

The Center for Mathematical Studies in Economics and Management Science
Northwestern University, Evanston, Illinois 60201

Discussion Paper No. 590

MULTISTAGE GAMES WITH COMMUNICATION

by

Roger B. Myerson

April 1984

revised February 1985

(Former title: "Sequential Correlated Equilibria of Multistage Games.")

Abstract. This paper considers multistage games with communication mechanisms that can be implemented by a central mediator. In a communication equilibrium, no player expects ex ante to gain by manipulating his reports or actions. A sequential communication equilibrium is a communication equilibrium with a conditional probability system under which no player could ever expect to gain by manipulation, even after zero-probability events. Codominated actions are defined. It is shown that a communication equilibrium is a sequential communication equilibrium if and only if it never uses codominated actions. Predominant communication equilibria are defined by iterative elimination of codominated actions and are shown to exist.

Acknowledgements. This research was supported by fellowships from the John Simon Guggenheim Memorial Foundation and the Alfred P. Sloane Foundation.

MULTISTAGE GAMES WITH COMMUNICATION

by
Roger B. Myerson

1. Introduction

When people try to exchange information or coordinate their activities in a social system, they are constrained by the need to give each other correct incentives. Individuals cannot be expected to testify against themselves, or to exert efforts for which they will not be rewarded. The goal of this paper is to develop a precise and tractable characterization of what rational individuals can achieve with communication, subject to such incentive constraints, in a dynamic social system.

In static or one-stage games, the limits of what rational players can accomplish with communication has been characterized by the concepts of correlated equilibrium, incentive compatibility, and the revelation principle (see Aumann [1974], and Myerson [1979, 1982, 1983]).^{1/} In dynamic or multistage games, even without communication, the problem of characterizing rational behavior is much more complicated, because of difficulties in defining how a rational player should behave after observing an event that had zero probability. Selten's [1975] concept of perfect equilibrium and Kreps' and Wilson's [1982] concept of sequential equilibrium have been proposed as characterizations of rational behavior in multistage games without communication. This paper combines ideas from these two strands in the

^{1/} The term correlated equilibrium was first used by Aumann [1974] in the context of pure moral-hazard problems, while the term Bayesian incentive compatible mechanism was originally used in the case of pure adverse-selection problems. For cases involving both adverse selection and moral hazard, Myerson [1982, 1983] continued to use the term "Bayesian incentive-compatible mechanism" while Forges [1984] used the term communication equilibrium. This paper follows Forges's terminology, because it is simpler and better indicates the game-theoretic context of these ideas.

literature, to define a concept of sequential communication equilibrium for multistage games with communication.

For one-stage games, let a (direct) communication mechanism be a centralized communication system with the following structure: first, each player makes a confidential report about all of his private information to some central mediator (which may be either a trustworthy person, or a specially designed machine); then the mediator computes a recommended action for each player, as a (possibly random) function of these reports; and then the mediator confidentially tells each player what is his recommended action. A communication equilibrium (or incentive-compatible communication mechanism) is any direct communication mechanism such that it would always be rational for each player to report honestly and choose his action obediently to the mediator, if all other players were expected to also be honest and obedient.

The revelation principle asserts that, for any (centralized or decentralized) communication system, and for any Nash equilibrium of strategies that the players might use in this communication system, there exists an equivalent communication equilibrium which always yields the same outcomes. That is, there is no loss of generality in assuming that the players use a centralized communication system, in which the players tell all of their information to a mediator, who in turn tells each player only what is necessary to guide his action. When the mediator knows all the players' information, he can simply tell them to do whatever they would have done under any other system. On the other hand, the more information that a player gets, the harder it may be to prevent him from finding ways to gain by disobeying the mediator. Thus, the only information that the mediator should transmit to a player is the name of the action or move that is recommended for him. If a

communication mechanism is a communication equilibrium, then a rational player should not need any further information to persuade him to obey the recommendation. So the set of these direct communication equilibria completely characterizes what the players can achieve by communication.

For multistage games, the same principle applies but matters are more complicated, for several reasons. In each stage, there may be new information available to each player that the mediator cannot prevent him from getting. For example, if players can observe each others' actions, then each player's information in the second stage must include his knowledge of what the other players did in the first stage. Our model of the game must include a description of all such unpreventable information flows. Furthermore, we must recognize that each player in later stages will remember what the mediator has told him in earlier stages.

To minimize a player's information in early stages of a multistage game, the mediator should not tell the player about what his recommended actions will be in later stages. Thus, at the beginning of the game, a mediator should not to simply tell each player what strategy to use thereafter. Instead, a new round of communication between players and mediator will be necessary in every stage of the game. In each stage, the mediator should ask each player to report all of his new information, and then the mediator should tell each player his recommended action or move for that stage only. Because we cannot assume that all communication occurs at the beginning of the game, the normal form is not an adequate representation of a multistage game with communication. This issue is discussed in Section 2.

Selten [1975] and Kreps and Wilson [1982] have shown that the simple definition of Nash equilibrium is not sufficiently restrictive to characterize rational behavior in multistage games without communication, and they have

proposed concepts of perfect equilibrium and sequential equilibrium to solve this problem. In Nash equilibrium, it is only required that there should be no event with positive probability in which a player would expect to gain by changing his strategy. In sequential equilibrium, it is required that there should be no events at all (even with zero probability) in which a player would expect to gain by changing his strategy. This may seem to be a fine point, but it can be very important, as is shown in Section 3. In multistage games, the probabilities of events in later stages may depend on the players' strategies in earlier stages. Thus, events that have zero probability in one equilibrium may have positive probability in another equilibrium, and so may not be negligible a priori. To characterize rational beliefs after zero-probability events, we need the theory of conditional probability systems developed in Section 5.

Three solution concepts for multistage games with communication are developed in this paper: communication equilibrium (in Section 4), sequential communication equilibrium (in Section 6), and predominant communication equilibrium (in Section 8). The first two of these concepts are analogous to the concepts of Nash equilibrium and sequential equilibrium for games without communication. The main result in this paper is the characterization of sequential communication equilibria in terms of codominated actions, which are defined in Section 7. We show that a communication equilibrium is a sequential communication equilibrium if and only if it would never recommend a codominated action to any player who has not lied to the mediator. This concept of codomination is closely related to the more familiar concept of domination of actions, discussed by Luce and Raiffa [1957]. Thus, this result connects two important ideas in noncooperative game theory: sequential rationality of equilibria and elimination of dominated actions. The process

of iterative elimination of codominated actions is considered, in Section 8, to develop the concept of predominant communication equilibrium, which extends Luce and Raiffa's [1957] concept of wide or iterative domination.


2. Insufficiency of the normal form

We begin by considering some illustrative examples. We represent our examples in extensive form, using standard notation and terminology (see Luce and Raiffa [1957], Chapter 3). The extensive form is slightly different from the multistage form that is introduced in Section 4 of this paper, but the translation between these two forms of representation can easily be made.


[Insert Figures 1 and 2 here.]

Consider now Examples 1 and 2. In Example 1, player 1 first chooses one of the three actions: t , m , or b . If he chooses m or b , then player 2 must choose either l or r . The dotted curve indicates that player 2 does not know whether 1 chose m or b . The payoffs to players 1 and 2 respectively are shown at the right ends of the game tree. (All payoffs are in units of von Neumann-Morgenstern utility functions).

Example 2 differs from Example 1 only in that player 1's three-branch decision node has been changed to a sequence of two-branch decision nodes. That is, player 1 first chooses either t or $-t$; and if he chooses $-t$ then he must next choose either m or b . Notice that these two examples have equivalent normal forms, as shown below.



Example 1
Figure 1



Example 2
Figure 2

		Player 2	
		2,2	2,2
Player 1	5,1	0,0	
	0,0	1,5	

Thus, the Nash equilibria of these two games are the same (if the players cannot communicate) because the set of Nash equilibria depend only on the normal form.

However, the difference between Examples 1 and 2 becomes significant when the players can communicate with a mediator. To be specific, consider a communication mechanism that assigns probability $1/2$ to each of the outcomes (m,ℓ) and (b,r) (after $-t$), so that each player gets an expected payoff of 3. In Example 1, this mechanism is not a communication equilibrium, because player 1 would never want to use his action b , which is dominated by t . In Example 2, however, this is feasible as a communication equilibrium. To implement it, a mediator should first recommend to player 1 that he choose the action $-t$ at the first node. Then after player 1 has made his choice between t and $-t$, the mediator should toss a fair coin. If it is heads then he should recommend that 1 choose m and 2 choose ℓ ; if it is tails then the mediator should recommend that 1 choose b and 2 choose r . Neither player could ever expect to gain by disobeying the mediator's recommendations, if the other player is expected to obey them. In particular, player 1 gets expected utility $3 = .5(5) + .5(1)$ from choosing $-t$ at his first decision node, whereas he would only get 2 if he chose t .

The key to implementing this mechanism in Example 2 is that player 1 must not learn whether m or b is recommended for him until after it is too late for him to select t . In Example 2, there is a point in time when m and b are still available as options for player 1 but t is not available; whereas in Example 1 there is no such point in time. (In Example 2, if the first node represents a decision to be made on Monday, and the second node represents a decision to be made on Wednesday, then the mediator's coin could be tossed on Tuesday.) If all communication had to occur before the beginning of the game then this distinction would not matter. But under the assumption that the players can communicate with each other or with a mediator at any point in the game, the set of communication equilibria is strictly larger in Example 2 than in Example 1, even though these two examples have equivalent normal forms.

Since von Neumann and Morgenstern [1944], game theorists have preferred to study games in normal (or strategic) form, rather than in the conceptually more complicated multistage or extensive form. There was no harm in doing so as long as the solution concept being applied was Nash equilibrium, because the Nash equilibria of an extensive game are the same as the Nash equilibria of its normal form representation. Thus it is disturbing to discover that, if communication equilibria are the solutions that we want to compute, then it is not sufficient to study the normal form; we must consider the extensive dynamic structure of the game.

Since changing our solution concept from Nash equilibrium to communication equilibrium is so analytically costly, it is important to understand what we gain by it. When we say that the players in a game can communicate freely with each other, we are saying that they have a wide range of actions available that affect each other's information but do not affect payoffs. They can send each other messages in any language; they can toss


coins or spin roulette wheels and observe the outcomes; they can even (as suggested by Aumann [1974]) build a machine or hire a mediator to send each of them confidential messages that are generated from any joint probability distribution. In principle, one could try to list all of these possibilities for communication as part of the explicit structure of the game, and then study its Nash equilibria. But the resulting game (with infinitely many options to toss coins, send messages, etc.) would be overwhelmingly complicated. By the revelation principle, we know that any equilibrium of this game with explicit communication possibilities is equivalent to some communication equilibrium of the original game, in which the communication possibilities are not explicitly modelled. Thus, when we change our solution concept from Nash equilibrium to communication equilibrium, we gain the right to simplify our model of the game by omitting the details of how players communicate and coordinate their actions.

3. Perfect equilibria and communication

As Selten [1975] has argued, the concept of Nash equilibrium is too weak to be an exact characterization of rational behavior in dynamic games. For a simple example to illustrate how irrational Nash equilibria can arise, consider Example 3.

[Insert Figure 3 here.]

In this example, (b,c) is a Nash equilibrium because player 1 would




Example 3

Figure 3

prefer b (giving him 1) over a (giving him 0) if he expected player 2 to choose c ; and player 2 could never gain from switching from c to d if player 1 were sure to choose b . However, (b,c) is an imperfect equilibrium, in the sense of Selten [1975], because if player 1 did choose a then player 2 would know that d was strictly better for him than c , and so player 1 should expect 2 to choose d after a . Since d after a leads to a better outcome for player 1 than c , we conclude that the unique perfect equilibrium for this game is (a,d) .

The problem with the concept of Nash equilibrium is that it permits players to behave irrationally in events that have zero probability. Although this may sound innocuous, it is not, because the events in a game that have zero probability are determined endogenously by the equilibrium strategies, so events of zero probability cannot be dismissed a priori (as they are in probability theory). In the Nash equilibrium (b,c) for this example, player 2 behaves irrationally in the zero-probability event that player 1 chooses a , and the event that player 1 chooses a has zero probability because player 2 is expected to play irrationally in this event. To rule out this kind of bizarre logic, we need stronger concepts of equilibrium that require rationality of all players in all possible events, not just the positive-probability events.

Selten's [1975] concepts of subgame perfect and trembling-hand perfect equilibrium, and Kreps and Wilson's [1982] concept of sequential equilibrium are three such stronger concepts of equilibrium that have been offered in the literature. However, when players can communicate, these concepts may eliminate too many Nash equilibria, as Example 4 illustrates.



Example 4

Figure 4

[Insert Figure 4 here.]

Example 4 differs from Example 3 in that there are two more players (3 and 4) who have actions to choose only if player 1 chooses a and player 2 chooses c . The dashed curve at right indicates that, when player 4 chooses between g and h , he does not know whether player 3 chose e or f .

After a and c , players 3 and 4 are in a subgame for which the unique equilibrium is for both players to use randomized strategies in which each action has probability $1/2$. When they use these strategies, player 2 gets an expected utility of 2 after a and c , just as in Example 3. It then follows that the only subgame-perfect equilibrium for Example 4 is the one in which player 1 chooses a and player 2 chooses d , with resulting payoffs $(2,3,0,0)$ for players 1 through 4. This $(2,3,0,0)$ outcome is also the unique trembling-hand perfect equilibrium and the unique sequential equilibrium for this example (since these are stronger solution concepts).

Now consider the following Nash equilibrium: player 1 chooses b , giving outcome $(1,9,0,0)$, because player 2 plans to choose c after a , player 3 plans to choose e and player 4 plans to randomly choose either g or h with equal probability if a and c occur. This equilibrium is not perfect or sequential because player 4's randomized strategy is not his best response to e . However, if the players can communicate then it is possible to make this imperfect Nash equilibrium into a perfect sequential equilibrium.

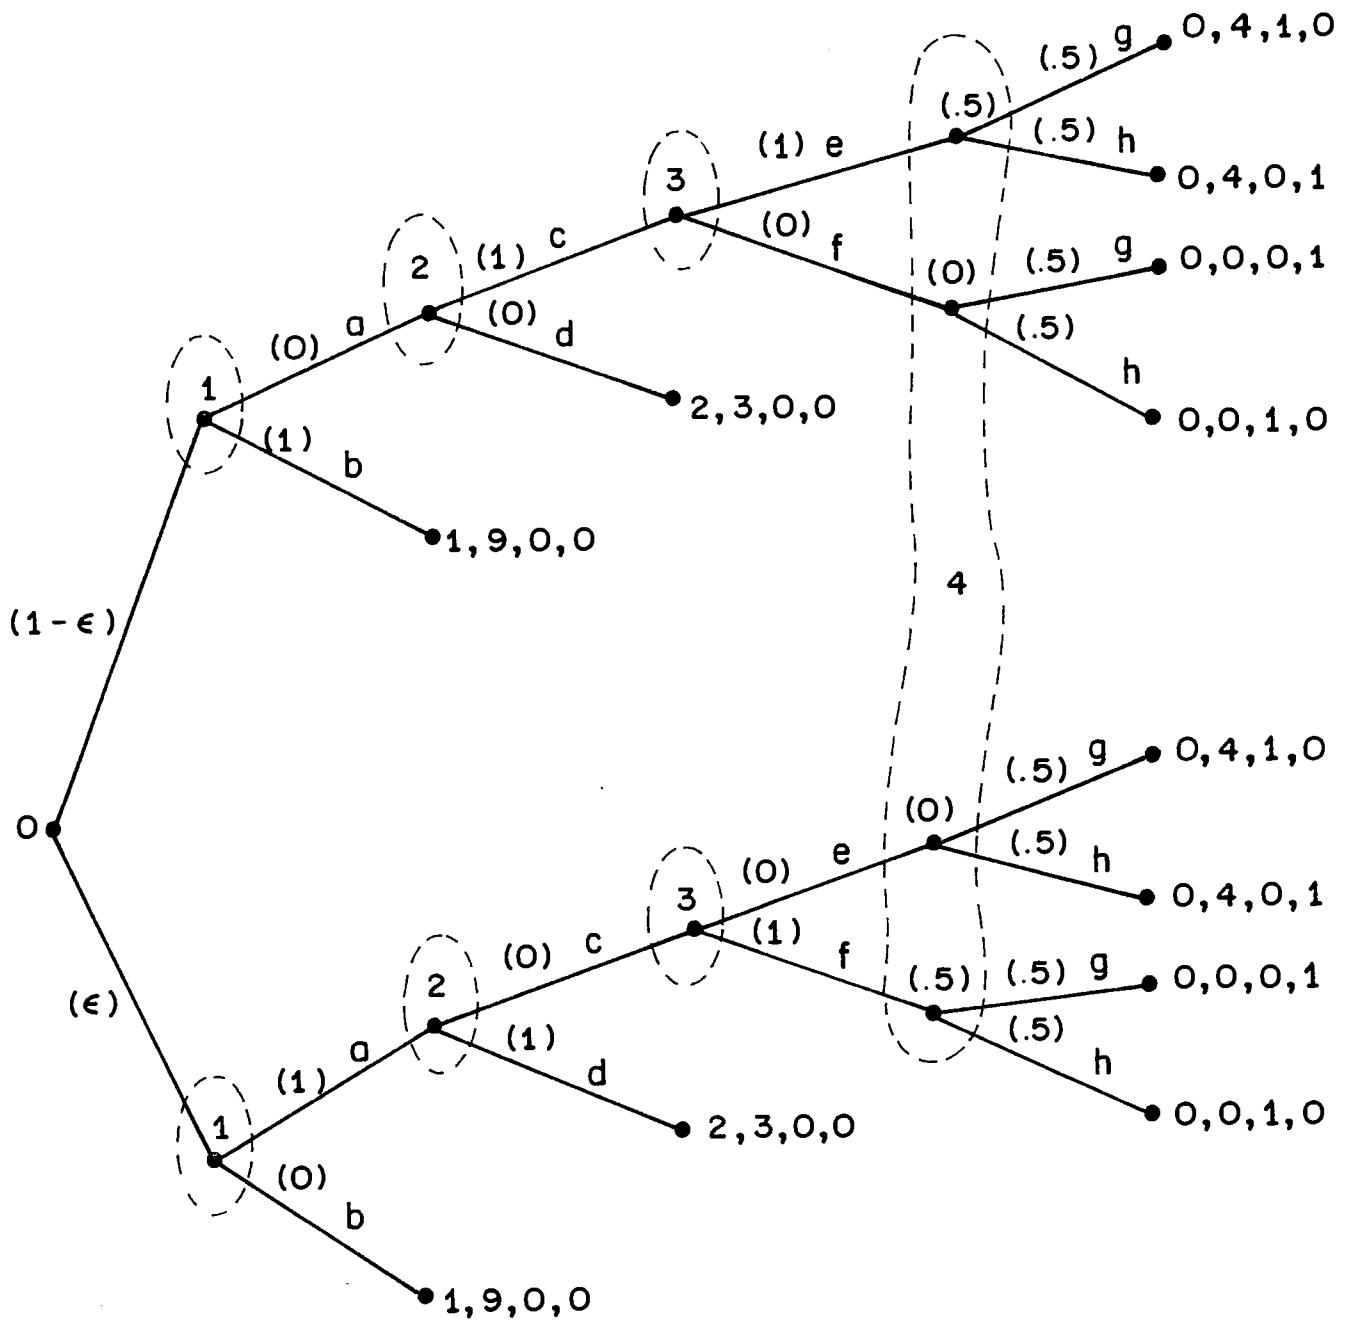


Figure 5

[Insert Figure 5 here.]

The essential idea is to convert Example 4 into Figure 5 by adding a payoff-irrelevant random event which is observed by players 1, 2, and 3, but is not observed by player 4. The event is either "Up", with probability $1 - \epsilon$, or "Down", with probability ϵ . In either case, the game after the random event is exactly as in Example 4, except that the actions of the three players who observe the event can be correlated with it.

The numbers in parentheses are probabilities forming a sequential equilibrium for this modified example. If the initial event is Up, then player 1 chooses b, 2 plans to choose c, and 3 plans to choose e if the opportunity to act arises. If the initial event is Down, then 1 chooses a, 2 chooses d, and 3 plans to choose f. Player 4 plans to randomly choose either action with equal probability. It is straightforward to check that these actions are rational for each of the first three players given the others' plans: if 3 would choose e then 2 would prefer c, and so 1 prefers b; if 3 would choose f then 2 prefers d and 1 prefers a; and player 3 is willing to choose either e or f if 4 is randomizing equally.

So, as before, it remains to show why player 4 might rationally choose to randomize his action. If player 4 found himself in a position to act, then he would know that either player 1 or player 2 must have made a mistake. Either the initial event was Up and 1 mistakenly chose a, in which case the play would be at the top node of player 4's information set, or the initial event was Down and 2 mistakenly chose c, in which case the play would be at the bottom node of 4's information set. Then it is consistent with the rules of

rational inference to suppose that player 4 might assign equal probability to these top and bottom nodes, and zero probability to the two middle nodes (as indicated in the figure) if he found himself in a position to act. Even if ϵ is very small, these beliefs are not irrational, as player 4 might believe that player 1 is much less likely to make a mistake than player 2. With these beliefs, player 4 would get an expected utility of .5 from either action, so he is willing to randomize. Now, as we let ϵ go to zero, the sequential equilibrium shown for Figure 5 gives the outcome (1,9,0,0) with probability one.

In Section 6, we develop a concept of sequential communication equilibrium, for multistage games with communication. When we consider Example 4 as a game with communication, we implicitly recognize that the players can transform the structure of information to that of Figure 5, by asking a mediator to do the initial randomization and communicate the results to players 1, 2, and 3. Thus, the imperfect equilibrium that gives outcome (1,9,0,0) should be (and is) a sequential communication equilibrium for Example 4, even though it is not a sequential equilibrium in the sense of Kreps and Wilson [1982]. In general, the set of sequential communication equilibria that are also Nash equilibria may be strictly larger than the set of Kreps-Wilson sequential equilibria. As Example 4 illustrates, even if communication is not actually needed to implement some given communication equilibrium (so that it is also a Nash equilibrium), the possibility of communication may make the equilibrium sequentially rational where it otherwise would not have been.

4. Basic definitions

In this paper, we analyze a general model of dynamic multistage games, in a form that is somewhat different from Kuhn's [1953] definition of the extensive form.

We let $N = \{1, \dots, n\}$ denote the set of players. We assume that the play of the game occurs in K sequential stages, which are numbered from 1 (first) to K (the last stage). We may refer to the end of the game, after all active play is finished, as stage $K+1$.

The overall structure of each stage is as follows. First, each player observes some signal, which may depend probabilistically on the actions and signals of earlier stages. Then the players have an opportunity to communicate with a mediator. Finally, each player must choose some action among the actions that are feasible for him.

For any player i and any stage k , we let C_i^k denote the set of actions that player i can choose among in stage k . Suppose that R_i^k denotes the set of possible signals that player i can observe at the beginning of stage k . Since we are assuming that the players have perfect recall, the additional information available to player i at the beginning of stage k that was not available at the beginning of stage $k-1$ is described by a point in T_i^k , where we let

$$T_i^1 = R_i^1, \quad T_i^{K+1} = C_i^K, \quad \text{and}$$

$$T_i^k = C_i^{k-1} \times R_i^k, \quad \forall k \in \{2, \dots, K\}.$$

(We define $T_i^{K+1} = C_i^K$ here because, at the end of the game, player i knows what action he chose at the last stage, in addition to everything that he knew

before that.)

We let T denote the subset of $\prod_{i=1}^n \prod_{k=1}^{K+1} T_i^k$ that consists of all

possible outcomes of the game. That is, T is the set of all possible states of all players' information at the end of the game. (We exclude from T impossible combinations of information states, such as when two players disagree about an event that they both have perfectly observed.) Given any vector $t = ((t_i^k)_{i=1}^n)_{k=1}^{K+1}$ in T , we may denote various subvectors of t as follows:

$$t^k = (t_1^k, \dots, t_n^k), \quad t_i = (t_i^1, \dots, t_i^{K+1}),$$

$$t^{\leq k} = (t^1, \dots, t^k), \quad t_i^{\leq k} = (t_i^1, \dots, t_i^k).$$

(Read $t^{\leq k}$ as "t up to k.") The sets of all such subvectors are denoted

$$T^k = \{t^k \mid t \in T\}, \quad T_i = \{t_i \mid t \in T\},$$

$$T^{\leq k} = \{t^{\leq k} \mid t \in T\}, \quad T_i^{\leq k} = \{t_i^{\leq k} \mid t \in T\}.$$

Thus, $T_i^{\leq k}$ is the set of all possible states of player i 's information at the beginning of stage k ; and $T^{\leq k}$ is the set of all possible states of all players' joint information at the beginning of stage k . We may refer to $T_i^{\leq k}$ as the set of possible types for player i at stage k .

Other related notation that we shall use is:

$$T_{-i}^{\leq k}(t_i^{\leq k}) = \{t_{-i}^{\leq k} = (t_j^{\leq k})_{j \neq i} \mid t^{\leq k} \in T^{\leq k}\} \quad \forall t_i^{\leq k} \in T_i^{\leq k},$$

$$T^{>k}(t^{\leq k}) = \{t^{>k} = (t^{k+1}, \dots, t^{K+1}) \mid t \in T\} \quad \forall t^{\leq k} \in T^{\leq k}.$$

That is, $T_{-i}^{\leq k}(t_i^{\leq k})$ is the set of all possible types for all players other than i at the beginning of stage k , if $t_i^{\leq k}$ is i 's type. When $t_i^{\leq k}$ in $T_i^{\leq k}$ and $t_{-i}^{\leq k}$ in $T_{-i}^{\leq k}(t_i^{\leq k})$ are given, then $t^{\leq k}$ is the vector

in $T^{\leq k}$ that is formed by merging these two vectors. We may write

$$t^{\leq k} = (t_{-i}^{\leq k}, t_i^{\leq k}).$$

Similarly, we let

$$C^k = \prod_{i=1}^n C_i^k, \quad C_i = \prod_{k=1}^K C_i^k, \quad C = \prod_{i=1}^n \prod_{k=1}^K C_i^k, \quad C_i^{\leq k} = \prod_{\ell=1}^k C_i^\ell.$$

Generally, $c^k = (c_1^k, \dots, c_n^k)$ denotes a vector in C^k , and $c_i = (c_i^1, \dots, c_i^K)$ denotes a vector in C_i , and so on. When a vector c in C is given, then c_i^k is the (i,k) -component of c , and $c_i^{\leq k}$ is the subvector (c_i^1, \dots, c_i^k) , and so on.

We let $p^1(t^1)$ denote the probability that t^1 in T^1 will be the state of all players' information at the beginning of the first stage of the game. For any $k \geq 1$, we let $p^{k+1}(t^{k+1} | c^k, t^{\leq k})$ denote the probability that t^{k+1} in T^{k+1} will be the vector of new information for the players at the beginning of stage $k+1$, if c^k in C^k is the vector of actions at stage k and $t^{\leq k}$ in $T^{\leq k}$ is the state of all information that the players have learned through the beginning of stage k . Of course, if $t^{k+1} = ((b_i^k, r_i^{k+1})_{i=1}^n)$ and $b_i^k \neq c_i^k$ for some i , then $p^{k+1}(t^{k+1} | c^k, t^{\leq k}) = 0$, since each player i knows his past action c_i^k as well as his new signal r_i^{k+1} in stage $k+1$. Similarly, at the end of the game,

$$p^{K+1}(t^{K+1} | c^K, t^{\leq K}) = \begin{cases} 1 & \text{if } t^{K+1} = c^K, \\ 0 & \text{if } t^{K+1} \neq c^K. \end{cases}$$

To justify our interpretation of T as the subset of $\prod_{i=1}^n \prod_{k=1}^{K+1} T_i^k$ that consists of all possible outcomes of the game, we assume that t is in T if and only if there exists some c in C such that

$$p^1(t^1) \cdot \prod_{k=1}^K p^{k+1}(t^{k+1} | c^k, t^{\leq k}) > 0.$$

Of course, we also have

$$\sum_{t^{k+1} \in T^{k+1}} p^{k+1}(t^{k+1} | c^k, t^{\leq k}) = 1, \quad \forall k \geq 1, \quad \forall c^k \in C^k, \quad \forall t^{\leq k} \in T^{\leq k};$$

and

$$\sum_{t^1 \in T^1} p^1(t^1) = 1.$$

The preferences of player i are characterized by a von Neumann-Morgenstern utility function on T ,

$$u_i: T \rightarrow \mathbb{R}.$$

Notice that the final information states in T record the action-choices in C as well as the signals in the sets R_i^k , so each player's payoff can depend on all signals and actions. These structures

$$\Gamma = (((C_i^k)_{k=1}^K, (T_i^k)_{k=1}^{K+1}, u_i)_{i=1}^n, T, (p^k)_{k=1}^{K+1})$$

complete our description of the multistage game Γ . We assume throughout this paper that the sets C_i^k and T_i^k are all finite.

(To keep the notation from being even more complicated, we have assumed that the set of actions available to a player at any given stage is independent of his type. However, none of the results in this paper depends on this assumption. Suppose instead that for each player i , the set of feasible actions were some function $\hat{C}_i^k(t_i^{\leq k})$ of his type $t_i^{\leq k}$. Given any such game, it is straightforward to construct an equivalent game in which the sets of feasible actions are independent of type. For any player i , at any stage k , if the sets $\hat{C}_i^k(t_i^{\leq k})$ all have the same number of actions, then it is only necessary to relabel the actions using the same set of labels C_i^k for all types; and there does not need to be any significance attached to the way in which actions for different types are identified in this common set of

labels. If the sets $\hat{C}_i^k(t_i^{\leq k})$ have different numbers of actions then we can make the numbers of actions equal by adding irrelevant duplicates of existing actions, with the same effect on all payoffs and observations, in the sets that have fewer actions.)

Suppose now that a mediator is helping the players to coordinate their actions. In each stage k , the mediator first asks each player i to report his new information in T_i^k , and then the mediator recommends some action in C_i^k to each player i . We assume that all players communicate confidentially with the mediator, so that no player directly observes the reports or recommendations of the other players.

Such a mediator should constrain each player i to choose his reports so that the reports up to each stage k form a vector in $T_i^{\leq k}$, which is the set of possible information states for player i up to the beginning of stage k .

(Recall that $T_i^{\leq k}$ may be a proper subset of $\prod_{\ell=1}^k T_i^\ell$.) Any sequence of reports outside of $T_i^{\leq k}$ would obviously include lies. (To effectively constrain a player to send reports that remain within this set, the mediator could first designate one member of the set and then announce that any report that is not in the set -- including the "report" of total silence -- will be interpreted as actually meaning this designated report.)

Thus, the set of possible sequences of reports that the mediator could get from the players in the first k stages is

$$\hat{T}^{\leq k} = \prod_{i=1}^n T_i^{\leq k}.$$

Notice that $T^{\leq k}$ may be a proper subset of $\hat{T}^{\leq k}$. If so, the mediator cannot constrain the players to send joint report-sequences in $T^{\leq k}$ without sometimes conveying information to them about each others' reports. We also let

$$\hat{T} = \hat{T}^{\leq K+1} = \prod_{i=1}^n T_i.$$

We let F denote the set of feedback rules that such a mediator could use to determine the recommended actions in each period, as a function of the given reports. Formally,

$$F = \{f = (f^k)_{k=1}^K \mid f^k: \hat{T}^{\leq k} \rightarrow C^k, \forall k\}.$$

So if the mediator uses the feedback rule f in F , then $f_i^k(t^{\leq k})$ is the recommended action for player i at stage k when $t^{\leq k}$ is the history of reports from the players to the mediator. (Here $f^k(t^{\leq k}) = (f_1^k(t^{\leq k}), \dots, f_n^k(t^{\leq k}))$.) We may write

$$c = f(t)$$

iff $c_i^k = f_i^k(t^{\leq k})$ for every player i and every stage k .

For any f in F and t in T , we let $P(t|f)$ denote the probability that t will be the final state of all players' information if the players coordinate their actions according to the rule f , being honest and obedient to the mediator at every stage. Thus,

$$P(t|f) = p^1(t^1) \cdot \prod_{k=1}^K p^{k+1}(t^{k+1} | f^k(t^{\leq k}), t^{\leq k}).$$

Notice that, by definition of T , for every t in T there exists some f in F such that $P(t|f) > 0$. We let $U_i(f)$ denote the expected utility payoff for player i if rule f is used, so that

$$U_i(f) = \sum_{t \in T} P(t|f) u_i(t).$$

Suppose now that a mediator is helping the players to coordinate their actions according to the rule f . In each stage k , the mediator first asks each player i to report the state of his new information in T_i^k , and then the mediator recommends to each player i that he should choose the action in C_i^k that is designated by f_i^k for the reported information states. We assume that all players communicate confidentially with the mediator, so that no player directly observes the reports or recommendations of the other players.

Any player can manipulate such a feedback rule by lying to the mediator or disobeying his recommendation. In stage k , player i could choose his report in T_i^k as any function τ_i^k of the mediator's past recommendations (in $C_i^{\leq k-1}$) and of i 's information from playing the game (in $T_i^{\leq k}$). Player i 's action in stage k could be any function γ_i^k of the mediator's recommendations through stage k (in $C_i^{\leq k}$) and of i 's information from playing (in $T_i^{\leq k}$). A manipulative strategy for player i is any pair $(\gamma_i, \tau_i) = ((\gamma_i^k)_{k=1}^K, (\tau_i^k)_{k=1}^K)$ where, for each k , γ_i^k maps $C_i^{\leq k} \times T_i^{\leq k}$ into T_i^k , τ_i^k maps $C_i^{\leq k-1} \times T_i^{\leq k}$ into C_i^k (so τ_i^1 maps T_i^1 into T_i^1), and τ_i could never send a sequence of reports outside of $T_i^{\leq K}$. We let M_i denote the set of all manipulative strategies for player i . That is,

$$(\gamma_i, \tau_i) \in M_i \quad \text{iff} \quad \gamma_i = (\gamma_i^k)_{k=1}^K, \quad \tau_i = (\tau_i^k)_{k=1}^K,$$

$$\gamma_i^k: C_i^{\leq k} \times T_i^{\leq k} \rightarrow C_i^k \quad \text{and} \quad \tau_i^k: C_i^{\leq k-1} \times T_i^{\leq k} \rightarrow T_i^k, \quad \forall k,$$

and

$$(\tau_i^\ell(c_i^{\leq \ell-1}, t_i^{\leq \ell}))_{\ell=1}^K \in T_i^{\leq K}, \quad \forall (c_i, t_i) \in C_i \times T_i.$$

For any (γ_i, τ_i) in M_i , we let $f \circ (\gamma_i, \tau_i)$ represent the feedback rule that would actually be implemented if the mediator tried to implement f but

player i used the manipulative strategy (γ_i, τ_i) while all other players were honest and obedient to the mediator. Thus, $f \circ (\gamma_i, \tau_i)$ is a feedback rule such that $c = (f \circ (\gamma_i, \tau_i))(t)$ iff there exist s_i in T_i and d_i in C_i such that, for all k ,

$$s_i^k = \tau_i^k(d_i^{\leq k-1}, t_i^{\leq k}), \quad d_i^k = f_i^k(t_{-i}^{\leq k}, s_i^{\leq k}), \quad c_i^k = \gamma_i^k(d_i^{\leq k}, t_i^{\leq k}),$$

and $c_j^k = f_j^k(t_{-j}^{\leq k}, s_j^{\leq k}) \quad \forall j \neq i.$

(Here, s_i^k would be player i 's report to the mediator and d_i^k would be the mediator's recommendation to player i in stage k .) We may call $f \circ (\gamma_i, \tau_i)$ the effective transformation of f by (γ_i, τ_i) .

For greatest generality, we allow that a mediator could be instructed to choose a feedback rule at random, according to any probability distribution over F , so that the players may not know which rule is being used. Thus, we say that a communication mechanism for the n players is any probability distribution μ over the set of decision rules F , where $\mu(f)$ denotes the probability that the mediator will use rule f .

For a communication mechanism to be feasible, it is necessary that no player should expect to gain by manipulating it when the others are not manipulating. Otherwise, the assumption that all players are participating honestly and obediently in the communication mechanism would be a self-denying prophecy. Thus, we say that a communication mechanism μ is a communication equilibrium iff, for every player i and every manipulative strategy (γ_i, τ_i) in M_i ,

$$(4.1) \quad \sum_{f \in F} \mu(f) U_i(f) \geq \sum_{f \in F} \mu(f) U_i(f \circ (\gamma_i, \tau_i)).$$

Notice that these incentive constraints (4.1) are linear inequalities in μ . Thus, the set of communication equilibria is a closed and convex subset of the set of communication mechanisms. The sets M_i and F may be enormously large in a multistage game, as a result of the combinatorial complexity that inevitably arises in dynamic games with communication, in which each player's action can be a function of anything that any player has observed earlier. In practice, therefore, the incentive constraints (4.1) may not be computationally tractable. Nevertheless, it is helpful to know that the set of communication equilibria has the simple mathematical structure of a compact polyhedron, defined by a finite collection of linear inequalities. The set of Nash equilibria of a game has no such structural simplicity.

The set of communication equilibria includes all of the Nash equilibria of the game, including irrational equilibria such as (b,c) in Example 3 above. Thus, we need a more restrictive solution concept to accurately characterize rational behavior in multistage games with communication. Ex ante, by (4.1), no player can expect to gain by planning to manipulate in a communication equilibrium. Furthermore, this implies that no player could expect to gain by manipulating after any event that is observable by him and that has positive probability of occurring (in the equilibrium). To strengthen our solution concept, we need to require that no player should ever expect to gain by manipulating after any possible event that is observable by him, including events that have zero probability in equilibrium. To make this restriction, we must first develop a theory of rational beliefs conditional on events of zero probability. In Section 5, we review and extend ideas of Kreps and Wilson [1982], to develop such a theory. Then, in Section 6, we return to define sequential communication equilibria for multistage games.

5. Conditional probability systems

To develop a theory of conditional probability systems, let us consider any nonempty finite set Ω . We may interpret Ω as the set of possible "states of the world." We let $\Delta(\Omega)$ denote the set of all probability distributions on Ω .

Given any distribution μ in $\Delta(\Omega)$, if $X \subseteq \Omega$ then $\mu(X)$ is the probability of the event X under the distribution μ . Suppose that a rational individual's beliefs about the unknown state in Ω were as given by the distribution μ , but he has now just received the additional information that the actual state is in the set Z , where $\Omega \supseteq Z \neq \emptyset$. If $\mu(Z) > 0$, then his conditional probability of the event X given Z , denoted $\mu(X|Z)$, should now be

$$(5.1) \quad \mu(X|Z) = \mu(X \cap Z) / \mu(Z).$$

On the other hand, if $\mu(Z) = 0$, then $\mu(X|Z)$ is not defined by the probability distribution μ . To define all conditional probabilities, we must construct a complete conditional probability system.

A conditional probability system is any function μ that specifies a nonnegative number $\mu(X|Z)$ for every X and Z such that $X \subseteq \Omega$ and $\emptyset \neq Z \subseteq \Omega$, and that satisfies the following three properties, for every X , Y , and Z such that $X \subseteq \Omega$, $Y \subseteq \Omega$, and $\emptyset \neq Z \subseteq \Omega$:

$$(5.2) \quad \text{if } X \cap Y = \emptyset \text{ then } \mu(X \cup Y|Z) = \mu(X|Z) + \mu(Y|Z);$$

$$(5.3) \quad \mu(\Omega|Z) = \mu(Z|Z) = 1;$$

$$(5.4) \quad \text{if } X \subseteq Y \subseteq Z \text{ and } Y \neq \emptyset \text{ then } \mu(X|Z) = \mu(X|Y) \mu(Y|Z).$$

Conditions (5.2) and (5.3) assert that $\mu(\cdot|Z)$ is a probability distribution over Ω that puts all probability weight on the given set Z . Condition (5.4) asserts that the conditional probabilities given Y are consistent with the conditional probabilities given Z . Notice that, if $\mu(Y|Z) > 0$ then the formula in (5.4) becomes

$$\mu(X|Y) = \mu(X|Z)/\mu(Y|Z);$$

so this equation asserts that the probability that an individual would assign to event X if event Y were known is equal to the probability that he would compute for X by Bayes formula if he learned that Y occurred when he already knew Z .

We let $\Delta^*(\Omega)$ denote the set of all conditional probability systems on Ω . Given any probability distribution η in $\Delta(\Omega)$, we say that a conditional probability system μ in $\Delta^*(\Omega)$ is an extension of η iff

$$\mu(X|\Omega) = \eta(X), \quad \forall X \subseteq \Omega.$$

One way to construct a conditional probability system on Ω is to start with a probability distribution that assigns positive probability to every point in Ω . If $\mu(Z) > 0$ for every nonempty set Z , then the conditional probabilities that are defined by equation (5.1) do satisfy (5.2)-(5.4), as is straightforward to check. In fact, every conditional probability system on Ω can be characterized as the limit of conditional probability systems that are constructed in this way.

Theorem 1. μ is a conditional probability system on Ω if and only if there exists a sequence of probability distributions $\{\eta^j\}_{j=1}^\infty$ such that

$$\eta^j(\{\omega\}) > 0, \quad \forall j, \quad \forall \omega \in \Omega; \quad \text{and}$$

$$\mu(X|Z) = \lim_{j \rightarrow \infty} \eta^j(X \cap Z) / \eta^j(Z), \quad \forall X, \quad \forall Z \neq \emptyset.$$

Proof. The proof is deferred to Section 9.

Theorem 1 explains the role of small mistakes or "trembling hands" in Selten's [1975] theory of perfect equilibrium. Every outcome has positive probability in each equilibrium of Selten's perturbed games, with small probabilities of mistakes. Then as the probabilities of mistakes go to zero, a limit of these equilibria generates a conditional probability system on the set of outcomes of the game, as in Theorem 1. Each agent's strategy in the limiting equilibrium is rational for him conditional on any event that he may observe in the game, if his beliefs in that event are as specified by the conditional probability system.

In a game with communication, there is no clear reason why one should assume that players' mistakes must be stochastically independent of each other. Thus we omit here the assumption that players tremble independently, which Kreps and Wilson [1982] used to restrict the class of permissible beliefs in their definition of sequential equilibrium. Nevertheless, Theorem 1 implies that, whenever we speak of conditional probability systems in the next section, we could equivalently speak of limits of perturbed games with small probabilities of players' mistakes, in terms similar to those of Selten [1975] and Kreps and Wilson [1982].

6. Sequential communication equilibria

To develop the concept of sequential communication equilibrium, we first need more definitions.

Let $P(t^{>k} | f, t^{<k})$ denote the probability that the information revealed in the game after stage k would be as in $t^{>k}$ if the players used the feedback rule f . That is,

$$P(t^{>k} | f, t^{<k}) = \prod_{\lambda=k}^K p^{\lambda+1}(t^{\lambda+1} | f^{\lambda}(t^{<\lambda}), t^{<\lambda}).$$

Let $U_i(f | t^{<k})$ denote the conditional expected utility for player i from rule f , given that $t^{<k}$ is the state of all players' information in stage k , so that

$$U_i(f | t^{<k}) = \sum_{t^{>k} \in T^{>k}(t^{<k})} P(t^{>k} | f, t^{<k}) u_i(t).$$

(Here $t = (t^{<k}, t^{>k})$.)

We let M_i^k denote the set of manipulative strategies for player i in which he is honest and obedient in all stages before stage k , so that

$$M_i^k = \left\{ (\gamma_i, \tau_i) \in M_i \mid \begin{array}{l} \forall \lambda < k, \forall c, \forall t, \gamma_i^{\lambda}(c_i^{<\lambda}, t_i^{<\lambda}) = c_i^{\lambda} \\ \text{and } \tau_i^{\lambda}(c_i^{<\lambda-1}, t_i^{<\lambda}) = t_i^{\lambda} \end{array} \right\}.$$

We let M_i^{*k} denote the set of manipulative strategies for player i in which he is honest and obedient in all stages before stage k , and he also reports his type honestly at the beginning of stage k ,

$$M_i^{*k} = \{ (\gamma_i, \tau_i) \in M_i^k \mid \forall c, \forall t, \tau_i^k(c_i^{<k-1}, t_i^{<k}) = t_i^k \}.$$

Thus, if player i decided to begin manipulating for the first time at the beginning of stage k , then he would have to choose a manipulative strategy in M_i^k . If he decided to begin manipulating for the first time at the end of stage k , after he has reported his information in T_i^k honestly, but before he chooses his action in C_i^k , then he would have to choose a manipulative strategy in M_i^{*k} .

We want to guarantee that it should always be rational for each player to obey the mediator's recommendations, even if the player has mistakenly disobeyed (or trembled) in the past. However, there are some cases in which a particular action could never be rationally chosen by a player. For example, if the mediator in Example 3 ever asked player 2 to use his dominated action "c", then player 2 would certainly prefer to disobey. In general, it may be necessary to impose some restrictions on the sequences of actions that the mediator can recommend to each player.

To understand the need for such restrictions, we must reconsider the argument for the revelation principle. There is no loss of generality in assuming that each player reports all of his new information at each stage and then receives in return only the recommendation of an action, because such communication systems maximize the mediator's information and minimize the player's information (and hence minimize his opportunities to find profitable ways to cheat). So, without loss of generality, we can assume that the vocabulary in which player i can report to the mediator at stage k is a subset of T_i^k , and the vocabulary in which the mediator speaks to player i is a subset of C_i^k . In fact, we argued in Section 4 that the mediator could require that player i must send his reports so that, at each stage k , the sequence of past reports should form a sequence in $T_i^{\leq k}$; but no smaller set of possible reports could be uninformatively specified, because the player's true information

could be anywhere in this set. Now we must ask, can we assume without any loss of generality that the set of possible recommendations that the mediator can send to player i at stage k must always be all of C_i^k ? Unfortunately, the answer to this question may be No, if we want to require that honesty and obedience should be rational for every player in any event that he can perceive. The problem is that, when we decrease the set of possible recommendations that the mediator could send to player i , we decrease the set of possible events that i could perceive, and this may make it easier to guarantee that honesty and obedience is rational in all such events. Thus, for maximum generality, we must allow that the mediator might, in some circumstances, restrict the set of actions that he could possibly recommend to a player, if the result of this restriction is to decrease the set of events that are considered possible for the player to perceive.

Since the goal of these restrictions on the mediator is only to reduce the set of events that a player could perceive, there is no loss of generality in assuming that the range of recommendations that the mediator can send to player i at stage k depends only on the communications between the mediator and player i up to stage k . Thus, we let a mediation range Q be a function that specifies, for each player i , each stage k , each type $t_i^{\leq k}$ in $T_i^{\leq k}$, and each $c_i^{\leq k-1}$ in $C_i^{\leq k-1}$, a set $Q(c_i^{\leq k-1}, t_i^{\leq k})$ such that

$$Q(c_i^{\leq k-1}, t_i^{\leq k}) \subseteq C_i^k.$$

We interpret $Q(c_i^{\leq k-1}, t_i^{\leq k})$ as the set of all actions that the mediator could possibly recommend to player i at the end of stage k if $t_i^{\leq k}$ were the vector of past reports from player i and $c_i^{\leq k-1}$ were the vector of past recommendations to player i . Given such a mediation range Q , we define

$$Q_i = \{(c_i, t_i) \in C_i \times T_i \mid c_i^k \in Q(c_i^{\leq k-1}, t_i^{\leq k}), \forall k\},$$

$$Q_i^k = \{(c_i^{\leq k-1}, t_i^{\leq k}) \mid (c_i, t_i) \in Q_i\},$$

$$Q_i^{*k} = \{(c_i^{\leq k}, t_i^{\leq k}) \mid (c_i, t_i) \in Q_i\},$$

and

$$G(Q) = \{f \in F \mid f_i^k(t^{\leq k}) \in Q(f_i^{\leq k-1}(t^{\leq k-1}), t_i^{\leq k}), \forall i, \forall k, \forall t \in \hat{T}\}.$$

(Here $f_i^{\leq k}(t^{\leq k}) = c_i^{\leq k}$ iff $c_i^{\ell} = f_i^{\ell}(t^{\leq \ell})$ for every $\ell \leq k$.)) Thus, Q_i is the set of all possible sequences of recommendations and type-reports for player i when the mediator is restricted by the mediation range Q .

Similarly, Q_i^k [or Q_i^{*k}] is the set of all possible sequences of type-reports and recommendations for player i up to the beginning [or end] of stage k if the mediator is restricted by Q . $G(Q)$ is the set of all feedback rules that satisfy the restriction imposed by the mediation range Q for all permissible sequences of reports from the players. (Notice that t ranges over \hat{T} in the definition of $G(Q)$.)

We are now ready to state the main definition of this paper.

A communication mechanism μ in $\Delta(F)$ is a sequential communication equilibrium (or, more fully, a sequentially rational communication equilibrium) iff there exists a mediation range Q and a conditional probability system $\bar{\mu}$ in $\Delta^*(G(Q) \times T)$ such that:

$$(6.1) \quad \bar{\mu}(f, t) = \mu(f) P(t|f), \quad \forall f \in G(Q), \forall t \in T;$$

$$(6.2) \quad \bar{\mu}(f, t|f^{\leq k}, t^{\leq k}) = \bar{\mu}(f|f^{\leq k}, t^{\leq k}) P(t^{\leq k}|f, t^{\leq k}), \quad \forall f \in G(Q), \forall t \in T, \forall k;$$

$$\begin{aligned}
 (6.3) \quad & \sum_{t_{-i}^{\leq k} \in T_{-i}^{\leq k}(t_i^{\leq k})} \sum_{f \in G(Q)} \bar{\mu}(f, t_{-i}^{\leq k} | c_i^{\leq k-1}, t_i^{\leq k}) U_i(f | t^{\leq k}) \\
 & \geq \sum_{t_{-i}^{\leq k} \in T_{-i}^{\leq k}(t_i^{\leq k})} \sum_{f \in G(Q)} \bar{\mu}(f, t_{-i}^{\leq k} | c_i^{\leq k-1}, t_i^{\leq k}) U_i(f \circ (\gamma_i, \tau_i) | t^{\leq k}), \\
 & \forall i, \forall k, \forall (c_i^{\leq k-1}, t_i^{\leq k}) \in Q_i^k, \forall (\gamma_i, \tau_i) \in M_i^k;
 \end{aligned}$$

and

$$\begin{aligned}
 (6.4) \quad & \sum_{t_{-i}^{\leq k} \in T_{-i}^{\leq k}(t_i^{\leq k})} \sum_{f \in G(Q)} \bar{\mu}(f, t_{-i}^{\leq k} | c_i^{\leq k}, t_i^{\leq k}) U_i(f | t^{\leq k}) \\
 & \geq \sum_{t_{-i}^{\leq k} \in T_{-i}^{\leq k}(t_i^{\leq k})} \sum_{f \in G(Q)} \bar{\mu}(f, t_{-i}^{\leq k} | c_i^{\leq k}, t_i^{\leq k}) U_i(f \circ (\gamma_i, \tau_i) | t^{\leq k}), \\
 & \forall i, \forall k, \forall (c_i^{\leq k}, t_i^{\leq k}) \in Q_i^{*k}, \forall (\gamma_i, \tau_i) \in M_i^{*k}.
 \end{aligned}$$

(Here we use the following notation:

$$\bar{\mu}(f, t) = \bar{\mu}(\{(f, t)\} | G(Q) \times T);$$

$$\bar{\mu}(f, t | f^{\leq k}, t^{\leq k}) = \bar{\mu}(\{(f, t)\} | \{(g, s) | \forall \lambda \leq k, g^\lambda = f^\lambda, s^\lambda = t^\lambda\});$$

$$\bar{\mu}(f | f^{\leq k}, t^{\leq k}) = \bar{\mu}(\{f\} \times T | \{(g, s) | \forall \lambda \leq k, g^\lambda = f^\lambda, s^\lambda = t^\lambda\});$$

and

$$\bar{\mu}(f, t_{-i}^{\leq k} | c_i^{\leq \lambda}, t_i^{\leq k}) =$$

$$= \bar{\mu}(\{(f, s) | s_{-i}^{\leq k} = t_{-i}^{\leq k}\} | \{(g, s) | (g_i^m(s^{\leq m}))_{m=1}^\lambda = c_i^{\leq \lambda}, s_i^{\leq k} = t_i^{\leq k}\});$$

where each indicated set is a subset of $G(Q) \times T$.)

Condition (6.1) asserts that the conditional probability system $\bar{\mu}$ is consistent with μ and the given dynamics of the game in all events that have positive probability ex ante. Notice that (6.1) implies that $\mu(f) = 0$ for any f that is not in $G(Q)$, since

$$\sum_{f \in G(Q)} \mu(f) = \sum_{f \in G(Q)} \sum_{t \in T} \bar{\mu}(f, t) = 1.$$

In any stage k , the players jointly know the state $t^{\leq k}$ and the mediator knows the components of the feedback rule $f^{\leq k} = (f^{\lambda})_{\lambda=1}^k$ that he has been using. Condition (6.2) asserts that, given all this information in stage k , it is expected that the future states $t^{>k}$ will depend on the feedback rule as required by the given dynamics of the game (when all players are honest and obedient). Condition (6.2) also implies that future actions by the players are not expected to influence the feedback rule.

Condition (6.3) contains the informational or adverse-selection incentive constraints for each stage. It asserts that, at the beginning of each stage k , given any history of recommendations and observations in Q_1^k , player i should not expect to gain by starting to use a manipulative strategy of lying and disobedience if he has never lied to the mediator before. Similarly, condition (6.4) contains the strategic or moral-hazard incentive constraints for each stage. It asserts that, at the end of stage k , given any history of recommendations and observations in Q_1^{*k} , player i should not expect to gain by starting to use a manipulative strategy if he has never lied before. Thus, conditions (6.3) and (6.4) guarantee that, for any event that player i could perceive in the course of the game when the mediator restricts his recommendations to those permitted by Q , it is always rational for player i to be honest and obedient if he has never lied before. That is, if the mediator

restricts himself to feedback rules in $G(Q)$, the beliefs specified by $\bar{\mu}$ guarantee that no player i would ever want to disobey or lie to the mediator in the sequential communication equilibrium, even after zero-probability events in which some players (not including i) have lied and some players (possibly including i) have disobeyed the mediator in the past.

In the definition of sequential communication equilibrium, we have not bothered to specify what a player would do or believe if he found that he had (accidentally) lied to the mediator at some previous stage. We can ignore such situations because no player could ever observe anything that would prove to him that another player had lied. Dishonesty is fundamentally different from disobedience in this respect, because we allow that players can directly observe each other's actions (in that t_i^k may depend on c_i^{k-1}). That is, if player j at stage k observed that player i used c_i^{k-1} at the previous stage, and if c_i^{k-1} is never in the mediation range, then player j would know that player i had disobeyed the mediator. Under these circumstances, the future behavior and beliefs of player i would be relevant to player j and so must be described in a sequential communication equilibrium. On the other hand, there is no event that a player could perceive that could only be explained by some other player having lied, because players do not directly observe each other's reports to the mediator. Thus, there is nothing to prevent us from assuming that every player always assigns probability zero to the event that any other players have lied to the mediator. Under this assumption, no player ever cares about what a dishonest player would do or believe.

This begs the question of whether we could get a larger set of sequentially rational communication equilibria if we allowed players to assign positive probability to the event that others have lied to the mediator.

Fortunately, by the revelation principle, this set would not be any larger. Given any mechanism in which a player lies to the mediator with positive probability after some event, there is an equivalent mechanism in which the player does not lie and the mediator makes recommendations exactly as if the player had lied in the given mechanism.

Existence of sequential communication equilibria is easy to verify (for finite games) because any sequential equilibrium in the sense of Kreps and Wilson [1982] is a sequential communication equilibrium. Also, it is easy to check that (6.1)-(6.4) imply (4.1), so every sequential communication equilibrium is indeed a communication equilibrium.

7. Codominated actions

The definition of sequential communication equilibrium in Section 6 is much more complicated than the definition of communication equilibrium in Section 4. In this section we show that the set of sequential communication equilibria of a game may be actually quite easy to characterize, using a new concept of codomination, which is closely related to more familiar notions of domination of strategies in games.

Let B be a correspondence that specifies sets $B(t_i^{\leq k})$ such that

$$(7.1) \quad B(t_i^{\leq k}) \subseteq C_i^k, \quad \forall i, \forall k, \forall t_i^{\leq k} \in T_i^{\leq k}.$$

Given any such correspondence we define

$$(7.2) \quad E^k(B) = \{f \in F \mid \forall i, \forall \ell > k, \forall t_i^{\leq \ell} \in \hat{T}_i^{\leq \ell}, f_i^\ell(t_i^{\leq \ell}) \notin B(t_i^{\leq \ell})\}.$$

That is, $E^k(B)$ is the set of all feedback rules that would never recommend actions in $B(t_i^{\leq \ell})$ to any type $t_i^{\leq \ell}$ in of any player i at any stage ℓ that is after stage k , even if players other than i lied. (Notice the $\hat{T}^{\leq \ell}$ in (7.2).)

For any c_i^k in C_i^k and any $s_i^{\leq k}$ in $T_i^{\leq k}$, we define

$$\phi^k(c_i^k, s_i^{\leq k}) = \{(f, t^{\leq k}) \in F \times T^{\leq k} \mid t_i^{\leq k} = s_i^{\leq k}, f_i^k(t^{\leq k}) = c_i^k\}.$$

That is, $\phi^k(c_i^k, s_i^{\leq k})$ is the set of pairs consisting of a feedback rule and a stage- k state such that player i 's type is $s_i^{\leq k}$, and player i 's recommended action is c_i^k . We let $\phi^k(B)$ denote the union of all sets $\phi^k(c_i^k, t_i^{\leq k})$ over all $i, t_i^{\leq k}$, and c_i^k in $B(t_i^{\leq k})$. That is

$$\phi^k(B) = \{(f, t^{\leq k}) \in F \times T^{\leq k} \mid \exists i \text{ such that } f_i^k(t^{\leq k}) \in B(t_i^{\leq k})\}.$$

We say that B is a (sequential) codomination correspondence iff, for every stage k and every probability distribution π in $\Delta(F \times T^{\leq k})$, if

$$(7.3) \quad \pi(E^k(B) \times T^{\leq k}) = 1 \quad \text{and} \quad \pi(\phi^k(B)) > 0$$

then there exists some player i , some $t_i^{\leq k}$ in $T_i^{\leq k}$, some c_i^k in $B(t_i^{\leq k})$, and some (γ_i, τ_i) in M_i^{*k} such that

$$(7.4) \quad \sum_{(f, t^{\leq k}) \in \phi^k(c_i^k, t_i^{\leq k})} \pi(f, t^{\leq k}) (U_i(f \mid t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i) \mid t^{\leq k})) < 0.$$

That is, if there is a positive probability that some players may be asked to use codominated actions in stage k , but no players would ever be asked to use codominated actions after stage k , then at least one player should expect to be able to gain by planning to manipulate in the event that he is told to use an action in B . (A related concept of weak codomination is introduced

elsewhere, by Myerson [1985]. The concept of codomination used here may be called sequential codomination whenever it is necessary to distinguish it from weak codomination.)

If B and \hat{B} are two codomination correspondences, then $B \cup \hat{B}$ is also a codomination correspondence. (Here $(B \cup \hat{B})(t_i^{\leq k}) = B(t_i^{\leq k}) \cup \hat{B}(t_i^{\leq k})$.) This is because, if μ satisfies (7.3) for $B \cup \hat{B}$ (in place of B), then μ must satisfy (7.3) for B or \hat{B} , so that μ must satisfy (7.4) in at least one case where c_i^k is in $B(t_i^{\leq k})$ or $\hat{B}(t_i^{\leq k})$.

With finite type sets and action sets, there can be at most finitely many codomination correspondences. So let D be the union of all codomination correspondences. Thus, D is the maximal codomination correspondence, containing all others. In general, we may say that an action c_i^k is codominated for type $t_i^{\leq k}$ of player i at stage k iff $c_i^k \in D(t_i^{\leq k})$.

When an action for player i is dominated in the sense of Nash [1951] or Luce and Raiffa [1957], it means that, in any mechanism that would ask i to use that action with positive probability, player i could expect to gain by planning to disobey after being told to use that action. The idea of codomination is that, in any mechanism that would recommend one or more codominated actions with positive probability, at least one player could expect to gain by planning to manipulate after being told to use a codominated action (but, for a different mechanism, it might be a different player or a different action). Thus, dominated actions are codominated.

Our main result is that a communication equilibrium (satisfying (4.1)) is sequentially rational if and only if the mediator would never recommend a codominated action to any player in any event. In our notation (from (7.2)), $E^0(D)$ denotes the set of all feedback rules that would never recommend a codominated action to any player.

Theorem 2. A communication equilibrium μ is a sequential communication equilibrium if and only if

$$(7.5) \quad \mu(E^0(D)) = 1.$$

Proof. The proof is deferred to Section 9.

By Theorem 2, once D is known, it is easy to check whether a mechanism μ is a sequential communication equilibrium, because it suffices to verify the ex ante incentive constraints (4.1) and the support condition (7.5). Equation (7.5) is satisfied if and only if μ assigns zero probability to every feedback rule outside of $E^0(D)$. Since (4.1) and (7.5) are both linear in μ , Theorem 2 implies that the set of sequential communication equilibria is convex.

Furthermore, when verifying that a mechanism is a communication equilibrium, it is actually unnecessary to check incentive constraints that involve disobedience to codominated actions. To express this result formally as a theorem, let L_i denote the set of all manipulative strategies for player i in which player i never uses a codominated action; that is,

$$L_i = \{(\gamma_i, \tau_i) \in M_i \mid \gamma_i^k(c_i^{\leq k-1}, t_i^{\leq k}) \notin D(t_i^{\leq k}), \quad \forall c_i \in C_i, \forall t_i \in T_i, \forall k\}.$$

Theorem 3. For any mechanism μ in $\Delta(F)$, if $\mu(E^0(D)) = 1$ and μ satisfies (4.1) for every player i and every (γ_i, τ_i) in L_i , then μ is a sequential communication equilibrium.

Proof. The proof is given in Section 9.

Let \bar{Q} be the mediation range consisting of all actions that are not codominated. That is

$$\bar{Q}(c_i^{\leq k-1}, t_i^{\leq k}) = C_i^k \setminus D(t_i^{\leq k}), \quad \forall i, \forall k, \quad \forall c_i^{\leq k-1} \in C_i^{\leq k-1}, \quad \forall t_i^{\leq k} \in T_i^{\leq k}.$$

Thus $G(\bar{Q}) = E^0(D)$. (Notice that $\bar{Q}(c_i^{\leq k-1}, t_i^{\leq k})$ actually depends only on the type $t_i^{\leq k}$, not on the history of past recommendations $c_i^{\leq k-1}$.) As in Section 6, \bar{Q}_i denotes the set of all possible histories of recommendations and type-reports for player i when the mediator is restricted to the mediation range \bar{Q} . In the proof of Theorem 2, we shall also show the following result.

Corollary 1. If conditions (6.1)-(6.4) can be satisfied with the mediation range Q for some sequential communication equilibrium, then $Q_i \subseteq \bar{Q}_i$ for every player i . Furthermore, conditions (6.1)-(6.4) can be satisfied with the mediation range \bar{Q} for every sequential communication equilibrium.

To get a practical method for finding codominated actions, we need a bit more notation. We let A^* be the set of all functions α that specify a nonnegative number $\alpha(\gamma_i, \tau_i | c_i^k, t_i^{\leq k})$ for every player i , stage k , (γ_i, τ_i) in M_i^{*k} , c_i^k in C_i^k , and $t_i^{\leq k}$ in $T_i^{\leq k}$. That is,

$$A^* = \prod_{k=1}^K \prod_{i=1}^n (M_i^{*k} \times C_i^k \times T_i^{\leq k}) \times \mathbb{R}_+.$$

We interpret $\alpha(\gamma_i, \tau_i | c_i^k, t_i^{\leq k})$ as a shadow price for the strategic incentive constraint (6.4) that says that player i should not expect to gain by beginning to use the manipulative strategy (γ_i, τ_i) when his informational type is $t_i^{\leq k}$ and the mediator has just recommended the action c_i^k .

We define the function $V^k: F \times T^{\leq k} \times A^* \rightarrow \mathbb{R}$ by

$$(7.6) \quad V^k(f, t^{\leq k}, \alpha) = \sum_{i=1}^n \sum_{(\gamma_i, \tau_i) \in M_i^{*k}} \alpha(\gamma_i, \tau_i | f_i^k(t^{\leq k}), t_i^{\leq k}) (U_i(f | t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i) | t^{\leq k})).$$

That is, $V^k(f, t^{\leq k}, \alpha)$ is a weighted sum of the contributions that f can make to the satisfaction of the incentive constraints (6.4) at $t^{\leq k}$. We may refer to $V^k(f, t^{\leq k}, \alpha)$ as the aggregate incentive value of the feedback rule f , k at the information state $t^{\leq k}$, with respect to the shadow prices α .

We say that (B, α) is a (sequential) codomination system iff α is in A^* , B is a correspondence satisfying (7.1) such that

$$(7.7) \quad \forall i, \forall k, \forall t_i^{\leq k} \in T_i^{\leq k}, \forall c_i^k \in C_i^k, \forall (\gamma_i, \tau_i) \in M_i^{*k},$$

$$\text{if } c_i^k \notin B(t_i^{\leq k}) \text{ then } \alpha(\gamma_i, \tau_i | c_i^k, t_i^{\leq k}) = 0;$$

and

$$(7.8) \quad \forall i, \forall k, \forall t_i^{\leq k} \in T_i^{\leq k}, \forall f \in E^k(D),$$

$$\text{if } f_i^k(t_i^{\leq k}) \in B(t_i^{\leq k}) \text{ then } V^k(f, t_i^{\leq k}, \alpha) < 0.$$

Condition (7.7) asserts that the α shadow prices are positive only for potential manipulations beginning when a player is asked to use a codominated action. Condition (7.8) asserts that, if the rule f would recommend an action in B to some player in state $t^{\leq k}$ but f would never recommend any actions in B after stage k , then the aggregate incentive value of f at $t^{\leq k}$ must be strictly negative.

Theorem 4. B is a codomination correspondence if and only if there exists some α in A^* such that (B, α) is a codomination system.

Proof. The proof is given in Section 9.

This result gives us a method to computationally show that actions are codominated. One notable aspect of this result is that it involves only the strategic incentive constraints (6.4); the informational incentive

constraints (6.3) after the first stage do not play any role. To understand this asymmetry, notice that it is easy to construct communication equilibria in which every player would be willing to report honestly in every information state. For example, any Kreps-Wilson sequential equilibrium of the game without communication can be reinterpreted as such a communication equilibrium, in which the mediator simply recommends to each player that he should use the actions designated for him in the Kreps-Wilson equilibrium on the basis of his own reported information. Since each player's reports affect only his own recommendations, he could never gain by lying. On the other hand, there are recommendations which a player would be unwilling to obey, no matter what communication mechanism is being implemented. For example, we have remarked that player 2 would not be willing to obey a recommendation to use his dominated action "c" in Example 3. Thus, there are codominated recommendations to which obedience would be impossible to motivate; but there are no information states in which honesty would be impossible to motivate.

To show that codomination actually eliminates more actions than other concepts of domination, consider Example 5, defined as follows. (This example is adapted from an example proposed by J. Farrell.) In the first stage, player 1 chooses either w_1 or x_1 . In the second stage, players 2 and 3 are informed as to what player 1's first-stage choice was, and then each player i in $\{2,3\}$ can choose x_i , y_i , or z_i . If player 1 chose w_1 in the first stage, then the final payoffs are $(u_1, u_2, u_3) = (1, 9, 9)$, no matter what happens in the second stage. If player 1 chose x_1 in the first stage, then the final payoffs depend on the actions of players 2 and 3 in the second stage as follows:

Player 3

After x_1 :

		x_3	y_3	z_3
Player 2	x_2	2, 1, 1	0, 2, 0	0, 2, 0
	y_2	0, 0, 2	0, 3, 0	0, 0, 3
	z_2	0, 0, 2	0, 0, 3	0, 3, 0

There are no dominated actions in this game, in the sense of Nash [1951], but y_i and z_i are codominated for $i=2$ and $i=3$ in the second stage after player 1 chooses x_1 . Furthermore, if $y_2, y_3, z_2,$ and z_3 are codominated, then w_1 is codominated for player 1 in the first stage (since he gets a payoff of 2 from x_1 followed by x_2 and x_3 , whereas he gets a payoff of 1 from w_1). Thus, the unique sequential communication equilibrium has player 1 choosing x_1 in the first stage, and players 2 and 3 choosing (x_2, x_3) in the second stage. There are other Nash equilibria, in which player 1 chooses w_1 in the first stage because players 2 and 3 would each choose y_i or z_i if he chose x_1 , but these are not sequential communication equilibria.

To check that y_i and z_i are codominated actions for players 2 and 3 after player 1 chooses x_1 , let $\alpha(x_i | y_i, "x_1") = \alpha(x_i | z_i, "x_1") = 1$ for $i = 2$ and $i = 3$. (Here, $\alpha(x_i | y_i, "x_1")$ is the shadow price for the incentive constraint that player i should not expect to gain in stage 2 by choosing x_i if y_i is recommended, when his type is $t_i^{\leq 2} = "x_1"$; that is, he knows that player 1 chose x_1 .) Then the aggregate incentive values equal -1 at each of the eight cells other than (x_2, x_3) in the payoff matrix after x_1 . For example, at (y_2, y_3) we get $V^2 = 0 + (3 - 2) + (0 - 2) = -1$, and at (x_2, z_3) we get $V^2 = 0 + 0 + (0 - 1) = -1$.

It is worth remarking here that there are no codominated actions in Example 4. Thus, although the sets of sequential communication equilibria and

subgame-perfect equilibria coincide for Example 5, they do not coincide for Example 4.

8. Predominant actions and equilibria

In Example 5, we first established that actions y_2, z_2, y_3, z_3 are codominated actions for players 2 and 3 at stage 2. Then, having eliminated all actions after x_1 except x_2 and x_3 , we could conclude that w_1 is codominated (even dominated, in fact) at stage 1. Without the elimination of the y_i and z_i actions at the second stage, no clear comparison between x_1 and w_1 could have been made.

Consider now Example 6, which differs from Example 5 only in that player 1 now has a third action y_1 available to him in stage 1. If player 1 chooses action y_1 then players 2 and 3 get the same observation at stage 2 as if x_1 had been chosen; that is, players 2 and 3 at stage 2 cannot distinguish between actions x_1 and y_1 (but they can distinguish w_1). The final payoffs (u_1, u_2, u_3) if player 1 chooses y_1 depend on the actions of players 2 and 3 as follows

		Player 3		
		x_3	y_3	z_3
<u>After y_1:</u>				
	x_2	0,0,0	0,0,0	0,0,0
Player 2	y_2	0,0,0	0,1,1	0,0,0
	z_2	0,0,0	0,0,0	0,0,0

In Example 6, action y_1 is dominated by w_1 for player 1, since w_1 gives him a payoff of 1 for sure and y_1 always gives him a payoff of 0. Thus y_1 is

a codominated action. Unfortunately, y_1 is also the only codominated action in Example 6. In fact, there is a sequential equilibrium in which player 1 chooses w_1 because he expects that players 2 and 3 would choose y_2 and y_3 if he chose x_1 or y_1 ; and players 2 and 3 would be willing to choose y_2 and y_3 because, if they saw that player 1 did not choose w_1 then they would assign probability one to the event that he chose y_1 .

Thus, identifying codominated actions at stage 2 can help to identify more codominated actions at stage 1, as in Example 5; but, as in Example 6, identifying codominated actions at stage 1 does not help to identify more codominated actions at stage 2. In general, codominated actions are identified by an analysis that begins with the last stage of the game and works backwards. In terms of the definition of codomination, this backwards analysis is derived from the appearance of $E^k(B)$ in condition (7.3). More fundamentally, it arises in our theory because players at any stage must always assign probability one to the event that everyone will obey the mediator at all future stages. Thus, if the mediator can never recommend codominated actions, then all players are sure that no codominated actions will be used in later stages. On the other hand, after an event of probability zero, players at some given stage may assign positive probability to the event that some players disobeyed the mediator and used codominated actions at earlier stages in the game. This asymmetry between backward and forward perceptions explains intuitively why we work backwards through the game, but not forwards, in the process of identifying codominated actions.

It is natural, however to impose some additional restrictions on the beliefs that rational players may have about past actions. Because no player would ever rationally choose a codominated action, one might suppose that a rational player, given any history of recommendations and observations at any

stage, should not assign positive probability to the event that any players have chosen codominated actions in the past, unless there is no other explanation for his observations. (Of course, if a player has directly observed players choosing codominated actions, then he must assign probability one to this event.) Such a restriction will generally decrease the set of "rational" communication equilibria and increase the set of actions that could never be rationally chosen by any given type of a player. Using this larger set of irrational actions, we may then impose further similar restrictions on players' beliefs about the past; and thus we may continue iteratively until no further "irrational" actions can be identified. We now develop a formal model of such iterative elimination of irrational actions.

Given any correspondence B that satisfies (7.1), let $S(B)$ be the set of all outcomes of the game that could have positive probability when no player ever uses an action in B . That is

$$S(B) = \{t \in T \mid \exists f \in E^0(B) \text{ such that } P(t|f) > 0 \}.$$

Analogous to earlier notation, we may also define

$$S^{\leq k}(B) = \{t^{\leq k} \mid t \in S(B)\}, \quad S_i^{\leq k}(B) = \{t_i^{\leq k} \mid t \in S(B)\}, \quad \hat{S}^{\leq k}(B) = \times_{j \in N} S_j^{\leq k}(B).$$

So $t^{\leq k} \in S^{\leq k}(B)$ iff $t^{\leq k}$ is a vector of players' types that could actually occur in stage k if all players always avoid actions in B .

We now inductively define, for any positive integer m , the correspondence D_m^* and the concept of m -iterative codomination. To begin, let

$$D_1^* = D.$$

For any $m > 1$ and any B satisfying (7.1), we say that B is an m -iterative codomination correspondence iff, for every stage k and every π in $\Delta(F \times T^{\leq k})$,

if $\pi(E_{m-1}^{*k}(B) \times S^{\leq k}(D_{m-1}^*)) = 1$ and $\pi(\phi^k(B)) > 0$

then there exists some player i , some $t_i^{\leq k}$ in $S_i^{\leq k}(D_{m-1}^*)$, some c_i^k in $B(t_i^{\leq k})$, and some (γ_i, τ_i) in M_i^{*k} such that

$$\sum_{(f, t^{\leq k}) \in \phi^k(c_i^k, t_i^{\leq k})} \pi(f, t^{\leq k})(U_i(f | t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i) | t^{\leq k})) < 0;$$

where $E_{m-1}^{*k}(B) = \{f \in E^0(D_{m-1}^*) \mid \forall i, \forall \ell > k, \forall t^{\leq \ell} \in \hat{S}^{\leq \ell}(D_{m-1}^*), f_i^\ell(t^{\leq \ell}) \notin B(t^{\leq \ell})\}$.

That is, B must satisfy the same conditions as in the definition of a codomination correspondence, except that we ignore all states that are impossible when all players avoid actions in D_{m-1}^* . It is easy to check that any union of m -iterative codomination correspondences is also an m -iterative codomination correspondence. Thus, we inductively define D_m^* to be the maximal m -iterative codomination correspondence.

It is straightforward to verify that, for every m

$$D_{m-1}^* \subseteq D_m^* \quad \text{and} \quad S(D_{m-1}^*) \supseteq S(D_m^*).$$

Thus, since there are only finitely many actions, there must exist some m and some correspondence D_∞^* such that

$$D_m^* = D_{m+1}^* = D_{m+2}^* = \dots = D_\infty^*.$$

We say that an action c_i^k in C_i^k is (sequentially) predominant for type $t_i^{\leq k}$ of player i at stage k iff c_i^k is not in D_∞^* . Similarly, we say that a communication equilibrium μ is (sequentially) predominant iff the mediator would always recommend only predominant actions. That is, μ is sequentially predominant iff $\mu(E^0(D_\infty^*)) = 1$. (A related concept of weak predominance is introduced elsewhere, by Myerson [1985]. The concept of predominance defined

here may be called sequential prodominance whenever it is necessary to distinguish it from weak predominance.)

Since $D \subseteq D_{\infty}^*$ and $E^0(D) \supseteq E^0(D_{\infty}^*)$, Theorem 2 implies that any predominant communication equilibrium is a sequential communication equilibrium.

Given any multistage game, its codominant residue is the game that remains when we eliminate all codominated actions from the feasible set of each type of each player at each stage of the game. (Thus, the codominant residue is generally a game in which the set of feasible actions for a player may depend on his type. However, as noted parenthetically in Section 4, all of the results in this paper can be extended to such games.) Similarly, let the m-iterative residue of a game be the game that remains when all m-iteratively codominated actions are eliminated. It is straightforward to check that the m-iterative residue is the codominant residue of the (m-1)-iterative residue.

In the codominant residue, the mediator is restricted to mechanism such that $\mu(E^0(D)) = 1$, and each player can only use manipulative strategies in L_1 . Thus, reinterpreting Theorems 2 and 3, the set of sequential communication equilibria of a multistage game is just the set of all communication equilibria of its codominant residue. More generally, the set of communication equilibria of the m-iterative residue is the set of all sequential communication equilibria of the (m-1)-iterative residue of the game. But the set of sequential communication equilibria of any finite multistage game is a nonempty subset of its communication equilibria. Thus, by induction, for any m, the set of communication equilibria of the m-iterative residue is a nonempty subset of the communication equilibria of the original game. For some sufficiently large m^* , the predominant

communication equilibria are the communication equilibria of the m^* -iterative residue. The Nash equilibria of this residue are, of course, a nonempty subset of its communication equilibria. Thus, we have derived the following general existence theorem.

Theorem 5. The set of predominant communication equilibria is nonempty and includes at least one Nash (communicationless) equilibrium.

In Example 6, only y_1 is codominated. Thus, the codominant residue of Example 6 is just the game in Example 5. Since everything except x_1 , x_2 , and x_3 is codominated in Example 5, (x_1, x_2, x_3) is the unique predominant equilibrium in Example 6.

9. Proofs.

Theorem 1. μ is a conditional probability system on Ω if and only if there exists a sequence of probability distributions $\{\eta^j\}_{j=1}^{\infty}$ such that

$$\eta^j(\{\omega\}) > 0, \quad \forall j, \quad \forall \omega \in \Omega; \quad \text{and}$$
$$\mu(X|Z) = \lim_{j \rightarrow \infty} \eta^j(X \cap Z) / \eta^j(Z), \quad \forall X, \quad \forall Z \neq \emptyset.$$

Proof of Theorem 1.

Suppose first that there exists a sequence of probability distributions $\{\eta^j\}_{j=1}^{\infty}$ such that $\eta^j(\{\omega\}) > 0$ for every j and ω in Ω , and

$$\mu(X|Z) = \lim_{j \rightarrow \infty} \eta^j(X \cap Z) / \eta^j(Z), \quad \forall X, \quad \forall Z \neq \emptyset.$$

Then (5.2)-(5.4) can be checked as follows. If $X \cap Y = \emptyset$ then

$$\begin{aligned}
 \mu(X \cap Y | Z) &= \lim_{j \rightarrow \infty} \eta^j((X \cap Y) \cap Z) / \eta^j(Z) \\
 &= \lim_{j \rightarrow \infty} \eta^j(X \cap Z) / \eta^j(Z) + \lim_{j \rightarrow \infty} \eta^j(Y \cap Z) / \eta^j(Z) \\
 &= \mu(X | Z) + \mu(Y | Z).
 \end{aligned}$$

$$\begin{aligned}
 1 &= \lim_{j \rightarrow \infty} \eta^j(Z) / \eta^j(Z) = \mu(Z | Z) \\
 &= \lim_{j \rightarrow \infty} \eta^j(\Omega \cap Z) / \eta^j(Z) = \mu(\Omega | Z).
 \end{aligned}$$

If $X \subseteq Y \subseteq Z$ and $Y \neq \emptyset$ then

$$\begin{aligned}
 \mu(X | Z) &= \lim_{j \rightarrow \infty} \eta^j(X \cap Z) / \eta^j(Z) = \\
 &= \lim_{j \rightarrow \infty} ((\eta^j(X) / \eta^j(Y)) (\eta^j(Y) / \eta^j(Z))) \\
 &= (\lim_{j \rightarrow \infty} \eta^j(X \cap Y) / \eta^j(Y)) (\lim_{j \rightarrow \infty} \eta^j(Y \cap Z) / \eta^j(Z)) \\
 &= \mu(X | Y) \mu(Y | Z).
 \end{aligned}$$

Thus, μ is a conditional probability system.

Conversely, suppose now that μ is a conditional probability system. We construct $\{\eta^j\}_{j=1}^{\infty}$ as follows. Let $W_0 = \Omega$ and then inductively define W_h for $h > 1$ by

$$W_h = \{\omega \in W_{h-1} \mid \mu(\{\omega\} | W_{h-1}) = 0\}.$$

Since these sets are strictly decreasing in size and Ω is finite, there exists some H such that $W_H \neq \emptyset$ and $W_{H+1} = \emptyset$. For any $X \subseteq \Omega$, let

$$\eta^j(X) = \left(\frac{j-1}{j - (1/j)^H} \right) \sum_{h=0}^H \mu(X|W_h) \left(\frac{1}{j} \right)^h.$$

Each η^j is a probability distribution on Ω , giving positive probability to every point. Given any sets X and Z such that $Z \neq \emptyset$, let g be the highest number such that $Z \subseteq W_g$. Then $\mu(Z|W_g) > 0$ and $\mu(Z|W_h) = 0$ for every $h < g$. Thus, using (5.4),

$$\begin{aligned} \mu(X|Z) &= \mu(X \cap Z|Z) = \frac{\mu(X \cap Z|W_g)}{\mu(Z|W_g)} = \\ &= \lim_{j \rightarrow \infty} \frac{\sum_{h=g}^H \mu(X \cap Z|W_h) \left(\frac{1}{j} \right)^h}{\sum_{h=g}^H \mu(Z|W_h) \left(\frac{1}{j} \right)^h} = \lim_{j \rightarrow \infty} \frac{\eta^j(X \cap Z)}{\eta^j(Z)}. \end{aligned} \quad \text{Q.E.D.}$$

We prove Theorem 4 and a series of lemmas before proving Theorem 2.

Theorem 4. B is a codomination correspondence if and only if there exists some α in A^* such that (B, α) is a codomination system.

Proof of Theorem 4

Let B be any correspondence satisfying (7.1). Let $E^k(B)$ and $\phi^k(B)$ be as defined in Section 7. Let

$$\hat{E}^k(B) = (E^k(B) \times T^{\leq k}) \cap \phi^k(B)$$

Thus $(f, t^{\leq k}) \in \hat{E}^k(B)$ if f would never recommend any actions in B after stage k , but f does recommend an action in $B(t_i^{\leq k})$ for some player i in stage k . By

the separating hyperplane theorem, the following two propositions are equivalent.

(1) There does not exist any probability distribution π such that $\pi(\hat{E}^k(B)) = 1$ and

$$\sum_{(f,t^{\leq k}) \in \phi^k(c_i^k, t_i^{\leq k})} \pi(f,t^{\leq k}) (U_i(f|t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i)|t^{\leq k})) \geq 0$$

$$\forall i, \forall t_i^{\leq k} \in T_i^{\leq k}, \forall c_i^k \in B(t_i^{\leq k}), \forall (\gamma_i, \tau_i) \in M_i^{*k}.$$

(2) There exist nonnegative numbers $\alpha(\gamma_i, \tau_i | c_i^k, s_i^{\leq k})$ ($\forall i, \forall s_i^{\leq k} \in T_i^{\leq k}, \forall c_i^k \in B(s_i^{\leq k}), \forall (\gamma_i, \tau_i) \in M_i^{*k}$) such that, for every $(f, t^{\leq k})$ in $\hat{E}^k(B)$

$$\sum_{\{i | f_i^k(t^{\leq k}) \in B(t_i^{\leq k})\}} \sum_{(\gamma_i, \tau_i) \in M_i^{*k}} (\alpha(\gamma_i, \tau_i | f_i^k(t^{\leq k}), t^{\leq k}) \cdot (U_i(f|t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i)|t^{\leq k}))) < 0.$$

It is straightforward to check that (B, α) is a codomination system if and only if (B, α) satisfies (2) for every stage k . Similarly, B is a codomination correspondence if and only if B satisfies (1) for every stage in k . Q.E.D.

Lemma 1. Suppose (B, α) is a codomination system and $(Q, \bar{\mu})$ satisfies the conditions (6.2)-(6.4) for a sequential communication equilibrium. For any player i , any stage k , and any (c_i, t_i) in Q_i , c_i^k is not in $B(t_i^{\leq k})$.

Proof. If the lemma fails, then let k be the maximal number such that

$$\{(f, t) \in G(Q) \times T \mid \exists i \text{ such that } f_i^k(t^{\leq k}) \in B(t_i^{\leq k})\} \neq \emptyset.$$

Let X denote this nonempty set. Then

$$\begin{aligned}
 & \sum_i \sum_{t_i^{\leq k}} \sum_{c_i^{\leq k}} \sum_{(\gamma_i, \tau_i) \in M_i^{*k}} (\bar{\mu}(c_i^{\leq k}, t_i^{\leq k} | X) \alpha(\gamma_i, \tau_i | c_i^k, t_i^{\leq k})) \\
 & \cdot \sum_{t_{-i}^{\leq k}} \sum_f \bar{\mu}(f, t^{\leq k} | c_i^{\leq k}, t_i^{\leq k}) (U_i(f | t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i) | t^{\leq k})) = \\
 & = \sum_{(f, t) \in X} \bar{\mu}(f, t | X) v^k(f, t^{\leq k}, \alpha) < 0.
 \end{aligned}$$

We use the maximality of k to guarantee that any f in $G(Q)$ is in $E^k(D)$; then (7.5) gives us the strict negativity result. Thus, at least one strategic constraint in (6.4) must be violated. Q.E.D.

Lemma 2. There exist \hat{B} , \hat{Q} , and $\bar{\sigma}$ such that (6.2)-(6.4) are satisfied (when $\bar{\mu} = \bar{\sigma}$ and $Q = \hat{Q}$) and

$$\hat{Q}_i = \{(c_i, t_i) \in C_i \times T_i \mid c_i^k \notin \hat{B}(t_i^{\leq k}) \ \forall k\}, \ \forall i.$$

Proof. For any ε between 0 and 1, we can construct an ε -perturbed game (in the sense of Selten [1975]) from our game Γ as follows. In each stage and every possible state, each player has an independent ε probability of trembling, in which case every action is equally likely. By the general existence theorem of Nash equilibria for finite games, an equilibrium exists for each ε -perturbed game. Since all the action sets are finite, there exists some correspondence \hat{B} and an infinite sequence $\{\varepsilon_j\}_{j=1}^{\infty}$ such that $\lim_{j \rightarrow \infty} \varepsilon_j = 0$, each $\varepsilon_j > 0$, and, for each j , there exists an equilibrium of the ε_j -perturbed game such that, for every player i , every stage k , and every $t_i^{\leq k}$ in $T_i^{\leq k}$, $\hat{B}(t_i^{\leq k})$ is the set of all actions to which player i gives zero probability in stage k if he is not trembling and his information state is $t_i^{\leq k}$. Let \hat{Q} be derived from \hat{D} as in the statement of the lemma.

Suppose that, before the game begins, a mediator offers to perform all the independent randomizations planned by all the players for all stages and all states in their equilibrium strategies for the ε_j -perturbed game. After performing these randomizations, the mediator will have generated a feedback rule in $G(\hat{Q})$. Let σ_j be the probability distribution over $G(\hat{Q}) \times T$ when the mediator selects his feedback rule in this way and then implements it in the ε_j -perturbed game. (So each player is assumed to obey the mediator with independent probability $1 - \varepsilon_j$ in every stage, and to tremble uniformly over his actions with probability ε_j .) Because each $f_i^k(t_i^{\leq k})$ is chosen independently by the mediator, with full support over $C_i^k \setminus \hat{D}(t_i^{\leq k})$, and because every action has positive probability in a tremble, the probability distribution of σ_j has full support over $G(\hat{Q}) \times T$.

Choosing a subsequence if necessary (to guarantee that all conditional probabilities converge), we can let $\bar{\sigma}$ be the conditional probability system on $G(\hat{Q}) \times T$ generated by $\{\sigma_j\}_{j=1}^{\infty}$ (as in Theorem 1). Then (6.2) is satisfied (when $\bar{\mu} = \bar{\sigma}$) because the probability of any trembles in stage k or thereafter goes to zero as $\varepsilon_j \rightarrow 0$, given any information about the mediator and the players up to the beginning of stage k . The incentive constraints (6.3) and (6.4) are satisfied because each σ_j represents an equilibrium of the ε_j -perturbed game, in which every information state occurs with positive probability. Q.E.D.

Lemma 3. There exist B , Q , and $\bar{\mu}$ such that $\bar{\mu}$ is a conditional probability system on $G(Q) \times T$, (6.2)-(6.4) are satisfied by Q and $\bar{\mu}$, B is a codomination correspondence, and

$$Q_i = \{(c_i, t_i) \in C_i \times T_i \mid c_i^k \notin B(t_i^{\leq k}) \ \forall k\}, \ \forall i.$$

Proof. Let \hat{B} , σ_j , and $\bar{\sigma}$ be as in Lemma 2 and its proof. If \hat{B} is a codomination correspondence, then we are done; so suppose that it is not. Let $B^0 = \hat{B}$. Then there must exist some finite sequence of correspondences (B^1, \dots, B^H) , such that B^H is a codomination correspondence

$$B^h(t_i^{\leq k}) \subseteq B^{h-1}(t_i^{\leq k}) \quad \forall i, \forall k, \forall t_i^{\leq k} \in T_i^{\leq k}, \forall h \in \{1, \dots, H\},$$

and, for every h and every k , there exists some probability distribution π^{hk} such that $\pi^{hk}(E^k(B^{h-1}) \times T_i^{\leq k}) = 1$ and, for every player i and every $t_i^{\leq k}$ in $T_i^{\leq k}$,

$$\pi^{hk}(\phi^k(c_i^k, t_i^{\leq k})) > 0 \quad \forall c_i^k \in B^{h-1}(t_i^{\leq k}) \setminus B^h(t_i^{\leq k}),$$

$$\pi^{hk}(\phi^k(c_i^k, t_i^{\leq k})) = 0 \quad \forall c_i^k \in B^h(t_i^{\leq k}), \quad \text{and}$$

$$\sum_{(f, t^{\leq k}) \in \phi^k(c_i^k, t_i^{\leq k})} \pi^{hk}(f, t^{\leq k}) (U_i(f | t^{\leq k}) - U_i(f \circ (\gamma_i, \tau_i) | t^{\leq k})) \geq 0$$

$$\forall (\gamma_i, \tau_i) \in M_i^{*k}, \quad \forall c_i^k \in B^{h-1}(t_i^{\leq k}).$$

These B^h and π^{hk} are constructed inductively as follows. If B^{h-1} is not a codomination correspondence, then there exists some stage ℓ and some probability distribution π such that condition (7.3) is satisfied (for $B = B^h$ and $k = \ell$) but condition (7.4) (the violation of some incentive constraint) does not hold. Then let $\pi^{h\ell} = \pi$ and let

$$B^h(t_i^{\leq \ell}) = \{c_i^\ell \in B^{h-1}(t_i^{\leq \ell}) \mid \mu(\phi^\ell(c_i^\ell, t_i^{\leq \ell})) = 0\}.$$

For every stage $k \neq \ell$, let

$$\pi^{\text{hk}}(f, t^{\leq k}) = \bar{\sigma}(f, t^{\leq k} | G(\hat{Q}) \times T), \quad \forall f, \quad \forall t^{\leq k},$$

and

$$B^h(t_i^{\leq k}) = B^{h-1}(t_i^{\leq k}), \quad \forall i, \quad \forall t_i^{\leq k} \in T_i^{\leq k}.$$

In this construction, the nonnegative integer

$$\sum_i \sum_k \sum_{t_i^{\leq k}} |B^h(t_i^{\leq k})|$$

strictly decreases with every increase in h . Thus, the construction must eventually terminate at a codomination correspondence. (Notice that the correspondence that always selects the empty set is a codomination correspondence.)

Let $B = B^H$, and let Q be derived from B so as to satisfy the equation in the statement of Lemma 3. It now remains for us to construct $\bar{\mu}$.

Let $\hat{G}^k(\hat{Q}) = \{f \in F \mid \exists g \in G(\hat{Q}) \text{ such that } \forall \ell < k \ f^\ell = g^\ell\}$. That is, $\hat{G}^k(\hat{Q})$ is the set of all feedback rules that look like rules in $G(\hat{Q})$ before stage k .

We can extend each π^{hk} to be a probability distribution over $G(Q) \times T$ by letting

$$\pi^{\text{hk}}(f, t) = \pi^{\text{hk}}(f, t^{\leq k}) P(t^{>k} | f, t^{\leq k}).$$

So π^{hk} is consistent with the given dynamics of the game after stage k .

Notice that we did not use any properties of f^ℓ for $\ell < k$ in our construction of the $\pi^{\text{hk}}(f, t)$ numbers. Thus, we can assume with no loss of generality that

$$\pi^{hk}(f,t) = 0 \text{ if } f \notin \hat{G}^k(\hat{Q}),$$

and

$$\pi^{hk}(f,t) = \pi^{hk}(g,t) \text{ if } f \in \hat{G}^k(\hat{Q}), \quad g \in \hat{G}^k(\hat{Q}), \text{ and } f^\lambda = g^\lambda, \quad \forall \lambda > k.$$

Thus, for the components of the feedback rule before stage k , π^{hk} gives the same marginal distribution as the uniform distribution on $G(\hat{Q})$.

By renumbering the $\{\sigma_j\}$ sequence if necessary, we can assume that there exists some J such that, for every $j \geq J$,

$$\sigma_j(f,t) \geq 1/j, \quad \forall f \in G(\hat{Q}), \quad \forall t \in T.$$

(Recall that each σ_j has full support over $G(\hat{Q}) \times T$.) For each j , let δ_j be a probability distribution over $G(\hat{Q}) \times T$ defined so that

$$\delta_j(f,t) = \sigma_j(f,t) + \sum_{k=1}^K \sum_{h=1}^H (1/j)^{(k-1)H+h+1} \pi^{hk}(f,t).$$

Given any event that has positive probability in σ_j , the conditional probabilities generated by δ_j approach those generated by σ_j as j becomes large. In any stage k , if the event that the mediator has used $f^{\leq k}$ and the players have learned $t^{\leq k}$ has positive probability under δ_j but has zero probability under σ_j , then the conditional probabilities generated by δ_j given this event must approach (as $j \rightarrow \infty$) the conditional probabilities generated by the first $\pi^{h\lambda}$ distribution in which this event has positive probability. (Here "first" $\pi^{h\lambda}$ means the lowest possible λ , and then lowest h given this λ .) Furthermore, the first $\pi^{h\lambda}$ that gives $(f^{\leq k}, t^{\leq k})$ positive probability must satisfy $\lambda \leq k$ if σ_j gives zero probability to $(f^{\leq k}, t^{\leq k})$. (This is because, when $\lambda > k$, $\pi^{h\lambda}$ gives positive probability to $(f^{\leq k}, t^{\leq k})$ only if σ_j does also.)

In any event that has positive probability under all δ_j , the conditions (6.2)-(6.4) are satisfied by the limiting conditional probability system that is generated as $j \rightarrow \infty$. The dynamics-consistency condition (6.2) is satisfied because each π^{hl} condition satisfies it in every stage $k \leq l$. The incentive constraints are satisfied because, if π^{hl} is the first to give positive probability to the event that i observes $(c_i^{\leq k}, t_i^{\leq k})$ [or $(c_i^{\leq k-1}, t_i^{\leq k})$] through the end of [or the beginning of] stage k , then $l \leq k$ and $c_i^l \in B^{h-1}(t_i^{\leq l}) \setminus B^h(t_i^{\leq l})$; and π^{hl} was constructed so that no player could expect to gain by manipulating in any event that he could observe with positive probability after getting such a recommendation c_i^l in state $t_i^{\leq l}$.

If c_i^k is in $B(t_i^{\leq k})$ (that is, $B^H(t_i^{\leq k})$) then there is a positive probability under every δ_j that c_i^k will be recommended at stage k to type $t_i^{\leq k}$ of player i , after some history of past recommendations $c_i^{\leq k-1}$, but not necessarily after every $c_i^{\leq k-1}$ such that $(c_i^{\leq k}, t_i^{\leq k}) \in Q_i^{*k}$. We now need to perturb the $\{\delta_j\}$ sequence slightly so that every history $(c_i^{\leq k}, t_i^{\leq k})$ in Q_i^{*k} should have positive probability, without losing incentive compatibility in the limit. Furthermore, these perturbed distributions should have full support over $G(Q) \times T$, so as to generate, in the limit, a conditional probability system on $G(Q) \times T$ satisfying the dynamic consistency condition (6.2).

We need some further notation. Let

$$\Psi_j^k(t^{\leq k}) = \{f^k \mid \exists (g,r) \text{ such that } g^k = f^k, r^{\leq k} = t^{\leq k}, \text{ and } \delta_j(g,r) > 0\},$$

$$\Psi_j^{\leq k}(t^{\leq k}) = \prod_{\lambda=1}^{k-1} \Psi_j^\lambda(t^{\leq \lambda}),$$

$$\Phi^k(f) = \{g \in F \mid \forall l \geq k, g^l = f^l\},$$

$$\xi_j^k(f,t) = \begin{cases} \sum_{g \in \Phi^k(f)} \delta_j(g,t) / |\Psi_j^{<k}(t^{\leq k})| & \text{if } f^{<k} \in \Psi_j^{<k}(t^{\leq k}), \\ 0 & \text{if } f^{<k} \notin \Psi_j^{<k}(t^{\leq k}). \end{cases}$$

Notice that $\xi_j^1(f,t) = \delta_j(f,t)$. Let $Z = KH + H + 1$, and let

$$\lambda_j(f,t) = \sum_{\ell=1}^{K+1} (1/j)^{(\ell-1)Z} \xi_j^\ell(f,t) + \sum_{m=1}^{K+1} (1/j)^{KZ+m} P(t^{>m} | f, t^{\leq m}).$$

Finally, for any (f,t) in $G(Q) \times T$, let

$$\eta_j(f,t) = \lambda_j(f,t) / \left(\sum_{g \in G(Q)} \sum_{s \in T} \lambda_j(g,s) \right).$$

For each j , η_j is a probability distribution with full support over $G(Q) \times T$ (because we let $P(t^{>m} | f, t^{\leq m}) = 1$ when $m = K + 1$, in the definition of $\lambda_j(f,t)$). Choosing a subsequence if necessary, let $\bar{\mu}$ be the conditional probability system on $G(Q) \times T$ generated by η_j as $j \rightarrow \infty$, in the sense of Theorem 1. We now show that $\bar{\mu}$ satisfies (6.2)-(6.4).

$\Psi_j^{<k}(t^{\leq k})$ is the set of feedback rules such that, in every stage before k , the mediator uses a feedback rule that has positive probability under δ_j , given the current state. Thus, ξ_j^k can be interpreted as follows. Suppose that there are two mediators: one manifest and the other subliminal. The players are coordinated by the subliminal mediator throughout the game according to the δ_j distribution. In stage k and thereafter, the manifest mediator uses the same feedback rule as the subliminal mediator. Before stage k , the manifest mediator has no influence over the players. For each $\ell < k$, the manifest mediator selects his feedback function for stage ℓ at random (uniformly) from $\Psi_j^k(t^{\leq \ell})$, if $t^{\leq \ell}$ is the current state of the players'

information. Then ξ_j^k is the joint distribution of the manifest mediator's feedback rule and the outcome of the game, under this (unusual) procedure.

Z is greater than the highest power of $(1/j)$ that is used in the definition of δ_j . Thus, for any event $X \subseteq G(Q) \times T$, the conditional probabilities $\bar{\mu}(\cdot|X)$ are completely determined by the term in the definition of λ_j that has the lowest power of $(1/j)$, among all terms that are positive for at least one point in X. Any event $(f^{\leq k}, t^{\leq k})$ that has positive probability under ξ_j^ℓ for some $\ell > k$ must also have positive probability under ξ_j^k , because of the definition of $\Psi_j^k(t^{\leq k})$. Thus, if $f^\ell \in \Psi_j^\ell(t^{\leq \ell})$ for every $\ell < k$, for all j sufficiently large, then the conditional probabilities $\bar{\mu}(\cdot|f^{\leq k}, t^{\leq k})$ are determined by the sequence $\{\xi_j^\ell\}_{j=1}^\infty$ for some $\ell < k$; and thus the conditional probabilities $\bar{\mu}(\cdot|f^{\leq k}, t^{\leq k})$ will be consistent with the dynamics of the game as required by (6.2), since the conditional probabilities generated by ξ_j^ℓ as $j \rightarrow \infty$ are consistent with these dynamics after stage ℓ (when the manifest and subliminal mediators coincide). On the other hand, if $f^\ell \notin \Psi_j^\ell(t^{\leq \ell})$ for some $\ell < k$, for all j sufficiently large, then let m be the lowest number such that there exists some $t^{>k}$ and $f^{>k}$ such that $P(t^{>m}|f, t^{\leq m}) > 0$. This m will not be greater than k, and the conditional probabilities $\bar{\mu}(\cdot|f^{\leq k}, t^{\leq k})$ will be completely determined by the term for this m in the second summation in the definition of λ_j . Thus, (6.2) is satisfied in this case as well.

Given any player i and any $(c_i^{\leq k}, t_i^{\leq k})$ in Q_i^{*k} [or $(c_i^{\leq k}, t_i^{\leq k+1})$ in Q_i^{k+1}] let ℓ be the highest number such that the recommendation c_i^ℓ would have probability zero under the δ_j distributions, for all j sufficiently large, after a recommendation-history of $c_i^{\leq \ell-1}$ and an information type $t_i^{\leq \ell}$. Clearly $\ell < k$. Then player i's conditional beliefs $\bar{\mu}(\cdot|c_i^{\leq k}, t_i^{\leq k})$ [or $\bar{\mu}(\cdot|c_i^{\leq k}, t_i^{\leq k+1})$] are completely determined by the $\{\xi_j^\ell\}_{j=1}^\infty$ sequence. Thus, player i believes

that, although the manifest mediator diverged from the subliminal mediator before stage l , the conditional probabilities that are generated by the δ_j distributions as $j \rightarrow \infty$ accurately characterize the current state of all information and the behavior of all individuals from stage l on. Since the relevant incentive constraints were satisfied by the beliefs generated by the δ_j sequence, $\bar{\mu}$ satisfies these constraints (6.3) and (6.4) as well.

Q.E.D.

Lemma 4. Let B and Q be as in Lemma 3, and let D and \bar{Q} be as defined in Section 7. Then $D = \bar{D}$ and $Q = \bar{Q}$.

Proof. By Lemma 1, mediation ranges that are consistent with (6.2)-(6.4) cannot include any codominated actions. By Lemma 3, Q is consistent with (6.2)-(6.4) and, in every state, allows all the actions that are not in the codomination correspondence B . Thus, B is equal to D , the maximal codomination correspondence, and Q is equal to \bar{Q} .

Q.E.D.

Theorem 2. A communication equilibrium μ is a sequential communication equilibrium if and only if $\mu(E^0(D)) = 1$.

Proof of Theorem 2 and Corollary 1.

If μ is a sequential communication equilibrium, then there is some mediation range Q on which (6.2)-(6.4) can be satisfied and such that $\mu(G(Q)) = 1$. But by Lemma 1, $Q \subseteq \bar{Q}$, because Q must avoid all codominated actions, and so $\mu(G(\bar{Q})) = 1$. (Recall $G(\bar{Q}) = E^0(D)$.)

Conversely, suppose that μ is a communication equilibrium and $\mu(G(\bar{Q})) = 1$. Let η_j be as in the proof of Lemma 3. Then we can define a sequence of probability distributions ρ_j with full support over $G(\bar{Q}) \times T$ by letting

$$\rho_j(f,t) = (1 - (1/j)) \mu(f) P(t|f) + (1/j) \eta_j(f,t), \quad \forall f \in G(\bar{Q}), \quad \forall t \in T.$$

Then let $\bar{\mu}$ be the conditional probability system generated by ρ_j as $j \rightarrow \infty$. For any event X , the conditional probabilities $\bar{\mu}(\cdot|X)$ are the same as under $\mu \cdot P$ if X has positive probability under $\mu \cdot P$; otherwise the conditional probabilities are the same as those generated by the η_j as $j \rightarrow \infty$. But $\mu \cdot P$ is consistent with the given dynamics of the game, and satisfies the incentive constraints in all events that have positive probability, since μ is a communication equilibrium. As shown in the proof of Lemma 3, the η_j generate conditional probabilities that satisfy (6.2)-(6.4) with the mediation range \bar{Q} . Thus, (6.2)-(6.4) are satisfied by $\bar{\mu}$ and \bar{Q} , and (6.1) follows immediately from the definition of ρ_j . Q.E.D.

Theorem 3. If $\mu(E^0(D)) = 1$ and μ satisfies (4.1) for every i and every (γ_i, τ_i) in L_i , then μ is a sequential communication equilibrium.

Proof of Theorem 3.

Suppose that, contrary to the theorem, $\mu(E^0(D)) = 1$ and μ satisfies (4.1) for every (γ_i, τ_i) in L_i , but μ is not a communication equilibrium. Then there exists some player i who could expect (ex ante) to gain by using some manipulative strategy that uses codominated actions with positive probability. Among the (finitely many) manipulative strategies that are optimal for player i against μ , in terms of his ex ante expected utility, we can choose (γ_i, τ_i) so that, for the lowest possible number k , γ_i would never select a codominated action after stage k . Let ν be the mechanism that effectively results when player i manipulates μ by (γ_i, τ_i) , so that

$$\nu(f) = \mu(\{g \mid g \circ (\gamma_i, \tau_i) = f\}).$$

Let π be the probability distribution over $F \times T^{\leq k}$ induced by ν , so that

$$\pi(f, t^{\leq k}) = \nu(f) \sum_{t > k} P(t|f).$$

Notice that $\pi(E^k(D) \times T^{\leq k}) = 1$, since no player is being told to use codominated actions after stage k . Also, $\pi(\phi^k(D)) > 0$, because player i must be using codominated actions with positive probability in stage k . (If player i were using codominated actions in stage k with zero probability only, then we could have found another manipulative strategy that is also optimal for i against μ and that never uses codominated actions after stage $k-1$, contradicting the minimality of k .) Thus, by definition of codomination, there must be some player who can expect to gain by planning to manipulate against π (or, equivalently, against ν) after he gets a recommendation to use a codominated action at stage k . But that player must be player i , because no other player ever gets a codominated recommendation under π . But if player i could expect to gain by planning to manipulate against π , then (γ_i, τ_i) could not have been an optimal manipulation against μ . This contradiction proves the theorem. Q.E.D.

References

- R. J. Aumann [1974], "Subjectivity and Correlation in Randomized Strategies," Journal of Mathematical Economics 1, 67-96.
- F. Forges [1984], "An Approach to Communication Equilibria," CORE discussion paper No. 8435, Université Catholique de Louvain.
- H. Kuhn [1953], "Extensive Games and the Problem of Information," in Contributions to the Theory of Games 2, H. Kuhn and A. Tucker eds. Princeton: Princeton University Press, 193-216.
- D. M. Kreps and R. Wilson [1982], "Sequential Equilibria," Econometrica 50, 863-894.
- R. D. Luce and H. Raiffa [1957], Games and Decisions. New York: John Wiley and Sons.
- R. B. Myerson [1979], "Incentive Compatibility and the Bargaining Problem," Econometrica 47, 61-73.
- R. B. Myerson [1982], "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," Journal of Mathematical Economics 10, 67-81.
- R. B. Myerson [1983], "Bayesian Equilibrium and Incentive Compatibility: An Introduction," Northwestern University discussion paper, to appear in Social Goals and Social Organization, L. Hurwicz, D. Schmeidler and H. Sonnenschein, eds. Princeton: Princeton University Press.
- R. B. Myerson [1985], "Acceptable and Predominant Correlated Equilibria," Center for Mathematical Studies discussion paper No. 591, Northwestern University.
- J. F. Nash [1951], "Noncooperative Games," Annals of Mathematics 54, 286-295.
- J. von Neumann and O. Morgenstern [1944], Theory of Games and Economic Behavior. Princeton: Princeton University Press.
- R. Selten [1975], "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory 4, 25-55.