finite number of players.

Proof. Let $\mathbf{P}$ be the element of $\Pi_1 \wedge \Pi_2$ that contains ω. Then we can write $\mathbf{P} = \cup_k \mathcal{P}^k$ where $\mathcal{P}^k$ are elements of Π_1. Since the event {agent 1's posterior belief is q_1} is common knowledge at ω, agent 1 must assign probability q_1 to event A at every partition element $\mathcal{P}^k$. So $q_1 = P(A \cap \mathcal{P}^k)/P(\mathcal{P}^k)$ for each k. This implies $q_1 \cdot P(\mathcal{P}^k) = P(A \cap \mathcal{P}^k)$. Summing over each of player 1's partition elements, we have $q_1 \sum_k P(\mathcal{P}^k) = \sum_k P(A \cap \mathcal{P}^k)$. Thus $q_1 \cdot P(\mathbf{P}) = P(A \cap \mathbf{P})$. But repeating the same line of logic for player 2, we obtain $q_2 \cdot P(\mathbf{P}) = P(A \cap \mathbf{P})$. So it must be that $q_1 = q_2$. ∎

The following example explains why it is important that players' posterior beliefs are common knowledge and not simply mutual knowledge.

EXAMPLE 1.5. Let $\Omega = \{1,2,3,4\}$ with a uniform prior, and define $\Pi_1 = \{\{1,2\},\{3,4\}\}$ and $\Pi_2 = \{\{1,2,3\},\{4\}\}$. Choose $A = \{1,4\}$ and $\omega = 2$. Then agent 1 assigns posterior probability $1/2$ to A while agent 2 assigns posterior probability $1/3$. Moreover, each agent knows one another's posterior probability. But Theorem 1.1 is not violated: agent 2 does not know that agent 1 knows his posterior probability to be $\frac{1}{3}$, so posterior beliefs are mutual knowledge but not common knowledge.

The starting hypothesis of Theorem 1.1—that individuals have common knowledge of one anothers' beliefs—is strong. Geanakoplos and Polemarchakis (1982) show that the same result obtains under a more realistic process: Communication of posterior beliefs converges to common knowledge of identical posterior beliefs, where this convergence occurs in fewer than $n_1 + n_2$ steps with n_i the size of agent i's partition.

EXAMPLE 1.6. Let $\Omega = \{1,2,3,4,5,6,7,8,9\}$ with all states equally likely. There are two agents, Bob and Carly, with information partitions

$$\Pi_B = \{\{1,2,3\},\{4,5,6\},\{7,8,9\}\}$$

and
$$\Pi_C = \{\{1,2,3,4\},\{5,6,7,8\},\{9\}\}.$$

Suppose the true state is $\omega = 1$, and the agents repeatedly communicate their beliefs about the event $A = \{3,4\}$.

Round 1: Bob's information partition reveals to him that the state belongs to $\{1,2,3\}$, so he assigns posterior probability $1/3$ to the event A. Carly's information partition reveals to her that the state belongs to $\{1,2,3,4\}$, so she assigns posterior probability $1/2$ to the event A. The two agents announce these posterior beliefs.

Round 2: That Bob assigns probability $1/3$ to A reveals to Carly that Bob was either informed that the state belongs to $\{1,2,3\}$ or informed that the state belongs to $\{4,5,6\}$. But Carly already knew in round 1 that these were the two partition elements that Bob might have been informed of (since she knew the state to be either 1, 2, 3, or 4), and so there is no information for her in this announcement. That Carly assigns probability $1/2$ to A reveals to Bob that agent 2 knows $\{1,2,3,4\}$, but again Bob knew this in round 1. So both agents' posterior beliefs are unchanged. They again announce $1/3$ and $1/2$.

And now something interesting happens. That Bob sticks to his original belief of $1/3$ tells Carly that Bob must have observed $\{1,2,3\}$. If instead Bob observed $\{4,5,6\}$, then upon hearing that Carly's belief was $1/2$ (and thus learning that Carly observed $\{1,2,3,4\}$), Bob would have deduced that the state was 4 with certainty, and hence revised his posterior belief of A to 1. So Carly now knows that the state is in $\{1,2,3\}$ and shares Bob's posterior belief, $\frac{1}{3}$. The two agents' beliefs have converged, and it is straightforward to show that these beliefs will not move after subsequent communication.

Although agents' beliefs must converge, the belief that they converge to need not be the belief that agents would have held had they pooled their information:

EXAMPLE 1.7. Let $\Omega = \{1,2,3,4\}$ with each state equally likely. Agents' partitions are given by $\Pi_1 = \{\{1,2\},\{3,4\}\}$ and $\Pi_2 = \{\{1,3\},\{2,4\}\}$. Let $\omega = 1$ and $A = \{1,4\}$. Both posteriors are $1/2$ and the process of belief revision converges in one step. But had agents shared their information, they would have learned that $\omega \in \{1,2\}$ and also $\omega \in \{1,3\}$, leading to a (common) posterior belief that A is true with probability 1.

1.3   The Email Game

Common knowledge assumptions appear frequently in analyses of strategic environments; for example, payoffs are assumed to be common knowledge in any complete-information game. Do strategic predictions made under an assumption of common knowledge approximately hold when we relax the assumption of common knowledge? Rubinstein (1989)'s email game shows that for one formalism of what "almost common knowledge" means, the answer is no: Strategic predictions can change discontinuously when we move from common knowledge to almost common knowledge.

In this game, two agents each choose an action from $\{A, B\}$. There are two possible payoff matrices indexed to $\{a, b\}$ (depicted below with a on the left and b on the right). The agents share a common prior that assigns probability $1 - p > \frac{1}{2}$ to the matrix indexed to a.

	A	B			A	B
A	M, M	$0, -L$		A	$0, 0$	$0, -L$
B	$-L, 0$	$0, 0$		B	$-L, 0$	M, M

We assume throughout that $L > M > 0$. Thus (A, A) yields higher payoffs for both agents when the payoff parameter is a while (B, B) yields higher payoffs when the payoff parameter is b. The action A is "safe," in that it never yields a negative payoff.

Communication Protocol. Both players have an automated email server, which is the only means by which the players can communicate. Agent 1 is informed of the payoff parameter. If (and only if) the parameter is b, agent 1's email server automatically sends an email to agent 2 announcing that the parameter is b. All emails are independently lost with probability $\varepsilon > 0$, so the agents' email servers are set up to automatically send back confirmations that emails have been received, and confirmations of confirmations, etc. Each agent i's type is the number of emails that agent i's computer sends, which is privately known to agent i.

In the special case $T_1 = T_2 = \infty$, there is common knowledge that the parameter is b. But if for example $T_1 = 2$, then agent 1 knows the parameter is b, and agent 1 knows that agent 2 knows that the parameter is b, but agent 1 does not know that agent 2 knows that agent 1 knows that agent 2 knows that the parameter is b. In general, so long as T_1 and T_2 are finite, then higher-order knowledge of parameter b must break down at some stage.

REMARK 1.2. In the partitional framework of Section 1.1, we would model this information environment as follows: The state space is

$$\Omega = \{(a, 0, 0), (b, 1, 0), (b, 1, 1), (b, 2, 1), (b, 2, 2), \dots\}$$

and agents' information partitions are given by

$$\Pi_1 = \{\{(a, 0, 0)\}, \{(b, 1, 0), (b, 1, 1)\}, \{(b, 2, 1), (b, 2, 2)\}, \dots\}$$
$$\Pi_2 = \{\{(a, 0, 0), (b, 1, 0)\}, \{(b, 1, 1), (b, 2, 1)\}, \dots\}$$

where, for example, $T_1 = 0$ reveals to player 1 the partition element $\{(a, 0, 0)\}$, while $T_2 = 0$ reveals to player 2 the partition element $\{(a, 0, 0), (b, 1, 0)\}$.

Proposition 1. *There is a unique Bayesian Nash equilibrium in which agent 1 plays A when the payoff parameter is a. In this equilibrium, both agents play A independently of the number of messages sent.*

Proof. Let $s_i : T_i \to \Delta(\{A, B\})$ denote player i's equilibrium strategy. By assumption, $s_1(0) = A$. We will show that also $s_2(0) = A$. Agent 2 of type $T_2 = 0$ knows that either agent 1's first message was never sent (the state is $(a, 0, 0)$), or agent 1's first message was sent but lost (the state is $(b, 1, 0)$). Unconditionally, the probabilities of these states are $(1 - p)$ and $p\varepsilon$. Conditional on $T_2 = 0$, agent 2 assigns a posterior probability of $\frac{1-p}{1-p+p\varepsilon}$ to $(a, 0, 0)$, a posterior probability of $\frac{p\varepsilon}{1-p+p\varepsilon}$ to $(b, 1, 0)$ and zero probability to all other states.

So agent 2's expected payoff from playing A is at least

$$M \cdot \left(\frac{1-p}{1-p+p\varepsilon} \right) + 0 \cdot \left(\frac{p\varepsilon}{1-p+p\varepsilon} \right) \tag{1.1}$$

while agent 2's expected payoff from playing B is no more than

$$(-L) \cdot \left(\frac{1-p}{1-p+p\varepsilon} \right) + M \cdot \left(\frac{p\varepsilon}{1-p+p\varepsilon} \right). \tag{1.2}$$

Since $1 - p > \frac{1}{2}$ and $L > M$ by assumption, (1.1) strictly exceeds (1.2), and so agent 2's strategy must satisfy $s_2(0) = A$.

Now suppose $s_i(T_i) = A$ for $i = 1, 2$ and all $T_i < t$. We'll argue that $s_1(t) = s_2(t) = A$. Suppose first that agent 1's computer sends t emails exactly, i.e., $T_1 = t$. Since agent 1's computer did not send a $(t + 1)$-th email, it must either be that agent 1's t-th message was lost (the state is $(b, t, t - 1)$), or that agent 1's t-th message was received, but its confirmation was lost (the state is (b, t, t)). Agent 2's posterior belief conditional on $T_1 = t$ then assigns probability $z := \frac{\varepsilon}{\varepsilon + (1-\varepsilon)\varepsilon} > \frac{1}{2}$ to $(b, t, t - 1)$ and probability $1 - z$ to (b, t, t). So the expected payoff to playing B is $z(-L) + (1 - z)(M) < 0$, while the payoff to playing A is zero. We conclude that agent 1's strategy must satisfy $s_1(t) = A$, with nearly identical reasoning yielding $s_2(t) = A$. ∎

This result shows a sharp discontinuity in strategic predictions at common knowledge. That is, (B, B) is an equilibrium when agents have common knowledge of the payoff parameter b, but fails to be an equilibrium when players have knowledge of b to arbitrarily high (finite) orders.

Whether this result is surprising depends on how natural we consider the relaxation of common knowledge to be. Rubinstein (1989) argues that "high T_i" is intuitively like common knowledge. Another view is that these are substantially different, since for arbitrarily small but strictly positive ε the informational model is the one described in Remark 1.2, but for $\varepsilon = 0$ (corresponding to common knowledge of the state) the set of states with positive ex-ante probability is $\Omega = \{(a, 0, 0), (b, \infty, \infty)\}$ and the agents' information partitions are complete. So there is a discontinuity in the informational environments as $\varepsilon \to 0$, and in this sense small ε may be quite unlike $\varepsilon = 0$.

1.4   (Common) p-Belief

We now consider an alternative approach to formalizing almost common knowledge, which defines common "almost-knowledge" in contrast to the above

"almost-common" knowledge.

DEFINITION 1.9. *For any $p \in [0,1]$, say that agent i p-believes A at ω if $P(A \mid \Pi_i(\omega)) \geq p$. The set of states at which agent i p-believes A is*

$$\mathcal{B}_i^p(A) = \{\omega : P(A \mid \Pi_i(\omega)) \geq p\}.$$

REMARK 1.3. Is the case $p = 1$ equivalent to knowledge? Suppose $\Omega = \{1, 2, 3\}$ and the prior is $P = (0, 1/2, 1/2)$. Agent 1's partition is $\{\{1,2\}, \{3\}\}$ while agent 2's partition is $\{\{1\}, \{2\}, \{3\}\}$. The state is $\omega = 2$. Then according to Definition 1.1, agent 2 knows $\{2\}$ but agent 1 does not, while according to Definition 1.9, both agents have 1-belief of $\{2\}$. Whether knowledge and 1-belief represent distinct modes of understanding is an interesting philosophical question, but we will not have more to say on it here.

The following construction of *common p-belief*, due to Monderer and Samet (1989), is parallel to Definition 1.3 for common knowledge.

DEFINITION 1.10 (Common p-Belief). *For any $p \in [0,1]$ and event $A \subseteq \Omega$, define $\mathscr{A}^1 = \bigcap_{i \in \mathcal{I}} \mathcal{B}_i^p(A)$ to be the set of states at which every agent p-believes A to be true, and recursively define $\mathscr{A}^k = \bigcap_{i \in \mathcal{I}} \mathcal{B}_i^p(\mathscr{A}^{k-1})$ for every $k \geq 2$. Then A is common p-belief at the set of states $\mathscr{A}^\infty = \cap_{n \geq 1} \mathscr{A}^n$.*

We can also define common p-belief by generalizing the definition of an evident event (Definition 1.7) to events that are evident p-belief.

DEFINITION 1.11. *For any $p \in [0,1]$, the event $A \subseteq \Omega$ is evident p-belief if $A \subseteq \bigcap_{i \in \mathcal{I}} \mathcal{B}_i^p(A)$.*

DEFINITION 1.12. *For any $p \in [0,1]$, the event $A \subseteq \Omega$ is common p-belief at ω if there exists an evident p-belief event E such that*

$$\omega \in E \subseteq \bigcap_{i \in \mathcal{I}} \mathcal{B}_i^p(A).$$

Definitions 1.10 and 1.12 are introduced in Monderer and Samet (1989) and shown to be equivalent.

EXERCISE 1.7 (G*). *Consider the email game of Rubinstein (1989). Let P denote the common prior on Ω (as defined in Remark 1.2), and define $\mathscr{C}^p$ to be the event that agents have common p-belief in parameter b. For each $\varepsilon \geq 0$, let*

$$\overline{p}(\varepsilon) = \sup_{p \in [0,1]} \{p : P(\mathscr{C}^p) > 0\}$$

be the supremum of the set of values of p such that $\mathscr{C}^p$ has positive ex-ante probability. Is $\overline{p}(0)$ equal to the limit of $\overline{p}(\varepsilon)$ as $\varepsilon \to 0$? Discuss your answer.

1.5 General State Spaces

To show that the preceding insights do not require assumption of a finite state space, we now briefly discuss two generalizations of these ideas. In each case, we begin with a probability space (Ω, Σ, P) where Ω is a set of states endowed with σ-algebra Σ, and $P : \Sigma \to [0,1]$ is a probability measure.

The first generalization. Let each information partition Π_i be a partition of Ω, where we require that each partition element is Σ-measurable and has strictly positive measure under P (see e.g., Monderer and Samet (1989)). All of the above definitions and proofs generalize as stated.

The second generalization. Alternatively, we might model each agent i's information as a sub σ-algebra of Σ, denoted by Π_i.[1] One foundation for this approach (which we will examine in detail in subsequent chapters) is that each agent i privately observes a random variable $X_i : \Omega \rightarrow \mathbb{R}$ that is measurable with respect to Σ. In this case, each agent i's σ-algebra is $\sigma(X_i)$, the σ-algebra generated by X_i, which is indeed coarser than Σ.

  The definition of knowledge can be extended as follows.

DEFINITION 1.13. *Agent i* knows *the event $A \in \Sigma$ to be true at ω if there exists some $B \in \Pi_i$ such that $\omega \in B \subseteq A$.*

  Common knowledge cannot in general be iteratively constructed (à la Definition 1.3) using this definition of K_i, since the set of states at which agent i knows A to be true may not be Σ-measurable. Nevertheless, similar to Definition 1.6, we can define $\bigwedge_{i \in \mathcal{I}} \Pi_i$ to be the finest common coarsening of the σ-algebras $\Pi_1, \ldots, \Pi_n$, and say that an event A is common knowledge at ω if there is an element A of $\bigwedge_{i \in \mathcal{I}} \Pi_i$ such that $\omega \in A$. We can also generalize Definition 1.8 as follows:

DEFINITION 1.14. *The event $A \in \Sigma$ is* evident *if $A \in \Pi_i$ for every $i \in \mathcal{I}$, i.e., A belongs to every agent's σ-algebra.*

DEFINITION 1.15. *The event $A \in \Sigma$ is common knowledge at state ω if there is an evident event E such that $\omega \in E$ and $E \subseteq A$.*

  Theorem 1.1 can also be generalized, although the previous proof does not extend (for example, there is no longer guaranteed to be a unique element of $\Pi_1 \wedge \Pi_2$ that contains ω).

Proposition 2. *Let $X \in \mathcal{L}^1(\Omega, \Sigma, P)$, and define $Y = \mathbb{E}(X \mid \Pi_1)$, $Z = \mathbb{E}(X \mid \Pi_2)$. If it is common knowledge that $Y = y$ and $Z = z$ at a state ω with strictly positive probability, then it must be that $y = z$.*

Proof. If it is common knowledge that $Y = y$ and $Z = z$, there must exist an event $E \in \Pi_1 \cap \Pi_2$, where Y takes the constant value y on E, and Z takes the constant value z on E. Let $\mathbb{1}_E$ denote the indicator variable that takes value 1 on E. Then

$$
\begin{aligned}
y \cdot P(E) &= \mathbb{E}(Y\mathbb{1}_E) \\
&= \mathbb{E}(X\mathbb{1}_E) \\
&= \mathbb{E}(Z\mathbb{1}_E) = z \cdot P(E)
\end{aligned}
$$

---

[1]That is, Π_i is a σ-algebra and $\Pi_i \subseteq \Sigma$.

using in the second and third equalities that Y and Z are conditional expectations of X. Since $P(E) > 0$ (by assumption that $\omega \in E$ has strictly positive probability), it follows that $y = z$ as desired. ∎

This result is in fact more general than Theorem 1.1, nesting the previous result as a special case when we choose X to be an indicator function on some set.

EXERCISE 1.8 (G*). *Generalize Proposition 2 by demonstrating that the conclusion still holds if we assume that there is a measurable set of states $B \subseteq \Omega$ with strictly positive probability, where at every $\omega \in B$ it is common knowledge that $Y = y$ and $Z = z$.*

1.6 Additional Exercises

EXERCISE 1.9 (G). *Two spies in an underground organization are stationed at remote locations. Each spy privately observes whether the coast is clear at their location. The spies share a common prior that the coast is clear at each location independently with probability 1/2.*

Communication protocol. The spies communicate by email with a third party electronic server at their home base. If and only if the coast is clear at a spy's location, that spy's computer will automatically send a message to the home base with the information that the coast is clear.

If the home base electronic server receives information from both spies indicating that the coast is clear, then it will automatically send a message to both spies indicating that it has received both messages. (Otherwise, it will send no messages.) As these are dangerous times, each message has only a $1 - \varepsilon$ chance of being received (again independent). If either spy receives a message from the home base, that spy will send a reply to the home base confirming receipt. The reply is lost with probability ε, independently of everything that's happened before. So on and so forth. Everything stated above is common knowledge.

Each spy observes the number of messages he has sent, and chooses an action in $\{A, B\}$. If the coast is clear at both locations, then payoffs are given by the right matrix below, and otherwise payoffs are given by the left matrix below.

	A	B
A	M, M	$0, -L$
B	$-L, 0$	$0, 0$

	A	B
A	$0, 0$	$0, -L$
B	$-L, 0$	M, M

The payoff parameters satisfy $L > 3M > 0$.

(a) *Prove the following analogue of Rubinstein (1989)'s result: Let $T_1 = \mathbb{Z}_+$ and $T_2 = \mathbb{Z}_+$ denote the two players' type spaces. There is a unique pure-strategy equilibrium in which both players choose A when the coast is **not** clear at their location, i.e. $s_1(0) = s_2(0) = A$. In this equilibrium, players choose A for any number of messages sent, i.e. $s_i(t) = A$ for both players i and all $t \in T_i$.*

(b) *Suppose instead that $L = 2$ while $M = 1$, and demonstrate that the result in Part (a) no longer holds by finding some $\varepsilon > 0$ and a pair of strategies (s_1, s_2) that constitute a pure-strategy Bayesian Nash equilibrium, where $s_1(0) = s_2(0) = A$ and $s_i(t) = B$ for some player i and type $t \in T_i$.*

EXERCISE 1.10 (G*). *Let $X \in \mathcal{L}^1(\Omega, \Sigma, P)$, and define $Y = \mathbb{E}(X \mid \Pi_1)$, $Z = \mathbb{E}(X \mid \Pi_2)$. Prove that if it is common knowledge that $Y \in A$ and $Z \in B$ at a state ω with strictly positive probability, then it must be that $A \cap B \neq \varnothing$.*