

Communication in Dynastic Repeated Games: 'Whitewashes' and 'Coverups'*

LUCA ANDERLINI
(Georgetown University)

ROGER LAGUNOFF
(Georgetown University)

June 2003

ABSTRACT. We ask whether communication can directly substitute for memory in *dynastic repeated games* in which short lived individuals care about the utility of their offspring who replace them in an infinitely repeated game. Each individual is unable to observe what happens before his entry in the game. Past information is therefore conveyed from one cohort to the next by means of communication.

When communication is costless and messages are sent simultaneously, communication mechanisms or *protocols* exist that sustain the same set of equilibrium payoffs as in the standard repeated game. When communication is costless but sequential, the incentives to “whitewash” the unobservable past history of play become pervasive. These incentives to whitewash can only be countered if some player serves as a “neutral historian” who verifies the truthfulness of others’ reports while remaining indifferent in the process. By contrast, when communication is sequential and (lexicographically) costly, all protocols admit only equilibria that sustain stage Nash equilibrium payoffs.

We also analyze a centralized communication protocol in which history leaves a “footprint” that can only be hidden by the current cohort by a unanimous “coverup”. We show that in this case the set of payoffs that are sustainable in equilibrium coincides with the weakly renegotiation proof payoffs of the standard repeated game.

JEL CLASSIFICATION: C72, C73, D82.

KEYWORDS: Dynastic Repeated Games, Communication, Whitewashing, Coverups.

ADDRESS FOR CORRESPONDENCE: Luca Anderlini, Department of Economics, Georgetown University, Washington DC 20057, USA. e-mail LA2@GEORGETOWN.EDU.

*We wish to thank an Associate Editor and an anonymous Referee of this journal as well as seminar participants at Arizona State, Columbia, Duke, Georgetown, Indiana, Montreal, Princeton, Rochester, Vanderbilt, VPI, the 2001 NSF/NBER Decentralization Conference, the Summer 2001 North American Econometric Society Meetings, and the Midwest Theory Conference, 2000, for useful comments and suggestions. All errors are our own.

“History is a pack of lies about events that never happened told by people who weren’t there.” - George Santayana

1. INTRODUCTION

1.1. Motivation

In any longstanding strategic relationship, history matters. The ability of the “players” to construct effective deterrents against “bad” behavior typically relies on accurate monitoring and recall of the history of play.

One chief interpretation of a long-term relationship is that of a stage game being repeated between “dynastic players” rather than between infinitely lived individuals. An infinitely repeated game is interpreted as an ongoing society populated by short lived individuals who care about the utility of their successors who replace them. Each successor then faces the same “types” of opponents as his predecessor.

Examples of repeated strategic interaction that would be modelled as dynastic repeated games abound. For example, in longstanding disputes between groups with competing claims (e.g., Catholics versus Protestants in Northern Ireland, Israelis versus Palestinians), the conflicts typically outlive any particular individual. Though the names of individuals involved change with time, the issues (payoffs) often remain the same. Other examples include electoral competition between political parties (e.g., Democrats versus Republicans) and strategic competition between firms. Firms, like political parties, are long lived organizations populated by short-lived managers, each of whom are periodically replaced. Putting agency issues aside, incentives may be structured so that each current manager acts in the long run interest of the firm, despite his relatively short tenure.

Since it seems unappealing to assume that any living individual observes something that takes place before he is “born,” a natural problem arises with dynastic games. It is well known that if the players do not have the means to condition their current actions on the history of play, equilibrium behavior changes dramatically. In the extreme case in which players have no knowledge of the past, strategic behavior can only depend on payoff relevant information (i.e., players must use so-called Markov strategies). When this happens and when the environment is stationary, then only repetitions of the stage game Nash equilibria are possible, even in an infinitely repeated game.¹

In a dynastic game, each new entrant cannot condition his behavior on history unless his “knowledge” of that history comes, directly or indirectly, from past participants. Often, that means that current players must rely on the historical accounts

¹Hence, Santayana’s other famous dictum: “Those who cannot remember the past are condemned to repeat it,” is quite literally true.

directly communicated by their predecessors.² This paper examines the properties of dynastic repeated games when participants do not observe history prior to their entry in to the game, and must therefore rely on accounts communicated by their predecessors.

For simplicity, we examine a model in which each member of a dynasty only lives one period. At the end of each period, individuals in the current cohort die, and are replaced by their successors in each dynasty. Each dynastic individual cares about his successor's utility as if it was his own discounted utility. Successors inherit the same preferences, but cannot observe prior behavior.

Since prior behavior is not directly observed, we assume that the only way current behavior can be linked to the past is through the reports of the previous cohort. We therefore augment the model to allow for messages to be sent at the end of each period from the current generation to the next. Communication is assumed to be publicly observed. Because the veracity of reports cannot be verified by neutral parties, the messages can also be manipulated. To see why incentives for manipulation may exist, suppose, for example, that two dynasties face off to play the Prisoners' Dilemma in Figure 1 below. Consider the Subgame Perfect equilibrium (SPE) which, for patient enough players, sustains perpetual mutual cooperation, (C, C) using "grim" "trigger" strategies. In this equilibrium, the dynastic players revert permanently to (D, D) if any defection is ever observed.

		Dynasty 2	
		C	D
Dynasty 1	C	2, 2	-2, 3
	D	3, -2	0, 0

Figure 1: Prisoners' Dilemma

Now suppose that at some date t , the date t member of Dynasty 1 defects by choosing " D ." Despite the fact that the individual in Dynasty 1 defected in the PD game, *both* individuals at date t may have an incentive to *whitewash* the defection by falsely reporting action (C, C) to the next generation. By lying, the current generation can insulate the next generation against the mutually destructive punishment phase. However, because lying precludes punishment, incentives for good behavior in the current stage are destroyed.

²For a useful perspective on the ways in which history is transmitted and collective memories are formed, see Pennebaker, Paez, and Rime (1997).

Unlike in standard communication (cheap talk) games,³ individuals within a dynasty value future payoffs in the same way. The potential incentive to misreport exists not because of payoff differences, but rather from a desire to protect future generations from the consequences of past deviations. By *whitewashing* deviations, the current generation has the chance to give their successors a clean slate to start the game. In this sense, the environment we analyze is reminiscent of repeated game models with renegotiation.⁴

Our interest, therefore, is in the extent to which history is accurately conveyed from one generation to another. How does potential manipulation of information across generations distinguish dynastic repeated play from the “full memory” model? To sensibly address this question, we adopt an implementation approach. We examine whether there exist useful *communication protocols*, i.e., mechanisms in which communication directly substitutes for memory.

We examine two models. The first, is a model of potential *whitewashing*. Information transmission constitutes cheap talk. Each individual may misrepresent the information unobservable to the next generation. Misrepresentation is typically costless to the sender. We examine whether or to what extent members of the current generation may whitewash the past in equilibrium. The second model examines the potential for *coverups*: The aggregation mechanism utilizes, to some extent, observable information. Consequently, members of the current cohort may attempt to hide information which might otherwise be observable to the next generation. Hiding information is difficult, and may require widespread agreement among the senders.

Despite the incentive to whitewash or to coverup detrimental histories, protocols that support full and honest communication exist. In the whitewashing model, if the communication protocol is decentralized and if messages are sent simultaneously then standard Nash implementation logic can be used to show that a Perfect Bayesian equilibrium (PBE) exists in which no whitewashing takes place. The idea is familiar: since reports contain redundant information, each individual’s report can be used to screen the veracity of others. Hence, if there are at least three players and all but one player’s messages agree, then the next generation uses the agreed upon message as the “official version” of history.⁵ If there are only two players, then the absence of an agreed upon message is treated as if a defection occurred. Hence, for any stage

³See Crawford and Sobel (1982) for a classic reference.

⁴See Farrell and Maskin (1989), Abreu, Pearce, and Stacchetti (1993), and Benoit and Krishna (1993).

⁵Of course, the original “cross-checking” argument goes back to Maskin (1999). Baliga, Corchon, and Sjöström (1997) use a similar type of mechanism in another model of cheap talk communication when there are three or more players. Similar types of mechanisms have also been used in repeated games with private monitoring and communication. See Ben-Porath and Kahneman (1996), Compte (1998), and Kandori and Matsushima (1998).

game the set of possible equilibrium strategies and payoffs is equivalent to that of the standard repeated game.

Ironically, the simultaneous moves protocol disciplines the players by instituting a coordination failure. For example, in our Prisoners' Dilemma example above, under the grim trigger strategy, both individuals in the present cohort would be better off by whitewashing a deviation, but neither can do it given the anticipated truthful message of the other. Clearly, if given the opportunity, one of the players would prefer to signal his intent to lie by moving first.⁶ Indeed, suppose that the messages are sent sequentially, and that members of Dynasty 1 communicate first. If the date t player 1 whitewashes his own defection, then it is clearly a best response for the date t player 2 to confirm the lie. Sequential moves therefore allows the players to break the "coordination failure" that prevented whitewashing in the simultaneous case. But since whitewashing will occur, the mutual cooperation equilibrium using grim trigger strategies cannot be sustained in the first place.

We characterize the PBE payoff set in any game when messages are sequential rather than simultaneous. A necessary condition for play to differ from the repetition of stage Nash actions in any PBE is that some player serves as "a neutral historian." The neutral historian is an individual who screens and verifies the truthfulness of reports of others, while remaining indifferent in the process. This necessity of a neutral historian rules out some types of equilibria. Nevertheless, a wide array of payoffs approximating the original payoffs of the repeated game are shown to be sustainable. It turns out that rectangular, "self generating" subsets of equilibrium payoff set are sustainable when communication is sequential.⁷ The analysis of the sequential communication protocol is important as a robustness check. If individuals are unable to commit themselves to the timing structure of the protocol, then some individual may break the simultaneous communication structure by attempting to "speak first."

Hence, on the one hand, our results are reassuring. Rectangular self generation is broad enough to include many if not most payoffs of interest in the full memory (non-dynastic) repeated game. On the other hand, it turns out that protocols with "neutral historians" are fragile. We show that these constructs fail when communication is no longer costless. When individuals weigh the complexity costs of the reporting strategies they use, then for any protocol, only the Nash equilibria of the stage game

⁶Lagunoff and Matsui (1997) analyze repeated coordination games in which the players move sequentially.

⁷Self generating payoff sets were first defined by Abreu, Pearce, and Stacchetti (1986) and Abreu, Pearce, and Stacchetti (1990) who used dynamic programming methods to characterize the equilibrium payoff set in a repeated game. Note also that we are somewhat abusing the meaning of the word "rectangular" here. What is required is that the lower contour of the self generating set is that of a rectangle. See the statement of Theorem 3 below.

are sustainable in the dynastic game. This is the case even when the actual payoffs from the stage game are lexicographically more important than the costs associated with a more complex reporting strategy.

In many instances, the assumption that the past history of play cannot be verified at all by the current cohort may be too extreme. It is easy to think of situations in which the past should be, at least in part observable, unless a concerted effort to hide it is made. The remnants of the Jewish Holocaust and the Stalin Purges all too quickly come to mind. To begin to address this type of set-up, we analyze a model of dynastic repeated game with a different set of assumptions about the communication protocol between one cohort and the next.

We assume history leaves a marker or “footprint” for the new generation. Efforts to manipulate information now entail effort to hide or *coverup* these footprints. We examine a protocol in which the truth can only be hidden by the current cohort when all individuals agree to the coverup. Somewhat counterintuitively, the difficulty in achieving consensus to unanimously coverup the truth may actually *increase* the incentive to hide it. We are able to show that, when communication is sequential, the set of sustainable payoffs coincides with the set of weakly renegotiation proof payoffs of the standard repeated game.⁸ The conclusion is that when the potential for *coverup* exists and when messages are sequenced, equilibria with strictly Pareto-ranked continuation payoffs cannot occur.

1.2. Outline

The material in the paper is divided into 4 further sections. In Section 2 we describe the model in detail. This includes briefly setting up the standard repeated game notation, and a complete description of the dynastic repeated game with communication. Section 3 is devoted to the analysis of our “whitewashing” model. We first analyze simultaneous and then sequential messages. We then move on to the case of lexicographic costs of more complex reporting strategies and show that only the stage game Nash equilibrium payoffs survive in this case. Section 4 is concerned with our model of “coverups.” After describing the model, we go on to show that only weakly renegotiation-proof equilibria are viable in this case. Section 5 concludes the paper with a brief discussion putting our results in the context of existing literature.

For ease of exposition, all proofs are confined to an appendix. In the numbering of equations, Lemmas, Theorems etc. a prefix of “A” means that the corresponding item is located in the Appendix.

⁸In the sense of Farrell and Maskin (1989).

2. THE MODEL

2.1. A Standard Repeated Game

We first describe a standard, n -player repeated game. We will then augment this structure to describe the dynastic repeated game with communication from one cohort to the next. The standard repeated game structure is of course familiar. We set it up below simply to establish the basic notation.

The stage game is described by the array $G = (S, u; I)$ where $I = \{1, \dots, n\}$ is the set of players, indexed by i . The n -fold cartesian product $S = \times_{i \in I} S_i$ is the set of pure action *profiles* $s = (s_1, \dots, s_n) \in S$, assumed to be finite. Stage game payoffs are defined by $u = (u_i)_{i \in I}$ where $u_i : S \rightarrow \mathbb{R}$ for each $i \in I$. Let $\sigma \in \Delta(S)$ denote a mixed action profile.⁹ The corresponding payoff to player i , denoted by U_i , is defined in the usual way: $U_i(\sigma) = \sum_s \sigma(s) u_i(s)$. Dropping the i subscript and writing $U(\sigma)$ gives the entire profile of payoffs. Finally, we let \mathcal{N} denote the set of Nash equilibria of the stage game.

In the repeated game, denote the behavior profile at time t by $\sigma(t) = (\sigma_1(t), \dots, \sigma_n(t))$. For $t \geq 1$, a period t *behavior history* (of length t) is an array $h^t \equiv (\sigma(0), \sigma(1), \dots, \sigma(t-1))$ of action profiles observed by time t . The null history is $h^0 = \emptyset$. Let $U_i(\sigma(t))$ denote the expected payoff at date t . The set of period t behavior histories is denoted by $H^t = \Delta(S)^t$. Let $H = \cup_{t=0}^{\infty} H^t$ denoting the collection of all (finite) behavior histories.¹⁰

The players' (for simplicity) common discount factor is denoted by $\delta \in (0, 1)$, so that for a given infinite history $h^\infty = (\sigma(0), \sigma(1), \dots)$, player i 's payoff in the repeated game is given by

$$V_i(h^\infty) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t U_i(\sigma(t)) \quad (1)$$

A behavior strategy in the repeated game is a map $f_i : H \rightarrow \Delta(S_i)$. Let $f = (f_1, \dots, f_n)$ denote the profile of strategies in the repeated game. Given any finite history h^t , the mixed action at date t is given by $f(h^t) = (f_1(h^t), \dots, f_n(h^t))$.

⁹At the expense of some extra notation and further manipulations we could consider *correlated* action profiles in the stage game without altering the nature of our results below.

¹⁰Therefore we are assuming that actual mixed strategies are observed. This simplifies our framework, and particularly notation, considerably. However, we later argue that none of the results in the paper depend on the assumption that mixed strategies are observed.

Given (1) the continuation payoff to i given strategy profile f after any history h^t follows the recursive equation

$$V_i(f|h^t) = (1 - \delta)U_i(f(h^t)) + \delta V_i(f|h^t, f(h^t)) \quad (2)$$

where $(h^t, f(h^t))$ denotes the period $t+1$ history given by the concatenation of history h^t and period $t+1$ behavior profile $f(h^t)$. Dropping the player subscript in equations (1) and (2) will denote the corresponding payoff profiles in the repeated game.

A *subgame perfect equilibrium* (SPE), f^* , for the repeated game is defined in the usual way: for each i , and each finite history h^t , and each strategy f_i for i , we require that $V_i(f^*|h^t) \geq V_i(f_i, f_{-i}^*|h^t)$.¹¹

We denote respectively by $\mathcal{F}(\delta)$ the set of SPE strategy profiles and by $\mathcal{E}(\delta)$ the set of SPE payoff profiles of the repeated game when the common discount factor is δ .

The standard model of repeated play we have just sketched out may be found in a myriad of sources. See, for example, Fudenberg and Maskin (1986) and the references contained therein. Hereafter, we refer to the standard repeated game model above as the *full memory repeated game*.

2.2. The Dynastic Repeated Game

Now assume that each $i \in I$ indexes an entire progeny of individuals. We refer to each of these as a *dynasty*. Individuals in each dynasty are assumed to live one period. At the end of each period t (the beginning of period $t+1$), a new individual from each dynasty — the date $(t+1)$ -lived individual — is born and replaces the date t lived individual in the same dynasty. Hence, $U_i(\sigma(t))$ now refers to the payoff received by the t -th individual in dynasty i . Each date t individual is altruistic in the sense that his payoff includes, as an additively separable argument, the utility of the $t+1$ -th individual from the same dynasty. The weight given to his own payoff is $1 - \delta$, while the weight given to his offspring (the $(t+1)$ -th individual) is δ . Therefore, in the dynastic repeated game, the long-run payoffs retain the same recursive structure given in equation (2). This observation is of course sufficient to show that, if all individuals in each dynasty can observe the past history of play, then the dynastic repeated game is in fact identical to the full memory repeated game described in Subsection 2.1. In fact, this full information dynastic repeated game is one extremely appealing interpretation/justification of the standard full memory repeated game model.

¹¹As is standard, here, and throughout the rest of the paper, a subscript of $-i$ indicates an array with the i -th element taken out.

2.3. Communication

If the t -th individuals in each dynasty cannot observe the history of play that took place before they are born, then their behavior cannot vary across distinct histories h^t . It follows immediately that, in the absence of communication from one cohort to the next, if all individuals in all dynasties are ignorant in this way, in each period of the dynastic repeated game only those payoffs that are Nash equilibria of the stage game can be attained.

The question then becomes: can communication substitute for memory? Assume that each time t a cohort *can* observe the action profile that takes place at t . Assume also that they have the chance to communicate with the next cohort. Can they credibly convey sufficient information to the $t + 1$ -th cohort to attain payoffs beyond the Nash equilibria of the stage game?

Of course, communication can take place in a variety of different ways. For instance, as we anticipated above, whether messages are sent simultaneously or sequentially will have an impact on the outcome of the game. We begin by defining a model of the communication between one cohort and the next in which the individuals in each dynasty speak simultaneously to the individuals in the next cohort. This will be modified in Subsection 3.2 to allow for sequential communication.

Let A_i denote a set of payoff-irrelevant communication actions for dynasty i , with $A = \times_i A_i$ being the set of profiles of such actions. These need not be related to the stage game itself, but are choices that collectively determine a message sent from one generation to the next. The messages that can possibly be sent to the next cohort is given by the set M . We assume that both A and M are invariant across time.¹² At each date t , let $a(t) = (a_1(t), \dots, a_n(t))$ denote the profile of communication actions, and let $m(t) \in M$ denote the message (or profile of messages) sent in period t .

Unless otherwise noted, we will assume that any message(s) transmitted to the next cohort is commonly observed by all members of the next cohort.¹³ A *communication protocol* in the dynastic repeated game is a list $C \equiv (A, M, \Phi)$ where $\Phi : A \rightarrow M$.

¹²It is worth emphasizing here that this assumption implies that “calendar time” cannot be “verified” by any cohort. For instance, the time t and the time $2t$ individuals could generate exactly the same message to be transmitted to the next cohort. In this case, calendar time would clearly “look the same” to the individuals in cohort $t + 1$ and in cohort $2t + 1$. A rather strong feature of this assumption, which may seem unappealing to some, is that any date t individuals could in fact generate a message for the next cohort which matches the null history \emptyset . In effect they would be telling their offspring that “the game has just started, you are the first to play.”

The qualitative nature of our results below would remain unaffected if somehow these “re-starting” messages were forbidden. In some cases (for instance Theorem 5 below), period $t = 0$ would have to be excluded from our characterization of equilibria which would only apply to periods $t \geq 1$.

¹³See our discussion of this assumption in the concluding section of the paper.

For $a = (a_1, \dots, a_n) \in A$, $\Phi(a)$ is the message sent to the next generation, after the payoff-relevant behavior occurs in that period. Some “natural” examples of features that communication protocols may satisfy are:

C^1 **Babbling:** $\Phi(a) = m$ for all $a \in A$. Clearly, the message m is uninformative.

C^2 **Dictatorial:** There exists a dynasty $i \in I$ such that $\Phi(a) = a_i$.

C^3 **Unanimity:** $A_i = M = H$, $\forall i$, and

$$\Phi(a) = \begin{cases} h & \text{if } a_i = h, \forall i \\ h^* & \text{otherwise.} \end{cases}$$

Here, a single version of history is sent if everyone agrees. Otherwise, a default history is reported.

C^4 **Decentralized Communication:** $M = \times_{i \in I} A_i$, $H \subseteq A_i$, $\forall i$, and $\Phi(m) = m$. In a decentralized communication protocol, everyone separately reports history to next generation.

While the first three examples fulfill an expository function, the last is a useful benchmark. It describes the least restrictive communication, and so it provides the most attractive environment for accurate transmission. We examine the case of decentralized communication protocols in detail below.

Every communication protocol identifies a dynastic repeated game with communication. A strategy for an individual in a dynasty is a pair consisting of an “action” strategy and a “communication” strategy. The former processes the messages received from the prior generation and determines current behavior in the stage game. The latter determines the individual’s communication action, which, via the communication protocol, determines the message conveyed to the next cohort of individuals.

We begin by defining action strategies. For simplicity, we examine action strategies that can be written as a single repeated game strategy describing the plans of all individuals within a dynasty. When $H \subseteq M$, there is little loss of generality with this assumption since an individual from, say, the t -th cohort, need only use that part of the repeated game strategy which follows histories h^t of length t . Let $g_i : M \rightarrow \Delta(S_i)$ denote an action strategy for dynasty i . Let $g = (g_1, \dots, g_n)$.

A communication strategy is a map μ_i from the prior generation’s messages and current (observed) actions to current messages. Formally, $\mu_i(m, \sigma)$ denotes a communication action $a_i \in A_i$ by an individual from dynasty i given that the prior

generation's message profile is m and that the current action profile is σ . The profile $(\mu_1(m, \sigma), \dots, \mu_n(m, \sigma))$ then maps to a message m' via Φ . This message m' is sent to the next generation. Let $\mu = (\mu_1, \dots, \mu_n)$. To summarize, date t individuals choose action profile $\sigma(t) = g(m(t-1))$ and take communication actions $a(t) = \mu(m(t-1), \sigma(t))$.

As with the full memory repeated game, something to start off play is needed. (In the full memory repeated game this is the empty history $h^0 = \emptyset$.) In the dynastic repeated game with communication we need to define which message the first (born at $t = 0$) cohort observes. Let this initial message be denoted by $m(-1) = h^0 = \emptyset$.

The pair (g, μ) describes all behavior in the dynastic repeated game with communication. An individual's dynamic payoff after receiving message m is expressed as $V_i(g, \mu | m)$. An individual's dynamic payoff after receiving message m and after action profile σ is expressed as $V_i(g, \mu | m, \sigma)$. In either case, an individual's payoff still follows the recursive form in (2). We can now define a Perfect Bayesian equilibrium (PBE) pair (g^*, μ^*) in the usual way: for each i , any μ_i and g_i , and for any m and any σ , $V_i(g^*, \mu^* | m) \geq V_i(g_{-i}^*, g_i, \mu^* | m)$ and $V_i(g^*, \mu^* | m, \sigma) \geq V_i(g^*, \mu_{-i}^*, \mu_i | m, \sigma)$.¹⁴

Given a communication protocol C and a common discount factor δ , we denote by $\mathcal{F}^C(\delta)$ the set of PBE, and by $\mathcal{E}^C(\delta)$ the set of PBE payoff profiles. Let C^1 denote the babbling protocol described above with $m = \emptyset$, and recall that C^4 denotes the decentralized protocol. Then it is easy to see that

$$\mathcal{E}^{C^1}(\delta) \subseteq \mathcal{E}^C(\delta) \subseteq \mathcal{E}^{C^4}(\delta) \quad (3)$$

for all δ and all C . Moreover, $\mathcal{E}^{C^1}(\delta)$ coincides with Nash equilibrium payoffs of the stage game. Clearly, if there is no communication, or when communication is uninformative, there is no hope of attaining anything beyond payoffs of the stage game. Conversely, if there are no restrictions on communication, then the largest possible payoff set can be sustained.

3. WHITEWASHING

In this section we focus on the case in which the current cohort has no access at all to any direct information about the past history of play. We examine first the case of simultaneous messages. Then we move on to the case in which the members of the current cohort speak sequentially to the next cohort. Finally, we turn to

¹⁴While there are no proper subgames beginning each date, all individuals have common knowledge regarding messages they receive. Hence, all Perfect Bayesian equilibria are Perfect *Public* equilibria of the dynastic game. Consequently, the definitions above make (partial) use of the "One-Shot Deviation Principle."

the case of sequential communication in which a more complex reporting strategy is lexicographically more costly than a simpler one.

3.1. *Decentralized Communication*

We now proceed to examine equilibrium behavior under the *decentralized communication protocol* defined in Section 2.3 above. Recall that in the decentralized communication protocol, all individuals in each cohort effectively report separately and simultaneously a history of play to the next generation, and all reports are commonly observed by all individuals in the next cohort. The messages are unrestricted in the sense that any finite history (of any length) can be conveyed by any individual (in other words $H \subseteq A_i$ for every i). Intuitively, a decentralized protocol corresponds to a world in which there is no attempt to collectively limit information, nor is there any direct trace left by the actual history of play. This is of course in contrast with the model of coverups that we will analyze in Section 4 below.

Our first characterization of the equilibrium set of dynastic games with communication tells us that in the case of decentralized simultaneous communication from one cohort to the next, the equilibrium set is the same as in the full memory standard repeated game. In other words, in this case *communication* can indeed *substitute for memory* in the dynastic repeated game. Those familiar with standard implementation logic will not find the following Theorem and its immediate Corollary very surprising.

It should be made clear at the outset that we are setting up Theorem 1 below as a benchmark case. Our task in the rest of this Section and in Section 4 will then be to show that this equivalence between the full memory standard repeated game and the dynastic repeated game is in fact not robust in more sense than one. Indeed Theorem 1 fails in a rather dramatic way when communication carries arbitrarily small costs and when the actual past history of play does leave a detectable trace.¹⁵

THEOREM 1: *Assume that the number of players n is at least 3. Fix any common discount factor δ and any SPE $f^* \in \mathcal{F}(\delta)$ of the full memory game.*

¹⁵Note also that we state Theorem 1 for the case of 3 or more players, ignoring the two-player case. This result can be generalized to the case of two players (a slightly weaker statement holds) for discount factors arbitrarily close to 1. Proceeding in this way saves a non-negligible amount of space. The proof of the result for the two-player case involves mimicking the “mutual minmax” argument used to prove Folk Theorems in two-player standard repeated games, and thus a substantial modification of the “cross-checking” argument that we outline below.

Since, as we just said, Theorem 1 largely plays the role of a straw man in what follows, we take the view that stating it only for the case of three or more players carries little or no cost for the sharpness of the overall message we are trying to convey.

Then, for any decentralized communication protocol C , there exists a PBE $(g^*, \mu^*) \in \mathcal{F}^C(\delta)$ of the dynastic game with communication protocol C that is equivalent to f^* in the following sense.

For each m such that $m = (h, \dots, h)$ for some $h \in H$ and for every $\sigma \in \Delta(s)$

$$g_i^*(m) = f_i^*(h) \quad \text{and} \quad \mu_i^*(m, \sigma) = (h, \sigma) \quad \forall i \quad (4)$$

In other words, provided that the messages sent by the previous cohort are all the same (equal to h), then the PBE (g^*, μ^*) prescribes the same actual behavior as the SPE f^* after h . Moreover, again provided that the messages sent by the previous cohort are all the same (equal to h), then in the PBE (g^*, μ^*) , any profile of current behavior (equilibrium or not) is truthfully reported to the next cohort by all i .

Finally, notice that (4) since the first cohort all receive the same (empty) message, obviously implies that the outcome path generated by (g^*, μ^*) is the same as the outcome path generated by f^* .

The proof of Theorem 1 is in the Appendix. As we mentioned above, it runs along familiar lines. The argument requires building into the equilibrium of the dynastic repeated game the correct incentives for truthful reporting by all individuals in each cohort.¹⁶ In turn, this of course requires a mechanism to detect and punish lies. Unilateral deviations from truthful reporting are easily identified when there are three dynasties or more. Since only single-player deviations from equilibrium ever need to be considered this is enough to induce truth-telling as required. Finally, since the argument constructs incentives for truthful reporting, taking as given the corresponding (action) incentives in the full memory game, Theorem 1 does not depend on our assumption that mixed strategies are observable.¹⁷ Indeed, we require only that *whatever* is observable in the full memory game can be observed (period-by-period) and reported in the dynastic game.

It is trivial that any equilibrium of the dynastic game with communication is also an equilibrium of the full memory game. Therefore, an immediate corollary of Theorem 1 is that the sets of equilibrium payoff profiles are the same in the two cases.

¹⁶Notice that all the equilibria that we construct in the paper are “truthful” in the sense that, without loss of generality, provided that $H \subseteq A_i$ for every i (so that the space of message actions is rich enough), we can take it to be the case that, in equilibrium, the members of each cohort report truthfully to the next generation. Truthfully, of course, can only mean that a given one-to-one map of action messages into histories is used throughout. This is a weak version of the so-called “revelation principle” that obviously holds in our model.

¹⁷For exactly the same reasons this is also true in all our other results below. See also footnote 10 above.

COROLLARY 1: *Assume that $n \geq 3$ and let any decentralized communication protocol C and any common discount factor δ be given. Then $\mathcal{E}^C(\delta) = \mathcal{E}(\delta)$.*

Before we turn to sequential communication, some remarks about Theorem 1 and Corollary 1 are in order. First, the “cross-checking” aspect of the mechanism we construct for three or more players in the proof is by no means new. Baliga, Corchon, and Sjostrom (1997) use this type of mechanism in a (static) model of communication. It is also reminiscent of communication mechanisms in repeated games with private monitoring. For example, Ben-Porath and Kahneman (1996) prove a Folk Theorem when public communication is admissible in a repeated game with private monitoring.¹⁸ Specifically, they show that the Folk Theorem applies in any private monitoring game in which individual’s behavior is (perfectly) observed by at least two others. Like ours, their proof also exploits a procedure whereby the deviator is identified as the one whose report fails to correspond to identical reports of at least two others.

Secondly, if we allow players in the dynastic game to (independently) *randomize* their communication actions, then Theorem 1 and Corollary 1 would have to be modified to take into account the fact that messages now become a potential coordination device. In particular, it is not hard to see that using random messages the players could achieve a randomization across different Nash equilibria of the stage game in a given period. This, of course, would alter (enlarge) the set of achievable long-run payoffs when δ is bounded away from 1.¹⁹ However the effect of randomized messages would become negligible as the discount factor approaches 1 since playing the different Nash equilibria in sequence through time would have the same effect on long-run payoffs. For simplicity, and because our main question here is whether the messages can substitute for memory, we focus on pure message strategies throughout the paper.

3.2. *Coordination Failure and Sequential Communication*

Clearly, equilibria in the full memory game represent the outer bound of what is possible in the dynastic repeated game. As we know from Theorem 1 and Corollary 1 decentralized communication with simultaneous choice of message actions achieves this bound. In a sense, this equivalence relies on a dependence on the simultaneity of messages, which is unsettling for at least two reasons. First of all, simultaneous messages may be simply not feasible.

Secondly, there may be incentives for the players to depart from any simultaneous communication protocol. Consider for instance the repeated PD game in Figure

¹⁸See also the papers by Compte (1998), and Kandori and Matsushima (1998).

¹⁹We are grateful to an anonymous referee for pointing out this fact.

1. A standard way of supporting mutual cooperation utilizes a joint punishment (e.g., permanent reversion to the unique equilibrium of the stage game) as a way to deter deviations from the path of perpetual cooperation. However, once a defection in behavior takes place, everyone, including the “injured parties,” would prefer to whitewash the history of defection. It is clear that the simultaneous structure of the communication protocol prevents any agent from signaling his intent to falsify the truth. If communication were sequential then the intent to “whitewash” could be relayed by one member of the current cohort to the other. Consequently, incentives to sequence messages may arise, especially if the timing of communication is not observed by subsequent cohorts.

In this Section we explore the consequences of assuming that message actions are taken *sequentially* by individuals in a given cohort.

To keep matters simple, we examine the simplest class of protocols with sequential choice of action messages. We consider the class of communication protocols that are decentralized in the sense that we specified above, but modified so that the individuals in each cohort choose their message actions one after the other.

DEFINITION 1: *We say that a communication protocol is a sequential decentralized communication protocol if it can be obtained as the following simple modification of a (simultaneous) communication protocol that is decentralized in the sense of protocol C^4 described above.*

There exists a permutation mapping $\theta : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ that describes the order in which the message actions are chosen.²⁰ In other words, first individual $\theta(1)$ chooses a message action $a_{\theta(1)} \in A_{\theta(1)}$. Immediately after, all other individuals in the cohort observe $a_{\theta(1)}$. Then individual $\theta(2)$ chooses a message action $a_{\theta(2)} \in A_{\theta(2)}$ which is then observed by all other individuals in the same cohort. The choice of action messages then continues in this fashion until individual $\theta(n)$ makes his choice.

The rest of the details of a sequential decentralized communication protocol are as in Example C^4 above.

With a sequential decentralized communication protocol, Theorem 1 no longer holds. Some equilibria of the full memory game are destroyed by the incentives to whitewash the past. For example, consider the stage game in Figure 2 below.

²⁰The actual permutation mapping θ is irrelevant for all our results below. Unless we specify otherwise in what follows we assume that θ is in fact the natural order so that $\theta(i) = i$.

		Dynasty 2		
		<i>L</i>	<i>M</i>	<i>R</i>
Dynasty 1	<i>T</i>	2, 2	0, 3	0, 0
	<i>M</i>	3, 0	1, 1	0, 0
	<i>B</i>	0, 0	0, 0	0, 0

Figure 2: A 3×3 Stage Game

In the full memory game, every payoff profile in the strictly individually rational set $\{(v_1, v_2) : v_i > 0, \forall i = 1, 2\}$ is sustainable as an SPE if δ is close enough to one.²¹

Now suppose that two dynasties play this game with sequential decentralized communication. We assert that any equilibrium in which a deviation is countered with permanent reversion to the worst Nash equilibrium (B, R) cannot be sustained. Consider the perpetual repetition of (T, L) each period. Suppose further that at some date t , the date t member of Dynasty 1 chooses to “cheat” by deviating to M . Now both members of generation t communicate sequentially to date $t + 1$ individuals. Dynasty 1 communicates first. If the date t player 1 whitewashes by lying about his defection then player 2 will confirm the lie *unless* player 2 can be made at least as well off by telling the truth about (M, L) taken in the current period. This means that in the continuation, player 2 must receive a payoff of at least 2 for telling the truth. However, if he receives a payoff of more than 2, then player 2 will always report “ (M, L) ” even when “ (T, L) ” was the true action taken.

Notice that since communication is cheap talk, the structure of the reporting “subgame” remains the same after every history.²² Hence, the set of continuation equilibria after the reporting stage must remain the same. Yet, the reporting “subgame” will typically have multiple “subgame perfect” equilibria, one for each possible history in the game. In this particular game, this means that player 2 must be *indifferent* between truthful reporting and whitewashing. Therefore, player 2’s continuation payoff in the putative continuation is 2 regardless of whether he reports “ (T, L) ” or “ (M, L) .”

²¹Notice that, as is often the case, the (lower) boundary of the individually rational set is not sustainable, except of course for the origin $(0, 0)$ which is a Nash equilibrium of the stage game.

²²We put “subgame” in quotes since the reporting stage is not literally a proper subgame of the dynastic game. Nevertheless one can refer to the subgame perfect equilibria of the extensive form reporting game, whose terminal payoffs are equilibrium continuations.




Figure 3

Moreover, since in the putative equilibrium, “(0, 0)” is the hypothesized continuation after a deviation, player 1 must have an incentive to truthfully report “(M, L)” after his own deviation. But since player 2 will truthfully report “(M, L)” should player 1 attempt to lie, the continuation payoff profile after the sequence of reports “((T, L), (M, L))” must be (0, 2). That is, if player 1 attempts to lie, and player 2 reports the truth, then player 2 must be indifferent between his reports, and player 1 must no better off than if he told the truth. In the latter case, player 1 cannot be strictly worse off since his lowest feasible payoff is 0, and so he too is indifferent between his reports.

The problem, however, is that (0, 2) is not a feasible SPE continuation payoff in the game. Hence, the mutual cooperation equilibrium using these particular punishments cannot be sustained. The relevant portion of the game tree in the message phase is represented in a schematic way in Figure 3 above.²³

²³Clearly, each player has 9 choices at each node in the complete version of the tree drawn in Figure 3 — one for each possible outcome of play at $t - 1$. We have only represented the two relevant ones purely for the sake of visual clarity.

3.3. Sequential Communication: The Necessity of a “Neutral Historian”

The above example demonstrates that, the potential for sequential moves to break the “coordination failure” that prevented whitewashing in the simultaneous case, is enough for whitewashing to occur in some cases. The example also demonstrates a general property of perfect equilibria when communication is sequential. To characterize this property, some further notation is required.

We will denote by a^{i-1} the $(i-1)$ -tuple of message actions chosen by individuals $1, \dots, i-1$ (by convention, set $a^0 = \emptyset$). In this way we can then write the strategy of individual i in the communication round as determining $a_i = \mu_i(m, \sigma, a^{i-1})$. In other words, i chooses his message action as a function of the message $m = (m_1, \dots, m_n)$ sent by the previous cohort, the current action profile σ , and the message actions a^{i-1} chosen by individuals 1 through to $i-1$ in the current cohort. In this way μ determines an entire path a of message actions that can be written as

$$a_1 = \mu_1(m, \sigma, a^0), a_2 = \mu_2(m, \sigma, a^1), \dots, a_i = \mu_i(m, \sigma, a^{i-1}), \dots, a_n = \mu_n(m, \sigma, a^{n-1})$$

In any sequential decentralized protocol, the choice of profile (g, μ) determines a reporting “subgame” in the communication round. This “subgame” can be viewed as an extensive form game of perfect information. The “terminal nodes” of this extensive form game are the equilibrium continuations that begin with the next generation’s play of the game. Since communication is cheap talk, the terminal nodes of this game do not vary with past reports and play. We write $v = V(g, \mu | a)$ to denote the payoff vector associated with the terminal node reached by path a given that the continuation of play is determined by (g, μ) . When the path has only been determined up to, say, player j ’s report, i.e., when the path is $a^j = (a_1, \dots, a_{j-1})$, then players’ strategies μ_{j+1}, \dots, μ_n are used to determine the terminal node. Since these strategies depend, in turn, on prior history (m, σ) , we must write $V(g, \mu | m, \sigma, a^j)$ to denote the terminal node reached by path a^j given history (m, σ) and strategies (g, μ) . In the result below, the *reporting subgame* refers to the induced extensive form game in the message phase whose terminal payoffs are equilibrium continuation payoffs.

THEOREM 2: (*Necessity of a “Neutral Historian”*) *Fix a sequential, decentralized protocol C , and fix a PBE, $(g, \mu) \in \mathcal{F}^C$. Let (m', σ') and (m'', σ'') denote the two prior histories of message and action profiles, respectively, each of which generate distinct message paths. Then, there is a path a , a player j , and a pair of distinct action messages, a'_j and a''_j , for that player such that $\mu_j(m', \sigma', a^{j-1}) = a'_j$ and $\mu_j(m'', \sigma'', a^{j-1}) = a''_j$ and*

$$V_j(g, \mu | m', \sigma', a^{j-1}, a'_j) = V_j(g, \mu | m'', \sigma'', a^{j-1}, a''_j)$$

In words, there exists some path a and some player (a “neutral historian”) who distinguishes between (m', σ') and (m'', σ'') by using two action messages that generate payoff equivalent continuations.

The result gives necessary conditions for the existence of multiple equilibria of the message game. This is of crucial importance since the message game must have multiple equilibria in order to successfully punish deviations. That is, since the next generation has no independent verification of the actual history, both the continuation from prescribed play and the continuations after deviations must all be equilibrium continuations of the reporting game. The older generation must be able to coordinate either on the original equilibrium path continuation, or on the punishment continuation. Generally, there will be at least $n + 1$ equilibria of the message game. One for the equilibrium prescription, and one each for each individual’s deviation from equilibrium.

However, because the reporting game cannot vary with prior history, the multiplicity of equilibria in the reporting game implies that subgame is nongeneric. Somewhere, there must be “ties,” in someone’s payoff, and these ties must have the structure described in Theorem 2. This structure implies that in every PBE (g^*, μ^*) , the reporting subgame must utilize a “neutral historian.”

Specifically, in any pair of equilibrium paths, there exists someone whose potential veto of one of the paths is governed by the later choice of a “neutral historian,” player j . The neutral historian is one who uses his indifference between two terminal payoffs (i.e., his “neutrality”) to influence the earlier messages of others. In essence, the neutral historian screens and verifies the truthfulness of others’ reports, while remaining indifferent in the process. An explicit construction that identifies the behavior of this neutral historian is given in the proof of Theorem 3 in the next Section.

3.4. Self Generation

Using a, by now standard, notion of self generating sets (Abreu 1988, Abreu, Pearce, and Stacchetti 1986, Abreu, Pearce, and Stacchetti 1990), one can construct equilibrium payoff sets for sequential protocols. Formally, a set \mathcal{V} is *self-generating* if for each $v \in \mathcal{V}$ there exists a $\sigma \in \Delta(S)$ and a map $w : \Delta(S) \rightarrow \mathcal{V}$ such that (i) $v = (1 - \delta)U(\sigma) + \delta w(\sigma)$, and (ii) $v_i \geq (1 - \delta)U(\sigma'_i, \sigma_{-i}) + \delta w_i(\sigma'_i, \sigma_{-i})$ for each i and each σ'_i . It is not difficult to show that all bounded self generating sets are contained in the SPE profiles, $\mathcal{E}(\delta)$, of the standard repeated game. Indeed, $\mathcal{E}(\delta)$ itself, as well as any stage Nash payoff vector are self generating sets. The following Theorem justifies our interest in a special type of self generating sets: those that have the lower boundary of a rectangle in \mathbb{R}^n . It asserts that any such self generating subset is sustainable in any sequential, decentralized protocol provided that the protocol allows

each individual to report histories and continuation payoffs to the next generation of players.

THEOREM 3: Fix any $0 < \delta < 1$, and in the full memory game let \mathcal{V} be a compact, self-generating set with the following property. For any i , and any $\hat{v} \in \mathcal{V}$, let $\underline{v}_i = \min v_i$ such that $(v_i, v_{-i}) \in \mathcal{V}$ for some v_{-i} . Then for every $\hat{v} \in \mathcal{V}$ we have that

$$(\underline{v}_i^*, \hat{v}_{-i}) \in \mathcal{V} \tag{5}$$

Then, for any sequential, decentralized communication protocol, C , in which $H \subseteq A_i$ for each i .

$$\mathcal{V} \subseteq \mathcal{E}^C(\delta)$$

The proof in the Appendix gives an explicit account of how the “neutral historian” is used to sustain payoffs in the self generating set. Roughly speaking, given any point in \mathcal{V} , a protocol is constructed in which player 2 (who speaks second), serves as the neutral historian. After player 1’s (who speaks first) report, player 2 is asked to “confirm” it or not. (The messages of all other players are ignored.) If player 2 confirms 1’s report, then play unfolds as dictated by a particular SPE of the full memory game. If on the other hand player 1 deviates from reporting the truth, this is treated as if he had behaviorally deviated from equilibrium play, and he is punished by being awarded the lowest possible payoff in \mathcal{V} . The equilibrium is also constructed so that player 2 is always *indifferent* between confirming 1’s report or not. Thus, exploiting the properties of \mathcal{V} , both 1 and 2 are given the correct incentives to report the truth and to serve as the “neutral historian” respectively.

The question of whether any particular subset of the equilibrium payoff set $\mathcal{E}(\delta)$ that has the lower boundary of a rectangle is self generating is open. However, particular examples of such self generating sets are not difficult to construct. Clearly, individual Nash equilibrium payoff profiles are degenerate self-generating rectangles. To see that more interesting self generating sets that have the lower boundary of a rectangle are common, consider once again the Prisoner’s Dilemma game in Figure 1 in the Introduction.

We claim that for $\delta \geq 4/7$ the rectangle $\{v : (1, 1) \leq v \leq (2, 2)\}$ is self generating. Notice that this includes mutual cooperation. While we don’t verify the property for all profiles in the set, we can easily do so for the extremal ones. First consider $(2, 2)$. Consider the following simple penal code. Any deviation from $(2, 2)$ by, say, a player 1, is followed by a one period punishment in which they play: $\sigma_1(C) = 4/5$, $\sigma_2(C) = 1/3$. Following this, the players revert to mutual cooperation. Further deviations by player 1 restart the punishment; deviations by player 2 are countered by the same punishments (switching the players’ roles).

We now verify that this is a SPE penal code when $\delta = 4/7$. In fact, by extending the length of the punishment, this type of penal code also works for discount factors greater than $4/7$ as well.²⁴

While the one shot punishment gives player 1 a stage payoff of $-1/3$, his dynamic payoff in the punishment phase is 1 since $1 = (1 - \delta)(-1/3) + \delta 2$ when $\delta = 4/7$. The payoff to player 2 in player 1's punishment phase is $(1 - \delta)2 + \delta 2 = 2$. Note that player 1 obtains payoff of unity from perpetual deviation. Hence, a (weak) best response is to submit to punishment in order to obtain 2 in the continuation. This penal code simultaneously verifies that the profiles $(2, 2)$, $(1, 2)$, and $(2, 1)$ are all SPE profiles (since punishment continuations are obviously SPE outcomes as well). To see that $(1, 1)$ is also a SPE profile, we construct a penal code which supports $(7/4, 7/4)$ as follows: play $\sigma_i(C) = .914$ (approximately) each period. Any deviation is met with one period reversion to the Nash equilibrium, after which time play resumes as before. This can be verified to be a SPE penal code. Since the value of the punishment profile is $(1, 1)$, it constitutes a SPE continuation. Other payoffs in the rectangle may be shown to be sustained more easily since the punishment continuation will generally be higher, and the one shot gain to deviation will generally be lower.²⁵

3.5. Costly Communication

Our discussion so far of sequential decentralized communication has yielded two insights. First of all the potential to whitewash *does* have an impact on the structure of equilibria of the repeated game. Theorem 1 no longer holds in this case. Some equilibria of the full memory game are not viable under sequential decentralized communication because they would leave one or more individuals with an incentive to whitewash the past after certain histories of play.

On the other hand, the logic of Theorem 3 demonstrates that whitewashing can be prevented, even if it makes all current and future generations better off, and even when individuals can signal their intentions to coordinate on the whitewashing of previous deviations. Since our examples suggest that most payoffs of interest (e.g., mutual cooperation in Prisoner's Dilemma) can be sustained by rectangular self generation if players are sufficiently patient, these payoffs are also sustainable by equilibria of sequential protocols. Our next step is to show that such equilibria are fragile.

We modify the dynastic game with sequential decentralized communication in the following way. We assume that an "infinitesimal" cost is associated with communication strategies that are more "complex". In other words, we assume that there are

²⁴Notice that we only establish $\delta = 4/7$ as a lower bound for self-generation when Nash reversion is the punishment.

²⁵The same logic that sustains $(7/4, 7/4)$ can be applied to verify the self generating property for other interior points in the rectangle.

costs attached to more complex communication strategies, but that they matter in the comparison between the payoffs that two strategies yield *only* if these two strategies yield actual payoffs (from actual play that is) that are the *same*. Using a standard term, we call these *lexicographic* costs of more complex communication strategies.

Once the model is modified to allow for lexicographic costs of more complex communication strategies, it behaves in a dramatically different way. Theorem 4 asserts that once complexity costs are taken into account the set of PBE of the dynastic game shrinks to \mathcal{N} — the set of Nash equilibria of the stage game.

We begin with a definition of what it means for a strategy to be more complex than another at the communication stage. We want to deem a communication strategy to be more complex than another if it prescribes communication actions that depend “more finely” on the history of play. To describe formally what we mean by more finely some extra notation is required.

Recall that with sequential decentralized communication the message action of individual i is denoted by $a_i = \mu_i(m, \sigma, a^{i-1})$, where $m = (m_1, \dots, m_n)$ is the message sent by the previous cohort, σ is the current action profile, and a^{i-1} is the profile of message actions chosen by individuals 1 through to $i - 1$ in the current cohort. Let \mathcal{M}_i be the set of all possible tuples (m, σ, a^{i-1}) .²⁶

Given a communication strategy μ_i we can of course identify the way in which μ_i *partitions* \mathcal{M}_i . We let this partition be denoted by $\mathcal{P}_i(\mu_i)$. The “cell” of $\mathcal{P}_i(\mu_i)$ that contains any given $(m, \sigma, a^{i-1}) \in \mathcal{M}_i$ is denoted by $\lambda_i(m, \sigma, a^{i-1})$ and is defined as follows.

$$\lambda_i(m, \sigma, a^{i-1}) = \{ (m', \sigma', a^{i-1'}) \in \mathcal{M}_i \mid \mu_i(m', \sigma', a^{i-1'}) = \mu_i(m, \sigma, a^{i-1}) \} \quad (6)$$

Lastly, as is standard, if the partition $\mathcal{P}_i(\mu_i)$ is *coarser* than the partition $\mathcal{P}_i(\mu'_i)$ we write $\mathcal{P}_i(\mu_i) \succ \mathcal{P}_i(\mu'_i)$.²⁷

We can now proceed with a formal definition of the assertion that a communication strategy is more complex than another.²⁸

²⁶So, to be precise we have that $\mathcal{M}_i = M \times \Delta(S) \times A^{i-1}$, where $A^{i-1} = A_1 \times \dots \times A_{i-1}$ if $i \geq 2$, and $A^{i-1} = \emptyset$ if $i = 1$.

²⁷Of course in this case we may also say that $\mathcal{P}_i(\mu'_i)$ is *finer* than $\mathcal{P}_i(\mu_i)$.

²⁸The notion of complexity embodied in Definition 2 below is related to the definition of complexity based on the number of states in an automaton needed to implement a strategy (Rubinstein 1986, Abreu and Rubinstein 1988, Aumann and Sorin 1989, Piccione 1992, Rubinstein and Piccione 1993, Chatterjee and Sabourian 2000, among others). However, it should be noted that the two are not the same. The reasons are two-fold. First of all, we do not restrict attention to strategies that are at all implementable by a finite automaton. In this sense the “domain” of Definition 2 is broader than the ones based on counting states in a finite automaton. Secondly, our definition below is “weaker”

DEFINITION 2: We say that communication strategy μ'_i is more complex than communication strategy μ_i if and only if $\mathcal{P}_i(\mu_i) \succ \mathcal{P}_i(\mu'_i)$.²⁹

As we anticipated above, we assume that whenever the payoffs stemming from the (repeated) stage-game are equal, communication strategies that are less complex in the sense of Definition 2 are preferred. The easiest way to include our assumption of lexicographic costs of more complex reporting strategies is to modify the definition of an PBE for the dynastic repeated game with sequential decentralized communication.

DEFINITION 3: Consider an PBE (g^*, μ^*) for the dynastic repeated game with sequential decentralized communication. We say that (g^*, μ^*) is robust to lexicographic complexity costs of communication if and only if for every dynasty i and every $(m, \sigma, a^{i-1}) \in \mathcal{M}_i$, there does not exist a communication strategy μ'_i such that

$$V_i(g^*, \mu'_i, \mu_{-i}^* | m, \sigma, a^{i-1}) = V_i(g^*, \mu^* | m, \sigma, a^{i-1}) \quad (7)$$

and μ'_i is more complex than μ_i in the sense of Definition 2.

Given a sequential decentralized communication protocol C and a common discount factor δ , the set of PBE that are robust to lexicographic complexity costs of communication will be denoted by $\tilde{\mathcal{F}}^C(\delta)$, while the corresponding set of payoffs will be denoted by $\tilde{\mathcal{E}}^C(\delta)$ throughout the rest of the paper.

Notice once again that Definition 3 embodies the idea that complexity costs of communication *only matter* if the payoffs from the (repeated) stage game are the same. A PBE is robust to lexicographic complexity costs of communication if, given the strategies of the others, no individual can choose a communication strategy that leaves his basic payoff unaffected but which has a lower degree of complexity than the equilibrium one.

than the automaton based ones in the following sense. Given two communication strategies that are implementable by a finite automaton, it is possible that one be less complex than the other in the sense that it requires fewer states, but that the two are *not comparable* in the sense of Definition 2 below. On the other hand, it is easy to check that in this case, if one strategy is less complex than the other in the sense used here, then it necessarily is less complex than the other in the sense of requiring fewer states. While counting states provides a complete order of strategies that are implementable by a finite automaton, Definition 2 only defines a partial order on this set. See also footnote 29 below.

²⁹Notice that the complexity of a communication strategy only defines a partial order on the set of communication strategies in the sense that clearly there exist pairs μ_i and μ'_i such that both $\mathcal{P}_i(\mu_i) \not\prec \mathcal{P}_i(\mu'_i)$ and $\mathcal{P}_i(\mu'_i) \not\prec \mathcal{P}_i(\mu_i)$ hold.

The idea is that players will not distinguish histories that have equal payoffs in the equilibrium continuation. Indeed, why would they expend energy to do otherwise? But if a player does not make fine distinctions between otherwise identical histories, then he must play the same way after each such history of play.

Notice that the way we have incorporated the role of complexity costs into the equilibrium notion for our model is in some strong sense the *weakest* possible one. If we modelled the complexity costs of communication to be even small but positive, their impact on the equilibrium set could not be smaller than in the lexicographic case we are considering here. In this sense, Theorem 4 below refers to the *limit* case in which complexity costs of communication have been shrunk to zero.

THEOREM 4: *Let any sequential decentralized communication protocol C and any common discount factor δ be given. Then any PBE (g^*, μ^*) that is robust to lexicographic complexity costs of communication has the following features. The action profile taken in any subgame $\sigma = g^*(m)$ is a Nash equilibrium of the stage game (in other words $g^*(m) \in \mathcal{N}$ for every $m \in M$). Moreover the action profile σ is the same in every period along the equilibrium path — except possibly in the first period.³⁰*

The proof of Theorem 4 is in the Appendix. A brief outline of the argument is as follows.³¹ We know from Theorem 2 that in order to generate distinct message paths in any PBE with sequential communication in the dynastic game it must be the case that at least one player is indifferent between two distinct action messages in terms of the continuation payoffs they generate. This immediately implies that at least one player can unilaterally deviate from this putative equilibrium and use a communication strategy that is less complex in the sense of Definition 2 and obtain the same continuation payoff. Hence, in any PBE of the dynastic game with sequential communication that is robust to lexicographic complexity costs according to Definition 3, the action messages of all players must be *the same* in every period.

Since the action messages chosen in equilibrium must be the same in every communication subgame, regardless of history, it now follows that the continuation payoffs to every individual cannot depend either on the current action profile or on the message received from the previous cohort. But then it follows immediately that the action profile σ chosen in every period cannot be anything other than a Nash equilibrium of the stage game.

³⁰Recall that each cohort cannot verify “calendar time” (see footnote 12 above). If, instead, we allowed calendar time to be verifiable it is relatively easy to see that different Nash equilibria of the stage game could be played in different periods in a PBE.

³¹We are grateful to an associate editor of this journal for pointing out that the proof of Theorem 4 could be shortened considerably by appealing to Theorem 2 in the way we do below.

The potential to whitewash is quite devastating when communication is sequential and decentralized and complexity costs of communication have even a lexicographic impact on payoffs. All deviations will be whitewashed by the current cohort. Continuation payoffs are therefore independent of current behavior, and only behavior that is equilibrium in a static sense will survive in any equilibrium of the dynastic game.

4. COVERUPS

So far we have analyzed the dynastic repeated game with communication protocols that ensured that *all* the information available to the current cohort is the result of message actions taken by the previous cohort. Of course, this is an extreme assumption. A more “realistic” view is that the information available to the current cohort is a mixture of the true history of play and of the communication behavior of the previous cohorts. The purpose of this section is to characterize equilibrium behavior in the dynastic repeated game under one such possible “mixed” communication protocol.

As we mentioned in Section 1 we examine a communication protocol in which the past history of play leaves a “footprint.” This footprint will be enough to reveal the true behavior of the previous cohort, unless the individuals in the current cohort *unanimously* agree to report a different history to the next generation. As we anticipated in Section 2 we call this a model of “coverups”.

To describe in detail the communication protocol with unanimous conspiracy to coverup, it is convenient to refer back to Example C^3 of Section 2.3 above. Essentially, we need to fill out the details of C^3 above: we need to specify the history h^* that is reported to the next generation in case of disagreement.

As in Example C^3 , let $A_i = M = H$. Now define $\Phi : M \times \Delta(S) \times A \rightarrow M$ as

$$\Phi(m, \sigma, a) = \begin{cases} h & \text{if } \exists h \text{ such that } h = a_i \forall i \\ (m, \sigma) & \text{otherwise} \end{cases} \quad (8)$$

Using (8) we can now proceed to define our communication protocol with unanimous conspiracy to coverup.

DEFINITION 4: *The dynastic repeated game with unanimous conspiracy to coverup is defined as follows.*

Individuals in each cohort choose their message actions sequentially as described in Section 3.2 above.

Consider a cohort that has received message m from the previous cohort. Assume that the individuals in the current cohort have chosen action profile σ . Let a be

the profile of action messages chosen by individuals in the current cohort. Then the message received by the next cohort is given by $m = \Phi(m, \sigma, a)$, where Φ is as in (8).

In other words, the message sent to the next cohort is equal to (m, σ) where m is the message of the previous cohort and σ is the *true* current action profile, *unless* all individuals in the current cohort choose *identical* action messages $a_1 = \dots = a_n = h$. In the latter case the message passed on to the next generation is h .

It turns out that the set of equilibrium payoffs of the dynastic repeated game with unanimous conspiracy to coverup is the same as the set of payoffs generated by equilibria of the full memory game that satisfy a restriction that has been analyzed before — namely those equilibria of the full memory game that are Weakly Renegotiation Proof (henceforth WRP) in the sense of Farrell and Maskin (1989). This is of independent interest since it tells us that Theorem 5 can be viewed as providing non-cooperative foundations for restricting attention to the set of WRP equilibria in a repeated game.

Before we proceed any further, for completeness we give a definition of those SPE that are WRP in the full memory game.

DEFINITION 5 [Farrell and Maskin (1989)]: *Consider the full memory game of Section 2. Let f^* denote an SPE of this game.*

We say that f^ is Weakly Renegotiation Proof if and only if it has the property that no continuation equilibrium is strictly Pareto-dominated by another continuation equilibrium.*

In other words, an SPE f^ of the full memory game is WRP if and only if there exist no pair of finite histories h and h' such that*

$$V_i(f^*|h) > V_i(f^*|h') \quad \forall i = 1, \dots, n \quad (9)$$

Throughout the rest of the paper, given a common discount factor δ the set of SPE strategy profiles of the full memory game that are WRP is denoted by $\mathcal{F}^R(\delta)$, while the corresponding set of payoff vectors is denoted by $\mathcal{E}^R(\delta)$.

We are now ready to state our last result.³²

³²We are grateful to an Associate Editor of this journal for suggesting that a previous version of Theorem 5 could be strengthened to yield the result that we now report below. See footnote 33 below.

THEOREM 5: *Let a communication protocol C with unanimous conspiracy to coverup and a common discount factor δ be given. Then $\mathcal{E}^C(\delta) = \mathcal{E}^R(\delta)$. In other words, the set of PBE payoffs in the dynastic repeated game with unanimous conspiracy to coverup is the same as the set of SPE payoffs of the full memory game that are WRP.*

The proof of Theorem 5 is in the Appendix. Intuitively, the argument that makes it hold runs along the following lines.

Start with $\mathcal{E}^C(\delta) \subseteq \mathcal{E}^R(\delta)$. Consider a reporting subgame of the dynastic repeated game, in which the players choose their actions sequentially, from 1 to n . Suppose that the equilibrium prescribes that all individuals report the action profile σ , but that the continuation payoffs associated with these reports are strictly Pareto-dominated by the continuation payoffs associated with another profile of message actions that are different from the “true” σ . Then using backwards induction (on the set of individuals, within the reporting “subgame”) it is possible to show that the true reporting behavior could not be an equilibrium in the first place. Since every individual can unilaterally trigger the true σ to be communicated to the next cohort, it is also possible to show that it cannot be the case that equilibrium behavior prescribes unanimous reporting of a “false” σ that is associated with continuation payoffs that are strictly dominated by the continuation payoffs associated with the true profile σ . In this way, it is possible to show that the equilibrium behavior in the reporting subgame cannot be associated with a profile of continuation payoffs that are dominated by the continuation payoffs associated with another profile of message actions. Hence no continuation equilibrium can be strictly Pareto-dominated by another continuation equilibrium.

To see that $\mathcal{E}^C(\delta) \supseteq \mathcal{E}^R(\delta)$ consider a reporting subgame of the dynastic repeated game and the following communication strategy profile which serves to implement any WRP equilibrium of the full memory game.³³

Player 1 starts by reporting the true history of play, and provided that there have been no deviations, all subsequent players also report the true history of play. If any deviations from truthful reporting have occurred (so that the true history is revealed through its “footprint”) then every subsequent player reports the true history of play. Now consider the possibility that, through a sequence of deviations, all players report some false history instead. Then, since the corresponding equilibrium of the full memory game is WRP there must be at least one player who (at least weakly) prefers the continuation payoffs associated with the true history of play. Since a single player disagreeing with the attempt to cover up the true history is sufficient to reveal it, this

³³This part of the statement and proof of Theorem 5 was suggested to us by an Associate Editor of this journal.

is enough to show that true reporting will in fact take place in equilibrium. Hence $\mathcal{E}^C(\delta) \supseteq \mathcal{E}^R(\delta)$.

5. CONCLUDING REMARKS

This paper examines play in dynastic repeated games.³⁴ Since each new generation cannot (perfectly) observe prior play, they must rely on messages of the prior generation. When these messages constitute cheap talk, then communication protocols must guard against whitewashing. When some prior information is available, then protocols must deter coverups.

Our results show that standard mechanism designs in the protocol can easily sustain all outcomes that were available in the full memory repeated game. Even when reports in the communication phase are sequenced, protocols which necessarily utilize some “neutral historian” exist to sustain most if not all outcomes of the full memory game.

However, our results also suggest that these equilibria are fragile. If individuals’ reports in any communication phase are sequenced, and if complexity matters even lexicographically, then only stage Nash equilibria can appear along the equilibrium path. In this world, the messages conveyed from one generation to the next are devoid of any real content.

Despite some similarities, the present model examines a very different type of communication than in typical sender-receiver, cheap talk models such as Crawford and Sobel (1982) and, more recently Krishna and Morgan (2001). The latter is representative of a more recent variety which, as in our model, features multiple senders of information. Yet, in all these models, difficulties in reporting incentives are due to different payoff functions between the sender and receiver. By contrast, in our model there are no payoff differences, at least between sender and receiver of the same dynasty. The incentive problems arise because of the requirement that the equilibria coordinate behavior on intertemporal sanctions. Sometimes these sanctions punish many or all individuals for the sins of one. This coordination on sanctions drives the necessary wedge between the senders and receivers of the hidden information.

Though no one in a cohort observes past history, the assumption of public messages means that all individuals of a given generation inherit the same “memory” from their predecessors. For this reason, the present model also bears some resemblance to repeated games with public monitoring. (Green and Porter (1984), Abreu, Pearce, and Stacchetti (1986), Abreu, Pearce, and Stacchetti (1990), and Fudenberg, Levine, and Maskin (1994), among others).

³⁴Related issues arise in two recent papers (Kobayashi 2003, Lagunoff and Matsui 2002) that model organizations populated by dynastic overlapping generations of players.

The public observation assumption is motivated by our desire to bias things as much as possible against whitewashing. Public observation allows the adoption of standard techniques from Nash implementation (see, for example, Jackson (1999) or Maskin and Sjöström (2002) and the references contained therein). For this reason, the sensitivity of such mechanisms to sequencing and complexity is somewhat unexpected.

Nevertheless, one could imagine dropping the assumption of public observability. In that event, we do not yet know what happens. With private, intra-dynastic communication and with only two dynasties, it is possible that folk theorem-like results similar to those reported in Kobayashi (2003) could be proved. However, with more than two dynasties, this is less clear.

With private communication, the model would bear closer resemblance to repeated games with private monitoring since each member of each cohort would possess a private version of the history of play that is not common knowledge across players.³⁵

On the one hand allowing private communication in the present set-up diminishes the scope for the use of “cross-checking” mechanisms to punish deviators. On the other hand, with private messages there seems to be a large degree of flexibility in defining out-of-equilibrium beliefs, which in turn suggests that a larger set of equilibrium outcomes may be sustainable.³⁶ It may well turn out to be the case that the role of out-of-equilibrium beliefs is, in fact, a key difference between dynastic games with private messages and repeated games with private monitoring.

The field of repeated games with private monitoring is an extremely important and currently active area of research in which general results have, by and large, eluded the best efforts of a formidable line-up of investigators. Though desirable, extending the model we have analyzed here to the general private-message case is evidently beyond the scope of the present paper.

³⁵See, for instance, Ben-Porath and Kahneman (1996), Sekiguchi (1997), Compte (1998), Kandori and Matsushima (1998), Mailath and Morris (1998), Bhaskar and van Damme (2002), Compte (2002a), Compte (2002b), Ely and Valimaki (2002), Kandori (2002), Mailath and Morris (2002), Matsushima (2002), and Piccione (2002). With few exceptions, this literature tends to examine outcomes of games that are “close” to those with public monitoring.

³⁶We are grateful to an anonymous Referee for pointing out a two-player example of out-of-equilibrium beliefs in which each player believes that the other player has received the same message with probability one. These out-of-equilibrium beliefs, when available, are helpful in sustaining many of the same equilibrium outcomes of the dynastic game with public communication.

APPENDIX

PROOF OF THEOREM 1: Let f^* denote any SPE in the full memory game.

We now define (g^*, μ^*) as follows. For each profile $m = (m_1, \dots, m_n)$, and each i ,

$$g_i^*(m) = \begin{cases} f_i^*(h) & \text{if } \exists h \text{ such that } m_j = h, \forall j \in J \subseteq I, \text{ with } |J| \geq n-1 \\ f_i^*(h^0) & \text{otherwise} \end{cases} \quad (\text{A.1})$$

and for each m , each σ , and each i ,

$$\mu_i^*(m, \sigma) = \begin{cases} (h, \sigma) & \text{if } \exists h \text{ such that } m_j = h, \forall j \in J \subseteq I, \text{ with } |J| \geq n-1 \\ (h^0, \sigma) & \text{otherwise} \end{cases} \quad (\text{A.2})$$

It is now straightforward to verify that (g^*, μ^*) is a PBE for the dynastic repeated game. The details are therefore omitted. Since the profile f^* was taken to be an arbitrary SPE of the full memory game, this is clearly enough to prove our claim. ■

PROOF OF COROLLARY 1: . It is immediate to check that $\mathcal{E}^C(\delta) \subseteq \mathcal{E}(\delta)$. The details of this claim are omitted. Since Theorem 1 obviously implies that $\mathcal{E}(\delta) \subseteq \mathcal{E}^C(\delta)$ the claim is proved. ■

PROOF OF THEOREM 2: Let (g, μ) denote any PBE. As hypothesized in the Theorem, let (m', σ') and (m'', σ'') , denote the prior histories of message and action profiles, respectively.

Now suppose that the Theorem is false. Then, for any path, a , and for every player j , we must have

EITHER

$$\mu_j(m', \sigma', a^{j-1}) = \mu_j(m'', \sigma'', a^{j-1}) \quad (\text{A.3})$$

OR

$$\begin{aligned} \mu_j(m', \sigma', a^{j-1}) &\neq \mu_j(m'', \sigma'', a^{j-1}), \text{ and} \\ V_j(g, \mu | m', \sigma', a^{j-1}, \mu_j(m', \sigma', a^{j-1})) &\neq V_j(g, \mu | m'', \sigma'', a^{j-1}, \mu_j(m'', \sigma'', a^{j-1})) \end{aligned} \quad (\text{A.4})$$

We use the following backward induction argument. Suppose, first, that either (A.3) or (A.4) hold for all paths and for player n . We now argue that (A.4) cannot hold for player n . To see this, observe that if (A.4) did indeed hold then, without loss of generality, we have

$$\hat{v}_n \equiv V_n(g, \mu | m', \sigma', a^{n-1}, \mu_n(m', \sigma', a^{n-1})) > V_n(g, \mu | m'', \sigma'', a^{n-1}, \mu_n(m'', \sigma'', a^{n-1})) \equiv v_n \quad (\text{A.5})$$

But since player n is the last mover to report, then at any node a^{n-1} , his preferences for \hat{v}_n over v_n cannot depend on prior histories (m', σ') and (m'', σ'') . That is, we can drop the notational dependence of V_n on (m', σ') and (m'', σ'') and rewrite (A.5) as

$$\hat{v}_n \equiv V_n(g, \mu | a^{n-1}, \mu_n(m', \sigma', a^{n-1})) > V_n(g, \mu | a^{n-1}, \mu_n(m'', \sigma'', a^{n-1})) \equiv v_n \quad (\text{A.6})$$

From (A.6), it is easy to see that $\mu_n(m'', \sigma'', a^{n-1})$ is not a best response, violating the equilibrium property of μ . Therefore, for player n , (A.3) must hold, i.e.,

$$\mu_n(m', \sigma', a^{n-1}) = \mu_n(m'', \sigma'', a^{n-1}) \quad (\text{A.7})$$

Notice that (A.7) implies that player n must play the same way on *every* path a , regardless of which prior history, (m', σ') or (m'', σ'') , occurred.

Now suppose that (A.3) holds for players $i + 1, \dots, n$. As was true for player n , these players must play exactly the same way on every path. Said another way, along any path all subsequent players fail to distinguish between (m', σ') and (m'', σ'') in their reporting behavior. But if strategies μ_{i+1}, \dots, μ_n fail to distinguish between (m', σ') and (m'', σ'') in these subgames, then this must also be true for player i . For if, instead, (A.4) held, i.e., if

$$V_i(g, \mu | m', \sigma', a^{i-1}, \mu_i(m', \sigma', a^{i-1})) \neq V_i(g, \mu | m'', \sigma'', a^{i-1}, \mu_i(m'', \sigma'', a^{i-1}))$$

then either $\mu_i(m', \sigma', a^{i-1})$ or $\mu_i(m'', \sigma'', a^{i-1})$ can no longer be a best response. Hence, player i must satisfy (A.3).

But now consider the incentives of player $i = 1$. Observe that, as hypothesized, (m', σ') and (m'', σ'') each generate distinct paths. Let a' denote the path following (m', σ') and let a'' denote the path following (m'', σ'') . Since equation (A.3) holds for all other players, $2, \dots, n$, it must be true that player 1 plays differently after each of (m', σ') and (m'', σ'') in order to distinguish a' from a'' . Specifically,

$$\mu_1(m', \sigma', a^0) \neq \mu_1(m'', \sigma'', a^0) \quad (\text{A.8})$$

That is, player 1 must have distinct choices after (m', σ') and (m'', σ'') since no other player distinguishes between the two histories. But (A.8) contradicts (A.3) for player 1. Since we have already established that player 1, as well as all other players cannot satisfy (A.4), it must be the case that player 1 violates both (A.3) and (A.4), and so we have obtained our contradiction. This concludes the proof. ■

The following Lemma will be used for the proof of Theorem 3.

LEMMA A.1: *Let \mathcal{V} be a self-generating closed set of long-run payoffs for the full memory game that satisfies (5).*

For any vector $v \in \mathcal{V}$, we let $\mathcal{Z}(v)$ be the set of strategy profiles that sustain the vector of long-run payoffs v as an SPE of the full memory game, with continuation payoffs that lie entirely in \mathcal{V} . For any vector $v \in \mathcal{V}$, we let f_v denote a generic element of $\mathcal{Z}(v)$. Also, for any vector $v \in \mathcal{V}$ we let $P_i(v)$ be the projection of v on the lower boundary of \mathcal{V} for player i . In other words, $P_i(v) = (v_1, \dots, v_{i-1}, \underline{v}_i, v_{i+1}, \dots, v_n)$.

Now consider an arbitrary $v^ \in \mathcal{V}$. Then there exists an $f^* \in \mathcal{Z}(v^*)$ with the following properties.*

For any history h^t , let $\sigma^(h^t)$ be the mixed profile of actions prescribed by f^* at time $t + 1$, conditional on history h^t taking place. Let also σ^i be any mixed action profile that agrees with*

$\sigma^*(h^t)$ on all components except for player i . That is σ^i satisfies $\sigma_j^i = \sigma_j^*(h^t)$ for every $j \neq i$ and $\sigma_i^i \neq \sigma_i^*(h^t)$.

Then, for any history h^t and for any i ,

$$V(f^*|h^t, \sigma^i) = P_i[V(f^*|h^t, \sigma^*(h^t))] \quad (\text{A.9})$$

Moreover, let $\hat{\sigma}$ be any mixed action profile that differs from $\sigma^*(h^t)$ on two or more components. Let \mathcal{D} be the set of players for which $\hat{\sigma}$ and $\sigma^*(h^t)$ differ. Then, for any history h^t

$$V_i(f^*|h^t, \hat{\sigma}) = \begin{cases} \underline{v}_i & \text{if } i \in \mathcal{D} \\ V_i(f^*|h^t, \sigma^*(h^t)) & \text{otherwise} \end{cases} \quad (\text{A.10})$$

In other words, without loss of generality, we can take f^* to have the property that any unilateral deviation by player i is punished by giving i a continuation payoff of \underline{v}_i and leaving the continuation payoffs of all other players unchanged. Moreover, again without loss of generality, we can take f^* to have the property that any deviation by two or more players yields “bad” continuation payoffs for the deviating players only as in the right-hand side of (A.10).

PROOF: Let any $\tilde{f} \in \mathcal{Z}(v^*)$ be given. We now construct f^* with the desired property as a modification of \tilde{f} . The construction is recursive.

On $h^0 = \emptyset$, f^* prescribes the same behavior as \tilde{f} . So long as no player deviates from the outcome path prescribed by \tilde{f} , the prescriptions of f^* are the same as those of \tilde{f} .

Suppose now that some history h^t (on the equilibrium path of \tilde{f}) has taken place and that a deviation by player i only has occurred at time t (we ignore deviations by more than one player for the time being). Let σ^i be the mixed action profile played at t which includes i deviation. Let also $\tilde{\sigma}(h^t)$ be the equilibrium prescription of \tilde{f} after history h^t , and let $v = V(\tilde{f}|h^t, \tilde{\sigma}(h^t))$ be the associated continuation payoff. Then, after i 's deviation at t the prescriptions of f^* are the same as those of $f_{P_i(v)}|h^0$ (recall that, according to our notation, $f_{P_i(v)}$ denotes a strategy profile that sustains the vector of long-run payoffs $P_i(v)$, with continuation payoffs that lie entirely in \mathcal{V}). Notice that this implies that the continuation payoff vector implied by f^* after (h^t, σ^i) is $P_i(v)$.

So long as the prescriptions of $f_{P_i(v)}|h^0$ are observed after time t , the prescriptions of f^* remain as we have just described. Suppose now that a history $h^m = (h^t, \sigma^i, h^{m-t})$ (with $m > t$ and h^{m-t} on the equilibrium path of $f_{P_i(v)}$) has occurred and that at time m a deviation by player j takes place. Let σ^j be the mixed action profile played at m which includes j deviation. Let also $\sigma^i(h^{m-t})$ be the equilibrium prescription of $f_{P_i(v)}$ after history h^{m-t} , and let $b^i = V(f_{P_i(v)}|h^m, \sigma^i(h^{m-t}))$ be the associated continuation payoff vector. Then, after j 's deviation at m the prescriptions of f^* are the same as those of $f_{P_j(b^i)}|h^0$. Notice that this implies that the continuation payoff vector implied by f^* after (h^m, σ^j) is $P_j(b^i)$.

So long as the prescriptions of $f_{P_j(b^i)}$ are observed after time m , the prescriptions of f^* remain as we have just described. If a further deviation occurs then the players “switch” to a new “phase” in which the deviating player is pushed down to the lowest payoff available for him in \mathcal{V} by playing the appropriate SPE from then on, in a way completely analogous to the one we have just described.

Thus the description of f^* can be completed by recursing forward the construction we have given. The rest of the details are omitted.

Clearly, the profile of strategies f^* that we have constructed has the property described in (A.9) by construction. Evidently it is also the case that, by construction, all continuation payoff vectors of f^* lie in \mathcal{V} , as required.

We now show that f^* is an SPE strategy profile of the full memory repeated game. This is relatively straightforward to check since only one-shot single-player deviations need ever be considered. To verify that no such profitable deviations are possible, suppose that some history h^s has taken place, and let v_i be i 's continuation payoff according to f^* after h^s . Thus, if i at time s adheres to the prescription of f^* he receives a payoff of v_i . Notice that v_i is also i 's continuation payoff in the particular SPE that is being played in the “phase” that follows history h^s . If on the other hand he deviates in any way from what f^* prescribes he receives a payoff of \underline{v}_i . Since \underline{v}_i is the lowest continuation payoff that i can get in any of the SPE that are used in the construction of f^* above, it is clear that this must be sufficient to deter i from deviating from the prescriptions of f^* after h^s has taken place.

Hence, we have shown that an SPE f^* satisfying (A.9) exists as required. It remains to show that f^* can be made to satisfy (A.10) as well. However, this is completely straightforward once we know that an SPE satisfying (A.9) exists since deviations by two or more players can always be ignored when checking if a given strategy profile is an SPE. The details are omitted for the sake of brevity. ■

PROOF OF THEOREM 3: Fix a set \mathcal{V} satisfying the hypothesis of the Theorem. Now fix $v^* \in \mathcal{V}$. We must show that $v^* \in E^C(\delta)$ for any sequential protocol C with $H \subseteq A_i$. Without loss of generality, we consider the sequential protocol with the natural order: player 1 speaks first, player 2 speaks second, and so forth.

Let f^* be an SPE of the full memory game that sustains v^* as vector of long-run payoffs. Using Lemma A.1 we can assume without loss of generality that f^* has the properties described in (A.9) and (A.10), and that all its continuation payoffs lie in \mathcal{V} .

We now construct the pair (g^*, μ^*) that sustains the arbitrary payoff vector $v^* \in \mathcal{V}$ as a PBE. Loosely speaking our construction of (g^*, μ^*) runs along the following lines. Only the messages of players 1 and 2 are ever taken into account. Player 1 is asked to report the history of play, then player 2 is asked to “confirm” 1's report. If player 1 reports the truth, then player 2 confirms 1's report and play unfolds according to f^* . If, on the other hand 1 ever issues a false report, then player 2 does not confirm and reports 1's deviation from the truth. In this case the continuation of play unfolds as if player 1 had *behaviorally* deviated from f^* , using the punishments prescribed by f^* . Since for player 1's deviations f^* punishes 1 in a way that leaves 2 *indifferent*, player 2 always has the correct incentives to report 1's deviation from truthful reporting. Given the punishment for behavioral deviations built into f^* , 1 now also has the correct incentives to always report the true history of play.

Recall that we set $m(-1) = h^0 = \emptyset$. To set the system in motion, let $g^*(h^0) = f^*|h^0$. The rest of the equilibrium is constructed recursively forward in the following way. In every period $t \geq 0$, player 1 reports the truth in the sense that $\mu_1^* = (h^{t-1}, \sigma)$ where h^{t-1} is the history reported by player 2 in period $t-1$ and σ is the true mixed profile that was played by the current cohort.

In every period $t \geq 0$, player 2's depends on the veracity of the report of player 1 in the following way. If player 1's report is truthful (as defined above) then player 2 issues an identical report (h^{t-1}, σ) where h^{t-1} is the history reported by player 2 in period $t-1$ (h^0 if $t=0$) and σ is the true mixed

profile that was played by the current cohort. If on the other hand player 1's report is not truthful, player 2 issues a report (h^{t-1}, σ') where h^{t-1} is the history reported by player 2 at $t-1$, and σ' is a mixed action profile that is *different* from that reported by player 1, *and* that records a behavioral deviation by player 1 only in period t . In other words, $\sigma'_1 \neq f_1^*|h^{t-1}$ and $\sigma'_i = f_i^*|h^{t-1}$ for every $i = 2, \dots, n$. Notice that these two conditions can clearly always be satisfied simultaneously.

The reports of all players i with $i \geq 3$ (if $n \geq 3$) are ignored. Therefore we simply set them equal to a fixed message m_i regardless of the history of play.

The g^* component of the equilibrium is easy to describe. In period t , if the reports of players 1 and 2 are the *same*, then all players behave according to f^* , conditional on the reported (h^{t-1}, σ) . If, on the other hand, the reports of players 1 and 2 differ, then all players behave according to f^* conditional on $\hat{h}^t = (\hat{h}^{t-1}, \hat{\sigma})$ defined as follows. We set \hat{h}^{t-1} equal to the $t-1$ history reported by player 1. Moreover, we set $\hat{\sigma}_1$ — the first component of $\hat{\sigma}$ — equal to the report of player 2, and all other components $(\hat{\sigma}_2, \dots, \hat{\sigma}_n)$ equal to the report of player 1.

Clearly the proposed equilibrium yields a vector of long-run payoffs v^* as required. Of course, it remains to show that f^* is indeed a PBE of the repeated game with decentralized communication protocol C . We need to verify that no player ever has an incentive to unilaterally deviate in any period, either at the communication stage or at the behavior stage.

All players $i \geq 3$ (if any) clearly have no incentive to deviate in any period. Their messages are ignored, and hence they cannot gain by deviating at the communication stage. At the behavior stage, since f^* is an SPE of the full memory game, and histories are reported truthfully, no individual deviation can be profitable.

Consider now player 1, at the reporting stage after some history h^{t-1} has been reported by player 2 of the previous cohort, and the mixed profile σ has taken place in the current period. If he reports the truth (h^{t-1}, σ) as required, he receives a continuation payoff of corresponding to f^* , conditional on (h^{t-1}, σ) . If on the other hand he reports anything else, he receives a payoff of \underline{v}_i . Since all continuation payoffs of f^* lie in \mathcal{V} , this cannot be a profitable deviation by player 1. Of course, at the behavior stage player 1 has no incentive to deviate simply because f^* is an SPE of the full memory game and histories are reported truthfully.

Lastly, consider player 2 at the reporting stage after some history h^{t-1} has been reported by player 2 of the previous cohort, the mixed profile σ has taken place in the current period, and player 1 has reported some (possibly false) $(\tilde{h}^{t-1}, \tilde{\sigma})$. Clearly, using (A.9) and (A.10) and because of the way we have defined $\hat{\sigma}$ above, the continuation payoff of player 2 is the same regardless of his report. Hence he cannot profitably deviate at this stage. Again, at the behavior stage player 2 has no incentive to deviate simply because f^* is an SPE of the full memory game and histories are reported truthfully. This is clearly enough to conclude the proof. ■

PROOF OF THEOREM 4: Consider an PBE (g^*, μ^*) of the dynastic repeated game with sequential decentralized communication protocol C , and assume that (g^*, μ^*) is robust to lexicographic complexity costs of communication.

By Definition 3, of course (g^*, μ^*) must be a PBE of the dynastic repeated game with sequential decentralized communication protocol C . Hence Theorem 2 applies to tell us that if $a' \neq a''$ are two distinct action paths following any (m', σ') and (m'', σ'') respectively, then for some player i we must have that

$$V_i(g^*, \mu^* | m', \sigma', a^{i-1'}, a'_i) = V_i(g^*, \mu^* | m'', \sigma'', a^{i-1}'', a''_i) \tag{A.11}$$

Then, using (A.11) it is clear that we could find a strategy $\mu_i \neq \mu_i^*$ such that $\mu(m', \sigma', a^{i-1'}) = \mu(m'', \sigma'', a^{i-1''})$ and

$$V_i(g^*, \mu^* | m', \sigma, a^{i-1'}, \mu(m', \sigma', a^{i-1'})) = V_i(g^*, \mu^* | m'', \sigma'', a^{i-1''}, \mu(m'', \sigma'', a^{i-1''})) \quad (\text{A.12})$$

and $\mathcal{P}(\mu^*) \succ \mathcal{P}(\mu)$. Therefore, using Definitions 2 and 3 we conclude if (g^*, μ^*) is robust to lexicographic complexity costs of communication we must have that

$$\mu_i^*(m', \sigma', a^{i-1'}) = \mu_i^*(m'', \sigma'', a^{i-1''}) \quad \forall m', m'', \sigma', \sigma'', a^{i-1'}, a^{i-1''} \quad (\text{A.13})$$

Using (A.13) we can now define *the* sequence of message actions that every individual in every cohort (except possibly the first one) will take. Recursively forward from individual 1 we set

$$a_1^* = \mu_1^*(m, \sigma, \emptyset) \quad \forall m, \sigma \quad (\text{A.14})$$

and (letting $a^{i-1*} = (a_1^*, \dots, a_{i-1}^*)$, for $i = 2, \dots, n$)

$$a_i^* = \mu_i^*(m, \sigma, a^{i-1*}) \quad \forall m, \sigma \quad (\text{A.15})$$

Lastly, we let $m^* = (a_1^*, \dots, a_n^*)$. This is the message that every cohort will receive in any subgame of the PBE (g^*, μ^*) , except of course for the first cohort that will, by assumption, receive a message $m = \emptyset$.

It now follows from (A.14) and (A.15) that the continuation payoff to individual i after any message m has been received from the previous cohort can be written as a function of his choice σ_i as

$$(1 - \delta)U_i(\sigma_i, g_{-i}^*(m)) + \delta V_i(g^*, \mu^* | m^*) \quad (\text{A.16})$$

Since all cohorts, except for the first one, receive message m^* from the previous cohort, the statement of the theorem now follows immediately from (A.16). The rest of the details are omitted. ■

LEMMA A.2: *Let a communication protocol C with unanimous conspiracy to coverup and a common discount factor δ be given. Let also $\mathcal{E}^R(\delta)$ be as in Definition 5. Then $\mathcal{E}^C(\delta) \subseteq \mathcal{E}^R(\delta)$.*

PROOF: Fix δ . Suppose, by contradiction, that $v^* \in \mathcal{E}^C(\delta)$ while $v^* \notin \mathcal{E}^R(\delta)$. Since $v^* \notin \mathcal{E}^R(\delta)$ then for all f^* that sustain v^* in the full memory repeated game, there exists some pair of histories, h', h'' , such that

$$v' \equiv V(f^* | h') \gg V(f^* | h'') \equiv v'' \quad (\text{A.17})$$

Now let (g^*, μ^*) sustain v^* under protocol C with unanimous conspiracy to coverup. Clearly, since (A.17) must hold for *every* SPE of the full memory game that sustains v^* , we must have that for some pair (m', σ') and (m'', σ'') , corresponding to h' and h'' respectively, the following holds

$$v' = V(g^*, \mu^* | m', \sigma') \gg V(g^*, \mu^* | m'', \sigma'') = v'' \quad (\text{A.18})$$

To derive the contradiction, suppose now that (m'', σ'') has in fact occurred. We proceed to show that v'' cannot be the equilibrium continuation of the communication phase. To verify this, we proceed by induction. Consider the incentives of player n , when all others have reported $a'_i = (m', \sigma')$, $\forall i \neq n$. That is, all others have (falsely) reported prior path (m', σ') . According to the protocol for unanimous conspiracy to coverup, if n also reports $a'_n = (m', \sigma')$ then v' is attained. However, if player n vetoes v' by reporting any other a_n , then the true history (m'', σ'') is revealed, and so continuation v'' occurs. But then (A.18) immediately implies that n 's best response is to in fact report a'_n .

Proceeding by induction, using the same argument as for player n , it is now easy to show that every player i 's ($i > 1$) best response to all preceding players $j = 1, \dots, i - 1$ having chosen a'_j is in fact to report a'_i . Finally, consider the choice of player 1. Clearly, if he reports a'_1 (given the best responses of all other players) he achieves a payoff of v'_1 while if he reports a''_1 he gets a payoff of v''_1 . Hence, using (A.18) again, reporting a''_1 cannot be player 1's equilibrium behavior in the reporting subgame. Moreover, given the equilibrium strategies of the other players in the reporting subgame, it is clear that player 1 (by choosing a'_1) can achieve a continuation payoff of v'_1 . Hence, v'' cannot be a continuation equilibrium payoff vector of the reporting subgame, as is in fact required. This contradiction is clearly enough to establish the result. ■

LEMMA A.3: *Let a communication protocol C with unanimous conspiracy to coverup and a common discount factor δ be given. Let also $\mathcal{E}^R(\delta)$ be as in Definition 5. Then $\mathcal{E}^C(\delta) \supseteq \mathcal{E}^R(\delta)$.*

PROOF: Fix δ and any $v^* \in \mathcal{E}^R(\delta)$. Let f^* be any WRP SPE strategy profile that sustains v^* in the full memory game.

Notice that since f^* is a WRP SPE we know that there is *no* pair of finite histories h' and h'' for which (A.17) holds. In other words, for any pair h' and h''

$$\exists i \text{ such that } V_i(f^*|h') > V_i(f^*|h'') \Rightarrow \exists j \text{ such that } V_j(f^*|h') \leq V_j(f^*|h'') \quad (\text{A.19})$$

We now specify a pair (g^*, μ^*) that sustains v^* as a PBE payoff in the dynastic game with unanimous conspiracy to coverup. The action strategy profile g^* is the same as f^* in the full memory game. The communication strategy profile is as follows. Player 1 always reports the true history of play. That is μ^*_1 prescribes $a_1 = (m, \sigma)$ for every possible (m, σ) .

For each $i = 2, \dots, n$ we let μ^*_i be defined as

$$\mu^*_i(m, \sigma) = \begin{cases} h & \text{if } \exists h \text{ such that } a_j = h, \forall j < i \text{ and } V_i(g^*, \mu^* | m, \sigma) < V_i(g^*, \mu^* | h) \\ (m, \sigma) & \text{otherwise} \end{cases} \quad (\text{A.20})$$

In other words, i reports the true history of play unless all players before him have agreed on a particular report which induces a continuation payoff that he prefers to the continuation payoff following the true history of play.

It is easy to verify that the profile μ^* we have just defined yields an equilibrium path in which all players report the true history of play. Hence (g^*, μ^*) yields payoffs v^* in the dynastic game. Of course, we still need to verify that it constitutes a PBE of the dynastic game.

We need to check that no player has an incentive to deviate from μ^*_i in the communication stage. Consider first players $i = 2, \dots, n$. If we are in the case in which (A.20) prescribes the report

(m, σ) , then either two players before i have made conflicting reports, or all players before i have reported a history h that induces a continuation payoff which gives leaves i no better off than the continuation following the true history of play. By sticking to the prescription of μ_i^* player i obtains a continuation payoff of $V_i(g^*, \mu^* | m, \sigma)$. Any deviation will yield either a payoff of $V_i(g^*, \mu^* | m, \sigma)$ or of $V_i(g^*, \mu^* | h)$ (depending on whether the deviation is to h or to some other history, and on whether all subsequent players agree to h or not). Hence no profitable deviation is possible in this case. If we are in the case described by the top line of (A.20), then any deviation leads to a continuation payoff of $V_i(g^*, \mu^* | m, \sigma)$. On the other hand abiding by the prescription of μ_i^* leads to a payoff of $V_i(g^*, \mu^* | h)$ if all subsequent players agree to h , ad to a payoff of $V_i(g^*, \mu^* | m, \sigma)$ otherwise. Hence no profitable deviation is possible in this case.

It then remains to check that player 1 cannot profitable deviate from μ_1^* . Of course, this is only possible if he can report a false history that all subsequent players agree upon. However, because (A.19) holds this is impossible. Some subsequent player must prefer (at least weakly) the continuation payoff that follows the true history of play. Hence, according to the bottom line of (A.20) he will report the true history of play. This is clearly sufficient to prove the claim. ■

PROOF OF THEOREM 5: The claim is a direct consequence of Lemma A.2 and Lemma A.3. ■

REFERENCES

- ABREU, D. (1988): “On the Theory of Infinitely Repeated Games with Discounting,” *Econometrica*, 56, 383–396.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1986): “Optimal Cartel Equilibrium with Imperfect Monitoring,” *Journal of Economic Theory*, 39, 251–269.
- (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1993): “Renegotiation and Symmetry in Repeated Games,” *Journal of Economic Theory*, 60, 217–40.
- ABREU, D., AND A. RUBINSTEIN (1988): “The Structure of Nash Equilibrium in Repeated Games with Finite Automata,” *Econometrica*, 56, 1259–1281.
- AUMANN, R. J., AND S. SORIN (1989): “Cooperation and Bounded Recall,” *Games and Economic Behavior*, 1, 5–39.
- BALIGA, S., L. CORCHON, AND T. SJOSTROM (1997): “The Theory of Implementation When the Planner is a Player,” *Journal of Economic Theory*, 77, 15–33.
- BEN-PORATH, E., AND M. KAHNEMAN (1996): “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 70, 281–297.
- BENOIT, J.-P., AND V. KRISHNA (1993): “Renegotiation in Finitely Repeated Games,” *Econometrica*, 61, 303–23.

- BHASKAR, V., AND E. VAN DAMME (2002): “Moral Hazard and Private Monitoring,” *Journal of Economic Theory*, 102, 16–39.
- CHATTERJEE, K., AND H. SABOURIAN (2000): “Multiperson Bargaining and Strategic Complexity,” *Econometrica*, 68, 1491–1510.
- COMPTE, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66, 597–626.
- (2002a): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66, 597–626.
- (2002b): “On Failing to Cooperate when Monitoring is Private,” *Journal of Economic Theory*, 102, 151–188.
- CRAWFORD, V., AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50, 1431–1452.
- ELY, J. C., AND J. VALIMAKI (2002): “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory*, 102, 84–105.
- FARRELL, J. V., AND E. S. MASKIN (1989): “Renegotiation in Repeated Games,” *Games and Economic Behavior*, 1, 327–360.
- FUDENBERG, D., D. K. LEVINE, AND E. MASKIN (1994): “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997–1039.
- FUDENBERG, D., AND E. S. MASKIN (1986): “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54, 533–556.
- GREEN, E., AND R. PORTER (1984): “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, 52, 87–100.
- JACKSON, M. (1999): “A Crash Course in Implementation Theory,” Social Science Working Paper 1076, California Institute of Technology.
- KANDORI, M. (2002): “Introduction to Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 102, 1–15.
- KANDORI, M., AND H. MATSUSHIMA (1998): “Private Observation, Communication and Collusion,” *Econometrica*, 66, 627–652.
- KOBAYASHI, H. (2003): “Folk Theorem for Infinitely Repeated Games Played by Organizations with Short-Lived Members,” Osaka Prefecture University, mimeo.

- KRISHNA, V., AND J. MORGAN (2001): "A Model of Expertise," *Quarterly Journal of Economics*, 116, 747–775.
- LAGUNOFF, R., AND A. MATSUI (1997): "Asynchronous Choice in Repeated Coordination Games," *Econometrica*, 65, 1467–1477.
- LAGUNOFF, R., AND A. MATSUI (2002): "Organizations and Overlapping Generations Games: Memory, Communication, and Altruism," Georgetown University, mimeo.
- MAILATH, G., AND S. MORRIS (1998): "Repeated Games with Imperfect Private Monitoring: Notes on a Coordination Perspective," University of Pennsylvania, mimeo.
- (2002): "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory*, 102, 189–228.
- MASKIN, E., AND J. T. SJÖSTRÖM (2002): "Implementation Theory," in *Handbook of Social Choice and Welfare*. North-Holland.
- MASKIN, E. S. (1999): "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23–38.
- MATSUSHIMA, H. (2002): "The Folk Theorem with Private Monitoring," *Econometrica*, forthcoming.
- PENNEBAKER, J., D. PAEZ, AND B. RIME (eds.) (1997): *Collective Memory of Political Events*. New Jersey: Lawrence Erlbaum Associates.
- PICCIONE, M. (1992): "Finite Automata Equilibria with Discounting," *Journal of Economic Theory*, 56, 180–193.
- (2002): "The Repeated Prisoners Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102, 70–83.
- RUBINSTEIN, A. (1986): "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 39, 83–96.
- RUBINSTEIN, A., AND M. PICCIONE (1993): "Finite Automata Play a Repeated Extensive Game," *Journal of Economic Theory*, 61, 160–168.
- SEKIGUCHI, T. (1997): "Efficiency in Repeated Prisoners Dilemma with Private Monitoring," *Journal of Economic Theory*, 76, 345–361.