Correlation in the Multiplayer Electronic Mail Game

Peter A. Coles[*]   Ran Shorrer[†]

[*]Harvard Business School, pacoles@gmail.com

[†]Harvard University, rshorrer@hbs.edu

Correlation in the Multiplayer Electronic Mail Game[*]

Peter A. Coles and Ran Shorrer

Abstract

In variants of the Electronic Mail Game (Rubinstein, 1989) where two or more players communicate via multiple channels, the multiple channels can facilitate collective action via redundancy, the sending of the same message along multiple paths or else repeatedly along the same path (Chwe, 1995 and De Jaegher, 2011). This paper offers another explanation for how multiple channels may permit collective action: parties may be able to coordinate their actions when messages' arrivals at their destinations are sufficiently correlated events. Correlation serves to fill in information gaps that arise when players are uncertain of the source of message failure, effectively strengthening messages from one player. This asymmetry in message strength in turn permits cutoff equilibria, where players take action after receiving a minimum number of confirmations.

KEYWORDS: electronic mail game, signaling, stag hunt, coordination, communication

---

1 Introduction

A squadron of planes plans a strike on a powerful enemy, but only a simultaneous attack with the full force of the squadron will ensure success. The captain in command is tasked with identifying the precise moment when conditions merit attacking, and then orders the attack via radio messages to his fellow pilots. Radio messages are likely to reach their destination, though transmission failures are also possible. Hence, when pilots receive a message, they confirm receipt to the captain. If the captain receives some, but not all confirmations, should he nevertheless initiate the attack, hoping it was simply confirmations that were lost and not his original messages? And if the captain requires confirmations from all before attacking, how does an individual pilot know that the captain indeed received the needed confirmations? The captain may choose to reconfirm to the other pilots that all received the attack message, but does he then require confirmations of receipt of this message as well? It seems we are back where we started.

In our pilots story, agents wish to coordinate on mutually beneficial action, but if an individual agent undertakes action without the participation of others, he incurs costs. This, of course, is the description of a *stag hunt*, a classic strategic setting which has been studied as far back as in the writings of the philosophers Rousseau and Hume.[1] In a stag hunt, collective action is the payoff-dominant equilibrium. But in a seminal paper, Rubinstein (1989) demonstrates that when two agents have detailed yet imperfect information about the existence of a risky opportunity, coordinated action may be impossible.

In our paper, we examine how this logic extends to multiplayer settings where one informed agent serves as a "leader," relaying messages to and from the other parties. Our main finding is that in contrast to Rubinstein's result, groups with an informed leader may indeed be able to coordinate on action, and abilty to coordinate depends crucially on the degree of correlation in the messages sent by the leader.

Before detailing our results, we briefly review the celebrated coordinated attack problem from the computer science literature (see e.g. Gray, 1978), and Rubinstein's game-theoretic formalization of this problem, the *elec-*

---

[1]In the classic two player stag hunt, each player has two options: hunt the stag, or gather berries. Berries are filling, but not nearly as delicious as venison. If both players choose to hunt stag, their hunt will be successful. If only one hunts stag, he will fail in his efforts and go hungry, while his counterpart settles for berries. So while the (hunt, hunt) equilbrium payoff dominates the (gather, gather) equilibrium, hunt is a risky action if a player is not confident that his counterpart will also hunt.

tronic mail game (henceforth EMG).

Two generals sit on opposite sides of an enemy. When "the time is ripe," if both generals attack, victory is certain. However, if only one general attacks, his army suffers devastating losses. General 1 observes that the "time is ripe" and sends a messenger to General 2 proposing an attack. However, there is a small probability that the messenger will not make it through the enemy lines. Since General 1 cannot be assured of the messenger's success, General 2 sends a messenger back, confirming receipt. But now General 2 does not know if General 1 received the confirmation, so General 1 again sends a messenger, confirming the receipt of the confirmation. The generals send confirmations back and forth until some messenger fails to make it through. Because only a finite number of messages can be relayed before a messenger fails, the attack proposal will never be common knowledge.[2] For the proposal to be common knowledge, it must be that for all integers $n > 0$,

(everybody knows that)^n an attack has been proposed.

Rubinstein replaces the couriers with email messages and considers a setting where confirmations are sent automatically upon message receipt. Intuition suggests that if sufficiently many messages have been relayed, the generals should be able to coordinate on attacking. Yet Rubinstein shows that regardless of the number of messages received, the uncertainty prevents the generals from coordinating on attacking in any equilibrium. By contrast, with common knowledge that "the time is ripe," they may indeed coordinate their attack.

The EMG demonstrates that "almost common knowledge" of a game can yield vastly different outcomes from common knowledge. But in addition to drawing attention to this distinction, Rubinstein's example raises a related, practical question: When can parties coordinate on risky, yet mutually beneficial action when communication channels are imperfect?

As a practical matter, almost any communication channel is necessarily flawed with some positive probability. Hence, unsurprisingly, the study of communication via flawed networks is a broad research topic with a rich history, so we will work to quickly narrow our scope. One important line of research is "network reliability theory" (see e.g. Colbourn, 1987 or Shier, 1991). In these models, a communication network is represented by a graph, where nodes represent agents, and edges between nodes represent communication channels. Each edge is operational with some probability, and one may calculate the probability of any two nodes being connected via a path of operational

---

[2]See Aumann (1976) for the first formal, set-theoretic treatment of common knowledge.

channels. Using this framework, one may consider, for example, designing a network from flawed channels to optimize the likelihood of connectedness of all network nodes.

Chwe (1995) introduces the notion of "strategic reliability" which formally combines network reliability with game theoretic foundations. In Chwe's formulation, agents are strategic players who take actions based on information obtained via a network. With strategic reliability, the design goal shifts from maximizing connectedness (or other measures of network reliability) to maximizing the likelihood that agents communicating over a flawed network can achieve desired outcomes, via equilibria in the induced game. Rubinstein's EMG falls squarely into this category, and indeed pre-dates and provides inspiration for Chwe's work.

The goal of players in the EMG is to coordinate their actions in a stag hunt. The stag hunt is a particularly appropriate setting to study flawed communication channels. Coordination is desirable, but taking action alone is costly. Hence, reliability of communication is critical in motivating players to play the risky but payoff-dominant equilibrium rather than the safe equilibrium. Several papers examine extensions of the two-player game, focusing primarily on the consequences of making assumptions in Rubinstein's example more realistic.[3]

In recent years, an intriguing line of literature has emerged that investigates collective action in multiplayer settings with flawed communication channels. With multiple players, agents can interact in ways that are not available in the two player setting. A primary focus of this literature is to see if this expanded interaction facilitates collective action or else corroborates Rubinstein's no-coordination result.

Morris (2001) considers a setting with "locally public communication." Nature determines whether the state is good or bad; only in the good state

---

[3]Dulleck (2007), Strzalecki (2011), and Takamiya and Tanaka (2006) consider the EMG when players are boundedly rational. Dulleck (2007) finds that when "absent-minded" players are unsure of the number of messages received, coordination on attacking is possible. Strzalecki (2011) demonstrates that players capable of only a finite number of steps of reasoning may likewise coordinate. Takamiya and Tanaka (2006) find that players' mutual knowledge of the state relative to their mutual knowledge of rationality dictates their ability to coordinate. Another class of papers explores plausibly realistic adjustments to the messaging technology. Binmore and Samuelson (2001) examine a setting when messages are costly and voluntary. Morris (2001) introduces a finite time horizon and messages that take a random length of time to arrive. In Dimitri (2003), a third party learns the state of nature and informs both generals with noise. De Jaegher (2008) examines a model in which message confirmations are strategic, rather than automatic. In each of these settings the adjustment to the messaging technology may permit partial coordination in equilibrium.

is there an opportunity for collective action. In each round, a subset of the players (smaller than the threshold for successful collective action) meets in a private caucus. In each such meeting, players learn the state, and also the history of messages received, and this information becomes common knowledge within the group. But in each round, with positive probability communication ceases, and the group receives no information from the previous caucus. This setting confirms Rubinstein's result: there is no equilibrium where players take collective action when the state is good, despite potentially many rounds of communication.

De Jaegher (2011) considers a multiplayer setting where collective action is facilitated via redundancy, the transmission of the same information along multiple paths. In De Jaegher's formulation, one informed player learns the state of nature, and if the state of nature is good, communicates this news to all other players. Each message successfully reaches its destination with probability $1 - \varepsilon$. Each time a player receives a message, he learns the state of nature, as well as the message's path: an analogy is the path of recipients on a long email thread. After a finite number of rounds, all communication ends and players select their actions. In this model, information sets for players are extremely complex, and the imposition of a finite time horizon helps to ensure collective action.[4] Nevertheless, the model is quite general (and democratic) in the sense that possible channels of communication are comprehensive. De Jaegher finds that in this setting there are multiple equilibria. In some of these equilibria, players take action upon receiving only a small number of confirmations. Notably, in these equilibria outcomes improve when the time horizon is lengthened, in contrast to the finite versions of Rubinstein's example. The logic underpinning these findings is that players benefit from multiple channels, as these offer backup paths to gather information, should other channels fail.

This paper offers a different explanation for how multiple channels may permit collective action: parties may be able to coordinate their actions when failures of messages to arrive at their destinations are sufficiently correlated events.

In our setting, instead of two generals, there is a single general and

---

[4]Finiteness of the horizon can play an important role in coordination. For example, in Rubinstein's EMG, when the number of allowable rounds of confirmations is finite, players can partially coordinate in equilibrium. In an extreme case, if General 1's computer sends a message whenever the state is good, but General 2's computer may not send any confirmations, then there is an equilibrium where General 1 attacks whenever the state is good and General 2 attacks whenever he receives a message. The only miscoordination occurs when the sole message fails to get through.

$N-1$ lieutenants. If "the time is ripe," the general sends messages to each of the lieutenants, who send confirmation messages back to the general (but not to each other). Any message fails to reach its destination with some small but positive probability. If the general hears back from all the lieutenants, he reconfirms to all, while if some lieutenant has failed to confirm, communication is terminated. Lieutenants respond with confirmations, and so forth. Once communication ceases, each party must decide whether or not to attack. For an attack to be successful, collective action is required from all parties.[5] We refer to this environment as the Hub and Spoke Electronic Mail Game (Hub and Spoke EMG).

We first demonstrate that when message failures are independent, the non-action result from the two-player model persists: players are unable to coordinate even when a large number of confirmations are sent. The logic behind this result parallels that of the two-player game: each player anticipates the actions of the others when communication terminates, and each is sufficiently pessimistic about the number of confirmations received by her counterpart that they are unable to coordinate on action.

Suppose, however, that failure rates of messages sent by the general are correlated. Why might we expect this? Consider any environment where a leader is *broadcasting* a message to others, who respond with individual communiqués. In such settings, failure rates may be correlated because of a failure in the broadcast, or else because of the lieutenants' inabilities to interpret the common broadcast. Returning to our opening example, in a squadron of planes, the control tower communicates to planes via its radio, while each plane confirms receipt via its individual radio. Similarly, in a currency attack setting, a ringleader may make a veiled public announcement about vulnerability of a sovereign currency (suggesting an attack may be in order), while other investors confirm via individual channels their intent to similarly pile on. A hopeful leader of a rebellion facing a regime keen to block communication may put out a semi-public call to arms over the internet, and must monitor individual replies via a variety of channels to gauge the depth of support. In short, in any setting where a leader communicates with followers, it is natural to think that failures of outbound messages may be correlated.

When the general sends a message but receives only a subset of confirmations, correlation may provide useful information about which messages

---

[5]In the Appendix, we consider the case where collective action from only a subset of the lieutenants is necessary for victory (see Examples A1 and A2). Compared to the case where action from all agents is required, in these settings coordination is possible under a wider parameter range, since even after message failures agents can often proceed with confidence that enough other players will join in the attack.

were lost. If the degree of correlation across outbound message failures is high, then it is likely that messages were not lost en route to the lieutenants, but instead confirmations were lost upon return. This has the effect of creating a "strong" message from the general: the general only needs some messages to get through to be optimistic that messages made it through to the other lieutenants as well. The lieutenants, by contrast, send "weak" messages, in that when any of the lieutenants' messages fails to make it through, communication stops.

This asymmetry generates decision points that share characteristics with the environments of Dimitri (2004) and Coles (2005).[6] These models are modifications of Rubinstein's two player EMG with one simple adjustment: messages from General 1 to General 2 are assumed to fail with probability differing from that of messages sent in the other direction. When the failure rates are sufficiently different, the no-coordination result of the EMG is overturned. The general who sends the "stronger" message, rather than being pessimistic about his counterpart having received the final communiqué, is in fact optimistic. This asymmetry allows the generals to agree on where communication is likely to have broken down, which in turn allows the generals to coordinate their actions.

In our setting, the correlation and the many-to-one structure of the communication network generate a natural asymmetry that is similar to the pre-specified, technological asymmetry of Dimitri and of Coles. This asymmetry permits coordination via cutoff strategies, strategies where agents attack provided they have received a sufficient number of confirmations. Unlike in the two-player game, or in the Hub and Spoke EMG with independent message failures, with correlated message failures players are able to coordinate their actions when communication ceases, because due to the asymmetry in message strength, they have coincidental beliefs about the source of signal failure.

From an organizational standpoint, our finding will help anticipate the likelihood of coordination in multiplayer communication settings, and to the extent that selection of a signaling mechanism is possible, it suggests that choosing a mechanism that induces correlated failures may be advisible, even if the mechanism involves higher individual message failure rates than mechanisms with independent channels of communication. The results demonstrate that we may take advantage of multiple communication channels in a way that is not possible in a two player setting, and in a way that is distinct from the redundancy benefits focused on previously in the literature.

---

[6]These two papers independently demonstrate similar results.

2   The Hub and Spoke Electronic Mail Game

In this section we model an environment where one informed, central party communicates with peripheral parties over faulty communication channels.

There is a set $\mathcal{N}$ of players with $|\mathcal{N}| = N$, $N > 2$.[7] There are two possible states of the world, a and b. States a and b occur with probability $1 - q$ and q respectively, where $q < 1/2$. Thus, a is the more likely state. For each player, there are two possible actions, A and B. Players who play A receive payoff 0. If the state is b, then if all players coordinate on B, each receives payoff $M > 0$. But if the state is b and some, but not all of the players play B, the players who chose B receive a payoff of $-L$, regardless of the state, with $L > M$. Finally, if the state is a, any player who plays B receives a payoff of $-L$. Thus, it is risky to play B without confidence that the state is b, and that opponents will also all play B. Following the coordinated attack language, b can be interpreted as the state of nature where conditions are ripe for an attack, B is "attack" and A is "do not attack."[8]

The game unfolds as follows. Player 1, the "general," learns the state of nature, either a or b. If the state of nature is b, email messages are automatically sent to each of the $N - 1$ peripheral players, "the lieutenants," indexed by $i = 2 \ldots N$. However, message channels are flawed, and some of these messages will not get through. Let $p : \{0, 1\}^{N-1} \longrightarrow [0, 1]$ be the distribution over message outcomes. That is, $p(\mathbf{1})$ is the probability that all messages get through, $p(\langle 1, 1, \ldots, 1, 0 \rangle)$ is the probability that all messages get through except that to lieutenant N, etc. Note that we place no restrictions on p, so that correlated failure is permitted. Sections 3 and 4 will examine outcomes for specific instances of $p(\cdot)$.

Whenever a lieutenant receives a message, a confirmation message is automatically sent back to the general. Messages from lieutenants are independent and each fails with probability ε. Whenever the general receives confirmations from *all* lieutenants, the general reconfirms to all, and sends out reconfirmations whose success is again described by $p(\cdot)$, independent of outcomes in other rounds. If, however, the general receives only an incomplete subset of the $N-1$ possible confirmations, communication ends: confirmations are sent back to none.

Note that this information structure preserves the feature from the

---

[7]When $N = 2$ the setting simplifies to the usual EMG.

[8]Payoffs in state a depart slightly from Rubinstein's model. Here, playing A always yields payoff 0 (even in state a), and B ensures a payoff of $-L$ in state a (even if all counterparts also play B). These changes, adopted in the literature by Chwe (1995), Morris and Shin (1997) and others, simplify the exposition, but do not affect any of the core analysis.
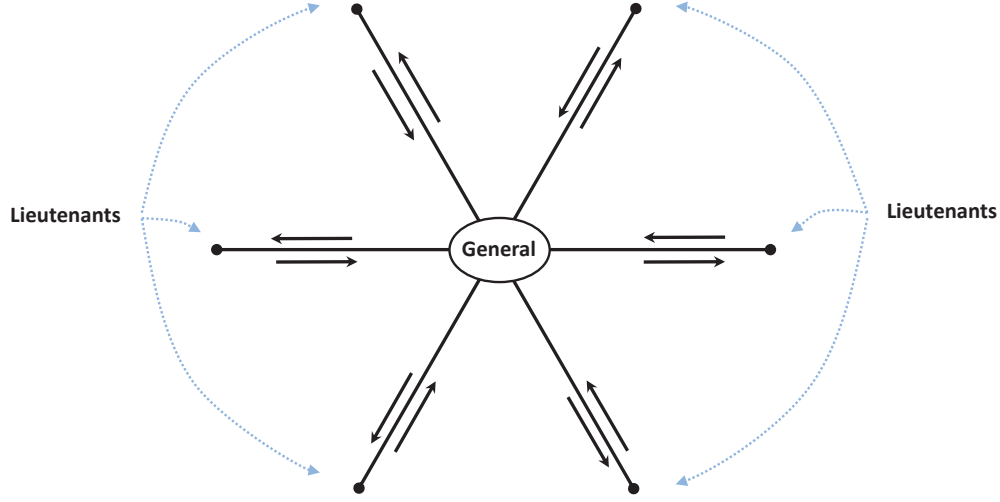
Figure 1: Messaging Structure in the Hub and Spoke Electronic Mail Game

original game that upon sending her Tth message, a player knows that all players have a minimum depth of knowledge of the state, yet do not have common knowledge. For example, when player $i \in \{2, \ldots, N\}$ has sent T_i messages, then $K_i(K_{Everyone})^{T_i-1}b$, where the notation $K_i(E)$ means that "i knows that event E has taken place."

Following this automatic process, the lieutenant's computer displays the number of messages T_i he has sent. A pure strategy $s_i : \mathbb{Z}^+ \to \{A, B\}$ for each lieutenant $i = 2 \ldots N$ is a mapping describing i's planned action when his machine has sent T_i messages (where $\mathbb{Z}^+$ indicates the non-negative integers). The computer of player 1, the general, displays both the number of *rounds* of messages T_1 his machine has sent, *and* the vector $\mathbf{m} \in \{0, 1\}^{N-1}$ of messages he receives in the final communication. For example, $\mathbf{m} = \langle 0, 1, 1, \ldots, 1 \rangle$ means the general received confirmations from all lieutenants except lieutenant 2. We let $m \equiv \mathbf{m} \cdot \mathbf{1}$ denote the number of messages received by the general in the final communication. A pure strategy $s_1 : \left(\mathbb{N} \times \{0, 1\}^{N-1}\right) \cup (0, *) \to \{A, B\}$ for the general is a mapping giving his action when his machine has sent T_1 rounds of messages and vector $\mathbf{m}$ of confirmations (with $m < N - 1$) triggered the communication breakdown. Note that $T_1 = 0$ is a special case. It means that the state of nature is a, so the general will never receive confirmations. His action under these circumstances is given by $s_1(0, *) \in \{A, B\}$.

We refer to this setting as the Hub and Spoke Electronic Mail Game (Hub and Spoke EMG).

3   Independent Message Failures

In this section, we suppose that failures of messages are independent. That is, in any round of messages sent out by the general to his lieutenants, any particular message will fail to reach its destination with probability ε, and this event is independent of the outcomes of the messages sent to the other lieutenants. Likewise, each message sent by a lieutenant fails to reach the general with probability ε, independent of other message outcomes. Under this specification of $p(\cdot)$ we have the following result.

Theorem 1. *In the Hub and Spoke EMG, there is a unique Nash equilibrium: all agents play A independently of the number of messages sent.*

The theorem states that, as in Rubinstein's two player EMG, players are unable to coordinate on playing B, even when it is valuable to do so, and when players have great depth of knowledge of the state.

Our proof relies on induction, but it is nevertheless instructive to see why there is no equilibrium in which agents use intuitive strategies that call for playing B after receipt of a threshold number of confirmations (the cutoff strategies referenced in the introduction).

Suppose lieutenant i's strategy calls for playing B when $T_i \geq \bar{T}_i$ is displayed on its machine (which means that the lieutenant has sent and received $\geq \bar{T}_i$ messages). What action should the general take if he has sent a $\bar{T}_i$th message, but fails to receive a confirmation from this lieutenant? He assigns the following conditional probabilities for T_i, the value displayed on the lieutenant's screen:

$$T_i = \begin{cases} \bar{T}_i - 1 & \text{with probability } \frac{1}{2-\varepsilon} \\ \bar{T}_i & \text{with probability } \frac{1-\varepsilon}{2-\varepsilon}. \end{cases}$$

Thus, it is more likely that the lieutenant did not receive the final message. Consequently, the general should take the safe action A. Since this logic applies to any of the lieutenants, it follows that the general's computer must display at least $\bar{T}_i + 1$ before the general will be confident enough to play B.

Now suppose that the general's strategy calls for playing B only if $T_1 \geq \bar{T}_1$ is displayed on its machine. What action should a lieutenant take if he has sent a $(\bar{T}_1 - 1)$th message to the general, but fails to receive a confirmation? If he is particularly optimistic and assumes that all other lieutenants have also received $\bar{T}_1 - 1$ messages and sent $\bar{T}_1 - 1$ confirmations, he assigns the following conditional probabilities to possibilities for T_1, the value displayed

on the general's screen:

$$T_1 = \begin{cases} \bar{T}_1 - 1 & \text{with probability } z \quad \equiv \frac{1-(1-\varepsilon)^{N-1}}{1-(1-\varepsilon)^{N-1}+\varepsilon(1-\varepsilon)^{N-1}} \\ \bar{T}_1 & \text{with probability } 1-z \end{cases}.$$

Note that $z > \frac{1}{2}$ (in particular, for small ε, $z \approx \frac{N-1}{N}$), so that the lieutenant should take the safe action A. In the less optimistic scenario that some lieutenant failed to send $T_1 - 1$ messages, then with certainty $T_1 < \bar{T}_1$, and the lieutenant should again take the safe action A. It follows that the lieutenant's computer must display at least $\bar{T}_1$ before he will be confident enough to play B.

These best response analyses can be interpreted as meaning that each player favors conservative action, and each would like to have a later cutoff than her counterpart. Hence, no cutoff point is acceptable for an equilibrium. By contrast, in the following section we will show that with correlated signals, this "battle for the later cutoff" is resolved, and players are content to attack after receiving a minimum number of confirmations.

4 Correlated Message Failures

In this section, we suppose that failures of messages sent by the general are correlated. We will demonstrate that with sufficiently high levels of correlation, players will be able to coordinate their actions in equilibrium by playing cutoff strategies.

For the Tth round of messages sent out by the general, let the success of messages be given by the random variable $\mathbf{X}^T = (X_2^T, \ldots, X_N^T)$, where $X_2^T, \ldots, X_N^T$ are Bernoulli variables whose values are distributed according to $p(\cdot)$. Let ρ_{ij} be the correlation between variables X_i^T and X_j^T (which does not depend on T). We continue to assume that each message sent by a lieutenant fails to reach the general with probability ε, independent of other message outcomes.

For player 1, we define $S_1^{\bar{T},\bar{m}}$, the *cutoff strategy* with *cutoff* $\bar{T}$ and *threshold* $\bar{m}$ as:

$$s_1^{\bar{T},\bar{m}}(T_1, \mathbf{m}) = \begin{cases} B & \text{for } T_1 \geq \bar{T} + 1 \\ B & \text{for } T_1 = \bar{T}, m \geq \bar{m} \\ A & \text{for } T_1 = \bar{T}, m < \bar{m} \\ A & \text{for } T_1 < \bar{T}. \end{cases}$$

For players $i \in \{2 \dots N\}$, we define $s_i^{\bar{T}}$, the *cutoff strategy* with *cutoff* $\bar{T}$ as:

$$s_i^{\bar{T}}(T_i) = \begin{cases} B & \text{for } T_i \geq \bar{T} \\ A & \text{for } T_i < \bar{T}. \end{cases}$$

A *cutoff equilibrium* is an equilibrium in which players employ cutoff strategies. For each lieutenant, a cutoff strategy simply specifies the minimum number of messages he must send in order to play B, while a cutoff strategy for the general specifies both the number of (rounds of) messages he must send, as well as the minimum number of confirmations he must receive in the final communication.

To prove our theorem, we assume the following (sufficient) condition, which stipulates that the probability of message failure be sufficiently small relative to the proportion $\frac{M}{M+L}$ (which reflects potential gains from collective action vs. losses from miscoordination.)

Condition 1. $\varepsilon < \frac{M}{M+L}$

We now have following result.

Theorem 2. *Fix ε, the failure probability for a message from any lieutenant to the general, and assume that condition 1 holds. Consider all distributions $p(\cdot)$ where, in expectation, any single message from the general fails with probability ε; that is, $E[X_i] = 1 - \varepsilon$. There exists $\rho \in (0, 1)$ such that for every $p(\cdot)$ with $\rho_{ij} \geq \rho$ for all i, j, we have that $(S_1^{\bar{T},1}, S_2^{\bar{T}}, S_3^{\bar{T}}, \dots, S_N^{\bar{T}})$ is an equilibrium in cutoff strategies for any $\bar{T} \geq 1$.*

Theorem 2 states that there is a multiplicity of equilibria in cutoff strategies, provided the correlation between failures of messages sent by the general is sufficiently high. In each of these equilibria, players use a common cutoff, and the threshold for the number of confirmations the general must receive in the final communication is 1. Observe that there is also a "babbling" equilibrium in which agents always play A, which can be interpreted as corresponding to $T = \infty$.

This result stands in contrast to that from the previous section. Why are the players now able to coordinate on action B after receiving a minimum number of confirmations, but were unable to do so when message failures were independent? The correlation generates a natural asymmetry that breaks the "battle for the later cutoff."

Consider the uncertainty lieutenant i faces when his machine has sent $T_i > 0$ messages. He knows that the communication breakdown could have

resulted from one or more messages not making it through to the general (including possibly the lieutenant's own message), and that it is much less likely that all $N - 1$ messages made it through and in fact it was the general's confirmation that was lost. Thus, in a sense i is sending a "weak" message: communication most likely ceased because he or one of the other lieutenants sent a message that failed to make it through to the general.

On the other hand, consider the general's decision when his machine has sent T_1 messages and he has received at least one confirmation in the final round. The general can be sure that at least one message made it through to a lieutenant. If the general's outgoing messages are (positively) correlated, conditional on a particular message arriving the general is more confident of his other messages' arrivals than he would have been if message failures were independent. With sufficiently high correlation, the general is highly confident that all of his messages arrived, and that the missing message(s) are instead the result of failures of confirmations from the lieutenants. In this sense, the general is sending a comparatively "strong" message, and prefers to use a cutoff *equal* to his counterparts'.

When correlation is high, there are no equilibria in cutoff strategies other than those described in Theorem 2. To see this, first note that strategy profiles where the general uses a different cutoff from the lieutenants can easily be ruled out, as one of the sides will be playing the risky action (tragically) early. There are also no cutoff equilibria where the general uses a threshold other than $\bar{m} = 1$; with sufficient correlation, the general's optimal action for $m > 1$ must be the same as the optimal action for $m = 1$.

The equilibria in Theorem 2 can be Pareto ranked; welfare is decreasing in cutoff. Each of these equilibria Pareto dominates the no-communication equilibrium ("$T = \infty$"). The best equilibrium is that where lieutenants require at least one message and the general at least one confirmation in order to play B.

While arguably an improvement over the no-communication result when messages are independent, valuable coordination in this environment is still partial compared to the payoff-dominant equilibrium when agents have common knowledge of the state. In our setting, message failures prior to the cutoff prompt all agents to play the safe action A. Separately, inopportune message failures near the cutoff can result in some players incurring costs as they play B, while others do not. However, with small ε, in the most efficient equilibrium in cutoff strategies, both of these occur with low probability. It can be shown, using an argument similar to one in Gray (1978), that in the class of communication protocols that involve a finite number of imperfect messages in expectation, no protocol can lead to perfect coordination in both states of

the world. In this sense, we can be satisfied with an outcome where agents coordinate on B when the state is b and on A when the state is a all but a small fraction of the time.

4.1 Examples

Example 1. Perfect Correlation

Suppose that when the general sends a message, it will reach all lieutenants with probability $1 - \varepsilon$ and will reach no lieutenants with probability ε. Messages from lieutenants independently fail to reach their destination each with probability ε. This scenario could arise when the source of messages from the general is subject to "system failure," as might happen, for example, if the general's radio breaks, or if a common message is sent out but is garbled and uniformly uninterpretable.

Since this is a special case of Theorem 2, we know that players may coordinate by using cutoff strategies. A lieutenant who has not received a message knows with certainty that his fellow lieutenants have also not received this confirmation. Conditional on communication breakdown, each lieutenant assigns high likelihood to one of their messages not making it through to the general, rather than all $N - 1$ messages making it through and the general's message not making it back. Hence, the lieutenants must indeed match the general's cutoff in equilibrium. Importantly, the general too can infer much about the information sets of the lieutenants, regardless of the number of messages he receives in the final round of communication. With a single confirmation, he can be sure that his message made it through to at least one player, and hence to all other players—a perfectly strong signal. Knowing his counterparts received the final message, he can safely play the risky action B, even when those final messages are just sufficient for the lieutenants to meet their requirements for playing B.

Example 2. Independent Messages with Chance of "System Failure"

Suppose when the general sends messages, there is a probability γ of "system failure" in which case no messages go out. If there is no system failure, messages are independent and fail with probability δ, where δ is chosen so that $E[X_i^T] = \varepsilon$ for all $T \geq 1$, $i \in \{2, \ldots, N\}$. By Theorem 2, as we increase the probability of system failure (up to ε) and correspondingly reduce δ, we will eventually enter a region where players can coordinate on attacking in equilibrium.

This result suggests that the general should consider connecting with his lieutenants via a "long-but-weak" common cable and then "short-and-strong" individual cables instead of via a "short-and-strong" common cable followed by "long-but-weak" individual cables. That is, the general should invest in a communication channel that gets a clear, common message as close to his followers as possible.

Example 3. An Odd Equilibrium

Theorem 2 characterizes the set of all cutoff equilibria when correlation is high. However, there are also less intuitive equilibria where the players coordinate on attacking only when receiving specific subsets of messages. We provide one such example.

As in Example 1, assume perfect correlation of failures in the general's outbound messages. Let $s_1(T, \mathbf{m}) = B$ iff T is odd and $m > 0$ or T is even and $m = 0$, and let $s_i(T) = B$ iff T is odd. We claim that these strategies constitute an equilibrium. First, as $\rho = 1$, lieutenants clearly must play A when T is even. When T is odd, lieutenants find it likely that at least one (but importantly, not all) of their signals was lost en route to the general, and hence play B.

As for the general, when T is odd and $m > 0$ he knows that all lieutenants will play B, so he should do so as well. When T is even and $m > 0$ he knows that all lieutenants will play A, so again he should do likewise. When $T > 0$ and $m = 0$ the general believes that the communication breakdown is most likely due to his message being lost and that lieutenants will play $s_i(T - 1)$. Hence, he should also play $s_i(T - 1)$.

5   Conclusion

In this paper we have examined collective action among multiple players who communicate over faulty channels. Our contribution is the illustration of a new way in which the multiplicity of channels may lead to coordination. Namely, when message failures are correlated, players may utilize the multiple channels to infer information about missing confirmations in a way that facilitates action. In our setting, this induced information mitigates pessimism about message receipts for one player at a key decision point and increases his confidence that he may safely take action. This breaks a battle of conservatism that would otherwise cause the game to unravel so that players never take risky action, even when it is mutually profitable to do so.

One limitation of our finding is that it holds only for specific regions of our parameters. In particular, we require that either the probability of messages breaking down is small, or alternatively that losses due to miscoordination are not too large relative to the gains from successful coordination (Condition 1). Furthermore, the degree of correlation that ensures the possibility of coordination depends on the number of players N.[9] These features of our result limit its applicability. How much correlation is necessary to permit coordination? The answer is that it depends.

The most important limitation of this finding is the specific nature of the network structure. We have assumed a single informed agent, the general, who communicates with other agents, but these agents do not communicate with each other. Furthermore, the general only reconfirms when receiving confirmations from all agents. These assumptions may be plausible in instances where there is a natural leader who seeks to coordinate a group of followers. The real advantage to using the hub and spoke model, however, is that it offers a convenient, tractible model for studying how agents may benefit from message failure correlation. By demonstrating the no-coordination result in a setting with independent messages, we are able to isolate correlation as the pivotal factor that leads to coordination. The result is robust in the sense that we do not rely on a specific distribution of message failures; any distribution with sufficiently high correlation permits coordination.

While we certainly do not purport (or believe) that correlated message failures universally facilitate collective action or otherwise increase welfare, we emphasize that our finding demonstrates their potential. Correlation should be acknowledged as one feature of a communication mechanism that may pivotally dictate feasibility of coordination. Since this seemingly subtle feature can mean the difference between reasonably efficient coordination and none at all, it underscores the need for a deeper understanding of what classes of communication protocols best facilitate coordination, and suggests opportunities for further research.

---

[9]Suppose we fix M, L, ε and select p such that the correlation between any two outbound messages from the general is fixed at $\bar{\rho} < 1$. As the number of players increases, one eventually arrives at a network structure where the only equilibrium is one in which all players select A. Hence, the degree of correlation necessary for coordination critically depends on N.

Appendix A: When Threshold for Successful Attack is $< N$.

In this appendix we investigate outcomes when the threshold participation level for collective action to be successful is less than the number of players N. We model this two different ways: one where the general seeks to coordinate separately with each lieutenant, and another where there are three players, but success is assured provided at least two attack. In each of these settings, partial coordination can be sustained in equilibrium. The intuition here is clear: when a big army undertakes a small operation so that only part of its power is needed, a lieutenant can count on at least partial support even when the communication is faulty. Additionally, as this fact reduces the risk from joining the operation, other lieutenants will likewise join, further reducing the likelihood of loss.

Example A1. Pairwise Coordination

Suppose that instead of the payoff structure from Section 2, now when the general chooses to play B, successful coordination with some of the lieutenants yields gains, while failed coordination with others incurs losses. Total payoff for the general is given by gains net of the losses.

We preserve the information structure, independent message distribution, and strategy sets of Section 3, altering only the payoffs. As in Section 2, after communication, each agent plays A or B. For each pair $(1, i \in \{2, \ldots, N\})$, payoffs are calculated as follows. An agent who plays A receives payoff 0. If both agents coordinate on B and the state is b, then each receives payoff M. If only one agent plays B when the state is b, that agent receives $-L$. Playing B when the state is a likewise yields $-L$. For each lieutenant i, this is all he will receive. The general's total payoff is the sum of the payoffs from each pairwise coordination effort. We refer to this setting as the Hub and Spoke Electronic Mail Game with Pairwise Coordination.

Notice that now the general may be willing to play B if he expects only a subset of the $N - 1$ peripheral players to play B. The lieutenants, however, will only play B if it is likely that the general will also play B.

To prove the result in this example, we assume the following (sufficient) condition.

Condition A1. $M(N - 2) - L > 0$

Condition A1 has the following interpretation: If the general manages to coordinate with all the lieutenants except one, his payoff is positive.

We have the following proposition:

Proposition A1. *In the Hub and Spoke EMG with Pairwise Coordination, for sufficiently small ε, $\exists \bar{m} < N - 1$ such that $(s_1^{\bar{T},\bar{m}}, s_2^{\bar{T}}, s_3^{\bar{T}}, \ldots, s_N^{\bar{T}})$ is an equilibrium in cutoff strategies for any $\bar{T} \geq 0$.*

Proposition A1 states that as in the Hub and Spoke EMG with correlated message failures, players may coordinate via cutoff strategies. But unlike with correlated messages where the threshold number of confirmations in the final communication for the general was 1, in this case the threshold may be higher, and is dictated by the number of lieutenants who, if they coordinate, will offset the losses from failed coordination attempts with the others.

Why can players coordinate on playing B, but not in the setting of Section 3? When communication stops, each lieutenant is uncertain if his final message made it through. However, he knows that it is highly likely that most of the final messages from the lieutenants did make it through to the general. For small ε, given that one message failed, it is unlikely that even two messages failed. Hence the lieutenants send a "strong" signal, because multiple signals need to break down for the general to miss his threshold. Upon sending T messages, each lieutenant i is confident that enough messages have gotten through to trigger an attack by the general, so the lieutenants are content with matching the general's cutoff.

Example A2. Threshold for Successful Attack $< N$

Let $N = 3$, and consider a model identical to that in Section 3, except that now only two agents need attack to ensure success. That is, assume each agent who attacks (plays B) receives payoff M if the state is good (b) and at least one other agent has attacked. An agent receives $-L$ if she alone plays B when the state is good, or if she plays B and the state is bad. Agents receive 0 from playing A.

We claim that for small ε and not-too-large L, the following cutoff strategies constitute an equilibrium:

$s_1(T_1, \mathbf{m}) = B$ iff $T_1 \geq 1$

$s_i(T_i) = B$ iff $T_i \geq 1$.

We begin by showing that s_i is a best response for the lieutenants. WLOG we inspect s_2. Clearly, if $T_2 \geq 1$, then $T_1 \geq 1$ and attacking is optimal. If $T_2 = 0$, then the conditional probability player 2 assigns to the state being good is small (namely $\frac{(1-q)\varepsilon}{q+(1-q)\varepsilon}$, which is less than 1/2 and approaches zero when epsilon approaches zero). Therefore, A is his best response.

We now inspect S_1. If $T_1 > 1$ or $T_1 = 1$ and $m > 0$, player 1 assigns probability 1 to a lieutenant playing B, so B is optimal. If $T_1 = 1$ and

$m = 0$ we have $Pr\{s_2(T_2) = s_3(T_3) = A \,|\, T_1, m\} = \frac{\varepsilon^2}{\varepsilon^2 + 2\varepsilon^2(1-\varepsilon) + (1-\varepsilon)^2\varepsilon^2}$ while
$Pr\{\exists i \text{ s.t. } s_i(T_i) = B \,|\, T_1, m\} = \frac{2\varepsilon^2(1-\varepsilon) + (1-\varepsilon)^2\varepsilon^2}{\varepsilon^2 + 2\varepsilon^2(1-\varepsilon) + (1-\varepsilon)^2\varepsilon^2}$. For small ε, the likelihood
of at least one other player attacking is roughly $3/4$. Therefore, provided L is
not too large, B is optimal. The case $T_1 = 0$ is trivial.

Appendix B: Proofs

Proof of Theorem 1

  Let $(s_1, \ldots, s_N)$ be a pure strategy Nash equilibrium. We will show by
induction on T that $s_1(T, \mathbf{m}) = s_i(T) = A \;\; \forall T, \mathbf{m}, i = 2 \ldots N$. Our proof also
establishes that even when considering mixed strategy equilibria, this is the
unique Nash equilibrium in the Hub and Spoke EMG.

  Observe first that in any equilibrium, we must have $s_1(0, *) = A$, since
A yields a payoff of zero for the general, while B yields $-L$.

  If lieutenant i has $T_i = 0$, then he did not receive any messages. It could
be that the state of nature is a and no message was sent, which occurs with
probability $1-q$, or that a message was sent and was lost en route, which occurs
with probability $q\varepsilon$. In the first case, by assumption, the general plays A. If
lieutenant i plays A, then regardless of what $s_1(1, \mathbf{m})$ is, his expected payoff
is 0. If he plays B, his payoff is at most $\frac{-L(1-q) + Mq\varepsilon}{1-q+q\varepsilon}$. Comparing these, we see
that it is strictly better for player i to play A; that is, $s_i(0) = A \; \forall i = 2 \ldots n+1$.

  Now suppose that we have shown that in equilibrium, for all $T < t$ we
have $s_1(T, \mathbf{m}) = s_i(T) = A \;\; \forall\; \mathbf{m}, i = 2 \ldots N$. Let us calculate the general's
optimal action when $T_1 = t$. In this case, the general has sent t messages, but
did not receive a confirmation from one or more lieutenants. Let i designate
one of these lieutenants. The general then assigns the following conditional
probabilities:

$$T_i = \begin{cases} t-1 & \text{with probability } \frac{1}{2-\varepsilon} \\ t & \text{with probability } \frac{1-\varepsilon}{2-\varepsilon} \end{cases}.$$

  When $T_i = t-1$, i plays A. Thus, if the general chooses B, he gets
at most $M\frac{1-\varepsilon}{2-\varepsilon} - L\frac{1}{2-\varepsilon}$. If he chooses A, he receives 0. Because $L > M$ and
$\frac{1}{2-\varepsilon} > \frac{1-\varepsilon}{2-\varepsilon}$, his best option is A. Thus, $s_1(t, \mathbf{m}) = A \;\; \forall\; \mathbf{m}$.

  We now calculate the optimal action for lieutenant i when $T_i = t$.
Suppose the lieutenant optimistically believes all the other lieutenants also
received t messages (and sent t confirmations). Then i is uncertain if the
general received all $N-1$ messages (so that $T_1 = t+1$), or if at least one
message failed to make it through (so that $T = t$). Player i assigns the

following conditional probabilities:

$$T_1 = \begin{cases} t & \text{with probability } z \\ t+1 & \text{with probability } 1-z \end{cases} \equiv \frac{1-(1-\varepsilon)^{N-1}}{1-(1-\varepsilon)^{N-1}+\varepsilon(1-\varepsilon)^{N-1}} .$$

When $T_1 = t$, 1 plays A. Thus, if i chooses B, his payoff is at most $-Lz + M(1-z)$. (If we include the less optimistic cases where the other lieutenants receive fewer than t messages then $\text{prob}(T_1 = t) > z$ and our payoff bound is lower still.) If he chooses A, his payoff is 0. Because $L > M$ and $z > \frac{1}{2}$, A is his best option. That is, $s_i(t) = A \ \forall \ i = 2 \dots N$.

Note that in our analysis each best response was unique, so that "always play A" is the unique equilibrium in the Hub and Spoke EMG. $\square$

Proof of Theorem 2

We start by stating some basic properties that will be useful throughout the proof. In the following four inequalities, random variable X_i is an indicator for the arrival a message from the general to lieutenant i (in any round of confirmations sent by the general.) At all periods, we have:

1. $P\{X_i = X_j = 1\} \geq \rho\varepsilon(1-\varepsilon) + (1-\varepsilon)^2$

2. $P\{X_i = X_j = 0\} \geq \rho\varepsilon(1-\varepsilon) + \varepsilon^2$

3. $P\{X_i \neq X_j\} \leq 2(1-\rho)\varepsilon(1-\varepsilon)$

4. $P(\mathbf{1}) + P(\mathbf{0}) = P\{X_2 = X_3 = \dots = X_N\} \geq 1 - P\{\exists i, j \ X_i \neq X_j\} \geq$
   $1 - \sum\limits_{i=3}^{N} P\{X_2 \neq X_i\} \geq 1 - 2(N-2)(1-\rho)\varepsilon(1-\varepsilon) \underset{\rho \to 1}{\to} 1$

We now turn to proving the theorem. We first show that the strategy of the general (player 1) is a best response to the strategies of the lieutenants.

First, $T_1 > \bar{T} \Rightarrow T_i \geq \bar{T}$, so in this case, B is player 1's best response. Similarly, $T_1 < \bar{T} \Rightarrow T_i < \bar{T}$, so A is 1's best response in this case.

When $T_1 = \bar{T}$ and $m = 0$, player 1 cannot be sure whether the lieutenants all received the $\bar{T}$th message and a confirmation failed to return (the "good" case, in which the player 1 should play B) or at least one lieutenant failed to receive a message (the "bad" case, in which the player 1 should play A). The probability of the bad case, conditional on the information known to 1, is at least $\frac{P(\mathbf{0})}{P(\mathbf{0})+P(\mathbf{1})\varepsilon^{N-1}}$, which converges to $\frac{1}{1+(1-\varepsilon)\varepsilon^{N-2}} > \frac{1}{2}$ as $\rho \to 1$. Hence, A is 1's best response for high enough ρ.

When $T_1 = \bar{T}$ and $m > 0$, player 1 again cannot be sure whether the lieutenants all received the $\bar{T}$th message and a confirmation failed to return

(the good case) or at least one lieutenant failed to receive a message (the bad case). Here, the probability of the good case is $\geq \frac{P(\mathbf{1})\varepsilon^{N-1-m}(1-\varepsilon)^m}{P(\mathbf{1})\varepsilon^{N-1-m}(1-\varepsilon)^m+1-P(\mathbf{0})-P(\mathbf{1})}$, and this expression approaches 1 as $\rho \to 1$, so player 1 should play B. Hence, the general's stated strategy is a best response to the strategies of the lieutenants.

We now turn to the strategies of the lieutenants. WLOG, we analyze lieutenant 2's strategy. Since $T_2 > \bar{T} \Rightarrow T_1 > \bar{T}$ and $T_i \geq \bar{T} \, \forall i$, we have that B is the best response in this case. Similarly, $T_2 < \bar{T} - 1 \Rightarrow T_i \leq \bar{T} - 1 \, \forall i$, so in this case A is 2's best response.

When $T_2 = \bar{T} - 1$, lieutenant 2 cannot be sure whether his fellow lieutenants all got a $\bar{T}$th message and not all confirmations were lost en route to the general (the good case), or else if some other lieutenant did not get a $\bar{T}$th message, or else if his fellow lieutenants all got a $\bar{T}$th message, but none of the confirmations made it back to the general (the bad cases). To create a (loose) lower bound on the probability of the bad case, we assume the "extreme" scenario in which 2 believes the general's $\bar{T}$th message was sent. The probability of the bad cases is at least $\frac{P(\mathbf{0})}{P(\mathbf{0})+[1-P(\mathbf{0})-P(\mathbf{1})]}$, which approaches 1 as $\rho \to 1$, so that A is 2's best response for large enough ρ.

Finally, when $T_2 = \bar{T}$ lieutenant 2 cannot be sure whether his fellow lieutenants all got a $\bar{T}$th message and at least one confirmation was lost en route to the general (the good case) or else if some lieutenant did not get a $\bar{T}$th mesage, or else if his fellow lieutenants all got a $\bar{T}$th message but all confirmations were lost en route to the general (the bad cases). The probability of the good case is at least

$$\frac{P(\mathbf{1})\cdot[(1-\varepsilon^{N-1})-(1-\varepsilon)^{N-1}+(1-\varepsilon)^{N-1}P(\mathbf{0})]}{P(\mathbf{1})\cdot[(1-\varepsilon^{N-1})-(1-\varepsilon)^{N-1}+(1-\varepsilon)^{N-1}P(\mathbf{0})]+[1-P(\mathbf{0})-P(\mathbf{1})]+P(\mathbf{1})\varepsilon^{N-1}} \xrightarrow[\rho\to 1]{}$$

$$\frac{(1-\varepsilon^{N-1})-(1-\varepsilon)^N}{(1-\varepsilon^{N-1})-(1-\varepsilon)^N+\varepsilon^{N-1}}.$$

The next series of expressions demonstrates that, provided Condition 1 holds, B is indeed 2's best response (for large enough ρ):

$$\alpha := \frac{(1-\varepsilon^{N-1})-(1-\varepsilon)^N}{(1-\varepsilon^{N-1})-(1-\varepsilon)^N+\varepsilon^{N-1}} = \frac{1-\varepsilon^{N-1}-(1-\varepsilon)^N}{1-(1-\varepsilon)^N} = 1 - \frac{\varepsilon^{N-1}}{1-(1-\varepsilon)^N} =$$

$$1 - \frac{\varepsilon^{N-1}}{\varepsilon[1+(1-\epsilon)+(1-\epsilon)^2+...+(1-\epsilon)^{N-1}]} \geq 1 - \epsilon^{N-2} \geq 1 - \epsilon \Rightarrow$$

$$\alpha M - (1-\alpha)L > 0 \iff 1 - \alpha < \frac{M}{M+L}.$$

So, as $1 - \alpha \leq \epsilon$ a sufficient condition is $\epsilon < \frac{M}{M+L}$.

Finally, note that the convergence was derived using only the bounds given in the beginning. Thus, we can find a $\rho < 1$ that suffices uniformly (independent of p). □

Proof of Proposition A1

We first find a threshold $\bar{m}$ such that $\forall\ T$, the cutoff strategy with cutoff T and threshold $\bar{m}$ is a best response for player 1 to the strategies of players $2 \ldots n + 1$.

Suppose the general's machine displays T_1, with $m < N - 1$ messages received in the final communication. If $T_1 \geq T + 1$, then the general can be sure that each lieutenant i has $T_i \geq T$, so that each will choose action B. Thus B is the best response. If $T_1 < T$, then each lieutenant player i has $T_i < T$, so that each will choose action A. Here, A is the best response.

Now suppose that $T_1 = T$. In this case, since 1 has received m confirmations, he is certain that at least m lieutenants have T. For each of the $(N - 1) - m$ remaining lieutenants, the general is unsure if his final communication was lost en route, or if it was received and the confirmation was lost. He assigns conditional probabilities $\frac{1}{2-\varepsilon}$ and $\frac{1-\varepsilon}{2-\varepsilon}$ to these two possibilities. Thus, playing B will yield payoff $mM + ((N - 1) - m)[-L\frac{1}{2-\varepsilon} + M\frac{1-\varepsilon}{2-\varepsilon}]$. This payoff is increasing and linear in m. When $m = 0$, the payoff is negative and by condition A1, the payoff is positive when $m = N - 2$. Hence, setting $\hat{m} \equiv (N - 1)\left[1 - \frac{(2-\varepsilon)M}{M+L}\right]$, the value of m for which this payoff is zero, we must have $\hat{m} \in (0, N - 2)$. Let $\bar{m}$ be $\lceil \hat{m} \rceil$. Then when $m \geq \bar{m}$, B is the best action, and when $m < \bar{m}$, the zero payoff action A is preferable.

Thus, cutoff strategy with cutoff T and threshold $\bar{m}$ is the best response to $S_i, i = 2 \ldots N$.

We now show that each lieutenant i's strategy is a best response to the strategies of the others (where the general uses a cutoff strategy with cutoff T and threshold $\bar{m}$.)

If $T_i \geq T + 1$, then $T_1 \geq T + 1$, in which case the general will play B. Thus, i's best response is B. If $T_i < t - 1$, then the general has $T_1 < t$ and will play A. Thus, A is the best response.

Suppose $T_i = T - 1$. Then the only way the general could have $T_1 = T$ is if i's $(T - 1)$th message made it through, all other players received a $(T - 1)$th message and sent successful confirmations to the general, and it was the confirmation sent from 1 to i that did not make it back. If player i is optimistic and assumes his fellow lieutenants all also received their $(T-1)$th message, he assigns conditional probabilities $z \equiv \frac{\varepsilon(1-\varepsilon)^{N-1}}{\varepsilon(1-\varepsilon)^{N-1}+1-(1-\varepsilon)^{N-1}}$ to $T_1 = T$ and $1 - z$ to $T_1 = T - 1$. When $T_1 = T - 1$, player 1 chooses A. Thus, player i choosing

B in this scenario yields a payoff of at most $Mz - L(1-z)$. But $z < \frac{1}{2}$ for any ε (and in the less optimistic scenarios, this conditional probability is lower still), so the safe, zero payoff choice A is preferred when $T_i = T - 1$.

Now suppose $T_i = T$. For lieutenant i to prefer to attack, he must find it sufficiently likely that the general either has $T_1 = t + 1$, or else has $T_1 = t$ and $m \geq \bar{m}$. As $\varepsilon \to 0$, the conditional probability that the general has $T_1 = T$ and $m \geq \bar{m}$ goes to one, since with high likelihood the source of failure stems from a single lieutenant's Tth confirmation failing to get through to the general. Hence, for sufficiently small ε, each lieutenant selects B when $T_i = T$, and the proof is complete. □

References

Aumann, Robert. 1976. "Agreeing to Disagree." *The Annals of Statistics*, 4: 1236–39.

Binmore, Ken, and Larry Samuelson. 2001. "Coordinated Action in the Electronic Mail Game." *Games and Economics Behavior*, 35: 6–30.

Chwe, Michael Suk-Young. 1995. "Strategic Reliability of Communication Networks." University of Chicago Department of Economics Working Paper.

Colbourn, Charles J. 1987. *The Combinatorics of Network Reliability.* New York, NY, USA:Oxford University Press, Inc.

Coles, Peter. 2005. "Essays on Signaling and Matching." Stanford University Department of Economics, Dissertation.

De Jaegher, Kris. 2008. "Efficient Communication in the Electronic Mail Game." *Games and Economics Behavior*, 63: 468–497.

De Jaegher, Kris. 2011. "Collective Action in the Multi-Player Electronic Mail Game." Utrecht University Department of Economics Working Paper.

Dimitri, Nicola. 2003. "Coordination in an Email Game without 'Almost Common Knowledge'." *Journal of Logic, Language and Information*, 12: 1–11.

Dimitri, Nicola. 2004. "Efficiency and equilibrium in the electronic mail game; the general case." *Theor. Comput. Sci.*, 314: 335–349.

Dulleck, Uwe. 2007. "The E-mail Game Revisted - Modeling Rough Inductive Reasoning." *International Game Theory Review*, 9: 323–339.

Gray, Jim. 1978. "Notes on Database Operating Systems." In *Lecture Notes in Computer Science. Operating Systems: An Advanced Course.* Vol. 66, , ed. R. Bayer, R. M. Graham and G. Seegmuller. Berlin: Springer-Verlag.

Morris, Stephen. 2001. "Faulty Communication: Some Variations on the Electronic Mail Game." *Advances in Theoretical Economics*, 1. Article 5.

Morris, Stephen, and Hyun Song Shin. 1997. "Approximate Common Knowledge and Co-ordination: Recent Lessons from Game Theory." *Journal of Logic, Language, and Information*, 6.

Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior Under 'Almost Common Knowledge'." *The American Economic Review*, 79: 385–391.

Shier, Douglas R. 1991. *Network reliability and algebraic structures.* New York, NY, USA:Clarendon Press.

Strzalecki, Tomasz. 2011. "Depth of Reasoning and Higher Order Beliefs." *Harvard University Working Paper.*

Takamiya, Koji, and Akira Tanaka. 2006. "Mutual Knowledge of Rationality in the Electronic Mail Game." *Osaka University Discussion Paper.*