

A “Super” Folk Theorem for Dynastic Repeated Games*

LUCA ANDERLINI
(Georgetown University)

DINO GERARDI
(Yale University)

ROGER LAGUNOFF
(Georgetown University)

September 2007

Abstract. We analyze *dynastic repeated games*. These are repeated games in which the stage game is played by successive generations of finitely-lived players with dynastic preferences. Each individual has preferences that replicate those of the infinitely-lived players of a standard discounted infinitely-repeated game. Individuals live one period and *do not* observe the history of play that takes place before their birth, but instead create social memory through *private messages* received from their immediate predecessors.

Under mild conditions, when players are sufficiently patient, *all feasible payoff vectors* (including those below the *minmax* of the stage game) can be sustained by Sequential Equilibria of the dynastic repeated game with private communication. In particular, the result applies to any stage game with $n \geq 4$ players for which the standard Folk Theorem yields a payoff set with a non-empty interior.

We are also able to characterize fully the conditions under which a Sequential Equilibrium of the dynastic repeated game can yield a payoff vector *not* sustainable as a Subgame Perfect Equilibrium of the standard repeated game. For this to be the case it must be that the players’ equilibrium beliefs violate a condition that we term “Inter-Generational Agreement.”

JEL CLASSIFICATION: C72, C73, D82.

KEYWORDS: Dynastic Repeated Games, Private Communication, Social Memory, Folk Theorem.

ADDRESS FOR CORRESPONDENCE: Luca Anderlini, LA2@GEORGETOWN.EDU, Department of Economics, Georgetown University, 37th and O Streets, Washington DC 20057, USA.

*A previous version of this paper was previously circulated as Anderlini, Gerardi, and Lagunoff (2004). We are grateful to Jeff Ely, Leonardo Felli, Navin Kartik, David Levine, Stephen Morris, Michele Piccione, Andrew Postlewaite, Lones Smith and to seminar audiences at Bocconi, Cambridge, CEPR-Guerzensee, Chicago, Columbia, Edinburgh, Essex, Georgetown, Leicester, LSE, Northwestern, Oxford, Rome (La Sapienza), Rutgers, SAET-Vigo, Stanford, SUNY-Albany, UCL, UC-San Diego, Venice and Yale for helpful feedback. All errors are our own.

1. Introduction

This paper establishes a “Super Folk Theorem” for *dynastic repeated games*. Dynastic repeated games are repeated games with periodic generational replacement. In these games, society is inhabited by sequences of generationally linked individuals. In each sequence or *dynasty*, a finitely lived decision maker is eventually replaced by a successor who has no direct knowledge of the past before his entry into the game. The members of a dynasty are linked by concerns for the future welfare of the group or organization. Examples include tribes, families, ethnic groups, or firms.¹ These ongoing organizations outlive any particular individual who occupies decision authority in the group at a particular point in time.

If all individuals in all dynasties could perfectly observe the past history of play, then the dynastic game and the standard repeated game are equivalent (in Subgame Perfect equilibria). However, because no individual person in a dynasty can directly observe events that occurred before his “birth,” he must rely on a private *messages* from his predecessor to form beliefs about the past history of play. This intergenerational communication forms a significant part of a society’s “social memory” at any given point in time.

Our main result indicates a stark difference between the dynastic repeated games and the standard repeated game model. Specifically, we find that in a broad class of games (which includes every $n \geq 4$ -player game for which the standard Folk Theorem yields a payoff set with a non-empty interior), as individuals become more and more altruistic (patient), *every* interior payoff vector in the convex hull of the payoffs of the stage game can be sustained by a Sequential Equilibrium (henceforth SE) of the dynastic repeated game with private communication.

The logic of this “Super Folk Theorem” relies on the possibility that individuals systematically misinterpret their predecessors’ messages indicating past deviations from within the dynasty. These individuals therefore remain unaware of the ensuing punishment phase, and hence are unable to protect themselves against it.

Even though our main result applies to infinite repetitions of a stage game with four players or more, we begin with an illustrative example involving a two-player game played twice. This is deliberately designed to be as simple as possible while still allowing the basic phenomenon underlying our main results to take place.

¹In the cases of firms, one may not think of managers in firms as “altruistic.” However, stock options and other incentive contracts help to align the manager’s interests with those of the firm’s.

Consider the following game played twice.

		Dynasty II		
		s_A	s_B	s_C
Dynasty I	s_A	4, 4	-5, 6	1, 3
	s_B	6, -5	3, 3	-10, 2
	s_C	3, 1	2, -10	0, 0

At the end of the first period, the two individuals who play the game in the first round are replaced by their successors, who play in the second period. There is no discounting: Each first period “father” cares about his “son’s” payoff as if it were his own. The second period players only care about their own payoffs.

Consider only pure strategies for now. It is clear that if the sons could observe the first period actions, then the only Subgame Perfect Equilibrium (henceforth SPE) is one in which (s_B, s_B) is played twice. Suppose, however, that a son cannot observe play in $t = 1$, but instead receives a private message from his father. We claim that there exists an SE in which (s_A, s_A) is played in the first round. In other words, a non-Nash profile is sustainable in the first period. This is so despite the fact that in the standard model the unique stage Nash equilibrium leaves no room for punishment in the second period.

The argument is as follows. Consider an SE in which two messages are used: m^* indicates a message of “no deviation,” while m^D indicates “someone’s father deviated.” The equilibrium prescribes that in the first period each father plays “ s_A .” He then sends m^* to his son if (s_A, s_A) was played, and sends m^D otherwise. Each son, upon observing the message from his father, plays s_B in the second round if he observes m^* and plays s_C if instead he sees m^D .

But how can each son justify playing “ s_C ” after observing the off-path message m^D ? For this to be the case, each son’s beliefs must be mismatched to the other son’s actions. In any SE, off-path beliefs must satisfy a well known consistency criterion. Namely, they must be the limit of a sequence of beliefs derived from a commonly shared theory of *mistakes or errors* (“trembles”). Hence, consider the following error structure in the first period actions of the fathers. Each father is believed to have erred in the action stage by playing s_B or s_C with probability $\varepsilon/4$ each. He is also believed to have sent the wrong message with probability ε . Given these error likelihoods, if a son receives message m^D he will believe that his opponent has received m^D or m^* with probability $1/2$ each. His best response in this case is “ s_C .” Since (s_C, s_C) provides an effective deterrent in the second period against deviation from (s_A, s_A) , each father will choose the prescribed action s_A in the first period.

The fact that s_C is a best response to “wrong” beliefs highlights how different the model is from standard cheap talk games. Messages are not simply “babbled” or garbled. Instead, the “mismatch” of beliefs to actions depends critically on the way the consistency criterion of SE actually works out. This criterion is much more than a technical artifact of SE. As

the example illustrates, the consistency criterion provides an important *social theory* of belief formation. Namely, it requires that all individuals have a complete and shared theory of the mistakes that might have caused deviations from equilibrium — a society’s *outlook* governing how “surprises” should be interpreted.

As in the example, the Super Folk Theorem utilizes this shared theory of mistakes. When an individual in the guilty dynasty receives an off-path message, he must weigh the possibility of a message error by his father against the possibility of his father’s truthful report of an earlier action deviation by the father or some other ancestor in the dynasty.

Another result establishes that the type of mismatch in beliefs in the Super Folk theorem is equivalent to saying that an SE violates *Inter-Generational Agreement*, a notion that relates end-of-period beliefs of the fathers to beginning-of-period beliefs of the sons. We are able to show that any SPE of the standard repeated game can be replicated as an SE of the dynastic repeated game that displays Inter-Generational Agreement. Conversely any SE of the dynastic repeated game that yields a payoff vector that is not sustainable as an SPE of the standard repeated game must violate Inter-Generational Agreement.

The outline of the rest of the paper is as follows. In Section 2 we lay down the notation and details of the canonical dynastic repeated game in which each cohort of individuals live only one period. The Theorems are proved for this canonical case, but we later discuss in our conclusions (Section 7) how they generalize to games with arbitrary (bounded) demographics. In Section 3 we define what constitutes an SE for the Dynastic Repeated Game. In Sections 4 and 5 we present our “Super” Folk Theorems and other results. Section 6 reviews some related literature, and Section 7 concludes.

For ease of exposition, and for reasons of space, no formal proofs appear in the main body of the paper. The main ingredients (public randomization devices, strategies, beliefs, and trembles) for the proof of our extended Folk Theorem (Theorem 1) appear in Appendix A. The complete proof of Theorem 3 appears in Appendix B. A technical addendum to the paper contains the rest of the formal proofs.²

2. The Model

The stage game is described by the array $G = (A, u, I)$ where $I = \{1, \dots, n\}$ is the set of players, indexed by i . The n -fold cartesian product $A = \times_{i \in I} A_i$ is the set of pure action profiles $a = (a_1, \dots, a_n) \in A$, assumed to be finite. Stage game payoffs are defined by $u = (u_1, \dots, u_n)$ where $u_i : A \rightarrow \mathbb{R}$ for each $i \in I$. Let $\sigma_i \in \Delta(A_i)$ denote a mixed strategy for i , with σ denoting the profile $(\sigma_1, \dots, \sigma_n)$.³ We extend the use of $u_i(\cdot)$ to mixed strategies in the usual way and hence we write $u_i(\sigma)$ for i ’s expected payoff given the profile σ . Dropping the i subscript and writing $u(\sigma)$ gives the entire profile of expected payoffs.

²The technical addendum is available at <http://www.georgetown.edu/faculty/la2/FolkAddendum.pdf>. In the numbering of equations, Lemmas etc. a prefix of “A” or “B” means that the item is located in the corresponding Appendix.

³As is standard, here and throughout the rest, given any finite set Z , we let $\Delta(Z)$ be the set of all probability distributions over Z .

Throughout the rest of the paper, we denote by V the convex hull of the set of payoff vectors from pure strategy profiles in G . We let $\text{int}V$ denote the (relative) interior of V .

Time is discrete and indexed by $t = 0, 1, 2, \dots$. In the dynastic repeated game, 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. We refer to each date t -lived individual in dynasty i simply as player $\langle i, t \rangle$.

The realized action profile at time t is denoted by a^t . The stage game payoff $u_i(a^t)$ now refers to the time t component of the payoff of player $\langle i, t \rangle$. Each player $\langle i, t \rangle$ is altruistic in the sense that his payoff includes the discounted sum of the direct payoffs of all future individuals in the same dynasty. The players’ (common for simplicity) discount factor is $\delta \in (0, 1)$. Player $\langle i, t \rangle$ gives weight $1 - \delta$ to $u_i(a^t)$, and weight $(1 - \delta)\delta^\tau$ to $u_i(a^{t+\tau})$ for every $\tau \geq 1$. So, the (dynastic) payoff to player $\langle i, t \rangle$ is $(1 - \delta) \sum_{\tau=0}^{\infty} \delta^\tau u_i(a^{t+\tau})$.

At the beginning of period t , player $\langle i, t \rangle$ receives a *private* message m_i^t from player $\langle i, t-1 \rangle$. He does not directly observe anything about the previous history of play. The finite set of possible messages m_i^t is denoted by M_i^t . We will return to the latter shortly. It should also be noted at this point that our results would survive intact if we allowed *public messages as well as private ones*. The equilibria we construct below would still be viable, with the public messages ignored.⁴

In each period, all players observe the realizations x^t and y^t of two public randomization devices \tilde{x}^t and \tilde{y}^t . The realization x^t is observed at the action stage, immediately after each player $\langle i, t \rangle$ observes m_i^t . The realization y^t is observed at the message stage, immediately before each player $\langle i, t \rangle$ sends message m_i^{t+1} .⁵ These devices are independent and i.i.d. across periods; we write \tilde{x} and \tilde{y} to indicate the random variables of which all the \tilde{x}^t s and \tilde{y}^t s are independent “copies.” We refer to \tilde{x} and \tilde{y} respectively as the action-stage correlation device and the message-stage correlation device. The random variables \tilde{x} and \tilde{y} have full support and take values in the *finite* sets X and Y respectively, so that their distributions are points in $\Delta(X)$ and $\Delta(Y)$ respectively.

To summarize, at the beginning of each period t , player $\langle i, t \rangle$ receives a private message m_i^t from player $\langle i, t-1 \rangle$. He then observes x^t and subsequently chooses a (mixed) action to play in G . After observing the realized action profile a^t , he observes y^t and then chooses a (mixed) message m_i^{t+1} to send to player $\langle i, t+1 \rangle$.

⁴Dealing explicitly with both private and public messages would be cumbersome, and make our results considerably less transparent. Analyzing the model with private messages only is the most economical way to put our main point across, and hence this is how we proceed.

⁵We return to the role of correlation devices in Section 7. Working with two separate correlation devices simplifies our arguments and, in our view, facilitates an intuitive understanding of our results. However, it should be noted that our main result and its proof survives literally unchanged if we assume that both correlation devices are simultaneously observable to all players in each cohort at the action stage. In this case, clearly a *single* correlation device would suffice.

The *action strategy* of player $\langle i, t \rangle$ is denoted by g_i^t . This determines the mixed strategy that player $\langle i, t \rangle$ plays in the stage game at time t . In particular g_i^t takes as input a message $m_i^t \in M_i^t$ and a value $x^t \in X$ and returns a mixed strategy $\sigma_i^t \in \Delta(A_i)$.⁶ We write $g_i^t(m_i^t, x^t)$ to indicate the mixed strategy $\sigma_i^t \in \Delta(A_i)$ that player $\langle i, t \rangle$ plays after observing the pair (m_i^t, x^t) .

The *message strategy* of player $\langle i, t \rangle$ is denoted by μ_i^t . It takes as inputs a message m_i^t , the realization x^t , the realized action profile a^t , the realized value y^t , and returns the probability distribution ϕ_i^t over messages $m_i^{t+1} \in M_i^{t+1}$. We write $\mu_i^t(m_i^t, x^t, a^t, y^t)$ to indicate the (mixed) message $\phi_i^t \in \Delta(M_i^{t+1})$ that player $\langle i, t \rangle$ sends player $\langle i, t + 1 \rangle$ after observing the quadruple (m_i^t, x^t, a^t, y^t) .

It is convenient to specify fully the finite message sets M_i^t at this point. The choice of message spaces is to a large extent arbitrary since enlarging message spaces does not shrink the set of SE payoffs, as Lemma T.3.1 demonstrates.⁷ We proceed in a way that allows comparability with the standard repeated game model. Define $h^t = (x^0, a^0, \dots, x^{t-1}, a^{t-1})$, with $h^0 = \emptyset$ and the set of all possible h^t 's denoted by H^t . From now on we take set of messages M_i^{t+1} available to each player $\langle i, t \rangle$ to send to player $\langle i, t + 1 \rangle$ to be equal to H^{t+1} .⁸

Because of our choice of message spaces $M_i^t = H^t$, formally the array $(g_i^0, g_i^1, \dots, g_i^t, \dots)$ is exactly the *same object* that defines a strategy for player i in the standard model where G is played infinitely many times by n infinitely lived players who in each period t observe the past history of play $h^t = (x^0, a^0, \dots, x^{t-1}, a^{t-1})$. This is immediate since each g_i^t takes as input an element of H^t and a value $x^t \in X$, and returns a mixed strategy $\sigma_i^t \in \Delta(A_i)$.

In denoting profiles and sub-profiles of strategies we use standard notational conventions. We let g_i denote the i -th dynasty profile $(g_i^0, g_i^1, \dots, g_i^t, \dots)$, while g^t indicates the time t profile (g_1^t, \dots, g_n^t) and g the entire profile of action strategies (g_1, \dots, g_n) . Similarly, we set $\mu_i = (\mu_i^0, \mu_i^1, \dots, \mu_i^t, \dots)$, as well as $\mu^t = (\mu_1^t, \dots, \mu_n^t)$ and $\mu = (\mu_1, \dots, \mu_n)$. Therefore, the pair (g, μ) entirely describes the behavior of all players in our model.

⁶It should be emphasized at this stage that nowhere in the paper we assume that mixed strategies are observable. Whatever (mutually independent) devices the players use to achieve a desired randomization among pure actions in G given m_i^t and a realization of \tilde{x}^t , remain unobservable to other players.

⁷This is because we can “replicate” any SE of the dynastic repeated game with smaller message spaces as an SE of the dynastic repeated game with larger message spaces by mapping each message in the smaller set to a finite set of messages in the larger message space. A choice of message in the smaller message space corresponds to a (uniform) randomization over the entire corresponding set in the larger message space. A player receiving one of the randomized messages in the larger message space acts just like the player who receives the corresponding message in the smaller message set.

⁸Observe that our choice of message spaces implies that at $t = 0$ all players $\langle i \in I, 0 \rangle$ receive the “null” message $m_i^0 = \emptyset$.

Notice also that we are excluding the relevant realized values y^τ from the message space of player $\langle i, t \rangle$. This is without loss of generality because, as we noted above, enlarging the message spaces does not shrink the set of SE payoffs.

3. Sequential Equilibrium

Our focus is the set of SE of the model. The notion of SE is a widely accepted benchmark. It is particularly well suited to dynastic games because it imposes a strong discipline on how memory gets transmitted through time. This is because information transmission in a dynastic game with private messages depends critically on the “interpretation” of off path events; and how they may be due to message “mistakes” or action ones. In an SE, a society must have a common outlook on the relative orders of magnitude of such mistakes.

While the original definition of SE in Kreps and Wilson (1982) does not readily apply to games with infinitely many players, only a minor adaptation of the concept is needed to apply it to our set-up.

To spell it out, we begin with the observation that the beliefs of player $\langle i, t \rangle$ can in fact be boiled down to a simpler object than one might expect at first sight, because of the structure of the dynastic repeated game. Upon receiving message m_i^t , in principle, we would have to define the beliefs of player $\langle i, t \rangle$ over the *entire set of possible past histories of play*. However, when player $\langle i, t \rangle$ is born, an entire cohort of new players replaces the $t - 1$ -th one, and hence the real history leaves no trace other than the messages (m_1^t, \dots, m_n^t) that have been sent to cohort t . It follows that, without loss of generality, after player $\langle i, t \rangle$ receives message m_i^t we can restrict attention to his beliefs over the $n - 1$ -tuple m_{-i}^t of messages received by other players in cohort t .⁹ This probability distribution, specifying the beginning-of-period beliefs of player $\langle i, t \rangle$, will be denoted by $\Phi_i^{tB}(m_i^t)$ throughout. When the dependence of this distribution on m_i^t can be omitted from the notation without causing any ambiguity we will write it as Φ_i^{tB} . The notation $\Phi_i^{tB}(\cdot)$ will indicate the entire array of possible probability distributions $\Phi_i^{tB}(m_i^t)$ as $m_i^t \in M_i^t$ varies.

Consider now the two classes of information sets at which player $\langle i, t \rangle$ is called upon to play: the first defined by a pair (m_i^t, x^t) when he has to select a mixed strategy σ_i^t , and the second defined by a quadruple (m_i^t, x^t, a^t, y^t) when he has to select a probability distribution ϕ_i^t over the set of messages H^{t+1} .

The same argument as above now suffices to show that at the (m_i^t, x^t) information set we can again restrict attention to the beliefs of player $\langle i, t \rangle$ over the $n - 1$ -tuple m_{-i}^t of messages received by other players in cohort t . Moreover, since all players observe the same x^t and this realization is independent of what happened in the past, player $\langle i, t \rangle$ beliefs over m_{-i}^t must be the same as when he originally received message m_i^t .

Finally, at the information set identified by the quadruple (m_i^t, x^t, a^t, y^t) , we can restrict attention to the beliefs of player $\langle i, t \rangle$ over the $n - 1$ -tuple m_{-i}^{t+1} of messages that the other players in cohort t are about to send to cohort $t + 1$. Just as before, since all players are replaced by a new cohort and time- t payoffs have already been realized, this is all that could ever

⁹It should be made clear that the beliefs of player $\langle i, t \rangle$ over m_{-i}^t will in fact depend on the relative likelihoods of the actual histories of play that could generate different $n - 1$ -tuples m_{-i}^t . What we are asserting here is simply that once we know the player’s beliefs over m_{-i}^t , we have all that is necessary to check that his behavior is optimal given his beliefs.

matter for the payoff to player $\langle i, t \rangle$ from this point on. This probability distribution, specifying the end-of-period beliefs of player $\langle i, t \rangle$, will be denoted by $\Phi_i^{tE}(m_i^t, x^t, a^t, y^t)$ throughout the rest of the paper and the technical addendum. When the dependence of this distribution on (m_i^t, x^t, a^t, y^t) can be omitted from the notation without causing any ambiguity we will write it as Φ_i^{tE} . The notation $\Phi_i^{tE}(\cdot)$ will indicate the entire array of possible probability distributions $\Phi_i^{tE}(m_i^t, x^t, a^t, y^t)$ as the quadruple (m_i^t, x^t, a^t, y^t) varies.

In the proofs of our results, we will also need to refer to the (revised) end-of-period beliefs of player $\langle i, t \rangle$ on the $n - 1$ -tuple of messages m_{-i}^t after he observes not only m_i^t , but also (x^t, a^t, y^t) . These will be indicated by $\Phi_i^{tR}(m_i^t, x^t, a^t, y^t)$, with the arguments omitted when this does not cause any ambiguity. The notation $\Phi_i^{tR}(\cdot)$ will indicate the entire array of possible probability distributions $\Phi_i^{tR}(m_i^t, x^t, a^t, y^t)$ as the quadruple (m_i^t, x^t, a^t, y^t) varies.

Throughout the rest of the paper we refer to the array $\Phi = \{\Phi_i^{tB}(\cdot), \Phi_i^{tE}(\cdot)\}_{t \geq 0, i \in I}$ as a *system of beliefs*. Following standard terminology we will also refer to a triple (g, μ, Φ) , a strategy profile and a system of beliefs, as an *assessment*. Also following standard terminology, we will say that an assessment (g, μ, Φ) is *consistent* if the system of beliefs Φ can be obtained (in the limit) using Bayes' rule from a sequence of completely mixed strategies that converges to (g, μ) . Since this is completely standard, for reasons of space we do not specify any further details here.

Definition 1. *Sequential Equilibrium:* An assessment (g, μ, Φ) constitutes an SE for the dynastic repeated game if and only if (g, μ, Φ) is consistent, and for every $i \in I$ and $t \geq 0$ strategy g_i^t is optimal for player $\langle i, t \rangle$ given beliefs $\Phi_i^{tB}(\cdot)$, and strategy μ_i^t is optimal for the same player given beliefs $\Phi_i^{tE}(\cdot)$.

4. A “Super” Folk Theorem

We anticipated that our “Super” Folk Theorem applies to a class of stage games that includes all games with $n \geq 4$ players for which the standard Folk Theorem yields a payoff set with a non-empty interior.

The class of stage games to which Theorem 1 below applies is larger — in a significant sense — than the one we just mentioned again. In essence, our result encompasses all stage games that *locally* satisfy the conditions we have outlined.

Before any formal definitions, an example will help bring out the point. Consider the following simple 4-player version of the Prisoners' Dilemma, which we will refer to as G^{PD} . Each player has two strategies labeled C and D . The payoffs to player i are $u_i(D, D, D, D) = 0$, $u_i(C, C, C, C) = 2$, $u_i(D, C_{-i}) = 3$, $u_i(C, D_{-i}) = -1$, $u_i(D, Z_{-i}) = 3$ and $u_i(C, Z_{-i}) = -1$, where Z_{-i} is any string of length 3 that contains at least one C and one D . Clearly D is a dominant strategy for every player in G^{PD} , and the minmax payoff vector is $(0, 0, 0, 0)$. Moreover, the payoff vector $(2, 2, 2, 2)$ strictly Pareto-dominates $(0, 0, 0, 0)$ and the convex hull of payoff vectors V^{PD} has dimension 4. So, all the hypotheses of the standard Folk Theorem are satisfied, and as a consequence our “Super” Folk Theorem applies to the dynastic repeated game when the stage game is G^{PD} .

Next, consider a new “augmented” version of G^{FD} , which we refer to as G^A , which is derived from G^{FD} by adding a third strategy H for each of the four players. The payoffs in G^A are as in G^{FD} whenever all players play either C or D , and all players get a payoff of 4 whenever one or more players play H . Clearly, the standard Folk Theorem does not yield any equilibrium payoff multiplicity in G^A . This is because each player’s minmax payoff in G^A is 4 — obtained by playing the dominant strategy H — and no vector of payoffs in V^A Pareto-dominates $(4, 4, 4, 4)$.

Our “Super” Folk Theorem *does* apply to the dynastic repeated game when the stage game is G^A . The reason is that, if we restrict attention to a subset of pure strategies for each player (those yielding G^{FD}), the hypotheses of the standard Folk Theorem apply and a payoff vector that strictly dominates the vector of minmax payoffs is in fact available. In short the conditions necessary for the standard Folk Theorem to yield a payoff set with a non-empty interior apply “locally” to the stage game G^A .

We now formalize these ideas in order to state formally our first result. Given a stage game $G = (A, u, I)$, we indicate by $\tilde{A} \subseteq A$ a typical set of pure action profiles with a *product structure*. In other words, we require that there exist an array $(\tilde{A}_1, \dots, \tilde{A}_n)$, with each \tilde{A}_i containing at least two pure strategies, and with $\tilde{A} = \times_{i \in I} \tilde{A}_i$. Given a product set \tilde{A} , we denote by $G^{\tilde{A}}$ the normal form game in which each player’s pure strategy set is given by \tilde{A}_i . Throughout, we refer to $G^{\tilde{A}}$ as a *component game* of G , and we denote the convex hull of payoffs in $G^{\tilde{A}}$ by $V(\tilde{A})$.

We are now ready to state our “Super” Folk Theorem.

Theorem 1. *Dynastic Folk Theorem:* *Let a stage game G with four or more players be given. Assume that G has a component game $G^{\tilde{A}}$ with the following two properties.¹⁰ First, there exists a payoff vector in $V(\tilde{A})$ that strictly Pareto-dominates the vector of minmax payoffs of $G^{\tilde{A}}$. Second, $V(\tilde{A})$ has dimension n . Then for every $v \in \text{int}V$ there exists a $\underline{\delta} \in (0, 1)$ such that $\delta > \underline{\delta}$ implies v is sustained by a SE with discount factor δ .*

The proof of Theorem 1 is constructive. In Appendix A we formally describe the randomization devices, the equilibrium strategies and the trembles that yield the equilibrium beliefs. The rest of the argument, deriving equilibrium beliefs from the trembles via Bayes’ rule and verifying sequential rationality, is available in Appendix A and in full detail in sections T.4 through T.7 of the technical addendum to the paper.¹¹

¹⁰Since $\tilde{A} \subseteq A$, the component game $G^{\tilde{A}}$ may be taken to be G itself.

¹¹It should also be noted that the actual statement that we prove is stronger than Theorem 1 above. This is reported below as Theorem A.1 in Appendix A. The statement of Theorem A.1 is stronger than that of Theorem 1 in two respects. First, it does not make explicit reference to the dimensionality of the payoff space of the component game $G^{\tilde{A}}$, but refers to the existence of certain vectors of payoffs in it. This makes explicit the fact that the dimensionality condition in the statement of Theorem 1 could be replaced by a weaker “Non-Equivalent Utilities” condition (Abreu, Dutta, and Smith, 1994) on the payoff space of $G^{\tilde{A}}$. Note the dimensionality restrictions can also be loosened in repeated games in continuous time — see

In certain respects, the SE that we construct to prove Theorem 1, mirrors the intuition of the Introductory Example. As in the example, a player $\langle i, t - 1 \rangle$ may want to communicate to his successor $\langle i, t \rangle$ that dynasty i is being punished for having deviated, but will be unable to do so in an effective way. This is not due to an inability to communicate due to a shortage of possible messages, even if the message space we use in the proof is smaller than the set of histories (see Footnote 7). Rather, the correct interpretation is that *in equilibrium* there is *no message* that player $\langle i, t \rangle$ might possibly *interpret* in the way that $\langle i, t - 1 \rangle$ would like.¹²

We now sketch the proof of our “Super” Folk Theorem. Consider a stage game $G = (A, u, I)$ and a product set $\tilde{A} \subseteq A$ yielding a component game $G^{\tilde{A}}$ satisfying the hypotheses of Theorem 1. Denote by $\omega_i(\tilde{A})$ the minmax payoff to player i in the component game $G^{\tilde{A}}$. Since $V(\tilde{A})$ has dimension n , we can find a $\hat{v} \in \text{int}V(\tilde{A})$ and an array of n payoff vectors $\bar{v}^1, \dots, \bar{v}^n$ in $V(\tilde{A})$ with the following properties. First, $\omega_i(\tilde{A}) < \bar{v}_i^i < \bar{v}_i^j$, and second $\bar{v}_i^i < \hat{v}_i$ for every $i \in I$ and every $j \neq i$.

For simplicity, assume that the payoff vectors $\bar{v}^1, \dots, \bar{v}^n$ can all be obtained from pure profiles of actions in \tilde{A} . Also for simplicity, assume that each of the payoffs $\omega_i(\tilde{A})$ can be obtained from some pure profile of actions in \tilde{A} .

The argument can be divided into two parts. First, we will argue that if δ is close to one, it is possible to sustain the payoff vector $\hat{v} \in V(\tilde{A})$ as an SE of the dynastic repeated game. Notice that $\hat{v} \in V(\tilde{A})$ could well already be *below the minmax* payoff in G for one or more players. We call this the “local” part of the argument. Second, we will argue that via a judicious use of the action-stage randomization device it is possible to go from the local argument to a “global” one and sustain every feasible payoff vector as required by the statement of the theorem.

The equilibrium path generated by the strategies we construct consists of $n + 1$ phases. We call the first one the *standard equilibrium phase*, the second one the *diversionary-1 equilibrium phase*, through to the *diversionary- n equilibrium phase*.

If all players $\langle i \in I, t \rangle$ receive message m^* then play is in the standard equilibrium phase. For simplicity again we proceed with our outline of the construction assuming that the equilibrium prescribes that the players $\langle i \in I, t \rangle$ play a pure action profile during the standard equilibrium phase, denoted by a' . The associated payoff vector is v' .

If all players $j \neq i$ in the t -th cohort receive message \check{m}^i , and player $\langle i, t \rangle$ receives any message $m^{i,\tau}$ in a finite set $\underline{M}(i, t) = \{m^{i,1}, \dots, m^{i,T}\} \subset M_i^t$, then play is in the diversionary- i

(Bergin, 2006). Second, in Theorem A.1 the vector of minmax payoffs for the component game $G^{\tilde{A}}$ is defined with the following added twist (see Definition A.1). When minmaxing player i , the other players are allowed to choose any *correlated* mixed strategies in $G^{\tilde{A}}$, while player i is only allowed to choose a best reply in \tilde{A}_i , without being able to condition on the same correlation device. In both cases, the weakened conditions of Theorem A.1 make it applicable to a wider class of stage games. We have refrained from including these extensions to Theorem 1 in the main body of the paper for reasons of space and of expository simplicity.

¹²At this point it is legitimate to wonder whether the concept of “neologism-proof” equilibrium (Farrell, 1993) has any impact on what we are saying here. While neologism-proofness in its current form does not apply to our model, we return to this point at some length in Section 7 below.

phase.¹³ Let \underline{a}^i be the vector of actions (pure, for simplicity) for which i receives $\omega_i(\tilde{A})$ — his minmax payoff in the component game $G^{\tilde{A}}$. During the diversionary- i equilibrium phase player $\langle i, t \rangle$ plays action \underline{a}_i^i . For all players $j \neq i$, let \check{a}_j^i be any action in \tilde{A}_j that is not equal to \underline{a}_j^i . Such action can always be found since, by assumption, each \tilde{A}_j contains at least two actions. During the diversionary- i equilibrium phase, any player $j \neq i$ plays action \check{a}_j^i . The (per-period) payoff vector associated with the diversionary- i equilibrium phase is denoted by \check{u}^i .

If in period t play is in any of the equilibrium phases we have just described, and no deviation occurs at the action stage, at the end of period t all players use the realization y^t of the message-stage randomization device to select the message to send to their successors. The possible realizations of \tilde{y}^t are $(y(0), y(1), \dots, y(n))$. The probability that $y^t = y(0)$ is $1 - \eta$ and the probability that $y^t = y(i)$ is η/n for every $i \in I$. Consider now the end of any period t in any equilibrium phase, and assume that no deviation has occurred. If $y^t = y(0)$ then all players $\langle i \in I, t \rangle$ send message m^* to their successors, so play in period $t + 1$ is in the standard equilibrium phase. If $y^t = y(i)$, then all players $j \neq i$ send message \check{m}^i to their successors and player $\langle i, t \rangle$ sends a (randomly selected) message $m^{i,\tau} \in \underline{M}(i, t)$ to player $\langle i, t + 1 \rangle$. So, in this case in period $t + 1$ play is in the diversionary- i equilibrium phase.

The profiles to be played in each diversionary- i equilibrium phase may of course be entirely determined by the need to differ from the component game minmax action profiles, so we have no degrees of freedom there. However, we are free to choose the profile a' in constructing the standard equilibrium phase. Recall that we take a' to be pure solely for expositional simplicity. Using the action-stage randomization device, clearly we could select correlated actions for the standard equilibrium phase that yield any payoff vector v' in $V(\tilde{A})$. Since $\hat{v} \in \text{int}V(\tilde{A})$ we can be sure that for some $v' \in V(\tilde{A})$ and some $\eta \in (0, 1)$

$$\hat{v} = (1 - \eta)v' + \frac{\eta}{n} \sum_{i=1}^n \check{u}^i \quad (1)$$

So that (modulo our expositional assumption that $v' = u(a')$) when play is in any equilibrium phase the expected (across all possible realizations of \tilde{y}^t) continuation payoff to any player $\langle i, t \rangle$ from the beginning of period $t + 1$ onward is \hat{v}_i .

The strategies we construct also define an off-path collection of n *punishment phases* and n *terminal phases*, one of each type for each dynasty i . In the punishment- i phase, in every period player i receives his component game minmax payoff $\omega_i(\tilde{A})$, and the phase lasts T periods. In the terminal- i phase in every period the players receive the vector of payoffs \bar{v}^i . The transition between any of the equilibrium phases and any of the punishment or terminal phases is akin to the benchmark construction in Fudenberg and Maskin (1986).

¹³In the formal proof the set of messages $\underline{M}(i, t)$ actually does depend on the time index t because not all messages are available for the first T periods of play. This is so in order to avoid any message space M_i^t having a cardinality that *exceeds* that of H^t .

In other words, a deviation by dynasty i during any of the equilibrium, any of the punishment phases or any of the terminal phases triggers the start (or re-start) of the punishment- i phase (deviations by two players or more are ignored). The terminal- i phase begins after play has been, without subsequent deviations, in the punishment- i phase for T periods. For an appropriately chosen (large) T , as δ approaches 1, with one critical exception, the inequalities in (iii) of the statement of Theorem 1 are used in much the same way as in Fudenberg and Maskin (1986) to guarantee that no player deviates from the prescriptions of the equilibrium strategies.

In Fudenberg and Maskin (1986), during the punishment- i phase player i plays a *myopic best response* to the actions of other players. Critically, this is not the case here. During the punishment- i phase dynasty i plays a best response to the strategy of others *restricted to the set of pure actions in \tilde{A}_i* . Clearly this could be very far away (in per-period payoff terms) from an unconstrained best reply chosen at will within A_i . To understand how this can happen in an SE we need to specify what message profiles mark the beginning of the punishment- i phase and what the players' beliefs are.

Suppose that player $\langle i, t \rangle$ deviates from the prescriptions of the equilibrium strategy and triggers the start of the punishment- i phase as of the beginning of period $t + 1$. Then all players $\langle j \in I, t \rangle$ send message $m^{i,T}$ to their successors. These messages are interpreted as telling all players $\langle j \neq i, t + 1 \rangle$ that the punishment- i phase has begun, and that there are T periods remaining. We return to the beliefs of player $\langle i, t + 1 \rangle$ from the guilty dynasty shortly. In the following period of the punishment- i phase all players $\langle j \in I, t + 1 \rangle$ send message $m^{i,T-1}$, then $m^{i,T-2}$ and so on, so that the index τ in a message $m^{i,\tau}$ is effectively interpreted (by dynasties $j \neq i$) as a “punishment clock,” counting down how many periods remain in the punishment- i phase.

The players' beliefs in the SE we construct are “correct” in all phases of play except for the beliefs of player $\langle i, t \rangle$ whenever play is in the punishment- i phase at time t . Upon receiving any message $m^{i,\tau} \in \underline{M}(i, t)$, player $\langle i, t \rangle$ believes that play is in the punishment- i phase with probability zero. Instead he believes that play is in the i -diversionary equilibrium phase with probability one. This is possible in an SE because player $\langle i, t \rangle$ conflates any message that might indicate the off-path punishment- i phase with the message indicating the on-path diversionary- i phase. It follows easily that player $\langle i, t \rangle$ does not want to deviate from the equilibrium strategies we have described when play is in the punishment- i phase.

Notice moreover that if at time t play is in the punishment- i phase, after the profile a^t is observed, player $\langle i, t \rangle$ will discover that play is in fact in the punishment- i phase, contrary to his beginning-of-period beliefs, even if a new deviation occurs at the action stage of period t . This is because, by construction, all players $\langle j \neq i, t \rangle$ play an action (namely \check{a}_j^i) in the diversionary- i equilibrium phase that is different from what they play in the punishment- i phase (namely \underline{a}_j^i). This, coupled with the assumption that $n \geq 4$ will ensure that $\langle i, t \rangle$ will discover the truth, and the identity of any deviator.¹⁴

¹⁴Clearly, if $\langle i, t \rangle$ could not be guaranteed to discover that play is in the punishment- i phase, or the identity

Therefore any player $\langle i, t \rangle$ who knows that in period $t + 1$ play will be in the punishment- i phase would like to “communicate effectively” to player $\langle i, t + 1 \rangle$ that play is in the punishment- i phase but is unable to do so, as in our discussion above concerning message spaces. After receiving $m^{i,\tau}$ and discovering that play is in the punishment- i phase, sending any message to player $\langle i, t + 1 \rangle$ that is not $m^{i,\tau-1}$ may cause him, or some of his successors, to deviate and hence to re-start the punishment- i phase.¹⁵ Sending $m^{i,\tau-1}$ will cause player $\langle i, t + 1 \rangle$ and his successors to play a best response among those that can be induced by any available message. This is because \underline{a}_i^i is in fact a best response to the minmax (in \tilde{A}_i) of the other players within the set \tilde{A}_i . Therefore, given the inequalities satisfied by \hat{v} and $\bar{v}^1, \dots, \bar{v}^n$, since T is sufficiently large, and δ is close enough to one, no profitable deviation is available to player $\langle i, t \rangle$.

The argument we have just outlined suffices to show that the payoff vector \hat{v} of the statement of the theorem can be sustained as an SE of the dynastic repeated game. We now argue that this fact can be used as a “local anchor” for our “global” argument that shows that *any* interior payoff vector can be sustained as an SE.

Fix any $v^* \in \text{int}V$ to be sustained in equilibrium. Since v^* is interior it is obvious that it can be expressed as $v^* = q\hat{v} + (1 - q)z$ for some $q \in (0, 1)$ and some $z \in V$. The “global” argument then consists of using the action-stage randomization device so that in each period with probability q play proceeds as in the construction above, while with probability $1 - q$ the (expected) payoff vector is z . The latter is achieved with action-stage strategies that do not depend on the messages received. A deviation by i from the (correlated) actions needed to implement z triggers the punishment- i phase. With one proviso to be discussed shortly, it is not hard to then verify that this is sufficient to keep all players from deviating at any point, and hence that v^* can be sustained as an SE payoff vector of the dynastic repeated game for δ sufficiently close to one.

The difficulty with the global argument we have outlined that needs some attention is easy to point out. The periods in which the action-stage randomization device tells the players to implement the payoff vector z cannot be counted as real punishment periods. They in fact stochastically interlace all phases of play, including any punishment- i phase. However, the length of effective punishment T has to be sufficiently large to deter deviations. The solution we adopt is to ensure that the punishment clock does not decrease in any period in which (with probability $1 - q$) the payoff vector z is implemented at the action stage. In effect, this makes the length of any punishment- i phase stochastic, governed by a punishment clock that counts down only with probability q in every period.

Theorem 1 assumes four dynasties or more. This is clearly essential to the construction

of any deviator at time t , then we could not construct strategies that guarantee that if $\langle j \neq i, t \rangle$ deviates during the punishment- i phase then play switches to the punishment- j phase, as required.

¹⁵Checking sequential rationality at the message stage takes a few more steps than may appear from our intuitive outline of the argument given here. This is because a deviation at the message stage may trigger *multiple deviations*; that is deviations at the action and/or message stage by more than one successor of any given player. The core of the argument dealing with this case is Lemma T.5.4 in the technical addendum.

we use in its proof. Is it essential to our results? One answer is that, under some conditions, our results can be extended to the case of three dynasties or more. In Anderlini, Gerardi, and Lagunoff (2004), Theorem 2, we prove

Theorem 2. *Dynastic Folk Theorem: Three Dynasties or More: Let any stage game G with three or more players be given. Assume that G is such that we can find two pure action profiles a^* and a' in A with*

$$u_i(a^*) > u_i(a') > u_i(a_i^*, a'_{-i}) \quad \forall i \in I \quad (2)$$

Then for every $v \in \text{int}V$ there exists a $\underline{\delta} \in (0, 1)$ such that $\delta > \underline{\delta}$ implies v is sustained by a SE with discount factor δ .

In the construction given in Anderlini, Gerardi, and Lagunoff (2004), the profile a' is used as a “common punishment.” The availability of a common punishment reduces the need from four dynasties or more to three or more. Intuitively this is because only the fact that a deviation has occurred needs to be discovered, as opposed to having to discover the identity of the deviator to administer personalized punishments, as in the construction used here. The construction used to prove a “Super” Folk Theorem for three dynasties or more also shows that, in a dynastic repeated game like the one set out here, the structure of SE can be radically different from the canonical one used to prove Folk Theorems in a standard repeated game (Fudenberg and Maskin (1986)). In particular, we construct SE of the dynastic repeated game in which some deviations trigger a permanent punishment phase in which a' as above is played (a' need not be a Nash equilibrium of G). The path of play can become trapped in a permanent punishment phase because of the mismatch in players’ beliefs to which we have referred several times above.

5. Inter-Generational (Dis)Agreement

Some of the SE of the dynastic repeated game we have identified in Theorem 1 above clearly do not correspond in any meaningful sense to any SPE of the standard repeated game. This is obvious if we consider an SE of the dynastic repeated game in which one or more players receive a payoff below their minmax value in the stage game G .

An obvious question to raise at this point is then the following. What is it that makes these SE viable? To put it another way, can we identify any properties of an SE of the dynastic repeated game which ensure that it must correspond in a meaningful sense to an SPE of the standard repeated game? The answer is yes, and this is what this section of the paper is devoted to.

The critical properties of an SE that we identify concern the players’ beliefs. These properties characterize entirely the set of SE yielding payoffs outside the set of SPE of the standard repeated game. Therefore, if one wanted to attempt to “refine away” the equilibria yielding payoffs outside the set of SPE, our results in this section pin down precisely which belief systems the proposed refinement would have to rule out.

The first of the properties we identify is that a player’s (revised) beliefs at the end of the period over the messages received by other players at the beginning of the period should be the *same* as at the beginning of the period. This is equation (3) below. The second is that the end-of-period beliefs of a player (over messages sent by his opponents) should be the *same* as the beginning-of-period beliefs of his successor (over messages received by his opponents). This is equation (4) below. In fact we are able to show that if this property holds for all players (and all information sets) in an SE of the dynastic repeated game, then this SE must be payoff-equivalent to some SPE of the standard repeated game. For want of a better term, when an SE of the dynastic repeated game has the two properties (of beliefs) that we just described informally, we will say that it displays *Inter-Generational Agreement*.

Definition 2. *Inter-Generational Agreement:* Let an SE triple (g, μ, Φ) of the dynastic repeated game be given.

We say that this SE displays *Inter-Generational Agreement* if and only if for every $i \in I$, $t \geq 0$, $m_i^t \in H^t$, $x^t \in X$, $a^t \in A$ and $y^t \in Y$ we have that

$$\Phi_i^{tR}(m_i^t, x^t, a^t, y^t) = \Phi_i^{tB}(m_i^t) \quad (3)$$

and for every m_i^{t+1} in the support of $\mu_i^t(m_i^t, x^t, a^t, y^t)$

$$\Phi_i^{tE}(m_i^t, x^t, a^t, y^t) = \Phi_i^{t+1B}(m_i^{t+1}) \quad (4)$$

We are now ready to state our last result.

Theorem 3. *SE of the Dynastic and SPE of the Standard Repeated Game:* Fix a stage game G , any $\delta \in (0, 1)$, and any \tilde{x} and \tilde{y} . Let (g, μ, Φ) be an SE of the dynastic repeated game. Assume that this SE displays *Inter-Generational Agreement* as in Definition 2. Let v^* be the vector of (dynastic) payoffs for $t = 0$ players in this SE. Then v^* is a SPE payoff profile of the standard repeated game with the same discount factor δ .

The proof of Theorem 3 is in Appendix B. Before proceeding with an intuitive outline, we state a remark on the implications of the theorem.

Clearly, Theorem 3 implies that if an SE payoff vector v is not sustainable as by a SPE, then it must be the case that no SE which sustains v in the dynastic repeated game displays *Inter-Generational Agreement*. Moreover, it is fairly straightforward to show that any SPE payoff profile can be sustained by an SE that satisfies *Intergenerational Agreement* (see Anderlini, Gerardi, and Lagunoff (2004), Theorem 1). Together with Theorem 3, this gives us a complete characterization in payoff terms of the relationship between the SE of the dynastic and the SPE of the standard repeated game as follows.

For any G and $\delta \in (0, 1)$, a payoff profile v is not sustained by a SPE if and only if any SE that sustains v in the dynastic repeated game violates *Inter-Generational Agreement*.

To streamline the exposition of the outline of the argument behind Theorem 3, make the following two simplifying assumptions. First, assume that at the message stage the players do not condition their behavior at the message-stage on the randomization device. The simplest way to fix ideas here is to consider a message-stage randomization device with a singleton Y (the set of possible realizations). Second, assume that all message-strategies μ_i^t are *pure*. In other words, even though they may randomize at the action stage, all players at the message stage send a single message, denoted $\mu_i^t(m_i^t, x^t, a^t)$, with probability one.¹⁶

Now consider an SE (g, μ, Φ) of the dynastic repeated game that satisfies Inter-Generational Agreement as in Definition 2. Fix any history of play $h^t = (x^0, a^0, \dots, x^{t-1}, a^{t-1})$. For each dynasty i , using the message strategies of players $\langle i, 0 \rangle$ through to $\langle i, t-1 \rangle$, we can now determine the message m_i^t that player $\langle i, t-1 \rangle$ will send to his successor, player $\langle i, t \rangle$. Denote this message by $m_i^t(h^t)$. Notice that $m_i^t(h^t)$ can be determined simply by recursing forward from period 0. Recall that at the beginning of period 0 all players $\langle i \in I, 0 \rangle$ receive message $m_i^0 = \emptyset$. Therefore, given $h^1 = (x^0, a^0)$, using μ_i^0 , we know $m_i^1(h^1)$. Now using μ_i^1 we can compute $m_i^2(h^2) = \mu_i^1(m_i^1(h^1), x^1, a^1)$, and so recursing forward the value of $m_i^t(h^t)$ can be worked out.

Because the SE (g, μ, Φ) satisfies Inter-Generational Agreement it must be the case that, after any actual history of play (on or off the equilibrium path) h^t , and therefore after receiving message $m_i^t(h^t)$, player $\langle i, t \rangle$ believes that his opponents have received messages $(m_1^t(h^t), \dots, m_{i-1}^t(h^t), m_{i+1}^t(h^t), \dots, m_n^t(h^t))$ with probability one.

To see why this is the case, we can recurse forward from period 0 again. Consider the end of period 0. Since all players in the $t = 0$ cohort receive message \emptyset , after observing (x^0, a^0) , player $\langle i, 0 \rangle$ *knows* that any player $\langle j \neq i, 0 \rangle$ is sending message $m_j^1(x^0, a^0) = \mu_j^0(\emptyset, x^0, a^0)$ to his successor player $\langle j, 1 \rangle$.

Equation (4) of Definition 2 guarantees that the beginning-of-period beliefs of player $\langle i, 1 \rangle$ must be the same as the end-of-period beliefs of player $\langle i, 0 \rangle$. So, at the beginning of period 1, player $\langle i, 1 \rangle$ believes with probability one that every player $\langle j \neq i, 1 \rangle$ has received message $m_j^1(x^0, a^0)$ as above.

Equation (3) of Definition 2 guarantees that player $\langle i, 1 \rangle$ will not revise his beginning-of-period beliefs during period 1. Therefore, after observing any (x^1, a^1) , player $\langle i, 1 \rangle$ still believes that every player $\langle j \neq i, 1 \rangle$ has received message $m_j^1(x^0, a^0)$ as above. But this, via the message strategies μ_j^1 implies that player $\langle i, 1 \rangle$ must believe with probability one that every player $\langle j, 1 \rangle$ sends message $m_j^2(h^2) = m_j^2(x^0, a^0, x^1, a^1) = \mu_j^1(m_j^1(h^1), x^1, a^1)$ to his successor $\langle j, 2 \rangle$. Continuing forward in this way until period t we can then see that the beginning-of-period beliefs of player $\langle i, t \rangle$ are as we claimed above.

Before we proceed to close the argument for Theorem 3, notice that *both* conditions of Definition 2 are necessary for our argument so far to be valid. Intuitively, the forward

¹⁶Abusing notation slightly, here and throughout, we will write $g_i^t(m_i^t, x^t) = a_i$ to mean that the distribution $g_i^t(m_i^t, x^t)$ assigns probability one to a_i . Similarly, we will write $\mu_i^t(m_i^t, x^t, a^t, y^t) = m_i^{t+1}$ to mean that the distribution $\mu_i^t(m_i^t, x^t, a^t, y^t)$ assigns probability one to m_i^{t+1} .

recursion argument we have outlined essentially ensures that the “correct” (because all its members receive a *given* message $m_i^0 = \emptyset$) beliefs of the first cohort “propagate forward” as follows. At the beginning of period $t = 1$ each player $\langle i \in I, 1 \rangle$ must have correct beliefs since the end-of-period beliefs of all players in period 0 are trivially correct, and equation (4) of Definition 2 tells us that the end-of-period beliefs must be the same as the beginning-of-periods beliefs of the next cohort. Now some pair (x^1, a^1) is observed by all players $\langle i \in I, 1 \rangle$. If this pair is consistent with their beginning-of-period beliefs, then clearly no player $\langle i \in I, 1 \rangle$ can possibly revise his beliefs on the messages received by others at the beginning of period 1. However, if (x^1, a^1) is *not* consistent with the beliefs of players $\langle i \in I, 1 \rangle$ and their action strategies, some players in the $t = 1$ cohort may be “tempted” to revise their beginning-of-period beliefs. This is because an observed “deviation” from what they expect to observe in period 1 can always be attributed to two distinct sources. It could be generated by an actual deviation at action stage of period 1, or it could be the result of one (or more) players in the $t = 0$ cohort having deviated at the message stage of period 0. Equation (3) of Definition 2 essentially requires that the $t = 1$ players should always interpret an “unexpected” pair (x^1, a^1) as an actual deviation at the action stage. The same applies to all subsequent periods. So, while equation (4) of Definition 2 ensures that the initially correct beliefs are passed on from one generation to the next, equation (3) of Definition 2 guarantees that actual deviations will be treated as such in the beliefs of players who observe them. The beliefs of players $\langle i \in I, 0 \rangle$ are correct and the end-of-period beliefs of any cohort are guaranteed to be the same as the beginning-of-period beliefs of the next cohort by equation (4). However, without equation (3) following an action deviation from the equilibrium path the end-of-periods beliefs of some players $\langle i, t \rangle$ could be incorrect, and be passed on to the next cohort intact.

Now recall that the punch-line of the forward recursion argument we have outlined is that if the SE (g, μ, Φ) satisfies Inter-Generational Agreement then we know that after any actual history of play h^t , player $\langle i, t \rangle$ believes that his opponents have received messages $(m_1^t(h^t), \dots, m_{i-1}^t(h^t), m_{i+1}^t(h^t), \dots, m_n^t(h^t))$ with probability one. To see how we can construct an SPE of the standard repeated game that is equivalent to the given SE, consider the strategies g_i^{t*} for the standard repeated game defined as $g_i^{t*}(h^t, x^t) = g_i^t(m_i(h^t), x^t)$. Clearly, these strategies implement the same payoff vector that is obtained in the given SE of the dynastic repeated game. Now suppose that the strategy profile g^* we have just constructed is not an SPE of the standard repeated game. Then, by the one-shot deviation principle we know that some player i in the standard repeated game would have an incentive to deviate in a single period t after some history of play h^t . However, given the property of beliefs in the SE (g, μ, Φ) with Inter-Generational Agreement that we have shown above, this implies that player $\langle i, t \rangle$ would have an incentive to deviate at the action stage in the dynastic repeated game. This of course contradicts the fact that (g, μ, Φ) is an SE of the dynastic repeated game. Hence the argument is complete. The proof of Theorem 3 that appears in Appendix B of course does not rely on the two simplifying assumptions we made here. However, modulo some additional notation and technical issues, the argument presented there runs along the same lines as the sketch we have given here.

6. Relation to the Literature

We do not attempt here to review the vast literature on repeated games.¹⁷ Instead, we point out three key ingredients of our model which, taken together, set this paper apart from previous contributions. First, we break infinitely-lived players into sequences of finitely-lived ones, each of whom has dynastic preferences. Second, the individual entrants have no memory of past play, and the past history of play leaves no tangible trace — only messages are available. Third, the messages within a dynasty from one generation of individuals to the next are private.

Via the first ingredient, the results of this paper can be related to those of overlapping generations games. Examples of the latter include Cremer (1986), Kandori (1992a), Salant (1991) and Smith (1992). In these papers there is no dynastic component to the players' payoffs and full memory (i.e., perfect observation of the past) is assumed.¹⁸ Consequently, the Folk Theorems for OLG games relate only to payoff profiles above all players' minmax values.

The second and third ingredients, bring out the relationship with other recent papers that study equilibria in dynastic environments when the full memory assumption is relaxed. Examples include Anderlini and Lagunoff (2005), Kobayashi (2003), and Lagunoff and Matsui (2004). Among these, Anderlini and Lagunoff (2005) is the closest and, in many ways, the most direct predecessor of the current paper. Anderlini and Lagunoff (2005) examines the same dynastic model when each player $\langle i, t \rangle$ receives a *public messages* from the player $\langle j \in I, t - 1 \rangle$ about the previous history of play. If the public messages from all player in the previous cohort are *simultaneous*, then a *Folk Theorem* in the sense of Fudenberg and Maskin (1986) can be obtained. If there are three or more players, all individually rational feasible payoffs can be sustained as an SE. Intuitively, this is because a version of a well known “cross-checking” argument that goes back to Maskin (1999) can be applied in this case. By contrast, the present paper studies the model in which private communication (within each dynasty) may occur. We show that the difference between purely public and possible private communication is potentially large. Equilibria that sustain payoffs below some dynasty's minmax exist, but they require inter-generational “disagreement.”

Kobayashi (2003) and Lagunoff and Matsui (2004) examine OLG games with a dynastic payoff component. As in Anderlini and Lagunoff (2005), these models assume entrants have no prior memory, and they also allow for communication across generations. Though substantive differences exist between each of the models, they both prove standard (for OLG games) Folk Theorems. Interestingly, both Folk Theorems make use of *intra*-generational disagreement of beliefs in the equilibrium continuations following deviations. Nevertheless, the constructed

¹⁷Mailath and Samuelson (2006) is a comprehensive source that includes classic as well as more recent developments.

¹⁸As he points out, the results in Kandori (1992a) generalize to less than full memory of the past, although some direct memory is required. Bhaskar (1998) examines a related OLG model with no dynastic payoffs and very little, albeit some, direct memory by entrants. He shows that very limited memory is enough to sustain optimal transfers in a 2-period consumption-loan smoothing OLG game.

equilibria in both papers leave no room for *inter*-generational disagreement at the message stage.

The role of *public* messages has been studied in other repeated game contexts such as in games with private monitoring. Models of Ben-Porath and Kahneman (1996), Compte (1998), and Kandori and Matsushima (1998) all examine communication in repeated games when players receive private signals of others’ past behavior. As in Anderlini and Lagunoff (2005), these papers exploit cross-checking arguments to sustain the truthful revelation of one’s private signal in each stage of the repeated game.

Recent works by Schotter and Sopher (2001a), Schotter and Sopher (2001b), and Chaudhuri, Schotter, and Sopher (2001) examine the role of communication in an experimental dynastic environment. These papers report on laboratory experiments designed to mimic the dynastic game. The general conclusion seems to be that the presence of private communication has a significant (if puzzling) effect, even in the full memory game.

It is also worth noting the similarity between the present model and games with imperfect recall.¹⁹ Each dynastic player could be viewed as an infinitely lived player with imperfect recall (e.g., the “absent-minded driver” with “multiple selves” in Piccione and Rubinstein (1997)) who can write messages to his future self at the end of each period.

By contrast, the present model is distinguishable from dynamic models that create memory from a tangible “piece” of history. For instance, Anderlini and Lagunoff (2005) and Anderlini, Gerardi, and Lagunoff (2007) extend the analysis to the case where history may leave a “footprint,” i.e, hard evidence of the past history of play. In particular, Anderlini, Gerardi, and Lagunoff (2007) examines the role of social memory in a dynastic repeated game with two dynasties when a (noisy) signal of the past history of play is available to the players.

Incomplete but hard evidence of the past history of play is also present in Johnson, Levine, and Pesendorfer (2001) and Kandori (1992b). In another instance, memory may be created from a tangible but intrinsically worthless asset such as fiat money. A number of contributions in monetary theory (e.g., Kocherlakota (1998), Kocherlakota and Wallace (1998), Wallace (2001), and Corbae, Temzelides, and Wright (2001)) have all shown, to varying degrees, the substitutability of money for memory. In fact, the role of money in creating “memory” is clarified by Aliprantis, Camera, and Puzzello (2007b) and Aliprantis, Camera, and Puzzello (2007a) who show that money is not needed when trade is periodically centralized since deviations can be deterred quickly without money. With sufficient lack of observability and decentralized trade, tangible assets such as money may again become useful.

7. Concluding Remarks

We posit a dynastic repeated game populated by one-period-lived individuals who rely on private messages from their predecessors to fathom the past. The set of equilibrium payoffs expands dramatically relative to the corresponding standard repeated game. Under extremely

¹⁹See the Special Issue of Games and Economic Behavior (1997) on Games with Imperfect Recall for extensive references.

mild conditions, as the dynastic players care more and more about the payoffs of their successors, *all* interior payoff vectors that are *feasible* in the stage game are sustainable in an SE of the dynastic repeated game.

We are able to characterize entirely, via a property of the players' beliefs, when an SE of the dynastic repeated game can yield a payoff vector not sustainable as an SPE of the standard repeated game: the SE in question must display a failure of what we have termed Inter-Generational Agreement.

The consistency condition of SE provides a coherent societal outlook rooted in the likelihood of mistakes and how such errors affect off-path beliefs. We find this discipline particularly compelling in the dynastic context. For one thing, it provides an intuitive model of how real world individuals might interpret off-path events. It also provides a natural language for distinguishing between what people say (i.e., message errors) and what they do (action errors). This paper shows how this dichotomy may give rise to systematically incorrect beliefs after certain off-path events. In turn, these beliefs sustain SE payoff profiles below a player's minmax.

We certainly recognize that there are other constructs that theorists are used to. The obvious one is that of "neologism-proofness" (Farrell, 1993, Mattehws, Okuno-Fujiwara, and Postlewaite, 1991, among others). As we mentioned earlier, at least in its current form, neologism-proofness does not apply to our framework. The reason is as follows. Roughly speaking, neologism-proofness builds into the solution concept the idea that in a sender-receiver game, provided the appropriate incentive-compatibility constraints are satisfied, a player's *exogenous type* (in the standard sense of a "payoff type") will be able to create a "neologism" (use an hitherto unused message) to distinguish himself from other types. The point is that in our dynastic game there are *no exogenous types* for any of the players. It would therefore be impossible to satisfy any standard form of incentive-compatibility constraints. The different "types" of each player in our dynastic repeated game are only distinguished by their beliefs, which in turn are determined by equilibrium strategies together with a complete theory of mistakes as in any SE. To see that the logic of neologism-proofness can be conceptually troublesome in our context, consider for instance the construction we use to prove Theorem 1 above. Suppose that some player $\langle i, t \rangle$ deviates so as to trigger the punishment- i phase. At the end of period t player $\langle i, t \rangle$ may want to communicate to player $\langle i, t + 1 \rangle$ that play is in the punishment- i phase so that he can play a best response to the actions of others in period $t + 1$. For a "neologism" to work at this point player $\langle i, t + 1 \rangle$ would have to believe it. He would have to believe what player $\langle i, t \rangle$ is saying: I have made a mistake, therefore respond appropriately to the punishment that follows (there are no exogenous types to which $\langle i, t \rangle$ can appeal in his "speech"). However, as always in a dynastic game, whether $\langle i, t + 1 \rangle$ believes or not what he is told, depends on the relative likelihood that he assigns to mistaken actions and mistaken messages; both are possible after all.²⁰ So, for the neologism to work it would have to be "trusted" more as a message than the

²⁰By contrast, in a sender-receiver game the sender communicates to the receiver his *exogenous type*, which

one prescribed in equilibrium. But, since there are no exogenous types to which to appeal, there do not seem to be compelling reasons for this to be the case. As with other possible refinements, our characterization of the new equilibria that appear in the dynastic repeated game in terms of Inter-Generational Agreement also seals the question of what bite a possible adaptation of neologism-proofness could have at a more general level. The new equilibria of the dynastic repeated game can be ruled out (without ruling out any of the traditional ones) if and only if beliefs that violate Inter-Generational Agreement can be ruled out. Whether this is the case or not is largely a matter of intuitive appeal.

While our results apply only to the actual formal model we have set forth, it is natural to ask which ones are essential and which ones are not. We have several remarks to make.

As we noted already, the absence of public messages alongside private ones is completely inessential to what we do here. Public messages could be added to our model without altering our results. It is always possible to replicate the SE of this model in another model with public messages as well; the public messages would be ignored by the players’ equilibrium strategies and beliefs.

We make explicit use of public randomization devices both at the action and at the message stage of the dynastic repeated game. While the use of two separate devices is not essential for our results (see footnote 5 above), whether the use of some public randomization device is necessary is a question for future research.²¹ In our constructions it is essential that the players should be able to correlate the messages they send to the next cohort. Without this it is hard to see how play could switch between the different phases on which our constructions depend.

We have stipulated a very specific set of “demographics” for our dynastic repeated game: all players live one period and are replaced by their successor at the end of their lives. Although the demographic structure of the model greatly simplifies the analysis, our results readily generalize in the following sense. Suppose that each dynasty is composed of a sequence of finitely lived individuals, each of whom can live any finite number of periods provided that there exists a uniform upper bound L on the length of each lifetime. Note that any overlap across dynasties that conforms to the L -boundedness assumption is possible. Using a technique known at least since Ellison (1994), one can show that our results extend to any L -bounded demographics. The idea involves constructing L interleaved “copies” of the equilibrium of the one-period lived demographics case.

To give an idea of how this works, consider the strategies and beliefs of individuals alive in periods $0, L, 2L, 3L$ and so on. These individuals “match” the strategies of the individuals alive in periods $0, 1, 2, 3$ and so on in the model with full replacement every period. Matching here means that when deciding how to play, the individuals alive at L will only consider

is chosen by Nature, and not by the sender himself.

²¹We have examples showing that even without any randomization devices it is possible to display SE that push one or more dynasties below their minmax in the stage game. Whether a “Super” Folk Theorem is available in this case is an open question at this point.

information concerning period 0, individuals alive at $2L$ will only consider information concerning periods 0 and L and so on, forward without bound. The same construction is used to match the strategies and beliefs of individuals alive in periods $1, L + 1, 2L + 1, 3L + 1$ and so on with those of the individuals alive in periods $0, 1, 2, 3$ and so on in the model with full replacement every period.

In this paper, we examine dynastic repeated games with 3 or more dynasties. It is not hard to construct examples of dynastic repeated games with two dynasties that admit SE in which the players' payoffs are below their minmax in the stage game. Thus, it seems that there is no definite need to have more than two dynasties to generate SE payoffs in the dynastic repeated game that are not sustainable as SPE of the corresponding standard repeated game. Whether and under what conditions a Folk Theorem like the one presented here is available for the case of two dynasties is an open question. We leave the characterization of the equilibrium set in this case for future work.

Lastly, our Folk Theorems for the dynastic repeated game show that, as δ approaches one, the set of SE payoffs includes “worse” vectors that push some (or even all) players below their minmax payoffs in the stage game. We do not have a full characterization of the SE payoffs for the dynastic repeated game when δ is bounded away from one. However, it is possible to construct examples showing that the set of SE payoffs includes vectors that Pareto-dominate those on the Pareto-frontier of the set of SPE payoffs in the standard repeated game when δ is bounded away from one. Intuitively, this is because some “bad” payoff vectors (pushing some players below the minmax) are sustainable in an SE when δ is bounded away from one. Thus, “harsher” punishments are available as continuation payoffs in the dynastic repeated game than in the standard repeated game. Using these punishments, higher payoffs are sustainable in equilibrium in the dynastic repeated game.

Appendix A: Proof of Theorem 1

A.1. A Stronger Statement

As we anticipated in the text, we show that Theorem 1 holds by proving a stronger statement, which readily implies it.

We begin with a definition. As we mentioned in footnote 11, this re-defines the benchmark payoffs as the minmax payoffs in the component game, but allowing for correlation among the minmaxing players while not allowing the player who is being minmaxed to take the correlation into account.

Definition A.1. *Restricted Correlated Minmax:* Consider a stage game $G = (A, u, I)$. Let a product set $\tilde{A} \subseteq A$ be given. Now let

$$\underline{\omega}_i(\tilde{A}) = \min_{z_{-i} \in \Delta(\tilde{A}_{-i})} \max_{a_i \in \tilde{A}_i} \sum_{a_{-i} \in \tilde{A}_{-i}} z_{-i}(a_{-i}) u_i(a_i, a_{-i}) \quad (\text{A.1})$$

where z_{-i} is any probability distribution over the finite set \tilde{A}_{-i} (not necessarily the product of independent marginals), and $z_{-i}(a_{-i})$ denotes the probability that z_{-i} assigns to the profile a_{-i} .

We then say that $\underline{\omega}_i(\tilde{A})$ is the restricted (to \tilde{A}) correlated minmax for i in G .

The statement that we will actually prove can now be made precise.

Theorem A.1. *Stronger Dynastic Folk Theorem:* Let a stage game $G = (A, u, I)$ with four or more players be given. Assume that G is such that we can find a product set $\tilde{A} \subseteq A$ and an array of $n + 1$ payoffs vectors $\hat{v}, \bar{v}^1, \dots, \bar{v}^n$ for which the following conditions hold.

- (i) For every $i \in I$, the set \tilde{A}_i contains at least two elements.
- (ii) $\hat{v} \in \text{int}V(\tilde{A})$, and $\bar{v}^i \in V(\tilde{A})$ for every $i \in I$.
- (iii) $\underline{\omega}_i(\tilde{A}) < \bar{v}_i^i < \bar{v}_i^j$ and $\bar{v}_i^i < \hat{v}_i$ for every $i \in I$ and every $j \neq i$.

Then for every $v \in \text{int}V$ there exists a $\underline{\delta} \in (0, 1)$ such that $\delta > \underline{\delta}$ implies v is sustained by SE for discount factor δ .

A.2. A Roadmap of the Proof

Our proof is constructive. It runs along the following lines. Given a $v^* \in \text{int}(V)$, we construct a randomization device \tilde{x} , a randomization device \tilde{y} , and an assessment (g, μ, Φ) , which implements the vector of payoffs v^* , and which for δ sufficiently large constitutes an SE of the dynastic repeated game. All the elements of our construction are defined independently of δ . The sequential rationality of the strategy profile given the postulated beliefs holds when δ is sufficiently close to one.

The basic logic follows along the familiar lines of other Folk Theorems. Namely, the equilibrium can be described in terms of phases and transitions. The phases, described informally in the main body of the paper, are given by the list

$$\{In, S, D^1, \dots, D^n, P^1, \dots, P^n, T^1, \dots, T^n, Z\}.$$

where In is the *initial phase*, S is the *standard equilibrium phase*, D^i , $i = 1, \dots, n$ is *diversionary- i phase*, P^i is i 's *punishment phase*, T^i is i 's *terminal phase*, and finally Z will be referred to as the *global phase*. The action and message strategies and randomization devices jointly facilitate transitions from one phase to another.

In what follows, $v^* \in \text{int}(V)$ is the vector of payoffs to be sustained as an SE as in the statement of Theorem A.1. Throughout the argument, \tilde{A} is a product set and \hat{v} and \bar{v}^1 through \bar{v}^n are vectors of payoffs as in the statement of Theorem A.1. Of course, these are fixed throughout the proof.

In Section A.4, we define formally the players' message spaces M_i^t , the randomization devices \tilde{x} and \tilde{y} , the strategy profile (g, μ) , and the players' beliefs. Throughout the argument, we assume that this message space for any player $\langle i, t - 1 \rangle$ consists of a set smaller than the set H^t . To work with “restricted” message spaces is sufficient to prove our claim because of Lemma T.3.1.

In Section A.5 we define the completely mixed strategies that generate the SE beliefs Φ . We also briefly outline the consistency and sequential rationality argument.

The rest of the details are relegated to the Technical Addendum. In Section T.4 of the Addendum we define formally the system of beliefs Φ . In Section T.5 we check that the assessment (g, μ, Φ) satisfies sequential rationality when δ is close to one. Finally in Section T.6 we verify the consistency of the equilibrium beliefs.

A.3. Messages and Randomization Devices

Before proceeding with a description of strategies, message spaces, beliefs, and devices, we first define the correlation probabilities and corresponding payoffs that each player will receive in each phases.

Definition A.2: Let $(a(1), \dots, a(\ell), \dots, a(\|A\|))$ be a list of all possible outcomes in G . Without loss of generality, assume that the first $\|\tilde{A}\| \leq \|A\|$ elements in this enumeration are the strategy profiles in the product set \tilde{A} . This enumeration will be taken as fixed throughout the rest of the argument.

To begin, we construct the correlation probabilities and the payoff profile that result in dynasty i 's punishment phase, P^i , for each i .

Remark A.1: From Definition A.1, for each $i \in I$ we can find a set of profiles of actions $\tilde{A}^i \subset \tilde{A}$ corresponding to a set of indices $(i_1, \dots, i_\ell, \dots, i_{\|\tilde{A}^i\|})$ in the enumeration of Definition A.2, and a set of positive weights $\{\underline{p}^i(a(i_\ell))\}_{\ell=1}^{\|\tilde{A}^i\|}$ adding up to one and such that

$$\underline{\omega}_i(\tilde{A}) = \max_{a_i \in \tilde{A}^i} \sum_{\ell=1}^{\|\tilde{A}^i\|} \underline{p}^i(a(i_\ell)) u_i(a_i, a_{-i}(i_\ell)) = \sum_{\ell=1}^{\|\tilde{A}^i\|} \underline{p}^i(a(i_\ell)) u_i(a(i_\ell)) \quad (\text{A.2})$$

Without loss of generality, we can take it to be the case that for every i_ℓ , $a_i(i_\ell)$ is the same action in \tilde{A}^i . We denote this by a_i^i so that $a_i(i_\ell) = a_i^i$ for $\ell = 1, \dots, \|\tilde{A}^i\|$.

For convenience, since \tilde{A} is fixed throughout the argument, in what follows we will use the following notation for the payoffs of each i corresponding to the weights $\{\underline{p}^j(a(j_\ell))\}_{\ell=1}^{\|\tilde{A}^j\|}$.

$$\underline{\omega}_i^j = \sum_{\ell=1}^{\|\tilde{A}^j\|} \underline{p}^j(a(j_\ell)) u_i(a(j_\ell)) \quad (\text{A.3})$$

The payoff $\underline{\omega}_i^j$ is what i receives in phase P^j . Of course, we have that $\underline{\omega}_i^i = \underline{\omega}_i(\tilde{A})$.

Next, we define correlation probabilities and payoff profiles resulting from the diversionary- i phases, D^i , for each i .

Definition A.3: Let \tilde{A}^i be as in Remark A.1. For each $i \in I$ and for each element $a(i_\ell)$ of \tilde{A}^i , construct a new action profile $\check{a}^i(i_\ell)$ as follows. For all $j \neq i$, set $\check{a}_j^i(i_\ell)$ to satisfy $\check{a}_j^i(i_\ell) \neq a_j(i_\ell)$ and $\check{a}_j^i(i_\ell) \in \tilde{A}_j$. Notice that this is always possible since, by assumption, \tilde{A}_j contains at least two elements for every $j \in I$. Finally, set $\check{a}_i^i(i_\ell) = a_i(i_\ell) = a_i^i$.

In what follows, for every i and j in I we will let

$$\check{\omega}_i^j = \sum_{\ell=1}^{\|\tilde{A}^j\|} \underline{p}^j(a(j_\ell)) u_i(\check{a}^j(j_\ell)) \quad (\text{A.4})$$

Payoff $\check{\omega}_i^j$ is what i receives in diversionary- j phase, D^j .

The correlation probabilities and payoffs resulting from terminal phases, T^i for each i , are as follows.

Remark A.2: Since each of the payoff vectors \bar{v}^j must only satisfy strong inequalities (see (iii) of the statement of the theorem), without loss of generality we can take it to be the case that $\bar{v}^j \in \text{int}(V(\tilde{A}))$, for each $j \in I$. It then follows that for every $j \in I$ we can find a set of positive weights $\{\bar{p}^j(a(\ell))\}_{\ell=1}^{\|\tilde{A}\|}$ adding up to one and such that for every $i \in I$

$$\bar{v}_i^j = \sum_{\ell=1}^{\|\tilde{A}\|} \bar{p}^j(a(\ell)) u_i(a(\ell)) \quad (\text{A.5})$$

Payoff \bar{v}_i^j is what i receives in terminal phase T^j .

Remark A.3: Since the payoff vector \hat{v} is in $\text{int}(V(\tilde{A}))$, we can find an $\eta \in (0, 1)$ and a set of positive weights $\{\hat{p}(a(\ell))\}_{\ell=1}^{\|\tilde{A}\|}$ adding up to one and such that for every $i \in I$

$$\hat{v}_i = (1 - \eta) \sum_{\ell=1}^{\|\tilde{A}\|} \hat{p}(a(\ell)) u_i(a(\ell)) + \frac{\eta}{n} \sum_{j=1}^n \check{u}_i^j \quad (\text{A.6})$$

From the construction in Remark A.3, the correlation probabilities and payoffs in the standard phase, S , are now defined.

Definition A.4: Let

$$\hat{v}_i = \sum_{\ell=1}^{\|\tilde{A}\|} \hat{p}(a(\ell)) u_i(a(\ell)) \quad (\text{A.7})$$

where the weights $\{\hat{p}(a(\ell))\}_{\ell=1}^{\|\tilde{A}\|}$ are as in Remark A.3. The payoff \hat{v}_i is what i receives in the standard phase S .

The correlation probabilities and payoffs corresponding to the global phase Z is defined as follows.

Remark A.4: Since the payoff vector v^* is in $\text{int}(V)$, we can find a $q \in (0, 1)$ and a set of positive weights $\{p^*(a(\ell))\}_{\ell=1}^{\|A\|}$ adding up to one and such that for every $i \in I$

$$v_i^* = (1 - q) \sum_{\ell=1}^{\|A\|} p^*(a(\ell)) u_i(a(\ell)) + q \hat{v}_i \quad (\text{A.8})$$

Of course, v^* is the payoff profile we wish to implement with a sequential equilibrium. The part of v^* that defines the payoff in the global phase Z is given by

$$z_i = \sum_{\ell=1}^{\|A\|} p^*(a(\ell)) u_i(a(\ell)). \quad (\text{A.9})$$

Finally, in the initial phase In , we have:

Remark A.5: Recall that by assumption the payoff vector v^* is in $\text{int}(V)$. Therefore, we can find a set of positive weights $\{p^0(a(\ell))\}_{\ell=1}^{\|A\|}$ adding up to one and such that for every $i \in I$

$$v_i^* = \sum_{\ell=1}^{\|A\|} p^0(a(\ell)) u_i(a(\ell)) \quad (\text{A.10})$$

Hence, the payoff in the initial phase In is the profile v^* that we wish to implement.

Definition A.5: Throughout the rest of the argument, we let T be an integer sufficiently large so as to guarantee that the following inequality is satisfied for all $i \in I$.

$$T(\bar{v}_i^i - \underline{\omega}_i^i) > \bar{u}_i - \underline{u}_i \quad (\text{A.11})$$

We now proceed to define message spaces and action randomization devices.

Definition A.6. *The Message Spaces:* As with payoffs and correlation probabilities, we can associate certain parts of the message space to each of four types of phases: S , P^i , D^i and T^i (of course, we must later construct equilibrium beliefs that correctly assign these messages to the various phases). There are no explicit messages indicating the initial phase In which is played only in period 0. There are also no explicit messages that indicate the global phase Z since the latter is exclusively signalled by the randomization devices. In particular, the global phase is reached with probability $(1 - q)$ each period, independently of the message.

First, let m^* denote the solitary message indicating the S phase.

Next, for each $j \in I$ and each $t = 1, \dots, T - 1$ let

$$\underline{M}(j, t) = \{\underline{m}^{j, T-t+1}, \underline{m}^{j, T-t+2}, \dots, \underline{m}^{j, T}\} \quad (\text{A.12})$$

and for every $t \geq T$ let

$$\underline{M}(j, t) = \underline{M}(j, T) = \{\underline{m}^{j, 1}, \dots, \underline{m}^{j, T}\} \quad (\text{A.13})$$

The set $\underline{M}(j, t)$ is the set of messages signalling to players $i \neq j$ the various stages of the P^j (punishment) phase. It is also the subset of messages received by player j indicating his own diversionary- j phase.

Now define $\overline{M} = \{\overline{m}^1, \dots, \overline{m}^n\}$. These are messages indicating to all players the terminal T^i phase for all i .

Next, define for each i the set $\check{M}_{-i} = \{\check{m}^1, \dots, \check{m}^{i-1}, \check{m}^{i+1}, \check{m}^n\}$. These are the messages indicating to player i the various diversionary- j phases, for $j \neq i$.

Putting all these sets together, recall that M_i^t denotes the set of messages that a player $\langle i, t-1 \rangle$ can send to player $\langle i, t \rangle$. For any $t = 1, \dots, T$ let

$$M_i^t = \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(1, t) \cup \dots \cup \underline{M}(n, t) \quad (\text{A.14})$$

For any $t \geq T + 1$ let

$$M_i^t = \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(1, t) \cup \dots \cup \underline{M}(n, t) \cup \overline{M} \quad (\text{A.15})$$

Definition A.7. *Action-stage Randomization Device:* Earlier we defined the action-correlation probabilities corresponding to each of the phases. Summarizing that information, we have:

- p^0 - initial phase In
- \hat{p} - standard phase S
- \underline{p}^i - diversionary- i phase D^i and punishment phase P^i
- \overline{p}^i - terminal phase T^i
- p^* - global phase Z

The action stage randomization device \tilde{x} therefore serves to indicate which of the above correlation probability distributions to use by the players. It is defined as follows.

The set X consists of $\|A\| \|\tilde{A}\|^{n+1} \prod_{i \in I} \|\tilde{A}^i\| + \|A\|^2$ elements. Let $(x(1), \dots, x(\kappa), \dots, x(\|X\|))$ be an enumeration of the elements of X , and let $\bar{\kappa} = \|A\| \|\tilde{A}\|^{n+1} \prod_{i \in I} \|\tilde{A}^i\|$.

Informally, the index $\bar{\kappa}$ separates the components of the action correlation device that define action play in the local phases, S , D^i , P^i and T^i from the global phase Z .

Formally, each of the first $\bar{\kappa}$ elements of X can be identified by a string of $1 + (n + 1) + n = 2n + 2$ indices as follows. With a slight abuse of notation, for $\kappa \leq \bar{\kappa}$, we will write $x(\kappa) = x(\ell_0, \hat{\ell}, \ell_1, \dots, \ell_n, 1_\ell, \dots, i_\ell, \dots, n_\ell)$ with ℓ_0 running from 1 to $\|A\|$, $\hat{\ell}$ and each of the indices ℓ_1 through ℓ_n running from 1 to $\|\tilde{A}\|$, and each of the n indices i_ℓ , each with ℓ running from 1 to $\|\tilde{A}^i\|$. Obviously, the last $\|X\| - \bar{\kappa}$ elements of X can

be identified by a pair of indices ℓ_{00} and ℓ^* both running from 1 to $\|A\|$. In this case, with a slight abuse of notation again, we will write $x(\kappa) = x(\ell_{00}, \ell^*)$.

To summarize, these indices are used to denote the action profiles in the support of the various phase-specific probability distributions. Hence, $\hat{\ell}$ indexes actions in the support of \hat{p} in the standard phase S , ℓ^* indexes actions in the global phase Z , ℓ_i is used for the T^i phase, and i_ℓ for both diversionary- i and punishment- i (resp., D^i and P^i) phases. Finally, ℓ_0 and ℓ_{00} are both used for the initial phase In .²²

We are now ready to aggregate the information over all phases in order to define the probability distribution governing the realization of \tilde{x} . For $\kappa \leq \bar{\kappa}$ let

$$\Pr(\tilde{x} = x(\ell_0, \hat{\ell}, \ell_1, \dots, \ell_n, 1_\ell, \dots, i_\ell, \dots, n_\ell)) = q \left[p^0(a(\ell_0)) \hat{p}(a(\hat{\ell})) \bar{p}^1(a(\ell_1)) \cdots \bar{p}^n(a(\ell_n)) \underline{p}^1(a(1_\ell)) \cdots \underline{p}^i(a(i_\ell)) \cdots \underline{p}^n(a(n_\ell)) \right] \quad (\text{A.16})$$

and for $\kappa = \bar{\kappa} + 1, \dots, \|X\|$ let

$$\Pr(\tilde{x} = x(\ell_{00}, \ell^*)) = (1 - q) \left[p^0(a(\ell_{00})) p^*(a(\ell^*)) \right] \quad (\text{A.17})$$

Definition A.8. Message-Stage Randomization Device: The message stage randomization device \tilde{y} is far simpler. It is defined as follows. The set Y consists of $n + 1$ elements, which we denote $(y(0), y(1), \dots, y(n))$. The random variable \tilde{y} takes value $y(0)$ with probability $1 - \eta$, and each of the other possible values with probability η/n (where η is defined in Remark A.3). The value $y(0)$ is the realization associated with the standard phase S and, for every i , $y(i)$ is the realization associated with the diversionary- i phase.

A.4. Strategies and Beliefs

In this Section, we describe the equilibrium action strategies, message strategies, and beliefs of an individual in each phase.

Henceforth, fix a dynasty i and a date t . We proceed to describe strategies and beliefs of individual $\langle i, t \rangle$ in each phase of the equilibrium. Also, we let k be any dynasty in I , and j be an element of I such that $j \neq i$. We will use the mnemonic BB to describe the *beginning of period beliefs* of individual $\langle i, t \rangle$, and EB to describe *end of period beliefs* of $\langle i, t \rangle$. Finally, we use the phrase “ i believes” to mean “puts probability one on the event...”

Definition A.9. The Initial Phase: The initial phase, In , only occurs in $t = 0$. Player $\langle i, 0 \rangle$ inherits null message \emptyset .

BB (beginning of period beliefs). His beliefs are trivial: $\langle i, 0 \rangle$ believes (with probability one) that all others received \emptyset .

Action strategy.

$$g_i^0(m_i^0, x^0) = \begin{cases} a_i(\ell_0) & \text{if } x^0 = x(\ell_0, \dots) \\ a_i(\ell_{00}) & \text{if } x^0 = x(\ell_{00}, \dots) \end{cases} \quad (\text{A.18})$$

²²Even though both indices ℓ_0 and ℓ_{00} are used to index the initial phase, the probability distribution does not vary across the local and global parts of x . The difference in ℓ_0 and ℓ_{00} is merely a notational convenience that allows us to group all the phase-specific randomization devices into one universal action randomization device.

Message strategy. Let $g^0(m^0, x^0) = (g_1^0(m_1^0, x^0), \dots, g_n^0(m_n^0, x^0))$, and define $g_{-k}^0(m^0, x^0)$ in the obvious way. Then

$$\mu_i^0(m_i^0, x^0, a^0, y^0) = \begin{cases} \check{m}^j & \text{if } a^0 = g^0(m^0, x^0) \quad \text{and } y^0 = y(j) \\ \underline{m}^{i,T} & \text{if } a^0 = g^0(m^0, x^0) \quad \text{and } y^0 = y(i) \\ \underline{m}^{k,T} & \text{if } a_{-k}^0 = g_{-k}^0(m^0, x^0) \quad \text{and } a_k^0 \neq g_k^0(m^0, x^0) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.19})$$

EB (end of period beliefs). If the realized profile is $g_i^0(m_i^0, x^0)$ and $y^0 = y(k)$, then play proceeds to phase D^k and $\langle i, 0 \rangle$ believes that every $\langle j, 0 \rangle$ with $j \neq k$ sends \check{m}^k and $\langle k, 0 \rangle$ sends \underline{m}^k (of course $\langle i, 0 \rangle$'s belief about $\langle k, 0 \rangle$ means effectively that $k \neq i$. This parenthetical remark applies to all the phases). In all other cases, $\langle i, 0 \rangle$ believes that others send the same message that he sends.²³

We now describe the equilibrium in the remaining phases. In each phase, at the beginning of each period, the players observe the message sent by their predecessors and the action correlation device $x^t = x(\kappa)$. Recall from Definition A.7 that if $\kappa \leq \bar{\kappa}$, then the game is in one of the local phases, S, D^k, P^k or T^k . If $\kappa > \bar{\kappa}$, then the game is in the global phase Z .

Definition A.10. The Standard phase: In the standard phase, S , player $\langle i, t \rangle$ inherits the message m^* .

BB. Player $\langle i, t \rangle$ believes that all others received m^* .

Action strategy.

$$g_i^t(m^*, x^t) = a_i(\hat{\ell}) \text{ whenever } x^t = x(\cdot, \hat{\ell}, \dots) \quad (\text{A.20})$$

Message strategy. Let²⁴

$$\mu_i^t(m^*, x^t, a^t, y^t) = \begin{cases} \check{m}^j & \text{if } x^t = x(\cdot, \hat{\ell}, \dots), a^t = a(\hat{\ell}) \quad \text{and } y^t = y(j) \\ \nu(\underline{M}(i, t+1)) & \text{if } x^t = x(\cdot, \hat{\ell}, \dots), a^t = a(\hat{\ell}) \quad \text{and } y^t = y(i) \\ \underline{m}^{k,T} & \text{if } x^t = x(\cdot, \hat{\ell}, \dots), a_{-k}^t = a_{-k}(\hat{\ell}) \quad \text{and } a_k^t \neq a_k(\hat{\ell}) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.21})$$

EB. If realized profile is $g^t(m^*, x^t)$ and $y^0 = y(k)$, then play proceeds to phase D^k and $\langle i, t \rangle$ believes that all individuals in all dynasties except k sends \check{m}^k while $\langle k, t \rangle$ randomizes uniformly over all messages in $\underline{M}(k, t+1)$. In all other cases, $\langle i, t \rangle$ believes that others send the same message that he sends.

Definition A.11. The Diversionary- j phase ($j \neq i$): In the Diversionary- j phase, D^j , with $j \neq i$, player $\langle i, t \rangle$ inherits the message \check{m}^j .

BB. Player $\langle i, t \rangle$ believes that the individual in dynasty j received a randomized message with support in $\underline{M}(j, t)$, and that all others received \check{m}^j .

Action strategy.

$$g_i^t(\check{m}^j, x^t) = \check{a}_i^j(j_\ell) \text{ whenever } x^t = x(\dots, j_\ell, \dots) \quad (\text{A.22})$$

²³Our description of beliefs, while complete, leaves out some of the formal notation which we fully provide in the external Technical Addendum.

²⁴Throughout the paper we adopt the following notational convention. Given any finite set, we denote by $\nu(\cdot)$ the uniform probability distribution over the set. So, if B is a finite set, $\nu(B)$ assigns probability $1/\|B\|$ to every element of B .

Message strategy. In the expression below, let j' be any element of I not equal to i . Then

$$\mu_i^t(\check{m}^j, x^t, a^t, y^t) = \begin{cases} \check{m}^{j'} & \text{if } x^t = x(\dots, j_\ell, \dots), a^t = \check{a}^j(j_\ell) & \text{and } y^t = y(j') \\ \nu(\underline{M}(i, t+1)) & \text{if } x^t = x(\dots, j_\ell, \dots), a^t = \check{a}^j(j_\ell) & \text{and } y^t = y(i) \\ \underline{m}^{k, T} & \text{if } x^t = x(\dots, j_\ell, \dots), a^t_{-k} = \check{a}^j_{-k}(j_\ell) & \text{and } a^t_k \neq \check{a}^j_k(j_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.23})$$

EB. If the realized profile is $(g^t_{-j}(\check{m}^j, x^t), a^j_i)$ and $y^0 = y(k)$, then play proceeds to phase D^k and $\langle i, t \rangle$ believes all individuals from dynasties $j' \neq k$ sends \check{m}^k while the individual in dynasty k randomizes uniformly over all messages in $\underline{M}(k, t+1)$. In all other cases, $\langle i, t \rangle$ believes others send the same message that he sends.

Definition A.12. The Diversionary- i AND Punishment- i phases: Significantly, these phases D^i and P^i are lumped together because $\langle i, t \rangle$ does not distinguish them in either his strategy or his beginning of period beliefs. In these phases, player $\langle i, t \rangle$ inherits a message $\underline{m}^{i, \tau} \in \underline{M}(i, t)$.

BB. Player $\langle i, t \rangle$ believes that all other dynasties received a \check{m}^i .

Action strategy.

$$g_i^t(\underline{m}^{i, \tau}, x^t) = a_i(i_\ell) = a_i^i \text{ for all } x^t \quad (\text{A.24})$$

Message strategy. Let

$$\mu_i^t(\underline{m}^{i, \tau}, x^t, a^t, y^t) = \begin{cases} \check{m}^j & \text{if } x^t = x(\dots, i_\ell, \dots), a^t = \check{a}^i(i_\ell) & \text{and } y^t = y(j) \\ \nu(\underline{M}(i, t+1)) & \text{if } x^t = x(\dots, i_\ell, \dots), a^t = \check{a}^i(i_\ell) & \text{and } y^t = y(i) \\ \underline{m}^{k, T} & \text{if } x^t = x(\dots, i_\ell, \dots), a^t_{-k} = \check{a}^i_{-k}(i_\ell) & \text{and } a^t_k \neq \check{a}^i_k(i_\ell) \\ \underline{m}^{k, T} & \text{if } x^t = x(\dots, i_\ell, \dots), a^t_{-k} = a^i_{-k}(i_\ell) & \text{and } a^t_k \neq a^i_k(i_\ell) \\ \underline{m}^{i, \tau-1} & \text{if } x^t = x(\dots, i_\ell, \dots) & \text{and } a^t = a(i_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.25})$$

where we set $\underline{m}^{i, 0} = \bar{m}^i$. Notice that player $\langle i, t \rangle$ may need to distinguish between the third and fourth cases of (A.25) since clearly they may be generated by different values of the index $k \in I$. To verify that this distinction is always feasible, recall that, by construction (see Definition A.3), $\check{a}_{-i}(i_\ell)$ differs from $a_{-i}(i_\ell)$ in every component, and that of course $n \geq 4$.

EB. If the realized profile is $(g^t_{-i}(\check{m}^i, x^t), a^i_i)$ and $y^0 = y(k)$, then play proceeds to phase D^k and $\langle i, t \rangle$ believes that individuals in dynasties $j \neq k$ send \check{m}^k while the individual in dynasty k randomizes uniformly over all messages in $\underline{M}(k, t+1)$. In all other cases, $\langle i, t \rangle$ believes that everyone sends the same message that he sends. In particular, notice that if the realized profile is $(g^t_{-i}(\underline{m}^{i, \tau}, x^t), a^i_i)$ then $\langle i, t \rangle$ believes that others will send either $\underline{m}^{i, \tau-1}$ if $\tau > 1$ or \bar{m}^i if $\tau = 1$. In other words, he discovers (after the fact) in this case that the correct phase was P^i rather than D^i .

Definition A.13. The Punishment- j phase, for all $j \neq i$: In the punishment- j phase, P^j , player $\langle i, t \rangle$ inherits the message $\underline{m}^{j, \tau} \in \underline{M}(j, t)$.

BB. Player $\langle i, t \rangle$ believes that all others received $\underline{m}^{j, \tau}$.

Action strategy.

$$g_i^t(\underline{m}^{j, \tau}, x^t) = a_i(j_\ell) \text{ whenever } x^t = x(\cdot, j_\ell, \dots) \quad (\text{A.26})$$

Message strategy. Let

$$\mu_i^t(\underline{m}^{j,\tau}, x^t, a^t, y^t) = \begin{cases} \underline{m}^{j,\tau-1} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot) & \text{and } a^t = a(j_\ell) \\ \underline{m}^{k,T} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a_{-k}^t = a_{-k}(j_\ell) & \text{and } a_k^t \neq a_k(j_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.27})$$

where we set $\underline{m}^{j,0} = \bar{m}^j$.

EB. Individual $\langle i, t \rangle$ believes that all send the same message that he sends.

Definition A.14. The Terminal- k' phase, for all $k' \in I$: In the terminal- k' phase, $T^{k'}$, player $\langle i, t \rangle$ inherits the message $\bar{m}^{k'}$.

BB. Player $\langle i, t \rangle$ believes that all others received $\bar{m}^{k'}$.

Action strategy.

$$g_i^t(\bar{m}^{k'}, x^t) = a_i(\ell_{k'}) \text{ whenever } x^t = x(\cdot, \ell_{k'}, \cdot \cdot \cdot) \quad (\text{A.28})$$

Message strategy. Let

$$\mu_i^t(\bar{m}^{k'}, x^t, a^t, y^t) = \begin{cases} \bar{m}^{k'} & \text{if } x^t = x(\cdot \cdot \cdot, \ell_{k'}, \cdot \cdot \cdot) & \text{and } a^t = a(\ell_{k'}) \\ \underline{m}^{k,T} & \text{if } x^t = x(\cdot \cdot \cdot, \ell_{k'}, \cdot \cdot \cdot), a_{-k}^t = a_{-k}(\ell_{k'}) & \text{and } a_k^t \neq a_k(\ell_{k'}) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.29})$$

EB. Once again, $\langle i, t \rangle$ believes that all send the same message that he sends.

Definition A.15. The Global phase: In the global phase, Z , player $\langle i, t \rangle$ observes that $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$, and he inherits whatever message was sent to him by his predecessor.

BB. Player $\langle i, t \rangle$'s belief depends on what belief he would have had if the phase had been local instead of global. This, in turn, depends on the message. So if he receives, for example, m^* , he inherits the beginning of period beliefs he would have had in phase S .

Action strategies. For all m_i^t , whenever $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$,

$$g_i^t(m_i^t, x^t) = a_i(\ell^*) \text{ whenever } x^t = x(\cdot, \ell^*) \quad (\text{A.30})$$

Message strategy. Let

$$\mu_i^t(m_i^t, x^t, a^t, y^t) = \begin{cases} m_i^t & \text{if } x^t = x(\cdot, \ell^*) & \text{and } a^t = a(\ell^*) \\ \underline{m}^{k,T} & \text{if } x^t = x(\cdot, \ell^*), a_{-k}^t = a_{-k}(\ell^*) & \text{and } a_k^t \neq a_k(\ell^*) \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.31})$$

EB. If $x^t = x(\cdot, \ell^*)$ and the realized profile is $a(\ell^*)$ then $\langle i, t \rangle$ believes that others send the same messages that they each received. In all other cases, $\langle i, t \rangle$ believes that others send the same message that he himself sends.

A.5. Trembles, Consistency, and Sequential Rationality

In this Section we lay out the necessary structure of trembles that satisfy the consistency requirement of Sequential equilibrium.

Definition A.16: Throughout this section we let ε denote a small positive number, which parameterizes the completely mixed strategies that we construct. It should be understood that our construction of beliefs involves the limit $\varepsilon \rightarrow 0$.

Definition A.17. *Completely Mixed Action Strategies:* Given ε , the completely mixed strategies for all players $\langle i, t \rangle$ at the action stage are denoted by $g_{i,\varepsilon}^t$ and are defined as follows.²⁵

After receiving a message $m \in \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$ and observing the realization x^t of the action-stage randomization device, any player $\langle i, t \rangle$ plays the action prescribed by the action-stage strategy described in Definition T.1.1 with probability $1 - \varepsilon^2(\|A\|_i - 1)$ and plays all other actions in A_i with probability ε^2 each.

After receiving any message $m \notin \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$ and observing the realization x^t of the action-stage randomization device, any player $\langle i, t \rangle$ plays the action prescribed by the action-stage strategy described in Definition T.1.1 with probability $1 - \varepsilon(\|A\|_i - 1)$ and plays all other actions in A_i with probability ε each.

Definition A.18. *Completely Mixed Message Strategies:* Given ε , the completely mixed strategies for all players $\langle i, t \rangle$ at the message stage are denoted by $\mu_{i,\varepsilon}^t$ and are defined as follows.

Player $\langle i, t \rangle$ sends the message prescribed by the message-stage strategy described in Definition T.1.2 with probability $1 - \varepsilon^{2n+1}(\|M_i^{t+1}\| - 1)$ and sends all other messages in M_i^{t+1} with probability ε^{2n+1} each.

Remark A.6. *Consistency:* Notice that deviations in the message stage are much less likely than deviations in the action stage. In particular, a single deviation in the message stage is infinitely less likely than n deviations in the action stage, one for each player. A consequence of this is that each player assigns probability one to the event that all other players receive messages that are matched to the same phase of the equilibrium as his. See Section T.6 in the Technical Addendum for details.

Remark A.7. *Sequential Rationality:* Given the structure of beliefs, the equilibrium incentives in the action stage are straightforward. For example, consider the action incentives of player $\langle i, t \rangle$ in phase S . If he keeps his prescribed action, he receives a payoff that converges to v_i^* as δ goes to one. Whereas if he deviates, he receives some payoff that converges to $q\bar{v}_i + (1-q)z_i < v_i^*$ as δ goes to one. Other incentive constraints follow a similarly standard logic.

Incentives at the message stage are more complicated and cannot be summarized in a short period of space. The key step is to show that player $\langle i, t \rangle$ sends his prescribed message in phase P^i . The fact that player $\langle i, t+1 \rangle$ assigns probability 0 in the action stage to the event that play is in phase P^i is critical. The full details are in Section T.5 in the Technical Addendum.

Appendix B: Proof of Theorem 3

B.1. Preliminaries

Definition B.1: Let a profile of message strategies μ be given. Fix an “augmented history” $\kappa^t = (x^0, a^0, y^0, \dots, x^{t-1}, a^{t-1}, y^{t-1})$. In other words, fix a history h^t , together with a sequence of realizations of the message-stage randomization device (y^0, \dots, y^{t-1}) . In what follows, κ^0 will denote the null history \emptyset , and for any $\tau \leq t$, κ^τ will denote the appropriate subset of κ^t .

For every $i \in I$ let $\mathcal{M}_i^0(m_i^0 | \kappa^0, \mu_i) = 1$. Then, recursively forward, define

$$\mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) = \sum_{m_i^{t-1} \in H^{t-1}} \mu_i^{t-1}(m_i^t | m_i^{t-1}, x^{t-1}, a^{t-1}, y^{t-1}) \mathcal{M}_i^{t-1}(m_i^{t-1} | \kappa^{t-1}, \mu_i) \quad (\text{B.1})$$

²⁵In the interest of brevity, we avoid an explicit distinction between the $t = 0$ players and all others. What follows can be interpreted as applying to all players re-defining m_i^0 to be equal to m^* for players $\langle i \in I, 0 \rangle$.

So that $\mathcal{M}_i^t(m_i^t|\kappa^t, \mu_i)$ is the probability that player $\langle i, t-1 \rangle$ sends message m_i^t given κ^t and the profile μ_i . We also let $\mathcal{M}_{-i}^t(m_{-i}^t|\kappa^t, \mu_{-i}) = \mathcal{M}_{-i}^t((m_i^t, \dots, m_{i-1}^t, m_{i+1}^t, \dots, m_n^t)|\kappa^t, \mu_{-i}) = \prod_{j \neq i} \mathcal{M}_j^t(m_j^t|\kappa^t, \mu_j)$.

Lemma B.1: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Fix any SE of the dynastic repeated game (g, μ, Φ) . Assume that it displays Inter-Generational Agreement as in Definition 2.

Let any augmented history κ^t as in Definition B.1 be given. Let also any $i \in I$ and any m_i^t such that $\mathcal{M}_i^t(m_i^t|\kappa^t, \mu_i) > 0$ be given.

Then for any m_{-i}^t

$$\Phi_i^{tB}(m_{-i}^t|m_i^t) = \mathcal{M}_{-i}^t(m_{-i}^t|\kappa^t, \mu_{-i}) \quad (\text{B.2})$$

Proof: We proceed by induction. Given the fixed κ^t , let $\kappa^0 = \emptyset$ and κ^τ with $\tau = 1, \dots, t$ be the augmented histories comprising the first three components (x^0, a^0, y^0) of κ^t , the first six components $(x^0, a^0, y^0, x^1, a^1, y^1)$ of κ^t and so on. First of all notice that setting $\tau = 0$ yields

$$\Phi_i^{1B}(m_{-i}^0|m_i^0) = \mathcal{M}_{-i}^1(m_{-i}^0|\kappa^0, \mu_{-i}) = 1 \quad (\text{B.3})$$

which is trivially true given that all players $\langle i \in I, 0 \rangle$ receive the null message by construction.

Our working hypothesis is now that the claim is true for an arbitrary $\tau - 1 < t - 1$, and our task is to show that it holds for τ .

Consider any message m_i^τ in $\text{Supp}(\mathcal{M}_i^\tau(\cdot|\kappa^\tau, \mu_i))$.²⁶ Then there must exist a message $m_i^{\tau-1}$ such that

$$\mu_i^{\tau-1}(m_i^\tau|m_i^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \mathcal{M}_i^{\tau-1}(m_i^{\tau-1}|\kappa^{\tau-1}, \mu_i) > 0 \quad (\text{B.4})$$

Therefore, using (4) we can write

$$\Phi_i^{\tau B}(m_{-i}^\tau|m_i^\tau) = \Phi_i^{\tau-1 E}(m_{-i}^\tau|m_i^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \quad (\text{B.5})$$

Notice that in any SE it must be the case that the right-hand side of (B.5) is equal to

$$\sum_{m_{-i}^{\tau-1}} \Phi_i^{\tau-1 R}(m_{-i}^{\tau-1}|m_i^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \left[\prod_{j \neq i} \mu_j^{\tau-1}(m_j^\tau|m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{B.6})$$

Using (3), we know that (B.6) is equal to

$$\sum_{m_{-i}^{\tau-1}} \Phi_i^{\tau-1 B}(m_{-i}^{\tau-1}|m_i^{\tau-1}) \left[\prod_{j \neq i} \mu_j^{\tau-1}(m_j^\tau|m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{B.7})$$

Our working hypothesis can now be used to assert that (B.7) is in turn equal to

$$\sum_{m_{-i}^{\tau-1}} \mathcal{M}_{-i}^{\tau-1}(m_{-i}^{\tau-1}|\kappa^{\tau-1}, \mu_{-i}) \left[\prod_{j \neq i} \mu_j^{\tau-1}(m_j^\tau|m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{B.8})$$

²⁶ $\text{Supp}(\cdot)$ denotes the support of a probability distribution.

Rearranging terms (B.8) we find that it can also be written as

$$\prod_{j \neq i} \left[\sum_{m_j^{\tau-1}} \mathcal{M}_j^{\tau-1}(m_j^{\tau-1} | \kappa^{\tau-1}, \mu_j) \mu_j^{\tau-1}(m_j^{\tau-1} | m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{B.9})$$

Using now (B.1), it is immediate that (B.9) is equal to

$$\prod_{j \neq i} \mathcal{M}_j^{\tau}(m_j^{\tau} | \kappa^{\tau}, \mu_j) = \mathcal{M}_{-i}^{\tau}(m_{-i}^{\tau} | \kappa^{\tau}, \mu_{-i}) \quad (\text{B.10})$$

and hence the claim is proved. ■

Definition B.2: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Fix any strategy profile, (g, μ) , for the dynastic repeated game.

Consider the standard repeated game with the same common discount factor δ , and with the following action-stage randomization device $\hat{\tilde{x}}$. The random variable $\hat{\tilde{x}}$ takes values in the finite set $Y \times X$ (the sets in which \tilde{y} and \tilde{x} take values respectively), and the probability that $\hat{\tilde{x}}$ is equal to $\hat{x} = (y, x)$ is $\Pr(\tilde{y} = y) \Pr(\tilde{x} = x)$. For notational convenience we will denote the realization \hat{x}^t of $\hat{\tilde{x}}^t$ by the pair (y^{t-1}, x^t) .

Recall that a history in the standard repeated game with randomization device $\hat{\tilde{x}}$ is an object of the type $h^t = (\hat{x}^0, a^0, \dots, \hat{x}^{t-1}, a^{t-1})$. Therefore, using our notational convention about time superscripts of the realizations of $\hat{\tilde{x}}^t$ we have that any pair (h^t, \hat{x}^t) can be written as a triple (y^{-1}, κ^t, x^t) , where κ^t corresponds to h^t in the obvious way.

We say that the strategy profile g^* for the standard repeated game with randomization device $\hat{\tilde{x}}$ is derived from the dynastic repeated game profile (g, μ) as above if and only if it is defined as follows.

$$g_i^{t*}(h^t, \hat{x}^t) = g_i^{t*}(y^{-1}, \kappa^t, x^t) = \sum_{m_i^t} \mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) g_i^t(m_i^t, x^t) \quad (\text{B.11})$$

Lemma B.2: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Consider any SE, (g, μ) , of the dynastic repeated game that displays Inter-Generational Agreement as in Definition 2.

Now consider the strategy profile g^* for the standard repeated game with randomization device $\hat{\tilde{x}}$ that is derived from (g, μ) as in Definition B.2.

Given g^* , fix any pair (h^t, \hat{x}^t) representing a history and realized randomization device for the standard repeated game. For any $a_{-i}^t \in A_{-i}$, let $\mathcal{P}_{g^* | h^t, \hat{x}^t}(a_{-i}^t)$ be the probability that the realized action profile for all players but i at time t is a_{-i}^t .

Given the pair (h^t, \hat{x}^t) , consider the corresponding triple (y^{-1}, κ^t, x^t) as in Definition B.2. Then

$$\mathcal{P}_{g^* | h^t, \hat{x}^t}(a_{-i}^t) = \prod_{j \neq i} \left\{ \sum_{m_j^t} \mathcal{M}_j^t(m_j^t | \kappa^t, \mu_j) g_j^t(a_j^t | m_j^t, x^t) \right\} \quad (\text{B.12})$$

Proof: The claim is a direct consequence of (B.11) of Definition B.2. ■

Lemma B.3: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Consider any SE (g, μ, Φ) of the dynastic repeated game that displays Inter-Generational Agreement as in Definition 2.

Fix any pair (h^t, \hat{x}^t) representing a history and realized randomization device for the standard repeated game. Given the pair (h^t, \hat{x}^t) , consider the corresponding triple (y^{-1}, κ^t, x^t) as in Definition B.2. Given the last two elements of this triple (κ^t, x^t) , now use (B.1) to find a message \bar{m}_i^t such that $\mathcal{M}_i^t(\bar{m}_i^t | \kappa^t, \mu_i) > 0$.

Finally, consider the following alternative action-stage and message-stage strategies $(\bar{g}_i^t, \bar{\mu}_i^t)$ for player $\langle i, t \rangle$. Whenever $m_i^t \neq \bar{m}_i^t$, set $\bar{g}_i^t = g_i^t$ and $\bar{\mu}_i^t = \mu_i^t$. Then define

$$\bar{g}_i^t(\bar{m}_i^t, x^t) = \sum_{m_i^t} \mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) g_i^t(m_i^t, x^t) \quad (\text{B.13})$$

and

$$\bar{\mu}_i^t(\bar{m}_i^t, x^t, a^t, y^t) = \sum_{m_i^t} \mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) \mu_i^t(m_i^t, x^t, a^t, y^t) \quad (\text{B.14})$$

Here and throughout the rest of the paper and the technical addendum, we denote by $v_i^t(g, \mu | m_i^t, x^t, \Phi_i^{tB})$ the continuation payoff to player $\langle i, t \rangle$ given the profile (g, μ) , after he has received message m_i^t has observed the realization x^t , and given that his beliefs over the $n-1$ -tuple m_{-i}^t are Φ_i^{tB} . See also our Point of Notation T.2.1. Then

$$v_i^t(g, \mu | \bar{m}_i^t, x^t, \Phi_i^{tB}) = v_i^t(\bar{g}_i^t, g_{-i}^{-t}, \bar{\mu}_i^t, \mu_{-i}^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) \quad (\text{B.15})$$

Proof: The claim is a direct consequence of Lemma B.1 and of (3) of Definition 2. The details are omitted for the sake of brevity. ■

B.2. Proof of the Theorem

Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Consider any SE triple (g, μ, Φ) for the dynastic repeated game. Assume that this SE displays Inter-Generational Agreement as in Definition 2.

Now consider the strategy profile g^* for the standard repeated game with common discount δ and randomization device \hat{x} that is derived from (g, μ) as in Definition B.2.

Since (g, μ) and g^* obviously give rise to the same payoff vector, to prove the claim it is enough to show that g^* is a SPE of the repeated game with δ and \hat{x} . By way of contradiction, suppose that it is not.

By the one-shot deviation principle this implies that there exist an i , an h^t , an \hat{x}^t and a σ_i such that

$$v_i(\sigma_i, g_i^{-t*}, g_{-i}^* | h^t, \hat{x}^t) > v_i(g^* | h^t, \hat{x}^t) \quad (\text{B.16})$$

Given the pair (h^t, \hat{x}^t) , consider the corresponding triple (y^{-1}, κ^t, x^t) as in Definition B.2. Given the last two elements of this triple (κ^t, x^t) , now use (B.1) to find a message \bar{m}_i^t such that $\mathcal{M}_i^t(\bar{m}_i^t | \kappa^t, \mu_i) > 0$.

Using Lemmas B.1 and B.2 we can now conclude that (B.16) implies that

$$v_i^t(\sigma_i, g_i^{-t}, g_{-i}, \bar{\mu}_i^t, \mu_{-i}^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) > v_i^t(\bar{g}_i^t, g_{-i}^{-t}, \bar{\mu}_i^t, \mu_{-i}^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) \quad (\text{B.17})$$

where σ_i is the profitable deviation identified in (B.16) and \bar{g}_i^t and $\bar{\mu}_i^t$ are the alternative action-stage and message-stage strategies of Lemma B.3.

However, using (B.15) of Lemma B.3, the inequality in (B.17) clearly implies that

$$v_i^t(\sigma_i, g_i^{-t}, g_{-i}, \bar{\mu}_i^t, \mu_i^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) > v_i^t(g, \mu | \bar{m}_i^t, x^t, \Phi_i^{tB}) \quad (\text{B.18})$$

But since (B.18) contradicts the fact that (g, μ, Φ) is an SE of the dynastic repeated game, the proof is now complete. ■

References

- ABREU, D., P. K. DUTTA, AND L. SMITH (1994): “The Folk Theorem for Repeated Games: A NEU Condition,” *Econometrica*, 62, 939–948.
- ALIPRANTIS, C., G. CAMERA, AND D. PUZZELLO (2007a): “Anonymous Markets and Monetary Trading,” *Journal of Monetary Economics*, forthcoming.
- (2007b): “Contagion Equilibria in a Monetary Model,” *Econometrica*, 75, 277–282.
- ANDERLINI, L., D. GERARDI, AND R. LAGUNOFF (2004): “The Folk Theorem in Dynastic Repeated Games,” Cowles Foundation Discussion Paper No 1490. <http://cowles.econ.yale.edu/P/cd/cfdpmain.htm>.
- ANDERLINI, L., D. GERARDI, AND R. LAGUNOFF (2007): “Social Memory and Evidence from the Past,” mimeo.
- ANDERLINI, L., AND R. LAGUNOFF (2005): “Communication in Dynastic Repeated Games: ‘Whitewashes’ and ‘Coverups,’” *Economic Theory*, 26, 265–299.
- BEN-PORATH, E., AND M. KAHNEMAN (1996): “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 70, 281–297.
- BERGIN, J. (2006): “The Folk Theorem Revisited,” *Economic Theory*, 27, 321–332.
- BHASKAR, V. (1998): “Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems,” *Review of Economic Studies*, 65, 135–149.
- CHAUDHURI, A., A. SCHOTTER, AND B. SOPHER (2001): “Talking Ourselves to Efficiency: Coordination in an Inter-generational Minimum Game with Private, Almost Common and Common Knowledge of Advice,” Working Paper 01-11, C.V. Starr Center for Applied Economics, New York University.
- COMPTE, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66, 597–626.
- CORBAE, D., T. TEMZELIDES, AND R. WRIGHT (2001): “Endogenous Matching and Money,” University of Pennsylvania, mimeo.
- CREMER, J. (1986): “Cooperation in Ongoing Organizations,” *Quarterly Journal of Economics*, 101, 33–49.

- ELLISON, G. (1994): "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61, 567–588.
- FARRELL, J. (1993): "Meaning and Credibility in Cheap Talk Games," *Games and Economic Behavior*, 5, 514–531.
- FUDENBERG, D., AND E. S. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533–556.
- GAMES AND ECONOMIC BEHAVIOR (1997): *Special Issue on Imperfect Recall*, vol. 20, no. 1. New-York: Academic Press.
- JOHNSON, P., D. K. LEVINE, AND W. PESENDORFER (2001): "Evolution and Information in a Gift Giving Game," *Journal of Economic Theory*, 100, 1–22.
- KANDORI, M. (1992a): "Repeated Games Played by Overlapping Generations of Players," *Review of Economic Studies*, 59, 81–92.
- (1992b): "Social Norms and Community Enforcement," *Review of Economic Studies*, 59, 63–80.
- 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.
- KOCHERLAKOTA, N. (1998): "Money is Memory," *Journal of Economic Theory*, 81, 232–251.
- KOCHERLAKOTA, N., AND N. WALLACE (1998): "Incomplete Record Keeping and Optimal Payout Arrangements," *Journal of Economic Theory*, 81, 272–289.
- KREPS, D. M., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.
- LAGUNOFF, R., AND A. MATSUI (2004): "Organizations and Overlapping Generations Games: Memory, Communication, and Altruism," *Review of Economic Design*, 8, 383–411.
- MAILATH, G. J., AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford University Press, *forthcoming*.
- MASKIN, E. S. (1999): "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23–38.
- MATTHEWS, S., M. OKUNO-FUJIWARA, AND A. POSTLEWAITE (1991): "Refining Cheap-Talk Equilibria," *Journal of Economic Theory*, 55, 247–273.

- PICCIONE, M., AND A. RUBINSTEIN (1997): “On the Interpretation of Decision Problems with Imperfect Recall,” *Games and Economic Behavior*, 20, 3–24.
- SALANT, D. (1991): “A Repeated Game with Finitely Lived Overlapping Generations of Players,” *Games and Economic Behavior*, 3, 244–259.
- SCHOTTER, A., AND B. SOPHER (2001a): “Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach,” New York University, mimeo.
- SCHOTTER, A., AND B. SOPHER (2001b): “Social Learning and Coordination Conventions in Inter-Generational Games: An Experimental Study,” New York University, mimeo.
- SMITH, L. (1992): “Folk Theorems in Overlapping Generations Games,” *Games and Economic Behavior*, 4, 426–449.
- WALLACE, N. (2001): “Whither Monetary Economics?,” *International Economic Review*, 42, 847–869.

A “Super” Folk Theorem for Dynastic Repeated Games: Technical Addendum

Luca Anderlini

Dino Gerardi

Roger Lagunoff

September 2007

T.1. A Compact Description of Strategies

Definition T.1.1. Action-Stage Strategies: Let k be an element of I , and j be an element of I not equal to i .

Let $(a(1), \dots, a(\|A\|))$ be the enumeration of the elements of A of Definition A.2 and consider the indexation of the elements of X in Definition A.7, according to whether $x(\kappa)$ has $\kappa \leq \bar{\kappa}$ or not.

Recall that at the beginning of period 0 all players $\langle i \in I, 0 \rangle$ receive message $m_i^0 = \emptyset$. For all players $\langle i \in I, 0 \rangle$ then define

$$g_i^0(m_i^0, x^0) = \begin{cases} a_i(\ell_0) & \text{if } x^0 = x(\ell_0, \dots) \\ a_i(\ell_{00}) & \text{if } x^0 = x(\ell_{00}, \dots) \end{cases} \quad (\text{T.1.1})$$

Now consider any player $\langle i, t \rangle$ with $t \geq 1$. It is convenient to distinguish between the two cases $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ and with $\kappa > \bar{\kappa}$.

For any $i \in I$ and $t \geq 1$ whenever $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ define^{T.1}

$$g_i^t(m_i^t, x^t) = \begin{cases} a_i(\hat{\ell}) & \text{if } m_i^t = m^* & \text{and } x^t = x(\cdot, \hat{\ell}, \dots) \\ \check{a}_i^j(j_\ell) & \text{if } m_i^t = \check{m}_j & \text{and } x^t = x(\dots, j_\ell, \dots) \\ a_i(\ell_k) & \text{if } m_i^t = \bar{m}^k & \text{and } x^t = x(\dots, \ell_k, \dots) \\ a_i(k_\ell) & \text{if } m_i^t \in \underline{M}(k, t) & \text{and } x^t = x(\dots, k_\ell, \dots) \end{cases} \quad (\text{T.1.2})$$

For any $i \in I$, $t \geq 1$ and m_i^t , whenever $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$ define

$$g_i^t(m_i^t, x^t) = a_i(\ell^*) \text{ if } x^t = x(\cdot, \ell^*) \quad (\text{T.1.3})$$

Definition T.1.2. Message-Stage Strategies: Let k be any element of I , and j be any element of I not equal to i .

We begin with period $t = 0$. Recall that $m_i^0 = \emptyset$ for all $i \in I$. Let also $g^0(m^0, x^0) = (g_1^0(m_1^0, x^0), \dots, g_n^0(m_n^0, x^0))$, and define $g_{-k}^0(m^0, x^0)$ in the obvious way.

We let

$$\mu_i^0(m_i^0, x^0, a^0, y^0) = \begin{cases} \check{m}^j & \text{if } a^0 = g^0(m^0, x^0) & \text{and } y^0 = y(j) \\ \bar{m}^{i,T} & \text{if } a^0 = g^0(m^0, x^0) & \text{and } y^0 = y(i) \\ \bar{m}^{k,T} & \text{if } a_{-k}^0 = g_{-k}^0(m^0, x^0) & \text{and } a_k^0 \neq g_k^0(m^0, x^0) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.4})$$

For the periods $t \geq 1$ it is convenient to distinguish between several cases. Assume first that $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$. Let

$$\mu_i^t(m_i^t, x^t, a^t, y^t) = \begin{cases} m_i^t & \text{if } x^t = x(\cdot, \ell^*) & \text{and } a^t = a(\ell^*) \\ \bar{m}^{k,T} & \text{if } x^t = x(\cdot, \ell^*), \ a_{-k}^t = a_{-k}(\ell^*) & \text{and } a_k^t \neq a_k(\ell^*) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.5})$$

^{T.1}Notice that the third case in (T.1.2) can only possibly apply when $t \geq T + 1$.

Now consider the case $x^t = x^t(\kappa)$ with $\kappa \leq \bar{\kappa}$. We divide this case into several subcases, according to which message player $\langle i, t \rangle$ has received. We begin with $m_i^t = m^*$. Let^{T.2}

$$\mu_i^t(m^*, x^t, a^t, y^t) = \begin{cases} \check{m}^j & \text{if } x^t = x(\cdot, \hat{\ell}, \cdot \cdot \cdot), a^t = a(\hat{\ell}) & \text{and } y^t = y(j) \\ \nu(\underline{M}(i, t + 1)) & \text{if } x^t = x(\cdot, \hat{\ell}, \cdot \cdot \cdot), a^t = a(\hat{\ell}) & \text{and } y^t = y(i) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdot, \hat{\ell}, \cdot \cdot \cdot), a_{-k}^t = a_{-k}(\hat{\ell}) & \text{and } a_k^t \neq a_k(\hat{\ell}) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.6})$$

Our next subcase of $\kappa \leq \bar{\kappa}$ is that of $m_i^t = \check{m}^j$. With the understanding that j' is any element of I not equal to i , we let

$$\mu_i^t(\check{m}^j, x^t, a^t, y^t) = \begin{cases} \check{m}^{j'} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a^t = \check{a}^j(j_\ell) & \text{and } y^t = y(j') \\ \nu(\underline{M}(i, t + 1)) & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a^t = \check{a}^j(j_\ell) & \text{and } y^t = y(i) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a_{-k}^t = \check{a}_{-k}^j(j_\ell) & \text{and } a_k^t \neq \check{a}_k^j(j_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.7})$$

Still assuming $\kappa \leq \bar{\kappa}$ we now deal with the subcase $m_i^t \in \underline{M}(i, t)$. For any $\underline{m}^{i, \tau} \in \underline{M}(i, t)$, we let

$$\mu_i^t(\underline{m}^{i, \tau}, x^t, a^t, y^t) = \begin{cases} \check{m}^j & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a^t = \check{a}^i(i_\ell) & \text{and } y^t = y(j) \\ \nu(\underline{M}(i, t + 1)) & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a^t = \check{a}^i(i_\ell) & \text{and } y^t = y(i) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a_{-k}^t = \check{a}_{-k}^i(i_\ell) & \text{and } a_k^t \neq \check{a}_k^i(i_\ell) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a_{-k}^t = a_{-k}^i(i_\ell) & \text{and } a_k^t \neq a_k^i(i_\ell) \\ \underline{m}^{i, \tau-1} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot) & \text{and } a^t = a(i_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.8})$$

where we set $\underline{m}^{i, 0} = \bar{m}^i$. Notice that player $\langle i, t \rangle$ may need to distinguish between the third and fourth cases of (T.1.8) since clearly they may be generated by different values of the index $k \in I$. To verify that this distinction is always feasible, recall that, by construction (see Definition A.3), $\check{a}_{-i}(i_\ell)$ differs from $a_{-i}(i_\ell)$ in every component, and that of course $n \geq 4$.

The next subcase of $\kappa \leq \bar{\kappa}$ we consider is that of $m_i^t \in \underline{M}(j, t)$. For any $\underline{m}^{j, \tau} \in \underline{M}(j, t)$, we let

$$\mu_i^t(\underline{m}^{j, \tau}, x^t, a^t, y^t) = \begin{cases} \underline{m}^{j, \tau-1} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot) & \text{and } a^t = a(j_\ell) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdot \cdot \cdot, j_\ell, \cdot \cdot \cdot), a_{-k}^t = a_{-k}(j_\ell) & \text{and } a_k^t \neq a_k(j_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.9})$$

where we set $\underline{m}^{j, 0} = \bar{m}^j$.

Finally, still assuming that $\kappa \leq \bar{\kappa}$, we consider the case in which $m_i^t = \bar{m}^{k'}$ for some $k' \in I$. We let

$$\mu_i^t(\bar{m}^{k'}, x^t, a^t, y^t) = \begin{cases} \bar{m}^{k'} & \text{if } x^t = x(\cdot \cdot \cdot, \ell_{k'}, \cdot \cdot \cdot) & \text{and } a^t = a(\ell_{k'}) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdot \cdot \cdot, \ell_{k'}, \cdot \cdot \cdot), a_{-k}^t = a_{-k}(\ell_{k'}) & \text{and } a_k^t \neq a_k(\ell_{k'}) \\ m^* & \text{otherwise} \end{cases} \quad (\text{T.1.10})$$

T.2. Notation

Point of Notation T.2.1: Abusing the notation we established for the standard repeated game, we adopt the following notation for continuation payoffs in the dynastic repeated game. Let an assessment (g, μ, Φ) be given.

^{T.2}Throughout the paper we adopt the following notational convention. Given any finite set, we denote by $\nu(\cdot)$ the uniform probability distribution over the set. So, if B is a finite set, $\nu(B)$ assigns probability $1/\|B\|$ to every element of B .

Recall that we denote by $v_i^t(g, \mu | m_i^t, x^t, \Phi_i^{tB})$ the continuation payoff to player $\langle i, t \rangle$ given the profile (g, μ) , after he has received message m_i^t , has observed the realization x^t , and given that his beliefs over the $n - 1$ -tuple m_{-i}^t are Φ_i^{tB} . In view of our discussion at the beginning of Section 3, it is clear that the only component of the system of beliefs Φ that is relevant to define this continuation payoff is in fact Φ_i^{tB} . Our discussion there also implies that the argument m_i^t is redundant once Φ_i^{tB} has been specified. We keep it in our notation since it helps streamline some of the arguments below.

We let $v_i^t(g, \mu | m_i^t, x^t, a^t, y^t, \Phi_i^{tE})$ denote the continuation payoff (viewed from the beginning of period $t + 1$) to player $\langle i, t \rangle$ given the profile (g, μ) , after he has received message m_i^t , has observed the triple (x^t, a^t, y^t) , and given that his beliefs over the $n - 1$ -tuple m_{-i}^{t+1} are given by Φ_i^{tE} . In view of our discussion at the beginning of Section 3, it is clear that once Φ_i^{tE} has been specified, the arguments (m_i^t, x^t, a^t, y^t) are redundant in determining the end-of-period continuation payoff to player $\langle i, t \rangle$. Whenever this does not cause any ambiguity (about Φ_i^{tE}) we will write $v_i^t(g, \mu | \Phi_i^{tE})$ instead of $v_i^t(g, \mu | m_i^t, x^t, a^t, y^t, \Phi_i^{tE})$.

As we noted in the text all continuation payoffs clearly depend on δ as well. To keep notation down this dependence will be omitted whenever possible.

Point of Notation T.2.2: We will abuse our notation for $\Phi_i^{tB}(\cdot)$, $\Phi_i^{tE}(\cdot)$ and $\Phi_i^{tR}(\cdot)$ slightly in the following way. We will allow events of interest and conditioning events to appear as arguments of Φ_i^{tB} , Φ_i^{tE} and Φ_i^{tR} , to indicate their probabilities under these distributions.

So, for instance when we write $\Phi_i^{tB}(m_{-i}^t = (z, \dots, z) | m_i^t) = c$ we mean that according to the beginning-of-period beliefs of player $\langle i, t \rangle$, after observing m_i^t , the probability that m_{-i}^t is equal to the $n - 1$ -tuple (z, \dots, z) is equal to c .

Point of Notation T.2.3: Whenever the profile (g, μ) is a profile of completely mixed strategies, the beliefs $\Phi_i^{tB}(\cdot)$, $\Phi_i^{tE}(\cdot)$ and $\Phi_i^{tR}(\cdot)$ are of course entirely determined by what player $\langle i, t \rangle$ observes and by (g, μ) using Bayes' rule. In this case, we will allow the pair (g, μ) to appear as a "conditioning event."

So, for instance, $\Phi_i^{tB}(m_{-i}^t | m_i^t, g, \mu)$ is the probability of the $n - 1$ -tuple m_{-i}^t , after m_i^t has been received, obtained from the completely mixed profile (g, μ) via Bayes' rule. Events may appear as arguments in this case as well, consistently with our Point of Notation T.2.2 above.

Moreover, since the completely mixed pair (g, μ) determines the probabilities of all events, concerning for instance histories, messages of previous cohorts and the like, we will use the notation \Pr to indicate such probabilities, using the pair (g, μ) as a conditioning event.

So, given any two events L and J , the notation $\Pr(L | J, g, \mu)$ will indicate the probability of event L , conditional on event J , as determined by the completely mixed pair (g, μ) via Bayes' rule.

T.3. A Preliminary Result

As we mentioned before, we work with message spaces that are smaller than the set H^t . We now proceed to show that this is without loss of generality.

Definition T.3.1: Consider the dynastic repeated game described in full in Section 2. Now consider the dynastic repeated game obtained from this when we restrict the message space of player $\langle i, t \rangle$ to be $M_i^{t+1} \subseteq H^{t+1}$, with all other details unchanged.

We call this the restricted dynastic repeated game with message spaces $\{M_i^t\}_{i \in I, t \geq 1}$. For any given $\delta \in (0, 1)$, \tilde{x} and \tilde{y} , we denote by $\mathcal{G}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1})$ the set of SE strategy profiles, while we write $\mathcal{E}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1})$ for the set of SE payoff profiles of this dynastic repeated game with restricted message spaces.

Lemma T.3.1: Let any $\delta \in (0, 1)$, \tilde{x} and \tilde{y} be given. Consider now any restricted dynastic repeated game with message spaces $\{M_i^t\}_{i \in I, t \geq 1}$. Then $\mathcal{E}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1}) \subseteq \mathcal{E}^D(\delta, \tilde{x}, \tilde{y})$.

Proof: Let a profile $(g^*, \mu^*) \in \mathcal{G}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1})$ with associated beliefs Φ^* be given. To prove the statement, we proceed to construct a new profile $(g^{**}, \mu^{**}) \in \mathcal{G}^D(\delta, \tilde{x}, \tilde{y})$ and associated beliefs Φ^{**} that are consistent with (g^{**}, μ^{**}) , and which gives every player the same payoff as (g^*, μ^*) .

Denote a generic element of M_i^t by z_i^t . Since $M_i^t \subseteq H^t$, we can partition H^t into $\|M_i^t\|$ non-empty mutually exclusive exhaustive subsets, and make each of these subsets correspond to an element z_i^t of M_i^t . In other words, we can find a map $\rho_i^t: M_i^t \rightarrow 2^{H^t}$ such that $\rho_i^t(z_i^t) \neq \emptyset$ for all $z_i^t \in M_i^t$, $\rho_i^t(z_i^t) \cap \rho_i^t(z_i^{t'}) = \emptyset$ whenever $z_i^t \neq z_i^{t'}$, and $\bigcup_{z_i^t \in M_i^t} \rho_i^t(z_i^t) = H^t$.

We can now describe how the profile (g^{**}, μ^{**}) is obtained from the given (g^*, μ^*) . We deal first with the action stage. For any player $\langle i, t \rangle$, and any $z_i^t \in M_i^t$, set

$$g_i^{t**}(m_i^t, x) = g_i^{t*}(z_i^t, x) \quad \forall m_i^t \in \rho_i^t(z_i^t) \quad (\text{T.3.1})$$

At the message stage, for any player $\langle i, t \rangle$, any (z_i^t, x^t, a^t, y^t) , any $m_i^t \in \rho_i^t(z_i^t)$, and any $z_i^{t+1} \in \text{Supp}(\mu_i^{t*}(z_i^t, x^t, a^t, y^t))$, set

$$\mu_i^{t**}(m_i^{t+1} | m_i^t, x^t, a^t, y^t) = \frac{1}{\|\rho_i^{t+1}(z_i^{t+1})\|} \mu_i^{t*}(z_i^{t+1} | z_i^t, x^t, a^t, y^t) \quad \forall m_i^{t+1} \in \rho_i^{t+1}(z_i^{t+1}) \quad (\text{T.3.2})$$

Next, we describe Φ^{**} , starting with the beginning-of-period beliefs. For any player $\langle i, t \rangle$, any $z_i^t \in M_i^t$ and any $z_{-i}^t \in M_{-i}^t$, set

$$\Phi_i^{tB**}(m_{-i}^t | m_i^t) = \frac{\Phi_i^{tB*}(z_{-i}^t | z_i^t)}{\prod_{j \neq i} \|\rho_j^t(z_j^t)\|} \quad \forall m_i^t \in \rho_i^t(z_i^t), \quad \forall m_{-i}^t \in \prod_{j \neq i} \rho_j^t(z_j^t) \quad (\text{T.3.3})$$

Similarly, concerning the end-of-period beliefs, for any player $\langle i, t \rangle$, any (z_i^t, x^t, a^t, y^t) and any $z_{-i}^{t+1} \in M_{-i}^{t+1}$, set

$$\begin{aligned} \Phi_i^{tE**}(m_{-i}^{t+1} | m_i^t, x^t, a^t, y^t) = \\ \frac{\Phi_i^{tE*}(z_{-i}^{t+1} | z_i^t, x^t, a^t, y^t)}{\prod_{j \neq i} \|\rho_j^{t+1}(z_j^{t+1})\|} \quad \forall m_i^t \in \rho_i^t(z_i^t), \quad \forall m_{-i}^{t+1} \in \prod_{j \neq i} \rho_j^{t+1}(z_j^{t+1}) \end{aligned} \quad (\text{T.3.4})$$

Since the profile (g^*, μ^*) is sequentially rational given Φ^* , it is immediate from (T.3.1), (T.3.2), (T.3.3) and (T.3.4) that the profile (g^{**}, μ^{**}) is sequentially rational given Φ^{**} , and we omit further details of the proof of this claim.

Of course, it remains to show that $(g^{**}, \mu^{**}, \Phi^{**})$ is a consistent assessment.

Let $(g_\varepsilon^*, \mu_\varepsilon^*)$ be parameterized completely mixed strategies which converge to (g^*, μ^*) and give rise, in the limit as $\varepsilon \rightarrow 0$, to beliefs Φ^* via Bayes’ rule.

Given any $\varepsilon > 0$, let $(g_\varepsilon^{**}, \mu_\varepsilon^{**})$ be a profile of completely mixed strategies obtained from $(g_\varepsilon^*, \mu_\varepsilon^*)$ exactly as in (T.3.1) and (T.3.2).

We start by verifying the consistency of the beginning-of-period beliefs. Observe that for any given $z^t = (z_i^t, z_{-i}^t)$, from (T.3.2) we know that whenever $m^t = (m_i^t, m_{-i}^t) \in \prod_{j \in I} \rho_j^t(z_j^t)$

$$\Pr(m_i^t, m_{-i}^t | g_\varepsilon^{**}, \mu_\varepsilon^{**}) = \frac{\Pr(z_i^t, z_{-i}^t | g_\varepsilon^*, \mu_\varepsilon^*)}{\prod_{j \in I} \|\rho_j^t(z_j^t)\|} \quad (\text{T.3.5})$$

Similarly, using (T.3.2) again we know that whenever $m_i^t \in \rho_i^t(z_i^t)$

$$\Pr(m_i^t | g_\varepsilon^{**}, \mu_\varepsilon^{**}) = \frac{\Pr(z_i^t | g_\varepsilon^*, \mu_\varepsilon^*)}{\|\rho_i^t(z_i^t)\|} \quad (\text{T.3.6})$$

Taking the ratio of (T.3.5) and (T.3.6) and taking the limit as $\varepsilon \rightarrow 0$ now yields that for any any $z_i^t \in M_i^t$ and any $z_{-i}^t \in M_{-i}^t$

$$\lim_{\varepsilon \rightarrow 0} \Phi_i^{tB^{**}}(m_{-i}^t | m_i^t, g_{\varepsilon}^{**}, \mu_{\varepsilon}^{**}) = \frac{\Phi_i^{tB^*}(z_{-i}^t | z_i^t)}{\prod_{j \neq i} \|\rho_j^t(z_j^t)\|} \quad \forall m_i^t \in \rho_i^t(z_i^t), \quad \forall m_{-i}^t \in \prod_{j \neq i} \rho_j^t(z_j^t) \quad (\text{T.3.7})$$

Hence we have shown that the beginning-of-period beliefs as in (T.3.3) are consistent with (g^{**}, μ^{**}) .

The proof that the end-of-period beliefs as in (T.3.4) are consistent with (g^{**}, μ^{**}) runs along exactly the same lines, and we omit the details. ■

T.4. Proof of Theorem A.1: Beliefs

Definition T.4.1. *Beginning-of-Period Beliefs:* Let k be any element of I , and j be any element of I not equal to i .

The beginning-of-period beliefs of all players $\langle i \in I, 0 \rangle$ are trivial. Of course, all players believe that all other players have received the null message $m_i^0 = \emptyset$.

The beginning-of-period beliefs $\Phi_i^{tB}(m_i^t)$ of any other player $\langle i, t \rangle$, depending on the message he receives from player $\langle i, t-1 \rangle$ are as follows^{T.3}

$$\begin{aligned} & \text{if } m_i^t = m^* \quad \text{then } m_{-i}^t = (m^*, \dots, m^*) \text{ with probability 1} \\ & \text{if } m_i^t = \check{m}^j \quad \text{then } \begin{cases} m_{-i-j}^t = (\check{m}^j, \dots, \check{m}^j) & \text{with pr. 1} \\ m_j^t \in \underline{M}(j, t) & \text{with pr. 1} \\ \Pr(m_j^t = \underline{m}^{j,\tau}) > 0 & \forall \underline{m}^{j,\tau} \in \underline{M}(j, t) \end{cases} \\ & \text{if } m_i^t = \underline{m}^{j,\tau} \quad \text{then } m_{-i}^t = (\underline{m}^{j,\tau}, \dots, \underline{m}^{j,\tau}) \text{ with probability 1} \\ & \text{if } m_i^t = \check{m}^i \quad \text{then } m_{-i}^t = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\ & \text{if } m_i^t = \bar{m}^k \quad \text{then } m_{-i}^t = (\bar{m}^k, \dots, \bar{m}^k) \text{ with probability 1} \end{aligned} \quad (\text{T.4.1})$$

Definition T.4.2. *End-of-Period Beliefs:* Let k be any element of I , and j be any element of I not equal to i .

We begin with period $t = 0$. Recall that $m_i^0 = \emptyset$ for all $i \in I$. As before, let also $g^0(m^0, x^0) = (g_1^0(m_1^0, x^0), \dots, g_n^0(m_n^0, x^0))$, and define $g_{-k}^0(m^0, x^0)$ in the obvious way.

Let $\Phi_i^{0E}(m_i^0, x^0, a^0, y^0)$ be as follows

$$\begin{aligned} & \text{if } a^0 = g^0(m^0, x^0) \text{ and } y^0 = y(j) \quad \text{then } m_{-i-j}^1 = (\check{m}^j, \dots, \check{m}^j), m_j^1 = \underline{m}^{j,T} \text{ with pr. 1} \\ & \text{if } a^0 = g^0(m^0, x^0) \text{ and } y^0 = y(i) \quad \text{then } m_{-i}^1 = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\ & \text{if } a_{-k}^0 = g_{-k}^0(m^0, x^0) \text{ and } a_k^0 \neq g_k^0(m_k^0, x^0) \quad \text{then } m_{-i}^1 = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with prob. 1} \\ & \text{otherwise} \quad \quad \quad m_{-i}^1 = (m^*, \dots, m^*) \text{ with probability 1} \end{aligned} \quad (\text{T.4.2})$$

Our next case is $t \geq 1$ and $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$. Let $x(\ell_{00}, \ell^*)$ denote the realization of x^t . For any

^{T.3}Notice that the second line of (T.4.1) does not fully specify the probability distribution over the component m_j^t of the beliefs of player $\langle i, t \rangle$. For the rest of the argument, what matters is only that all elements of $\underline{M}(j, t)$ have positive probability, and that no message outside this set has positive probability. The distribution can be computed using Bayes' rule from the equilibrium strategies described in Definitions T.1.1 and T.1.2 above. We omit the details for the sake of brevity.

player $\langle i, t \rangle$, let $\Phi_i^{tE}(m_i^t, x(\ell_{00}, \ell^*), a^t, y^t)$ be as follows^{T.4}

$$\begin{array}{ll}
\text{if } a^t = a(\ell^*) \text{ and } m_i^t = \check{m}^j & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^j, \dots, \check{m}^j) & \text{with pr. 1} \\ m_j^{t+1} \in \underline{M}(j, t) & \text{with pr. 1} \\ \Pr(m_j^{t+1} = \underline{m}^{j,\tau}) > 0 & \forall \underline{m}^{j,\tau} \in \underline{M}(j, t) \end{cases} \\
\text{if } a^t = a(\ell^*) \text{ and } m_i^t = \underline{m}^{j,\tau} & \text{then } m_{-i}^{t+1} = (\underline{m}^{j,\tau}, \dots, \underline{m}^{j,\tau}) \text{ with probability 1} \\
\text{if } a^t = a(\ell^*) \text{ and } m_i^t = \underline{m}^{i,\tau} & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a^t = a(\ell^*) \text{ and } m_i^t = \bar{m}^k & \text{then } m_{-i}^{t+1} = (\bar{m}^k, \dots, \bar{m}^k) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(\ell^*) \text{ and } a_k^t \neq a_k(\ell^*) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.4.3}$$

We divide the case of $t \geq 1$ and $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ into several subcases, according to which message player $\langle i, t \rangle$ has received. We begin with $m_i^t = m^*$. Let $x(\cdot, \hat{\ell}, \dots)$ denote the realization of x^t . For any player $\langle i, t \rangle$, with the understanding that $\underline{m}^{j,\tau}$ is a generic element of $\underline{M}(j, t+1)$, let $\Phi_i^{tE}(m^*, x(\cdot, \hat{\ell}, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = a(\hat{\ell}) \text{ and } y^t = y(j) & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^j, \dots, \check{m}^j) \\ m_j^{t+1} = \underline{m}^{j,\tau} \end{cases} \text{ with pr. } \frac{1}{\|\underline{M}(j, t+1)\|} \\
\text{if } a^t = a(\hat{\ell}) \text{ and } y^t = y(i) & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(\hat{\ell}) \text{ and } a_k^t \neq a_k(\hat{\ell}) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.4.4}$$

The next subcase is that of $m_i^t = \check{m}^j$. Let $x(\dots, j_\ell, \dots)$ denote the realization of x^t . With the understanding that j' is an element of I not equal to i and that $\underline{m}^{j',\tau}$ is a generic element of $\underline{M}(j', t+1)$, let $\Phi_i^{tE}(\check{m}^j, x(\dots, j_\ell, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = \check{a}^j(j_\ell) \text{ and } y^t = y(j') & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^{j'}, \dots, \check{m}^{j'}) \\ m_{j'}^{t+1} = \underline{m}^{j',\tau} \end{cases} \text{ with pr. } \frac{1}{\|\underline{M}(j', t+1)\|} \\
\text{if } a^t = \check{a}^j(j_\ell) \text{ and } y^t = y(i) & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a_{-k}^t = \check{a}_{-k}^j(j_\ell) \text{ and } a_k^t \neq \check{a}_k^j(j_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.4.5}$$

The next subcase is that of $m_i^t = \underline{m}^{i,\tau} \in \underline{M}(i, t)$. Let $x(\dots, i_\ell, \dots)$ denote the realization of x^t . With the understanding that $\underline{m}^{j,\tau}$ is a generic element of $\underline{M}(j, t+1)$, let $\Phi_i^{tE}(\underline{m}^{i,\tau}, x(\dots, i_\ell, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = \check{a}^i(i_\ell) \text{ and } y^t = y(j) & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^j, \dots, \check{m}^j) \\ m_j^{t+1} = \underline{m}^{j,\tau} \end{cases} \text{ with pr. } \frac{1}{\|\underline{M}(j, t+1)\|} \\
\text{if } a^t = \check{a}^i(i_\ell) \text{ and } y^t = y(i) & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a_{-k}^t = \check{a}_{-k}^i(i_\ell) \text{ and } a_k^t \neq \check{a}_k^i(i_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(i_\ell) \text{ and } a_k^t \neq a_k(i_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{if } a^t = a(i_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{i,\tau-1}, \dots, \underline{m}^{i,\tau-1}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.4.6}$$

where we set $\underline{m}^{i,0} = \bar{m}^i$.

The next subcase of $t \geq 1$ and $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ that we consider is that of $m_i^t = \underline{m}^{j,\tau} \in \underline{M}(j, t)$.

^{T.4}Similarly to (T.4.1), the first line of (T.4.3) does not fully specify the probability distribution over the component m_j^{t+1} of the beliefs of player $\langle i, t \rangle$. For the rest of the argument, what matters is only that all elements of $\underline{M}(j, t)$ have positive probability, and that no message outside this set has positive probability. The distribution can be computed using Bayes' rule from the equilibrium strategies described in Definitions T.1.1 and T.1.2 above. We omit the details for the sake of brevity.

Let $x(\dots, j_\ell, \dots)$ denote the realization of x^t . Let $\Phi_i^{tE}(\underline{m}^{j,\tau}, x(\dots, j_\ell, \dots), a^t, y^t)$ be as follows

$$\begin{aligned} &\text{if } a^t = a(j_\ell) && \text{then } m_{-i}^{t+1} = (\underline{m}^{j,\tau-1}, \dots, \underline{m}^{j,\tau-1}) \text{ with probability 1} \\ &\text{if } a_{-k}^t = a_{-k}(j_\ell) \text{ and } a_k^t \neq a_k(j_\ell) && \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\ &\text{otherwise} && m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1} \end{aligned} \quad (\text{T.4.7})$$

where we set $\underline{m}^{j,0} = \bar{m}^j$.

The final subcase to consider is that of $m_i^t = \bar{m}^{k'}$ for some $k' \in I$. Let $x(\dots, \ell_{k'}, \dots)$ denote the realization of x^t . Let $\Phi_i^{tE}(\bar{m}^{k'}, x(\dots, \ell_{k'}, \dots), a^t, y^t)$ be as follows

$$\begin{aligned} &\text{if } a^t = a(\ell_{k'}) && \text{then } m_{-i}^{t+1} = (\bar{m}^{k'}, \dots, \bar{m}^{k'}) \text{ with probability 1} \\ &\text{if } a_{-k}^t = a_{-k}(\ell_{k'}) \text{ and } a_k^t \neq a_k(\ell_{k'}) && \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\ &\text{otherwise} && m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1} \end{aligned} \quad (\text{T.4.8})$$

T.5. Proof of Theorem A.1: Sequential Rationality

Definition T.5.1: Let \mathcal{I}_i^{tE} denote the end-of-period- t collection of information sets that belong to player $\langle i, t \rangle$, with typical element \mathcal{I}_i^{tE} .

It is convenient to partition \mathcal{I}_i^{tE} into mutually disjoint exhaustive subsets on the basis of the associated beliefs of player $\langle i, t \rangle$. The fact that they exhaust \mathcal{I}_i^{tE} can be checked directly from Definition T.4.2 above.

Let $\mathcal{I}_i^{tE}(\ast) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to (m^*, \dots, m^*) with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\ast)$.

Let $\mathcal{I}_i^{tE}(\check{i}) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to $(\check{m}^i, \dots, \check{m}^i)$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\check{i})$.

For every $j \in I$ not equal to i , let $\mathcal{I}_i^{tE}(\check{j}, t) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i-j}^{t+1} is equal to $(\check{m}^j, \dots, \check{m}^j)$ with probability one, that $\Pr(m_j^{t+1} = \underline{m}^{j,\tau}) > 0 \forall \underline{m}^{j,\tau} \in \underline{M}(j, t)$, and that $\Pr(m_j^{t+1} \in \underline{M}(j, t)) = 1$.^{T.5} These beliefs will be denoted by $\Phi_i^{tE}(\check{j}, t)$.

For every $j \in I$ not equal to i , let $\mathcal{I}_i^{tE}(\check{j}, t+1) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i-j}^{t+1} is equal to $(\check{m}^j, \dots, \check{m}^j)$ with probability one, that $\Pr(m_j^{t+1} = \underline{m}^{j,\tau}) = \|\underline{M}(j, t+1)\|^{-1} \forall \underline{m}^{j,\tau} \in \underline{M}(j, t+1)$. These beliefs will be denoted by $\Phi_i^{tE}(\check{j}, t+1)$.

For every $k \in I$, let $\mathcal{I}_i^{tE}(\bar{k}) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to $(\bar{m}^k, \dots, \bar{m}^k)$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\bar{k})$.

For every $k \in I$, and every $\tau = \max\{T-t, 1\}, \dots, T$ let $\mathcal{I}_i^{tE}(\underline{k}, \tau) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to $(\underline{m}^{k,\tau}, \dots, \underline{m}^{k,\tau})$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\underline{k}, \tau)$.

Definition T.5.2: Let the strategy profile (g, μ) described in Definitions T.1.1 and T.1.2 be given. Fix a period t and an n -tuple of messages $m^{t+1} = (m_1^{t+1}, \dots, m_n^{t+1})$, with $m_k^{t+1} \in M_k^{t+1}$ for every $k \in I$.

Clearly, the profile (g, μ) together with m^{t+1} uniquely determine a probability distribution over action profiles over all future periods, beginning with $t+1$.

Therefore, we can define the expected discounted (from the beginning of period $t+1$) payoff to player $\langle i, t \rangle$, given (g, μ) and m^{t+1} in the obvious way. This will be denoted by $\check{v}_i^t(m^{t+1})$. Moreover, since they play a special role in some of the computations that follow, we reserve two pieces of notation for two particular instances of m^{t+1} . The expression $\check{v}_i^t(\ast)$ stands for $\check{v}_i^t(m^{t+1})$ when $m^{t+1} = (m^*, \dots, m^*)$. Moreover, for any $k \in I$, the expression $\check{v}_i^t(k, \tau)$ stands for $\check{v}_i^t(m^{t+1})$ when $m_{-k}^{t+1} = (\check{m}^k, \dots, \check{m}^k)$ and $m_k^{t+1} = \underline{m}^{k,\tau} \in \underline{M}(k, t+1)$.

^{T.5}See footnote T.4 above.

Lemma T.5.1: For any $i \in I$, any $k \in I$, any t , and any $\tau = \max\{T - t, 1\}, \dots, T$, we have that

$$\ddot{v}_i^t(*) = \frac{(1 - \delta) [q \hat{v}_i + (1 - q) z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad (\text{T.5.1})$$

and

$$\ddot{v}_i^t(k, \tau) = \frac{(1 - \delta) [q \check{u}_i^k + (1 - q) z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad (\text{T.5.2})$$

where $\ddot{v}_i^t(*)$ and $\ddot{v}_i^t(k, \tau)$ are as in Definition T.5.2, \hat{v}_i is as in (A.7), z_i is as in Remark A.4, v_i^* is as in the statement of the Theorem, and \check{u}_i^k is as in (A.4).

Proof: Assume first that $t \geq T$. Using Definitions T.1.1 and T.1.2 we can write $\ddot{v}_i^t(*)$ and $\ddot{v}_i^t(k, \tau)$ recursively as

$$\ddot{v}_i^t(*) = q \left\{ (1 - \delta) \hat{v}_i + \delta \left[(1 - \eta) \ddot{v}_i^{t+1}(*) + \frac{\eta}{n} \sum_{k' \in I} \sum_{\tau=1}^T \frac{\ddot{v}_i^{t+1}(k', \tau)}{T} \right] \right\} + \frac{\delta q v_i^*}{(1 - q) [(1 - \delta) z_i + \delta \ddot{v}_i^{t+1}(*)]} \quad (\text{T.5.3})$$

and

$$\ddot{v}_i^t(k, \tau) = q \left\{ (1 - \delta) \check{u}_i^k + \delta \left[(1 - \eta) \ddot{v}_i^{t+1}(*) + \frac{\eta}{n} \sum_{k' \in I} \sum_{\tau=1}^T \frac{\ddot{v}_i^{t+1}(k', \tau)}{T} \right] \right\} + \frac{\delta q v_i^*}{(1 - q) [(1 - \delta) z_i + \delta \ddot{v}_i^{t+1}(k, \tau)]} \quad (\text{T.5.4})$$

Since the strategy profile (g, μ) described in Definitions T.1.1 and T.1.2 is stationary for $t \geq T$, we immediately have that $\ddot{v}_i^t(*) = \ddot{v}_i^{t+1}(*)$ and, for any $k \in I$ and any $\tau = 1, \dots, T$, $\ddot{v}_i^t(k, \tau) = \ddot{v}_i^{t+1}(k, \tau)$. Hence we can solve (T.5.3) and (T.5.4) simultaneously for the $NT + 1$ variables $\ddot{v}_i^t(*)$ and $\ddot{v}_i^t(k, \tau)$ ($k \in I$ and $\tau = 1, \dots, T$). Using (A.8) this immediately gives (T.5.1) and (T.5.2), as required.

Proceeding by induction backwards from $t = T$, it is also immediate to verify that the statement holds for any $t < T$. The details are omitted for the sake of brevity. ■

Lemma T.5.2: Let the strategy profile (g, μ) and system of beliefs Φ described in Definitions T.1.1, T.1.2, T.4.1 and T.4.2 be given. Then the end-of-period continuation payoffs for any player $\langle i, t \rangle$ (discounted as of the beginning of period $t + 1$) at any information set $\mathcal{I}_i^t \in \mathcal{I}_i^{tE}$ (as categorized in Definition T.5.1) are as follows.^{T.6}

$$v_i^t(g, \mu | \Phi_i^{tE}(*)) = \frac{(1 - \delta) [q \hat{v}_i + (1 - q) z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad (\text{T.5.5})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\simeq i)) = \frac{(1 - \delta) [q \check{u}_i^i + (1 - q) z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad (\text{T.5.6})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\simeq j, t)) = v_i^t(g, \mu | \Phi_i^{tE}(\simeq j, t + 1)) = \frac{(1 - \delta) [q \check{u}_i^j + (1 - q) z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad \forall j \neq i \quad (\text{T.5.7})$$

^{T.6}See our Point of Notation T.2.1 above.

$$v_i^t(g, \mu | \Phi_i^{tE}(\bar{k})) = q \bar{v}_i^k + (1 - q) z_i \quad \forall k \in I \quad (\text{T.5.8})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\underline{k}, \tau)) = \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^\tau \right] [q \omega_i^k + (1 - q) z_i] + \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^\tau [q \bar{v}_i^k + (1 - q) z_i] \quad \forall k \in I \quad \forall \tau = \max\{T - t, 1\}, \dots, T \quad (\text{T.5.9})$$

where \hat{v}_i is as in (A.7), z_i is as in Remark A.4, v_i^* is as in the statement of the Theorem, \check{u}_i^k is as in (A.4), and ω_i^k is as in (A.3).

Proof: Equations (T.5.5), (T.5.6) and (T.5.7) are a direct consequence of Definition T.5.1 and Lemma T.5.1.

Equation (T.5.8) follows directly from Definition T.5.1 and the description of the profile (g, μ) in Definitions T.1.1 and T.1.2.

Using the notation established in Definition T.5.2, consider the quantity $\check{v}_i^t(\underline{m}^{k,\tau}, \dots, \underline{m}^{k,\tau})$. Given the strategies described in Definitions T.1.1 and T.1.2 it is evident that this quantity does not depend on t . Therefore, for any $k \in I$ and $\tau = \max\{T - t, 1\}, \dots, T$, we can let $\check{v}_i(\underline{k}, \tau) = \check{v}_i^t(\underline{m}^{k,\tau}, \dots, \underline{m}^{k,\tau})$, for all t . Clearly, using Definition T.5.1, we have that for all k, τ and t , $v_i^t(g, \mu | \Phi_i^{tE}(\underline{k}, \tau)) = \check{v}_i(\underline{k}, \tau)$.

From the description of (g, μ) in Definitions T.1.1 and T.1.2, for any $k \in I$ and for any $\tau = 2, \dots, T$, the quantity $\check{v}_i(\underline{k}, \tau)$ obeys a difference equation as follows.

$$\check{v}_i(\underline{k}, \tau) = q [(1 - \delta) \omega_i^k + \delta \check{v}_i(\underline{k}, \tau - 1)] + (1 - q) [(1 - \delta) z_i + \delta \check{v}_i(\underline{k}, \tau)] \quad (\text{T.5.10})$$

Using again Definitions T.1.1 and T.1.2, the terminal condition for (T.5.10) is

$$\check{v}_i(\underline{k}, 1) = q [(1 - \delta) \omega_i^k + \delta [q \bar{v}_i^k + (1 - q) z_i]] + (1 - q) [(1 - \delta) z_i + \delta \check{v}_i(\underline{k}, 1)] \quad (\text{T.5.11})$$

Solving (T.5.10) and imposing the terminal condition (T.5.11) now yields (T.5.9), as required. ■

Purely for expositional convenience, before completing the proof of sequential rationality at the message stage, we now proceed with the argument that establishes sequential rationality at the action stage.

Definition T.5.3: Recall that at the action stage, player $\langle i, t \rangle$ chooses an action after having received a message m_i^t and having observed a realization x^t of the randomization device \tilde{x}^t .

Let \mathcal{I}_i^{tB} denote period- t action-stage collection of information sets that belong to player $\langle i, t \rangle$, with typical element \mathcal{I}_i^{tB} . Clearly, each element of \mathcal{I}_i^{tB} is identified by a pair (m_i^t, x^t) .

It is convenient to partition \mathcal{I}_i^{tB} into mutually disjoint exhaustive subsets. The fact that they exhaust \mathcal{I}_i^{tB} can be checked directly from Definition T.4.1 above.

Let $\mathcal{I}_i^{tB}(\ast) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to (m^*, \dots, m^*) with probability one.^{T.7} These beliefs will be denoted by $\Phi_i^{tB}(\ast)$.

Let $\mathcal{I}_i^{tB}(\surd i) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to $(\check{m}^i, \dots, \check{m}^i)$ with probability one. These beliefs will be denoted by $\Phi_i^{tB}(\surd i)$.

For every $j \in I$ not equal to i , let $\mathcal{I}_i^{tB}(\surd j) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i-j}^t is equal to $(\check{m}^j, \dots, \check{m}^j)$ with probability one, that $\Pr(m_j^t = \underline{m}^{j,\tau}) > 0 \forall \underline{m}^{j,\tau} \in \underline{M}(j, t)$, and that $\Pr(m_j^t \in \underline{M}(j, t)) = 1$.^{T.8} These beliefs will be denoted by $\Phi_i^{tB}(\surd j)$.

^{T.7}In the interest of brevity, we avoid an explicit distinction between the $t = 0$ players and all others. What follows can be interpreted as applying to all players re-defining m_i^0 to be equal to m^* for players $\langle i \in I, 0 \rangle$.

^{T.8}See footnote T.3.

For every $j \in I$ not equal to i , and every $\tau = \max\{T-t+1, 1\}, \dots, T$ let $\mathcal{I}_i^{tB}(j, \tau) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to $(\underline{m}^{j, \tau}, \dots, \underline{m}^{j, \tau})$ with probability one. These beliefs will be denoted by $\Phi_i^{tB}(j, \tau)$.

For every $k \in I$, let $\mathcal{I}_i^{tB}(\bar{k}) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to $(\bar{m}^k, \dots, \bar{m}^k)$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\bar{k})$.

Lemma T.5.3: *There exists a $\delta \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ the action-stage strategies described in Definition T.1.1 are sequentially rational given the beliefs described in Definition T.4.1 for every player $\langle i, t \rangle$.*^{T.9}

Proof: Consider any information set $\mathcal{I}_i^{tB} \in \{\mathcal{I}_i^{tB}(\ast) \cup \mathcal{I}_i^{tB}(\simeq i) \cup \mathcal{I}_i^{tB}(\simeq j)\}$.^{T.10}

Using Definition T.1.1, Lemma T.5.2 and Definition T.5.3, it is immediate to check that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from following the action-stage strategies described in Definition T.1.1 at any of these information sets is

$$v_i^* = q\hat{v}_i + (1-q)z_i \quad (\text{T.5.12})$$

In the same way, it can be checked that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from deviating at any of these information sets is

$$q\bar{v}_i^i + (1-q)z_i \quad (\text{T.5.13})$$

Since by assumption $\hat{v}_i > \bar{v}_i^i$ this is of course sufficient to prove our claim for any information set $\mathcal{I}_i^{tB} \in \{\mathcal{I}_i^{tB}(\ast) \cup \mathcal{I}_i^{tB}(\simeq i) \cup \mathcal{I}_i^{tB}(\simeq j)\}$.

Now consider any information set \mathcal{I}_i^{tB} either in $\mathcal{I}_i^{tB}(j, \tau)$ or in $\mathcal{I}_i^{tB}(\bar{j})$ (with $j \neq i$).

Using Definition T.1.1, Lemma T.5.2 and Definition T.5.3, it is immediate to check that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from following the action-stage strategies described in Definition T.1.1 at any of these information sets is

$$q\bar{v}_i^j + (1-q)z_i \quad (\text{T.5.14})$$

In the same way, it can be checked that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from deviating at any of these information sets is exactly as in (T.5.13).

Since by assumption for any $j \neq i$ we have that $\bar{v}_i^j > \bar{v}_i^i$ this is of course sufficient to prove our claim for any of these information sets.

To conclude the proof of the lemma, we now consider any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(\bar{i})$. Using Definition T.1.1, Lemma T.5.2 and Definition T.5.3, it can be checked that the expected continuation payoff to player $\langle i, t \rangle$ from following the action-stage strategies described in Definition T.1.1 at any of these information sets is bounded below by

$$(1-\delta)\underline{u}_i + \delta [q\bar{v}_i^i + (1-q)z_i] \quad (\text{T.5.15})$$

^{T.9}It should be understood that we are, for now, taking it as given that each player $\langle i, t \rangle$ follows the prescriptions of the message-stage strategies described in Definition T.1.2. Of course, we have not demonstrated yet that this is in fact sequentially rational given the beliefs described in Definition T.4.2. We will come back to this immediately after the current lemma is proved.

^{T.10}See Definition T.5.3.

In the same way it can be readily seen that the expected continuation payoff to player $\langle i, t \rangle$ from deviating at any of these information sets is bounded above by

$$(1 - \delta)\bar{u}_i + \delta \left\{ \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] [q\omega_i^i + (1 - q)z_i] + \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T [q\bar{v}_i^i + (1 - q)z_i] \right\} \quad (\text{T.5.16})$$

The difference given by (T.5.15) minus (T.5.16) can be written as

$$(1 - \delta) \left\{ \frac{\delta q \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] (\bar{v}_i^i - \underline{\omega}_i^i)}{(1 - \delta)} - (\bar{u}_i - \underline{u}_i) \right\} \quad (\text{T.5.17})$$

Consider now the term inside the curly brackets in (T.5.17). We have that

$$\lim_{\delta \rightarrow 1} \frac{\delta q \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] (\bar{v}_i^i - \underline{\omega}_i^i)}{(1 - \delta)} - (\bar{u}_i - \underline{u}_i) = T(\bar{v}_i^i - \underline{\omega}_i^i) - (\bar{u}_i - \underline{u}_i) \quad (\text{T.5.18})$$

Using (A.11), we know that the quantity on the right-hand side of (T.5.18) is strictly positive. Hence we can conclude our claim is valid at any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(\bar{i})$. ■

Lemma T.5.4: *Consider the notation we established in Definition T.5.2. For any given t and $\tau = \max\{T - t, 1\}, \dots, T$ let $\check{v}_i^t(m, \underline{m}^{i,\tau})$ denote $\check{v}_i^t(m^{t+1})$ when the vector m^{t+1} has the i -th component equal to a generic $m \in M_i^{t+1}$ and $\underline{m}^{t+1} = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. As in the proof of Lemma T.5.2, let $\check{v}_i(i, \tau) = \check{v}_i^t(\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$.*

Then there exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ for every player $\langle i, t \rangle$, for every $m \in M_i^{t+1}$, and for every $\tau = \max\{T - t, 1\}, \dots, T$

$$\check{v}_i(i, \tau) \geq \check{v}_i^t(m, \underline{m}^{i,\tau}) \quad (\text{T.5.19})$$

Proof: We prove the claim for the case $t \geq T$. The treatment of $t < T$ has some completely non-essential complications due to the fact that the players' message spaces increase in size for the first T periods. The details are omitted for the sake of brevity.

We now introduce a new random variable \tilde{w} , independent of \tilde{x} and \tilde{y} (see Definitions A.7 and A.8), and uniformly distributed over the finite set $\{1, \dots, T\}$. This will be used in the rest of the proof of the lemma to keep track of the “private” randomization across messages that members of dynasty i may be required to perform (see Definition T.1.2). Just as we did for the action-stage and the message-stage randomization devices, we consider countably many independent “copies” of \tilde{w} , one for each time period, denoted by \tilde{w}^t , with typical realization w^t .

To keep track of all “future randomness” looking ahead for $t' = 1, 2, \dots$ periods from t , it will also be convenient to define the random vectors $\tilde{s}^{t,t'}$

$$\tilde{s}^{t,t'} = [(\tilde{x}^{t+1}, \tilde{y}^{t+1}, \tilde{w}^{t+1}), \dots, (\tilde{x}^{t+t'}, \tilde{y}^{t+t'}, \tilde{w}^{t+t'})] \quad (\text{T.5.20})$$

A typical realization of $\tilde{s}^{t,t'}$ will be denoted by $s^{t,t'} = [(x^{t+1}, y^{t+1}, w^{t+1}), \dots, (x^{t+t'}, y^{t+t'}, w^{t+t'})]$. The set of all possible realizations of $\tilde{s}^{t,t'}$ (which obviously does not depend on t) is denoted by $S^{t'}$.

Recall that the profile (g, μ) described in Definitions T.1.1 and T.1.2 is taken as given throughout. Now suppose that in period t , player $\langle i, t \rangle$ sends a generic message $m \in M_i^{t+1}$ and that $m_{-i}^{t+1} = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. Then, given any realization $s^{t,t'}$ we can compute the actual action profile played by all players $\langle k \in I, t+t' \rangle$. This will be denoted by $\mathbf{a}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})$. Similarly, we can compute the profile of messages $m_{-i}^{t+t'}$ received by all players $\langle j \neq i, t+t' \rangle$. This $n-1$ -tuple will be denoted by $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})$.

Recall that the messages received by all time- $t+t'$ players are the result of choices and random draws that take place on or before period $t+t'-1$. Therefore it is clear that if we are given two realizations $\hat{s}^{t,t'} = [s^{t,t'-1}, (\hat{x}^{t+t'}, \hat{y}^{t+t'}, \hat{w}^{t+t'})]$ and $\hat{\hat{s}}^{t,t'} = [s^{t,t'-1}, (\hat{\hat{x}}^{t+t'}, \hat{\hat{y}}^{t+t'}, \hat{\hat{w}}^{t+t'})]$, then it must be that

$$\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, \hat{s}^{t,t'}) = \mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, \hat{\hat{s}}^{t,t'}) \quad (\text{T.5.21})$$

Notice next that from the description of the profile (g, μ) in Definitions T.1.1 and T.1.2 it is also immediate to check that for any t' , any $m \in M_i^{t+1}$ and any realization $s^{t,t'}$ the message profile $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})$ can only take one out of two possible forms. Either we have $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'}) = (\bar{m}^i, \dots, \bar{m}^i)$ or it must be that $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'}) = (\underline{m}^{i,\tau'}, \dots, \underline{m}^{i,\tau'})$ for some $\tau' = 1, \dots, T$.

Lastly, notice that, given an arbitrary message $m \in M_i^{t+1}$ we can write

$$\ddot{v}_i^t(m, \underline{m}^{i,\tau}) = (1 - \delta) \sum_{t'=1}^{\infty} \delta^{t'-1} \sum_{s^{t,t'} \in S^{t'}} \Pr(\tilde{s}^{t,t'} = s^{t,t'}) u_i[\mathbf{a}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})] \quad (\text{T.5.22})$$

Since the strategies described in Definitions T.1.1 and T.1.2 are stationary for $t \geq T$, and the distribution of $\tilde{s}^{t,t'}$ is independent of t , it is evident from (T.5.22) that $\ddot{v}_i^t(m, \underline{m}^{i,\tau})$ does not depend on t . From now on we drop the superscript and write $\ddot{v}_i(m, \underline{m}^{i,\tau})$.

We now proceed with the proof of inequality (T.5.19) of the statement of the lemma. In order to do so, from now on we fix a particular $t = \hat{t}$, $m = \hat{m}$ and $\tau = \hat{\tau}$, and we prove (T.5.19) for these fixed values of t , m and τ . Since the lower bound on δ that we will find will clearly not depend on t , and since there are finitely many values that m and τ can take, this will be sufficient to prove the claim.

Inequality (T.5.19) in the statement of the lemma is trivially satisfied (as an equality) if $m = \underline{m}^{i,\tau}$. From now on assume that \hat{m} and $\hat{\tau}$ are such that $\hat{m} \neq \underline{m}^{i,\hat{\tau}}$.

Given any $t' = 1, 2, \dots$, we now partition the set of realizations $S^{t'}$ into five disjoint exhaustive subsets; $S_1^{t'}$, $S_2^{t'}$, $S_3^{t'}$, $S_4^{t'}$ and $S_5^{t'}$. This will allow us to decompose the right-hand side of (T.5.22) in a way that will make possible the comparison with (a similar decomposition of) the left-hand side of (T.5.19) as required to prove the lemma.

Let

$$S_1^{t'} = \{s^{\hat{t},t'} \mid \mathbf{m}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\underline{m}^{i,\tau'}, \dots, \underline{m}^{i,\tau'}) \text{ for some } \tau' = 1, \dots, \hat{\tau}\} \quad (\text{T.5.23})$$

and notice that if $t' \leq \hat{\tau}$ then $S_1^{t'} = S^{t'}$.

Assume now that $t' > \hat{\tau}$ and let

$$S_2^{t'} = \{s^{\hat{t},t'} \mid \mathbf{m}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i) \text{ and } u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \leq u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}))\} \quad (\text{T.5.24})$$

and

$$S_3^{t'} = \{s^{\hat{t},t'} \mid \mathbf{m}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i) \text{ and } u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) > u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}))\} \quad (\text{T.5.25})$$

Notice that if the first condition in (T.5.24) holds, then $\mathbf{m}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i)$. Therefore, $S_1^{t'}$ and $S_2^{t'}$ and $S_3^{t'}$ are disjoint.

Next, let any $s^{\hat{t},t''} \in S_3^{t''}$ with $t'' < t'$ be given and define

$$S_4^{t'}(s^{\hat{t},t''}) = \{s^{\hat{t},t'} \mid s^{\hat{t},t'} = (s^{\hat{t},t''}, s^{t'',t'}) \text{ for some } s^{t'',t'} \text{ and} \\ \|\{t \in (t'' + 1, \dots, t' - 1) \mid x^t = x(\kappa) \text{ with } \kappa \leq \bar{\kappa}\}\| \leq T - 1\} \quad (\text{T.5.26})$$

Now let

$$S_4^{t'} = \bigcup_{\substack{t' < t' \\ s^{\hat{t},t''} \in S_3^{t''}}} S_4^{t'}(s^{\hat{t},t''}) \quad (\text{T.5.27})$$

From the strategies described in Definitions T.1.1 and T.1.2 it can be checked that if $s^{\hat{t},t'} \in S_4^{t'}$ then $\mathbf{m}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\underline{m}^{i,\tau'}, \dots, \underline{m}^{i,\tau'})$ for some τ' and $\mathbf{m}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i)$. Therefore, it is clear that $S_4^{t'}$ is disjoint from $S_1^{t'}$, $S_2^{t'}$ and $S_3^{t'}$.

The last set in the partition of $S^{t'}$ is defined as the residual of the previous four.

$$S_5^{t'} = S^{t'} / \{S_1^{t'} \cup S_2^{t'} \cup S_3^{t'} \cup S_4^{t'}\} \quad (\text{T.5.28})$$

Using (T.5.22), we can now proceed to compare the two sides of inequality (T.5.19) of the statement of the lemma for the five distinct (conditional) cases $s^{\hat{t},t'} \in S_1^{t'}$ through $s^{\hat{t},t'} \in S_5^{t'}$. Notice first of all that when $s^{\hat{t},t'} \in S_2^{t'}$, we know immediately from (T.5.24) that there is nothing to prove.

We begin with $s^{\hat{t},t'} \in S_1^{t'}$. Notice first of all that if we fix any $\bar{s}^{\hat{t},t'} \in S_1^{t'}$, then it follows from (T.5.21) and (T.5.23) that any $s^{\hat{t},t'}$ of the form $[\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}]$ (where $\bar{s}^{\hat{t},t'-1}$ are the first $t' - 1$ triples of $\bar{s}^{\hat{t},t'}$) is in fact in $S_1^{t'}$.

Using, (T.5.23) and Definitions A.1, T.1.1 and T.1.2 we get

$$\sum_{s^{t'-1,t'} \in S^1} \Pr(\bar{s}^{t'-1,t'} = s^{t'-1,t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, [\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}])) = q\underline{\omega}_i^i + (1 - q)z_i \geq \\ \sum_{s^{t'-1,t'} \in S^1} \Pr(\bar{s}^{t'-1,t'} = s^{t'-1,t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, [\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}])) \quad (\text{T.5.29})$$

Therefore, since the $\bar{s}^{\hat{t},t'}$ that we fixed is an arbitrary element of $S_1^{t'}$, we can now conclude that

$$\sum_{s^{\hat{t},t'} \in S_1^{t'}} \Pr(\bar{s}^{\hat{t},t'} = s^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t},t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \geq \sum_{s^{\hat{t},t'} \in S_1^{t'}} \Pr(\bar{s}^{\hat{t},t'} = s^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t},t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \quad (\text{T.5.30})$$

Now fix any $\bar{s}^{\hat{t},t'} \in S_3^{t'}$. Using, (T.5.25), (T.5.26) and (T.5.27), and Definitions T.1.1 and T.1.2 we get that the difference given by

$$\Pr(\bar{s}^{\hat{t},t'} = \bar{s}^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, \bar{s}^{\hat{t},t'})) + \\ \sum_{t''=t'+1}^{\infty} \delta^{(t''-t')} \sum_{s^{\hat{t},t''} \in S_4^{t''}(\bar{s}^{\hat{t},t'})} \Pr(\bar{s}^{\hat{t},t''} = \bar{s}^{\hat{t},t''}) u_i(\mathbf{a}^{\hat{t}+t''}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, \bar{s}^{\hat{t},t''})) \quad (\text{T.5.31})$$

minus

$$\begin{aligned} & \Pr(\tilde{s}^{\hat{t},t'} = \bar{s}^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, \bar{s}^{\hat{t},t'})) + \\ & \sum_{t''=t'+1}^{\infty} \delta^{(t''-t')} \sum_{s^{\hat{t},t''} \in S_4^{t''}(\bar{s}^{\hat{t},t'})} \Pr(\tilde{s}^{\hat{t},t''} = \bar{s}^{\hat{t},t''}) u_i(\mathbf{a}^{\hat{t}+t''}(\hat{m}, \underline{m}^{i,\hat{\tau}}, \bar{s}^{\hat{t},t''})) \end{aligned} \quad (\text{T.5.32})$$

is greater or equal to

$$\Pr(\tilde{s}^{\hat{t},t'} = \bar{s}^{\hat{t},t'}) \left\{ \frac{\delta q \left[1 - \left(\frac{\delta q}{1 - \delta(1-q)} \right)^T \right] (\bar{v}_i^i - \underline{\omega}_i)}{(1-\delta)} - (\bar{u}_i - \underline{u}_i) \right\} \quad (\text{T.5.33})$$

Notice now that we know that the quantity in (T.5.33) is in fact positive for δ sufficiently close to 1. This is simply because the term in curly brackets in (T.5.33) is the same as the right-hand side of (T.5.18). Therefore, we have dealt with any $\bar{s}^{\hat{t},t'} \in S_3^{t'}$ and with all its relevant ‘‘successors’’ of the form $S_4^{t''}(\bar{s}^{\hat{t},t'})$. Since t' is arbitrary, by (T.5.27), this exhausts $S_3^{t'}$ and $S_4^{t'}$ for all possible values of t' .

Finally, we deal with $s^{\hat{t},t'} \in S_5^{t'}$. Notice first of all that if we fix any $\bar{s}^{\hat{t},t'} \in S_5^{t'}$, then it follows from (T.5.21) and (T.5.28) that any $s^{\hat{t},t'}$ of the form $[\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}]$ (where $\bar{s}^{\hat{t},t'-1}$ are the first $t' - 1$ triples of $\bar{s}^{\hat{t},t'}$) is in fact in $S_5^{t'}$.

Using, (T.5.28) and Definitions T.1.1 and T.1.2 we get

$$\begin{aligned} & \sum_{s^{t'-1,t'} \in S^1} \Pr(\tilde{s}^{t'-1,t'} = s^{t'-1,t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, [\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}])) = q\bar{v}_i^i + (1-q)z_i > \\ & q\underline{\omega}_i^i + (1-q)z_i \geq \sum_{s^{t'-1,t'} \in S^1} \Pr(\tilde{s}^{t'-1,t'} = s^{t'-1,t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, [\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}])) \end{aligned} \quad (\text{T.5.34})$$

Therefore, since the $\bar{s}^{\hat{t},t'}$ that we fixed is an arbitrary element of $S_5^{t'}$, we can now conclude that

$$\sum_{s^{\hat{t},t'} \in S_5^{t'}} \Pr(\tilde{s}^{\hat{t},t'} = s^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t},t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \geq \sum_{s^{\hat{t},t'} \in S_5^{t'}} \Pr(\tilde{s}^{\hat{t},t'} = s^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t},t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \quad (\text{T.5.35})$$

Hence, the proof of the lemma is now complete. ■

Remark T.5.1: Let the strategy profile (g, μ) described in Definitions T.1.1 and T.1.2 be given. Consider a player $\langle i, t \rangle$, and a realization of future uncertainty $s^{t,t'}$ as defined in the proof of Lemma T.5.4.

Let any message $m \in M_i^{t+1}$ be given, and fix any information set \mathcal{I}_i^{tE} and associated beliefs $\Phi_i^{tE}(\cdot)$.

It is then clear from Definitions T.1.1 and T.1.2 and T.5.1, that for any t' the action that player $\langle i, t \rangle$ expects player $\langle i, t + t' \rangle$ to take is uniquely determined by m , $s^{t,t'}$ and \mathcal{I}_i^{tE} .

For the rest of the argument we will denote this by $\mathbf{a}_i^{t+t'}(m, s^{t,t'}, \mathcal{I}_i^{tE})$.

Lemma T.5.5: There exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ the message-stage strategies described in Definition T.1.2 are sequentially rational given the beliefs described in Definition T.4.2 for every player $\langle i, t \rangle$.

Proof: Consider any information set $\mathcal{I}_i^{tE} \in \mathcal{I}_i^{tE}(\underline{i}, \tau)$, where $\mathcal{I}_i^{tE}(\underline{i}, \tau)$ is as in Definition T.5.1. It is then evident from Lemma T.5.4 and from the beliefs $\Phi_i^{tE}(\underline{i}, \tau)$ described in Definition T.5.1 that for δ sufficiently close to 1, the message strategies described in Definition T.1.2 are sequentially rational at any such information set.

From now on, consider any information set $\mathcal{I}_i^{tE} \notin \mathcal{I}_i^{tE}(\underline{i}, \tau)$. Let $m \in M_i^{t+1}$ be the message that player $\langle i, t \rangle$ should send according to the strategy μ_i^t , and let \hat{m} be any other message in M_i^{t+1} . Consider a particular realization $\bar{s}^{t,t'}$, and for any $t'' \in \{1, \dots, t' - 1\}$, let $\bar{s}^{t,t''}$ denote the first t'' triples of $\bar{s}^{t,t'}$.

Next, assume that $\mathbf{a}_i^{t+t'}(m, \bar{s}^{t,t'}, \mathcal{I}_i^{tE}) \neq \mathbf{a}_i^{t+t'}(\hat{m}, \bar{s}^{t,t'}, \mathcal{I}_i^{tE})$, and that either $t' = 1$, or alternatively that $\mathbf{a}_i^{t+t''}(m, \bar{s}^{t,t''}, \mathcal{I}_i^{tE}) = \mathbf{a}_i^{t+t''}(\hat{m}, \bar{s}^{t,t''}, \mathcal{I}_i^{tE})$ for every $t'' \in \{1, \dots, t' - 1\}$.

Clearly, in periods $\{t + 1, \dots, t' - 1\}$, conditional on $\bar{s}^{t,t'}$, the payoff to player $\langle i, t \rangle$ is unaffected by the deviation to \hat{m} . Now consider the payoff to player $\langle i, t \rangle$, conditional on $\bar{s}^{t,t'}$, from the beginning of period t' on, for simplicity discounted from the beginning of period t' . If player $\langle i, t \rangle$ sends message m as prescribed by μ_i^t , and δ is close enough to 1, the payoff in question is bounded below by

$$(1 - \delta)\underline{u}_i + \delta(q\bar{v}_i^i + (1 - q)z_i) \quad (\text{T.5.36})$$

Now consider the payoff to player $\langle i, t \rangle$ if he sends message \hat{m} , conditional on $\bar{s}^{t,t'}$, from the beginning of period t' on, for simplicity discounted from the beginning of period t' . In period t' the action played cannot yield him more than \bar{u}_i . From Lemma T.5.4, we know that, for δ close enough to 1, from the beginning of period $t' + 1$ the payoff is bounded above by $\bar{v}_i(\underline{i}, T)$. Hence, for δ close enough to 1, using (T.5.9) the payoff in question is bounded above by

$$\delta\bar{u}_i + (1 - \delta) \left\{ \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] [q\omega_i^i + (1 - q)z_i] + \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T [q\bar{v}_i^i + (1 - q)z_i] \right\} \quad (\text{T.5.37})$$

Notice now that the quantity in (T.5.36) is the same as the quantity in (T.5.15), and the quantity in (T.5.37) is in fact the same as the quantity in (T.5.16). Hence, exactly as in the proof of Lemma T.5.3, we know that, for δ sufficiently close to 1, the quantity in (T.5.36) is greater than the quantity in (T.5.37). This is clearly enough to conclude the proof. ■

T.6. Proof of Theorem A.1: Consistency of Beliefs

Remark T.6.1: Let $(g_\varepsilon, \mu_\varepsilon)$ be the completely mixed strategy profile of Definitions A.17 and A.18. It is then straightforward to check that as $\varepsilon \rightarrow 0$ the profile $(g_\varepsilon, \mu_\varepsilon)$ converges pointwise (in fact uniformly) to the equilibrium strategy profile described in Definitions T.1.1 and T.1.2, as required.

Lemma T.6.1: The strategy profile (g, μ) described in Definitions T.1.1 and T.1.2 and the beginning-of-period beliefs described in Definition T.4.1 are consistent.

Proof: When $t = 0$, there is nothing to prove. Assume $t \geq 1$. We consider two cases. First assume that player $\langle i, t \rangle$ receives message $m \in \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$. Clearly, this is on the equilibrium path generated by the profile of strategies (g, μ) described in Definitions T.1.1 and T.1.2. Therefore, consistency in this case simply requires checking that the beginning-of-period beliefs described in Definition T.4.1 are obtained via Bayes' rule from the profile (g, μ) . This is a routine exercise, and we omit the details.

Now assume that player $\langle i, t \rangle$ receives message $m \notin \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$. From Definition T.4.1 it is immediate to check that in this case player $\langle i, t \rangle$ assigns probability one to the event that $m_{-i}^t = (m, \dots, m)$. Given (g, μ) , this event may of course have been generated by several possible histories. Notice however, that the profile (g, μ) is such that a *single* deviation by one player at the action stage is sufficient to generate the message profile $m^t = (m, \dots, m)$. Therefore, upon observing $m \notin \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$ the probability that

$m_{-i}^t = (m, \dots, m)$ is an infinitesimal in ε of order no higher than 2. This needs to be compared with the probability that $m_{-i}^t \neq (m, \dots, m)$ and $m_i^t = m$. The latter event is impossible given the profile (g, μ) unless a deviation at the message stage has occurred at some point. Therefore its probability is an infinitesimal in ε of order no lower than $2n + 1$. This is obviously enough to prove the claim. ■

Lemma T.6.2: *The strategy profile (g, μ) described in Definitions T.1.1 and T.1.2 and the end-of-period beliefs described in Definition T.4.2 are consistent.*

Proof: The case $t = 0$ is trivial. Assume $t \geq 1$, and consider any player $\langle i, t \rangle$ after having observed (m_i^t, x^t, a^t, y^t) .

We deal first with the case in which $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$. Let $x(\ell_{00}, \ell^*)$ denote the realization x^t . In this case, the action-stage strategies described in Definition T.1.1 prescribe that every player $\langle k \in I, t \rangle$ should play $a_k^t(\ell^*)$. Therefore, if the observed action profile a^t is equal to $a(\ell^*)$, player $\langle i, t \rangle$ does not revise his beginning-of-period beliefs during period t . Hence consistency in this case follows immediately from the profile μ and from the consistency of beginning-of-period beliefs, which of course was proved in Lemma T.6.1. Notice now that if $a^t \neq a(\ell^*)$, then the message strategies described in Definition T.1.2 prescribe that each player $\langle k \in I, t \rangle$ should send a message that does not depend on the message m_k^t he received. Hence, in this case consistency is immediate from Definition T.4.2 and the profile μ .

We now turn to the case in which $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$. Here, it is necessary to consider several subcases, depending on the message m received by player $\langle i, t \rangle$. Assume first that $m \notin \check{M}_{-i} \cup \underline{M}(i, t)$. Then for any possible triple (x^t, a^t, y^t) we have that

$$\lim_{\varepsilon \rightarrow 0} \Pr(m_{-i}^t = (m, \dots, m) \mid m_i^t = m, x^t, a^t, g_\varepsilon, \mu_\varepsilon) = 1 \quad (\text{T.6.1})$$

To see this consider two sets of possibilities. First, $m = m^*$, $x^t = x(\cdot, \hat{\ell}, \dots)$, and $a^t = (a_1(\hat{\ell}), \dots, a_n(\hat{\ell}))$. Then play is as prescribed by the equilibrium path generated by the profile (g, μ) , and from Definitions T.1.1 and T.1.2 there is nothing more to prove. For all other possibilities, notice that the event $m^t = (m, \dots, m)$ is consistent with any a^t together with n deviations at the action stage of the second type described in Definition A.17. Therefore, for any a^t , the probability of $m^t = (m, \dots, m)$ and a^t is an infinitesimal in ε of order no higher than $2n$. On the other hand, from Definition A.18 it is immediate that the probability that $m_{-i}^t \neq (m, \dots, m)$ (since it requires at least one deviation at the message stage) is an infinitesimal in ε of order no lower than $2n + 1$. Hence (T.6.1) follows. From (T.6.1) it is a matter of routine to check the consistency of end-of-period beliefs from using the profile (g, μ) . We omit the details.

Still assuming that $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$, now consider the case $m = \check{m}_j \in \check{M}_{-i}$. In this case we can show that

$$\lim_{\varepsilon \rightarrow 0} \Pr(m_{-i-j}^t = (\check{m}_j, \dots, \check{m}_j) \text{ and } m_j^t \in \underline{M}(j, t) \mid m_i^t = \check{m}_j, x^t, a^t, g_\varepsilon, \mu_\varepsilon) = 1 \quad (\text{T.6.2})$$

using an argument completely analogous to the one we used for (T.6.1). The details are omitted. As in the previous case, from (T.6.2) it is a matter of routine to check the consistency of end-of-period beliefs from using the profile (g, μ) .

The last case remaining is $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ and $m = \underline{m}^{i, \tau}$. In this case we have that

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \Pr(m_{-i}^t = (\check{m}_i, \dots, \check{m}_i) \mid m_i^t = \underline{m}^{i, \tau}, x^t, a^t, g_\varepsilon, \mu_\varepsilon) + \\ \lim_{\varepsilon \rightarrow 0} \Pr(m_{-i}^t = (\underline{m}^{i, \tau}, \dots, \underline{m}^{i, \tau}) \mid m_i^t = \underline{m}^{i, \tau}, x^t, a^t, g_\varepsilon, \mu_\varepsilon) = 1 \end{aligned} \quad (\text{T.6.3})$$

Again, the argument is completely analogous to the one used for (T.6.1) and (T.6.2), and the details are omitted. Now take (T.6.3) as given and let $x^t = x(\dots, i_\ell, \dots)$.

Suppose next that $a_{-i}^t = \check{a}_{-i}^i(i_\ell)$. Then player $\langle i, t \rangle$ does not revise his beginning-of-period beliefs, and hence, using the profile μ and Lemma T.6.1 it is immediate to check that his end-of-period beliefs are consistent in this case.

Now suppose that for some $j \neq i$ we have that $a_j^t \neq \check{a}_j^i(i_\ell)$ and $a_{-i-j}^t = \check{a}_{-i-j}^i(i_\ell)$. Consistency of beliefs in this case requires showing that the first element in the sum in (T.6.3) is equal to 1. Of course given (T.6.3) it suffices to compare the probabilities of the two events $m_{-i}^t = (\check{m}_i, \dots, \check{m}_i)$ and $m_{-i}^t = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. The first is compatible with a single deviation at the action stage on the part of player $\langle j, t \rangle$. Therefore its probability is an infinitesimal in ε of order no higher than 2. The latter requires an action-stage deviation in some period $t' < t$ (order 2 in ε), and $n - 2$ action-stage deviations in period t (order 1 each). Hence, player $\langle i, t \rangle$ has consistent beliefs if he assigns probability 1 to $m_{-i}^t = (\check{m}_i, \dots, \check{m}_i)$. The consistency of his end-of-period beliefs can then be checked from the profile μ .

Now suppose that for some $j \neq i$ we have that $a_j^t \neq \check{a}_j^i(i_\ell)$ and $a_{-i-j}^t = a_{-i-j}^i(i_\ell)$. Consistency of beliefs in this case requires showing that the second element in the sum in (T.6.3) is equal to 1. Of course given (T.6.3) it suffices to compare the probabilities of the two events $m_{-i}^t = (\check{m}_i, \dots, \check{m}_i)$ and $m_{-i}^t = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. The first requires $(n - 2)$ deviations at the action-stage of period t , each of order 2 in ε . Since $n \geq 4$, this is therefore an infinitesimal in ε of order no lower than 4. The second is consistent with a deviation of order 2 in ε at the action-stage of some period $t' < t$, together with a deviation of order 1 in ε at the action stage of period t . Therefore its probability is an infinitesimal in ε of order no higher than 3. Hence, player $\langle i, t \rangle$ has consistent beliefs if he assigns probability 1 to $m_{-i}^t = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. The consistency of his end-of-period beliefs can then be checked from the profile μ . The same argument applies to show the consistency of his end-of-period beliefs when $a_{-i}^t = a_{-i}^i(i_\ell)$. We omit the details.

In all other possible cases for a^t , the messages sent by all players $\langle j \neq i, t \rangle$ do not in fact depend on a^t , provided that m_j^t is either \check{m}_i or $\underline{m}^{i,\tau}$. Given (T.6.3), the consistency of the end-of-period beliefs of player $\langle i, t \rangle$ can then be checked directly from the profile μ . ■

T.7. Proof of Theorem A.1

Given any $v^* \in \text{int}(V)$ and any $\delta \in (0, 1)$, using (A.10), (A.8) and the strategies and randomization devices described in Definitions A.7, A.8, T.1.1 and T.1.2 clearly implement the payoff vector v^* .

From Lemmas T.5.3 and T.5.5 we know that there exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ each strategy in the profile described in Definitions T.1.1 and T.1.2 is sequentially rational given the beliefs described in Definitions T.4.1 and T.4.2.

From Lemmas T.6.1 and T.6.2 we know that the strategy profile described in Definitions T.1.1 and T.1.2 and the beliefs described in Definitions T.4.1 and T.4.2 are consistent.

Hence, using Lemma T.3.1, the proof of Theorem A.1 is now complete. ■