Achieving Risky Coordination in the Electronic Mail Game

Anonymous ID: 6

April 24, 2013

1 Introduction

The electronic mail game [8] or "email game" is a coordination game with two possible states, a and b. In each state, players are trying to coordinate their actions so that both play A (in a) or B (in b). The first case happens more than half the time, and the second is not only less frequent but also carries a penalty if only one player chooses B. So in the absence of extra information, it makes sense to pick A.

However, in the email game, player 1 (and only player 1) has access to the ground truth about whether the state is a or b. Whenever the state is b, player 1 automatically sends a message to player 2, allowing them to try to coordinate on B. Messages have a chance of being dropped, so players send confirmation messages (and confirmations of confirmations) until a message is dropped. Surprisingly, this additional information does not help players coordinate on B: it is a strictly dominant strategy to ignore the messages and always play A.

In this paper, we present the proof from [8] which says that coordination on the risky action (B) is impossible. Then we present results from the literature which show how changing assumptions about the email game can get around the impossibility result. Finally, we analyze the email game under an alternative paradigm for reasoning about uncertainty, called possibility measures. This new way of looking at the original email game also sidesteps the impossibility result and allows coordination in the risky case.

2 Background

2.1 The Electronic Mail Game

The electronic mail game was first presented in 1989 by Rubinstein [8], and has since become a classic problem in coordination game theory. The game is similar to the coordinated attack problem [4] in distributed systems.

In the electronic mail game, two players (1 and 2) each have to choose between two actions, A and B. There are two states of the world, a and b, which give different payoff matrices for the actions, as shown in Figure 1. Essentially, both players get payoff M if they coordinate on A when the world state is a, or if they coordinate on B when the world state is b. If the players coordinate on the wrong action, they both get zero payoff. And if the players pick different actions, the player who picks B gets payoff $-L$, where $L > M$, while the A player gets payoff zero.

Furthermore, the state b happens with probability $p < 1/2$ and a happens with probability $1 - p$. The ground truth about the actual state of nature is available only to player 1. Notice that B is the "risky" action because not only does the state b happen less frequently, but there is a penalty for

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

(a) state a, probability $1-p$

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

(b) state b, probability p

Figure 1: Payoff matrices for the electronic mail game. Nature decides whether the state is a or b.

miscoordination, and that penalty only applies to the player who chose B. In order to coordinate on the same action, player 1 sends a message to player 2 whenever the state is b. However, any message can be dropped by the network with probability ϵ (after which there is no further communication between the players). Therefore, player 2 sends a confirmation message back to player 1. This can also be dropped, so player 1 confirms the receipt of the confirmation message; and so on.

The sending of messages happens automatically, outside of the players' control. The state or type of each player is determined by the number of messages he has sent to the other player (we say T_i is how many messages were sent by player i). Players are given access to T_i (and player 1 knows the actual world state), and must choose either A or B.

2.2 Common Knowledge

Common knowledge is discussed at length by Fagin et al in [3]. Common knowledge is defined so that "player i knows x", "player j knows that (i knows x)", "i knows that (j knows that (i knows x))", and all the infinite statements of this form must hold simultaneously. When a message is sent from one player to another, this will usually establish one more "level" of knowledge, but not common knowledge.

Let the set of these statements about x be called ϕ. The key point in the definition of common knowledge is the word "simultaneously". There must exist some point in time at which everyone came to know ϕ. So the knowledge must have been part of the common prior, or else at one instant in time, everyone transitioned to knowing the statements about x. As [3] explains,

> "...the transition from ϕ not being common knowledge to its being common knowledge must involve a *simultaneous* change in all relevant agents' knowledge."

In particular, if player i learns all of ϕ at time t_i, and j learns ϕ at time $t_j \neq t_i$ (without loss of generality, assume $t_j > t_i$), then player i would have to learn that j learns ϕ at time t_j. But i was supposed to learn all of ϕ at time t_i, so ϕ did not actually become common knowledge at t_i (nor, by a similar argument, at t_j). For common knowledge, we need $t_i = t_j$.

Application to the email game: One game closely related to the email game is the coordinated attack problem [4], where two generals are trying to coordinate an attack. One general sends an "attack!" message to the other, which may be dropped and requires a confirmation, and so on. By [6], whenever generals attack it must be common knowledge that they are attacking (assuming they do not attack in the absence of successful communication). In other words, high levels of knowledge are not sufficient to allow an attack in the coordinated attack problem. The generals must have full common knowledge to attack.

In terms of the email game, this means that players must have common knowledge that the state of the world is b before they can collaborate on B (rationalizably). However, it is not possible to achieve common knowledge solely by exchanging messages that have a chance of failure. This explains why in the original email game, collaborating on B does not happen, because it is not possible to achieve common knowledge. The proof of this result is presented next.

3 Analysis of the Electronic Mail Game

Let us define the strategy space of players to be $s_i(t)$ where t is the number of messages they have sent from their machine. We will prove by induction that irrespective of the value of t and the state of nature, it is a strictly dominant strategy to play A for both players. (This proof is based on [8].)

Base Case: If $T_2 = 0$, then player 2 knows this is either because the state of nature is a, or because player 1 sent a message which got dropped. The probability of the former is $1-p$ whereas the probability of the latter is $p \cdot \epsilon$. We also know that if the state of nature is a, then it is a dominant strategy for player 1 to play A, however, we are unsure about the optimal strategy for player 1 if the state of nature is b. From the perspective of player 2, the worst expected payoff if he plays A is

$$u_2(A) = \frac{(1-p)M + p\epsilon \cdot 0}{(1-p) + p\epsilon} \tag{1}$$

whereas the best expected payoff if he plays B is

$$u_2(B) = \frac{(1-p)(-L) + p\epsilon M}{(1-p) + p\epsilon} \tag{2}$$

Now, given that $0 < \epsilon < 1$, $0 < p < 1/2$, and $L > M$, we have from (1) and (2) that

$$\begin{aligned} u_2(B) = \frac{(1-p)(-L) + p\epsilon M}{(1-p) + p\epsilon} &< \frac{(1-p)(-M) + p\epsilon M}{(1-p) + p\epsilon} = \frac{(1-p-p\epsilon)(-M)}{(1-p) + p\epsilon} \\ &< \frac{(1-p-p\epsilon)M}{(1-p) + p\epsilon} = \frac{(1-p)M - p\epsilon M}{(1-p) + p\epsilon} < \frac{(1-p)M}{(1-p) + p\epsilon} = u_2(A) \end{aligned}$$

This means that $u_2(A) > u_2(B)$, and hence, it is a dominant strategy for player 2 to play A. So $s_2(0) = A$, irrespective of the state of nature. Moreover, note that this conclusion does not depend on any private information that player 2 holds except the fact that no message was received by player 2. So, player 1 can also do the same logical inference for player 2, in the case when he does not receive a confirmation from player 2, and decide to play A. In other words, $s_1(0) = A$.

Induction Hypothesis: For all values of $T_i < t$, it is a strictly dominant strategy to play A (for both players).

Induction Step: Let us see the game from the perspective of player 1 if he is of type $T_1 = t$. player 1 is uncertain whether $T_2 = t-1$ (player 1's outgoing message got lost) or $T_2 = t$ (the returning confirmation from player 2 got lost). Given that player 1 did not receive a confirmation for his t'th message, the probability that $T_2 = t-1$ is

$$p(T_2 = t-1) = \frac{\epsilon}{\epsilon + (1-\epsilon)\epsilon} = \frac{1}{2-\epsilon} > \frac{1}{2} \tag{3}$$

whereas the probability that $T_2 = t$ is

$$p(T_2 = t) = \frac{(1-\epsilon)\epsilon}{\epsilon + (1-\epsilon)\epsilon} = \frac{1-\epsilon}{2-\epsilon} < \frac{1}{2} \tag{4}$$

From (3) and (4), player 1 concludes that it is more likely that $T_2 = t - 1$ rather than $T_2 = t$ and hence, by the inductive hypothesis it is a strictly dominant strategy for player 2 to play A. Let the probability from (3) be called q (so that $q = \frac{1}{2-\epsilon}$). Then, for player 1 the expected payoff for playing A is 0 whereas the expected payoff for playing B is:

$$u_1(B) = q(-L) + (1-q)M < 0 \tag{5}$$

Since $L > M$ and $q > \frac{1}{2}$, playing A is a strictly dominant strategy. Thus, $s_1(t) = A$. A symmetric argument shows that $s_2(t) = A$. $\square$

4 Modifications to the Email Game

The above proof gives us a counter-intuitive result: irrespective of the number of messages exchanged between the two players, they will play action A when communication stops. However, it is important to note that this result is obtained under a certain set of assumptions. The most prominent one is that the process of message exchange between the players is involuntary, i.e. a reply is sent as soon as a message is received. Another strong assumption is that there is no cost associated with the exchange of messages and hence, players can afford to perform an unbounded number of message exchanges. In this section we see what happens when we change these assumptions.

Voluntary responses: First, let us restrict the number of messages an agent i will respond to, setting a maximum-response limit of $n_i \in \mathbb{N}$. This modification was considered in [1], where the authors proved that a new set of Nash equilibria arises. If the number of messages successfully exchanged is less than the predetermined message cap n_i, then both players will play A. However, if the number of messages exchanged is greater than or equal to n_i, then B would be played in equilibrium given that:

$$\epsilon \cdot M > (1-\epsilon) \cdot L \tag{6}$$

This is because in the last round of message exchange even if the player is unsure whether his previous message went through or not, equation (6) makes it profitable for him to trust that the message went through and hope that the other player would coordinate with him, and hence, play B in state b. Notice that the original email game discussed is a special case of this game with $n_i = \infty$. Furthermore, for this new set of equilibria, the expected payoff increases as we increase the values of n_i, but at the same time the probability that B will be played in equilibrium decreases as well.

Messages with costs: We obtain similar results by modifying the original email game to introduce a cost $\underline{c}$ to pay attention to the messages being received in the game. In this case, the cost $\underline{c}$ is such that $0 < \underline{c} < \overline{c} = \epsilon(p \cdot M - (1-p) \cdot L)$ sets a cap of n_i on the number of messages to which an agent "should" pay attention, by making the benefit from paying attention to one more confirmation strictly less than the cost of paying attention to it. In [1], the authors prove that under this new setting, the email game has an equilibrium in which A is played by both the players if the number of messages received n is less than n_i, and another equilibrium in which B is played by both the players if $n \geq n_i$.

Correlated failures: Now, let us look at a more realistic distributed system setting with multiple players (or agents). In this case, we have one centralized agent (the general) which observes the state of nature and sends out messages to N other agents in the network (lieutenants), using a broadcast message to coordinate on playing B in state of nature b. After receiving the message, the lieutenants send a confirmation of the receipt of this message and this message exchange continues until they have attained common knowledge, or stops if one or more messages are lost due to noise in the channel. We assume that the drop rate of messages is the same value ϵ in both directions.

The proof given in section 3 works in this case as well, which shows that we would never be able to attain common knowledge and play B in state of nature b. However, in this game, the messages from the general to the lieutenants are sent over a common channel which makes their drop rate correlated by a factor ρ, whereas the lieutenants' confirmations are sent over individual channels and hence, their drop rates are independent. In [2] the authors have shown that for any message drop rate ϵ, there exists a threshold correlation value ρ' such that for all $\rho \geq \rho'$, we have a Nash equilibrium where B is played by both the players when the number of messages exchanged is at least $\overline{T}$ (the cap on the number of messages that can be exchanged). Rather than giving a proof, we would like to present an intuitive argument. Assume the case when $\rho = 1$. In this case, if the general receives only one confirmation from a lieutenant, he can safely assume that his other messages also went through and the respective confirmations got lost, allowing him to collaborate on B in the state of nature b.

5 Analysis Under Possibility Measures

5.1 Definition of Possibility Measures

Possibility measures, described by Halpern [5], are an alternative way of reasoning about uncertainty. They are a replacement for probabilities, based on ideas from fuzzy logic. Like probabilities, possibility measures assign a number in $[0, 1]$ to each subset of the world set W (i.e. each event). Unlike probabilities, possibility measures cannot be added or multiplied in the usual way.

The probability of $U \cup V$ is the sum of the probabilities of each event (if $U \cap V = \emptyset$), because probabilities say that either U or V could happen. But possibility measures optimistically assume that the most likely outcome will happen. So the possibility of a union is whichever possibility is higher for the two arguments: $\mathrm{Poss}(U \cup V)$ is the maximum of $\mathrm{Poss}(U)$ and $\mathrm{Poss}(V)$.

This leads us to the main axioms of possibility theory, which are:

1. $\mathrm{Poss}(\emptyset) = 0$
2. $\mathrm{Poss}(W) = 1$, where W is the universal set
3. $\mathrm{Poss}(U \cup V) = \max(\mathrm{Poss}(U), \mathrm{Poss}(V))$ if U and V are disjoint

Another property of possibility measures from Chapter 3 of [5], which we state without proof, is

$$\mathrm{Poss}(U_1 \cap U_2 \mid U_3) = \min(\mathrm{Poss}(U_1 \mid U_2 \cap U_3), \mathrm{Poss}(U_2 \mid U_3))$$

In particular, if $U_3 = W$ and U_1 and U_2 are independent, then

$$\mathrm{Poss}(U_1 \cap U_2) = \min(\mathrm{Poss}(U_1), \mathrm{Poss}(U_2))$$

So we have seen how to calculate possibilities for unions and intersections of events. The only other tool we need is the ability to calculate expected values under possibility measures. From [7], we find that the expected value of a function $h : X \to [0,1]$ under a possibility measure Poss is

$$\mathop{\mathbb{E}}_{x\sim\mathrm{Poss}(x)}[h(x)] = \sup_{\alpha\in[0,1]} \Big[\alpha \wedge \mathrm{Poss}(A_\alpha)\Big] \qquad \text{where } A_\alpha = \{x \mid h(x) \geq \alpha\} \tag{7}$$

Here the function h can be from any set, but its range must be normalized to $[0,1]$. The operator $\wedge$ is simply the minimum function for our purposes (as in fuzzy logic). And $\sup$ is the supremum, which for finite subsets of $\mathbb{R}$ is just the maximum element of the set.

5.2 Applying Possibility Measures to the Email Game

Consider the original unmodified electronic mail game. Recall that when this game is analyzed with probability theory, there is a dominant strategy (for both players) to always play A, regardless of how many messages were exchanged. We will prove that under possibility measures, playing B is a dominant strategy when at least one message is sent.

Setup: Let the universal set of all events $W = \{a, b\}$ (the two states of nature). We let $\mathrm{Poss}(\{b\}) = p < 1/2$. Notice that $\{a\} \cup \{b\} = W$, so

$$\mathrm{Poss}(\{a\} \cup \{b\}) = \max(\mathrm{Poss}(\{a\}), \mathrm{Poss}(\{b\})) = \mathrm{Poss}(W) = 1$$

But we know that $\mathrm{Poss}(\{b\}) = p$, so it must be that $\mathrm{Poss}(\{a\}) = 1$.

Now, as per the email game, denote the possibility of a message getting dropped as $\mathrm{Poss}(drop) = \epsilon$. Then by a similar argument, the possibility of not getting dropped is $\mathrm{Poss}(\overline{drop}) = 1$.

Base case, player 1: Suppose player 1 has not received any messages. Player 1 knows the state of the world, and if that state is a, it is a strictly dominant strategy for him to play A. (Then, no messages will be sent and as we will see in the next case, player 2 will also play A.) However, if the state is b, player 1 must have sent one message to player 2 (which may or may not have gotten through), and that is covered by the "recursive" case below.

Base case, player 2: Suppose player 2 has not received any messages; we will compare the expected utility (under possibility measures) of playing A versus playing B. There are two reasons player 2 could receive no messages: either the state of the world is a, and player 1 never sent any messages; or the state of the world is in fact b, and player 1's message was dropped. The first case happens with possibility 1 (the possibility of the state being a), and the second case happens with possibility

$$\mathrm{Poss}(b \wedge drop) = \min(\mathrm{Poss}(b), \mathrm{Poss}(drop)) = \min(p, \epsilon)$$

Now we want to calculate the expected utility of player 2's actions (i.e. $u_2(A)$ and $u_2(B)$). To use the notion of expected values under possibility measures, we define a "reward" function $h_0 : Y \to \mathbb{R}$ where $Y = \{A, B\} \times \{a, b\}$. This function is defined (knowing player 1 will play A in a) as

$$\begin{aligned}
&h_0(A,a) = M && \text{(coordinating on } A \text{ with player 1)}\\
&h_0(A,b) = 0 && \text{(playing } A \text{ in state } b \text{ always gives } 0)\\
&h_0(B,a) = -L && \text{(since player 1 will play } A \text{ in state } a)\\
&h_0(B,b) = \begin{cases} -L & \text{if player 1 picks } A \\ M & \text{if player 1 picks } B \end{cases} && \text{(utility depends on player 1's choice)}
\end{aligned}$$

Next, define a normalization mapping $n : \mathbb{R} \to [0, 1]$ as a linear function so that $n(-L) = 0$ and $n(M) = 1$, and denote $n(0) = \delta$. Now we can define $h = n \circ h_0$ (i.e. $h(x) = n(h_0(x))$) as a normalized version of h_0. At this point, h is of the form useful for computing expected values, and we can find $u_2(A)$ and $u_2(B)$.

- First, we find $u_2(A)$, which is a possibility-measures expectation, so we compute the A_α sets from equation (7). The function h only maps its output to three critical points, 0 and δ and 1, so we need only consider these three values of α. And in fact, when restricted to values involving A, h only evalutes to 1 and δ. So the A_α sets will be

  $$A_1 = \{(A, a)\}, \quad A_\delta = A_1 \cup \{(A, b)\}$$

  and thus

  $$\begin{aligned} u_2(A) = \mathop{\mathbb{E}}_{x \sim \mathrm{Poss}(x)}[h(x)] &= \sup\{1 \wedge \mathrm{Poss}(A_1), \delta \wedge \mathrm{Poss}(A_\delta)\} \\ &= \sup\{1 \wedge 1, \delta \wedge \max(1, \min(p, \epsilon))\} = 1 \end{aligned}$$

  Note that this is a normalized utility value; the function n^{-1} could be used to map back to actual utilities ($n^{-1}(1) = M$), but we do not need to do so in this proof.

- Next, we find $u_2(B)$ assuming player 1 picks A in state of nature b. The only A_α set will be $A_0 = \{(B, a), (B, b)\}$, and then

  $$u_2(B|A) = \mathop{\mathbb{E}}_{x \sim \mathrm{Poss}(x)}[h(x)] = \sup\{0 \wedge \mathrm{Poss}(A_0)\} = \sup\{0 \wedge \min(1, \min(p, \epsilon))\} = 0$$

- Finally, we find $u_2(B)$ assuming player 1 picks B in state of nature b. The A_α sets will be

  $$A_1 = \{(B, b)\}, \quad A_0 = A_1 \cup \{(B, a)\}$$

  and thus

  $$\begin{aligned} u_2(B|B) = \mathop{\mathbb{E}}_{x \sim \mathrm{Poss}(x)}[h(x)] &= \sup\{1 \wedge \mathrm{Poss}(A_1), 0 \wedge \mathrm{Poss}(A_0)\} \\ &= \sup\{1 \wedge \min(p, \epsilon), 0 \wedge \max(\min(p, \epsilon), 1)\} \\ &= \sup\{\min(p, \epsilon), 0\} = \min(p, \epsilon) \end{aligned}$$

In other words, since we assume that $p, \epsilon < 1$, it is a strictly dominant strategy for player 2 to play A in the case when no messages are received. This fits with our intuition of the email game.

Recursive case: Suppose some player i has sent $t > 0$ messages and then received no response (i.e. $T_i = t$). At this stage, player i must know that the state of the world is b. Without loss of generality, let us assume that player i is player 1. (This is the trickier case. If player 1 sends one message, he is unsure whether player 2 receives it and knows that the state is b; but if player 2 sends a message, he knows player 1 already knows that the state is b regardless of whether the message gets through.)

Now, since there was no response to the t'th message, there are two possibilities: either 1's last message was dropped and $T_2 = t - 1$, or 1's last message went through but 2's response was dropped, so that $T_2 = t$. Under possibility measures, the possibility that $T_2 = t-1$ is just the possibility that 1's message was dropped: i.e. $\mathrm{Poss}(T_2 = t-1) = \mathrm{Poss}(drop) = \epsilon$. Meanwhile, the possibility that $T_2 = t$ is the possibility that 1's message went through and 2's was dropped, or in other words,

$$\mathrm{Poss}(T_2 = t) = \mathrm{Poss}(\overline{drop_1} \wedge drop_2) = \min(\mathrm{Poss}(\overline{drop}), \mathrm{Poss}(drop)) = \min(1, \epsilon) = \epsilon$$

So player 1 considers it equally possible (with possibility ϵ) that $T_2 = t$ and that $T_2 = t-1$. (This is in contrast to the proof in section 3, where $P(T_2 = t) < 1/2 < P(T_2 = t-1)$.) In other words, from player 1's perspective it is equally possible that 1's last message was dropped, or the message got through and 2's returning message was dropped.

Let us now consider the expected utility of player 1's actions, i.e. compare $u_1(A)$ and $u_1(B)$. Since the state of the world is known to be b, define a function $h_1 : Z \to \mathbb{R}$ which describes the utility for each action (where $Z = \{A, B\} \times \{A, B\}$ represents player 1's actions plus player 2's actions):

$$\begin{aligned}
h_1(A,A) &= 0 && \text{(coordinating on } A\text{, but state is } b\text{)} \\
h_1(A,B) &= 0 && \text{(miscoordinating, but choosing safe option } A\text{)} \\
h_1(B,A) &= -L && \text{(choosing } B \text{ is bad if player 2 goes } A\text{)} \\
h_1(B,B) &= M && \text{(coordinating on } B \text{ in state } b\text{)}
\end{aligned}$$

We define $h = n \circ h_1$ using the same normalization function n from before. Now, outputs of h_1 that were $-L$ become 0, 0 become δ, and M becomes 1. Using this h, we can compute $u_1(A)$ and $u_1(B)$ as follows:

- We find $u_1(A)$ (the utility of player 1 playing A) first. The only set A_α will be $A_\delta = \{(A, A), (A, B)\}$, and thus

$$u_1(A) = \mathop{\mathbb{E}}_{x \sim \mathrm{Poss}(x)}[h(x)] = \sup\left\{\delta \wedge \mathrm{Poss}(A_\delta)\right\} = \sup\{\delta \wedge \max(\epsilon, \epsilon)\} = \min(\delta, \epsilon)$$

- Next we find $u_1(B)$. If player 2 has sent $t' > 0$ messages, i.e. $T_2 > 0$, then this whole "recursive" case analysis applies to player 2 as well. The argument is symmetric (and deterministic), so player 1 and 2 should decide to either both go A or both go B (but we are assuming player 1 is going B here, so player 2 would be going B as well). However, if player 2 has sent 0 messages, then the first base case applies and player 2 will go A.

  Now we have two subcases depending on whether it is possible that $T_2 = 0$ (which can only happen if $T_1 = 1$).

  - If $T_1 = 1$, then it is possible that $T_2 = 0$ and player 2 could go A with possibility $\mathrm{Poss}(T_2 = 0) = \epsilon$. Then

$$A_1 = \{(B, B)\}, \quad A_0 = A_1 \cup \{(B, A)\}$$

    and thus

$$\begin{aligned}
u_1(B|T_1 = 1) = \mathop{\mathbb{E}}_{x \sim \mathrm{Poss}(x)}[h(x)] &= \sup\{1 \wedge \mathrm{Poss}(A_1), 0 \wedge \mathrm{Poss}(A_0)\} \\
&= \sup\{1 \wedge \epsilon, 0 \wedge \epsilon\} = \max(\epsilon, 0) = \epsilon
\end{aligned}$$

– If $T_1 > 1$, then player 1 and player 2 will either both play A or both play B. We are assuming player 1 plays B so player 2 will play B as well, and therefore, the only A_α set will be $A_1 = \{(B, B)\}$, and

$$u_1(B|T_1 > 1) = \mathop{\mathbb{E}}_{x\sim \mathrm{Poss}(x)}[h(x)] = \sup\{1 \wedge \mathrm{Poss}(A_1)\} = \sup\{1 \wedge 1\} = 1$$

Thus, player 1 stands to get utility $\min(\delta, \epsilon)$ from playing A, and either utility ϵ or 1 from playing B (depending on whether only one message, or at least two messages were sent). So in this case it is actually a (weakly) dominant strategy to play B.

Remarks: Player 1 will play A if the state of the world is a, and similarly, player 2 will play A if no messages are received. However, if the state of the world is b, player 1 sends a message and has a (weakly) dominant strategy to play B; meanwhile, player 2 has a strictly dominant strategy to play B if at least one message is received. If player 1 sends at least two messages then playing B becomes a strictly (not weakly) dominant strategy.

The dominance is also strict in all cases whenever $\epsilon < \delta$. We have some restrictions on δ: since $L > M$ and $M > 0$, we have that $1/2 < \delta < 1$. So if the communication channel is very noisy, with a drop rate of more than 50%, the players may be indifferent between A and B in the "recursive" case. But otherwise, with $\epsilon < \delta$, playing B is a strictly dominant strategy whenever player 1 sees the state as b or player 2 receives a message.

In short, this proof shows that coordination on the risky action B is possible in the original email game, when you analyze it under possibility measures. This is a very different result from the original impossibility proof, which showed that there is a strictly dominant strategy to always play A. But that proof relied on examining the game from the context of probability theory, not possibility measures. Possibility measures are optimistic, picking the event with highest possibility and hoping it will happen. So once each player has knowledge that could give him reason to switch to B, possibility theory supposes that he will do so, which gives the other player incentive to coordinate on B as well. □

6 Conclusion

In this paper we presented a detailed study of the email game, starting from the classical proof which shows that achieving common knowledge to coordinate on the risky action is impossible. We then showed different modifications which could be made to the game to overcome this impossibility result (such as placing a bound on the number of message exchanges that can take place). Finally, we analyzed the email game under a different method of reasoning about uncertainty, called possibility measures, where (roughly speaking) the most likely events are assumed to take place. The original impossibility proof relies on probabilities, so this alternative paradigm of reasoning is able to allow coordination on the risky action in an unmodified email game.

References

[1] K. Binmore and L. Samuelson. Coordinated action in the electronic mail game. *Games and Economic Behavior*, 35(1):6–30, 2001.

[2] P. A. Coles and R. Shorrer. Correlation in the Multiplayer Electronic Mail Game. *The BE Journal of Theoretical Economics*, 12(1), 2012.

[3] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. Common knowledge revisited. *Annals of Pure and Applied Logic*, 96(1):89–105, 1999.

[4] J. N. Gray. *Notes on Data Base Operating Systems*. Springer, 1978.

[5] J. Y. Halpern. *Reasoning about uncertainty*. 2003.

[6] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM (JACM)*, 37(3):549–587, 1990.

[7] J. M. Keller and B. Yan. Possibility expectation and its decision making algorithm. In *Fuzzy Systems, 1992., IEEE International Conference on*, pages 661–668. IEEE, 1992.

[8] A. Rubinstein. The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge". *The American Economic Review*, pages 385–391, 1989.