

Social Memory, Evidence, and Conflict*

LUCA ANDERLINI
(Georgetown University)

DINO GERARDI
(Yale University)

ROGER LAGUNOFF
(Georgetown University)

June 2009

Abstract. This paper examines an equilibrium model of *social memory* — a society’s vicarious beliefs about its past. We show that incorrect social memory is a key ingredient in creating and perpetuating destructive conflicts.

We analyze an infinite-horizon model in which two countries face off each period in a game of conflict characterized by the possibility of mutually destructive “all out war” that yields catastrophic consequences for both sides. Each country is inhabited by a dynastic sequence of individuals. Each individual cares about future individuals in the same country, and can communicate with the next generation of their countrymen using *private messages*. Social memory is based on these messages, and on physical evidence — a sequence of imperfectly informative public signals of past behavior. We find that if the future is sufficiently important for all individuals, then *regardless of the precision of physical evidence from the past* there is an equilibrium in which the two countries engage in all out war with arbitrarily high frequency, an outcome that cannot arise in the standard repeated game. In our construction, each new generation “repeats the mistakes” of its predecessors, leading to an endless cycle of destructive behavior.

Surprisingly, we find that *degrading* the quality of information that individuals have about current decisions may “improve” social memory. This in turn ensures that arbitrarily frequent all out wars cannot occur.

JEL CLASSIFICATION: C72, C79, D80, D83, D89.

KEYWORDS: Social Memory, Private Communication, Dynastic Games, Physical Evidence.

*An older version of this paper circulated under the title: “Social Memory and Evidence from the Past.” The authors are grateful to Narayana Kocherlakota and three anonymous referees for their helpful comments and suggestions, and to the National Science Foundation for financial support (Grant SES-0617789). We also thank Helena Poon, Leeat Yariv, and seminar participants at Collegio Carlo Alberto [Turin], Georgetown, Penn, Princeton, Vanderbilt, and the CET Conference, the JPET Conference, and North American Econometric Society Meetings for the many helpful suggestions received on earlier drafts of the paper. All errors are our own.

1. Introduction

“In 1989, the Serbs commemorated their defeat at the hands of the Turks in the Battle of Blackbird Field, in 1389, and it formed the starting point for the Balkans wars of the 1990s.” — Baumeister and Hastings (1997).

The “memories” that induced such a destructive war were obviously not direct. Instead, they were shaped by oral or written accounts of history that were passed on through the generations. This paper develops a model in which destructive behavior can arise from this type of indirect memory, which social scientists in various fields refer to as *social memory*. One commonly accepted definition, given by Crumley (2002), describes social memory as “the means by which information is transmitted among individuals and groups and from one generation to another. Not necessarily aware that they are doing so, individuals pass on their behaviors and attitudes to others in various contexts but especially through emotional and practical ties and in relationships among generations [...]”

The definition is obviously neutral about how social memory is used.¹ This paper develops a model that shows how social memory, specifically *incorrect* social memory in the face of contrary physical evidence, is instrumental in creating and perpetuating destructive conflict.

Clearly, not all conflicts are universally destructive. Some yield clear winners and losers. Indeed, there is a long history of strategic models of war and conflict in this vein dating back at least to Schelling (1960). Some recent models, for instance Schwarz and Sonin (2007), Jackson and Morelli (2008), Chassang and i Miguel (2008), and Yared (2009), are rooted in the theory of repeated and/or dynamic games. Others such as Fearon (1995) and Glaeser (2005) place conflict in a signalling context.² A typical characteristic of the literature is that there may be positive incentives for conflict in some states of the world by at least one social stratum that stands to gain from what is an overall destructive conflict with another group or society.

Our interest in this paper, however, is in conflicts that are known to have no winners at all; conflicts that are so destructive that existing theories have trouble explaining them. Hence,

¹Fields in which social memory is commonly studied include anthropology, psychology, and distributed cognition. See, for instance, Cattell and Climo (2002), Connerton (1989), Crumley (2002), Fentress and Wickham (1992), Pennebaker, Paez, and Rimè (1997), Rogers (1997) and Sutton (2005).

²Overall, the literature on war and conflict is immense and impossible to do justice here. However, see Yared (2009) for an excellent bibliography. In Section 5 we discuss those models of conflict that, like ours, have some element of intergenerational memory or communication.

standard repeated games are less useful for our purposes since they only sustain “moderately bad” outcomes, i.e., outcomes above the stage game minmax payoff. In this sense, repeated game models of war afford each participant a modicum of self protection. Similarly, signalling games are limited in that they require agents to assign large likelihood to objective states of the world in which the conflict is desirable from at least one country’s point of view. Yet in some conflicts, World War I being a case in point, decision makers are often aware at the outset that the hostilities are extremely undesirable for all concerned. Foley (2005) makes this argument plainly.

Rather than dismissing these models, the present paper attempts to expand them in a way that can accommodate a role for social memory in conflicts. We therefore posit a model that abstracts away from scenarios in which there are positive incentives for conflict.

Consider a game faced by two nations each period in which all out conflict is catastrophically bad for all involved in a way that everyone understands. For concreteness, we call these nations \mathcal{F} (France) and \mathcal{G} (Germany). These countries are engaged in a repeated *Game of Conflict*. In a Game of Conflict, there are three types of “peaceful” profiles, and one exceedingly “destructive” action profile. The peace profiles come in three varieties. One that is good for both sides (“cooperation”), one that is bad for one side and good for the other (“domination” by one country over the other), and a third one in which domination by one side over the other is reversed. The “destructive” action profile takes both countries below their “individually rational” (minmax) payoff levels in the stage game. We refer to this profile as “all out war.” Our main result demonstrates the existence of an equilibrium in which all out war occurs with arbitrarily high frequency.

Clearly, if these countries are thought of as unitary actors with perfect memory, then the aforementioned equilibrium could not occur since it would take each nation’s long run payoff below its stage minmax. Instead, we formulate a generational model based loosely on Crumley’s notion of social memory. Specifically, each nation in this conflict is a “dynasty,” i.e., a placeholder for a sequence of individual decision makers, each of whom is finitely lived but cares about what happens to future generations who inhabit the same nation-dynasty. Families, tribes, and ethnic groups are all ongoing entities with similar features. We initially focus attention on a canonical case of one-period lived individuals, however we later extend the results to any demographic structure in which individual lifetimes are uniformly bounded.

At the end of each period, each “stand-in” decision maker observes the realized action

profile, then chooses what and how much of his information to pass on to his successor in the dynasty by way of a private message. Each new entrant has no direct memory of the past, but nevertheless forms a belief about it from two possible sources. One is the message about the past — the written or oral historiography of the dynasty — received from his predecessor. The other source is the physical evidence — history’s “footprint” — in the form of an informative but imperfect public signal of past events. Social memory is therefore created over time by a combination of evidence and intergenerational communication within dynastic nations.

Our main result demonstrates that if each individual decision maker is sufficiently concerned about the future, then there exists a sequential equilibrium in the dynastic Game of Conflict with the following properties. (1) All out war occurs with arbitrarily high frequency. (2) Physical evidence is ignored - i.e., neither beliefs nor actions condition on evidence from the past. (3) Social memory following counterfactual histories may be incorrect.

The equilibrium roughly works as follows. On path, behavior cycles deterministically between long periods of war and (relatively) short periods of peace.³ Along this path, each side’s long run payoff is below its stage minmax. This means that even though a self-protective action is available to each nation during the war epoch, it is not chosen. Instead, the equilibrium prescribes that the all out war profile is chosen. In order to rationalize this choice, each decision maker believes that peace-time will be a lot worse for his country if he *fails* to wage war today.

To understand how such beliefs can occur, consider a deviation by, say \mathcal{F} , in which \mathcal{F} fails to wage war during the war epoch. In that case the equilibrium prescribes a punishment continuation in which, during the peaceful period, the nations shift from a cooperative peace to a hegemonic one in which \mathcal{G} dominates \mathcal{F} . This asymmetric peace continues for a short time until the next war epoch, at which point a long phase of all out war resumes.

But why should \mathcal{F} follow this prescription when the self-protective (minmax) action is available? It is at this point where the possibility of incorrect social memory is critical. Recall that the asymmetric peace is off-path, i.e., it occurs only after an initial deviation by a decision maker in \mathcal{F} . Furthermore, knowledge of the deviation must be communicated to future decision makers in each country. So, after receiving this out-of-equilibrium message

³Although the equilibrium is deterministic, the logic extends to a “sunspot” equilibrium in which a war epoch is brought about by some exogenous randomization device.

from his predecessor, the future entrant in \mathcal{F} must weigh two possibilities. On the one hand, the predecessor's message may be in error — that is, no such action deviation took place. On the other, the message is correct. A deviation did occur, indicating an action error by a predecessor in \mathcal{F} . Each of these possibilities has different implications. If the former is an order of magnitude less likely than the latter, then the current decision maker in \mathcal{F} suffers from the illusion that it will not be dominated in the future peace period. In other words, the entrant will not anticipate his own dynasty's punishment. Moreover, \mathcal{F} will hold this belief even if the evidence points to the contrary. In short, while \mathcal{F} “forgets” that its dynasty has deviated, \mathcal{G} does not. Moreover, in the equilibria we construct below, even once the leader of country \mathcal{F} realizes his mistake, he is unable to relay this information to \mathcal{F} 's future leaders who are then doomed to make the same mistake.

A few clarifications may be helpful at this point. First, the equilibrium construction does not depend on anyone's failure to understand the consequences of war. Everyone in this model correctly anticipates that it will be horrific.⁴ Rather, what matters is the potential failure of future decision makers to understand the true nature of the error that leads to an off-path message. Because of this, current decision makers are lead to the inescapable conclusion *on path* that war is necessary. This logic and its important consequences do not hold in standard repeated games with infinitely-lived players.⁵

Second, the logic is similar to but generally different from a related result that also utilizes dynastic communication. In Anderlini, Gerardi, and Lagunoff (2008) (AGL), it is shown that for certain stage games, any (interior) feasible payoff vector of the stage game can be sustained by a sequential equilibrium in a dynastic game with intra-dynastic communication if the participants place sufficient weight on future generations' payoffs. However, critically, even the least restrictive AGL results apply only to three or more dynasties. By contrast, the present result covers the important, and perhaps more common, case of bilateral conflict. The difference is not just superficial: the construction in AGL makes use of the “large numbers” of participants in a way that cannot be adapted to the current environment.⁶ More importantly,

⁴The argument is therefore consistent with arguments often made, for instance, about World War I that both German and French leaders were aware in advance of the casualties that the Battle of Verdun would entail - see Foley (2005).

⁵Nor do they hold in overlapping generations games with full memory (Bhaskar, 1998, Kandori, 1992a, Salant, 1991, Smith, 1992, among others).

⁶There are two results in AGL, one requiring three or more players, and another which applies to a larger class of stage games requiring four or more players. The first result requires three or more so that unilateral deviations from the equilibrium path can be identified. This is standard in repeated games. The second result

the role of evidence and the question of whether it plays any role in swaying beliefs of the participants is not even considered in AGL. The present paper shows that incorrect social memory can be perpetual, even in the face of contrary physical evidence. The logic relies on the fact that messages can in principle convey more information than any imperfect physical evidence. This is because they are sent after the current action profile is *observed*. It turns out that this is sufficient to make viable equilibria in which physical evidence is ignored, but the messages convey the “wrong” information to future individuals.

To underscore this last point, we examine a variation of the main model in which information is further degraded so that present actions of the rival country are not directly observed. All dynastic members observe the same imperfect evidence about the rival. We show that, in any pure strategy equilibrium, messages would be useless in this case, and so participants would be forced to confront the evidence. The perpetual state of frequent all out war sketched above cannot arise in this case.

The rest of the paper is organized as follows. We first set up the model in Section 2. In Sections 3 and 4 we present our main results. Section 5 examines related literature. We compare our work to other potential explanations of destructive conflicts. Finally, Section 6 concludes. For ease of exposition, some technical material has been relegated to an Appendix.

2. The Baseline Model

There are two countries — France (\mathcal{F}) and Germany (\mathcal{G}) — that face off in an infinitely repeated game. We first describe the stage game and posit an explicit demographic structure that defines each nation-dynasty. We then define the strategies, beliefs, and Sequential Equilibrium of this game. Social memory will be shown to be a well defined notion arising naturally from a standard (though perhaps under-appreciated) aspect of Sequential Equilibrium.

2.1. A Dynastic Game of Conflict

Though the class of stage games to which our analysis applies is broad, we focus attention on the following symmetric 3×3 Game of Conflict. Each country has three feasible actions, $\{C, D, W\}$. Restricting attention momentarily to the actions C and D , the game is a version of the Prisoners’ dilemma. That is, the payoff, from the point of view of the row player is $\pi(D, C) > \pi(C, C) > \pi(D, D) > \pi(C, D)$. The third action (W) available to each country is

in AGL requires four or more players (dynasties) in order to identify a deviation from among those $n - 1$ players who carry out the punishment of one player.

interpreted as “war.” War is strictly dominated by strategy D for both players: $\pi(D, \cdot) > \pi(W, \cdot)$. Both countries choosing W is an “all out war.” An all out war is worse for both players than the outcome (D, D) : $\pi(D, D) > \pi(W, W)$. The two key properties embedded in these assumptions are that (1) all out war pushes both sides below their individually rational (minmax) payoffs; (2) as with the standard Prisoner’s dilemma, mutual cooperation CC is the unique, symmetric efficient payoff. A numerical example of this game is given in (1) below.

	C	D	W	
C	2, 2	−1, 3	−25, 1	
D	3, −1	0, 0	−8, −5	(1)
W	1, −25	−5, −8	−10, −10	

Notice that in (1), if \mathcal{F} and \mathcal{G} were to face off indefinitely and were modeled as standard long-run players with perfect recall, then “very frequent” all out wars could not take place. This is because, as all out war occurs more and more frequently, the long-run average payoff of both countries approaches -10 . But both \mathcal{F} and \mathcal{G} can guarantee a payoff of at least -8 in every period by unilaterally choosing D .

Our results derive from the assumption that each of the two countries \mathcal{F} and \mathcal{G} identifies a *dynasty*. A dynasty is a placeholder for successive generations of individual decision makers. At any given time, the dynasty is inhabited by one such decision maker who cares about his own payoff and those of future generations within his own dynasty.

For ease of exposition we first restrict attention to a baseline model in which each individual in each dynasty lives one period. This simplifies the notation considerably since each period t can then be identified with a unique cohort. In Section 6, we discuss an extension to all games in which there is a uniform upper bound L on the length of life of all individuals in the model.

The individual decision makers who inhabit \mathcal{F} and \mathcal{G} in period t are denoted by \mathcal{F}^t and \mathcal{G}^t respectively. Let $a^t = (a_{\mathcal{F}}^t, a_{\mathcal{G}}^t)$ denote the action profile in period t , and the two countries’ per-period payoffs by $\pi_{\mathcal{F}}(a^t)$ and $\pi_{\mathcal{G}}(a^t)$, respectively.

Individuals care not only about their per-period payoffs, but about the long-run future payoff of their country as well. Note that since each country is populated by separate in-

dividuals through time, the degree by which future payoffs are taken into account can be interpreted as the degree of “altruism” that individuals exhibit towards future individuals in their own country. This is modeled using standard geometric weights, so that given a stream of per-period payoffs, the time t continuation payoffs to \mathcal{F}^t and \mathcal{G}^t , including both the present and the future components can be written as

$$\Pi_{\mathcal{F}}^t = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \pi_{\mathcal{F}}(a^{\tau}) \quad \text{and} \quad \Pi_{\mathcal{G}}^t = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \pi_{\mathcal{G}}(a^{\tau}) \quad (2)$$

In the dynastic game what each individual observes about the past history of play when he enters the game is critical. If the individuals \mathcal{F}^t and \mathcal{G}^t could only observe the actual history of action profiles, denoted by $h^t = (a^0, a^1, \dots, a^{t-1}) \in H^t$, then the model would be equivalent to the standard repeated game world of, say, Fudenberg and Maskin (1986);⁷ incorrect social memory would be ruled out by assumption, and no frequent all out wars could take place. Instead, when an individual (\mathcal{F}^t or \mathcal{G}^t) enters the game, he observes two things (in no particular order): (i) imperfect publicly available evidence of the past history of actions and (ii) a *private* message about the previous history from his predecessor, who was alive at $t - 1$ in his own country and whom he replaces.

Regarding (i) the publicly available evidence, its precise form does not matter for the results, however, for concreteness, it is in the form of a history of imperfect, but informative signals of the action history. Formally, let $s^t = (p^0, p^1, \dots, p^{t-1})$ denote the evidence observed by individuals \mathcal{F}^t and \mathcal{G}^t at the beginning of t . At the end of each period t , a signal p^t is conditionally realized given action profile a^t . Again, the parametric form is not important, but to fix ideas, let $p^t = a^t$ with probability $(1 + 8\gamma)/9$ (there are 9 possible action profiles), and is equal to any $\hat{a}^t \neq a^t$ with probability $(1 - \gamma)/9$. Here, γ parameterizes the precision of s^t ; when $\gamma = 0$ the observed evidence s^t contains no information about the previous history, and when $\gamma = 1$ the observed evidence s^t equals the true history of action profiles with probability one.

Note that we take the per-period signals p^t to be realized *once and for all* in each period t , in the sense that, for instance, all individuals from the end of $t = 0$ onwards, observe the *same* realized signal about what happened in period $t = 0$. Once this (imprecise) historical

⁷This is an immediate consequence of the “one-shot deviation principle.” See for instance the textbooks by Fudenberg and Tirole (1991) or Osborne and Rubinstein (1994).

footprint is set, it remains as given through time.

Regarding (ii) the intra-dynastic message, let m_i^t denote the private message received by the individual in dynasty $i = \mathcal{F}, \mathcal{G}$ at the start of period t . To start things off, m_i^0 is the null message, and for $t \geq 1$, message m_i^t is sent by the individual in dynasty i in cohort $t - 1$ to his dynastic successor in cohort t . The message space M_i^t may be, but need not be, the set of action histories h^t .⁸

The timing within each period is as follows. An individual from dynasty $i = \mathcal{F}, \mathcal{G}$, first observes the evidence s^t and his private message m_i^t . He then selects an action a_i^t . Then observes the action a_j^t , $j \neq i$ selected by the opposing individual at time t and subsequently, the updated evidence s^{t+1} . Once period t is over, the individual sends a private message m_i^{t+1} to his successor in the dynasty.

To summarize, in each period each country is inhabited by a representative one-period lived decision maker who cares about his own and future countrymen's payoffs. Upon entry, an individual observes an imprecise footprint of what took place prior to his entry. He also observes a private message concerning the past left by his predecessor. While alive, he observes directly the other country's decisions and how current action profiles generate a footprint for the future. Finally, he sends a message to his successor just before his own exit.

2.2. Equilibrium

In every period t , the t -action strategies determine the actions of individuals \mathcal{F}^t and \mathcal{G}^t as a function of what they have observed at the beginning of the period. Formally, for each dynasty $i = \mathcal{F}, \mathcal{G}$, the t -action strategy is expressed as $\alpha_i^t(m_i^t, s^t) = a_i^t$.⁹ The individual's t -message strategy determines the messages sent to his successor individual at the end of the period. His t -message strategy is expressed as $\mu_i^t(m_i^t, s^{t+1}, a^t) = m_i^{t+1}$. Taken together, α_i^t and μ_i^t make up the individual's full t -strategy $\mathbf{f}_i^t = (\alpha_i^t, \mu_i^t)$. A strategy profile \mathbf{f} refers to the array of all individuals in both countries in all periods.

One can distinguish between *beginning-of-period- t beliefs* and *end-of-period- t beliefs* of individuals \mathcal{F}^t and \mathcal{G}^t . For an individual t in dynasty i , the beginning-of-period belief

⁸Because our main results are constructive, there is considerable latitude in the possible choices of message spaces. The ones considered here seem the natural ones in many ways. All our results hold unchanged for arbitrary "sufficiently rich" message spaces.

⁹The main results are constructive, utilizing only pure strategies. Consequently, we restrict attention to pure strategies in the formal definitions.

$\mathbf{b}_i^t(m_i^t, s^t)$ is a probability distribution over the entire set of past action and message histories in $H^t \times (\times_{\tau=0}^t (M_{\mathcal{F}}^\tau \times M_{\mathcal{G}}^\tau))$, given his inherited message m_i^t and evidence s^t . The end-of-period belief $\mathbf{b}_i^t(m_i^t, s^{t+1}, a^t)$ describes conditional probability distribution over the set of end of period action and message histories in $H^{t+1} \times (\times_{\tau=0}^{t+1} (M_{\mathcal{F}}^\tau \times M_{\mathcal{G}}^\tau))$ given one's message, the evidence, and also the action profile taken in date t . As is standard, a *system of beliefs* \mathbf{b} refers to an array of beliefs of all individuals in all periods.

We refer to the pair (\mathbf{f}, \mathbf{b}) as an *equilibrium* of the model if and only if it constitutes a Sequential Equilibrium (Kreps and Wilson, 1982) of the dynastic game. Sequential Equilibrium is generally accepted as one of the (if not *the*) benchmark equilibrium concepts for dynamic games of incomplete information. Sequential Equilibrium requires that strategies of all individuals must be sequentially rational given the equilibrium beliefs. Just as importantly, it also requires that equilibrium beliefs \mathbf{b} must be recoverable as the limit of beliefs entirely determined by Bayes' rule using "fully mixed" perturbations of the equilibrium strategies. This is more than a technical condition. Intuitively, it means that all individuals share a complete philosophy about how other individuals' mistakes or trembles rationalize off-path events. Section 3 below demonstrates how the theory makes critical use of this "shared philosophy" to generate frequent all out wars.

3. Frequent All Out Wars: A Result

Our first and main result is that it is possible in equilibrium for the two countries \mathcal{F} and \mathcal{G} to be caught in a perpetual cycle of frequent all out wars.

Proposition 1. *Arbitrarily Frequent All Out Wars:* Fix any positive integer N . Then there exists a pair (\mathbf{f}, \mathbf{b}) which: (i) is an equilibrium for any δ sufficiently close to 1 and for any precision level $\gamma \in [0, 1)$, and (ii) is such that the action profile (W, W) is played $N - 1$ times out of any N consecutive periods.

Proposition 1 asserts that catastrophic conflict resulting in below-minmax long run payoffs can occur in equilibrium. The argument, given in three basic steps, follows along the lines of the intuition outlined in the Introduction. The rough idea is that a self-protective (minmax) action may not be taken by the current decision maker after a prior deviation by his dynasty if his belief about whether the deviation occurred is incorrect. Our task in the proof is to construct an equilibrium strategy profile $\mathbf{f} = (\boldsymbol{\alpha}, \boldsymbol{\mu})$ in which this is the case. This is done in Step 1. The external appearance of the equilibrium in Step 1 is standard. Namely, there are

three phases: an “equilibrium” phase, and a “punishment phase” for each $i = \mathcal{F}, \mathcal{G}$ following a unilateral deviation that came most recently from within dynasty i . However, because the equilibrium phase results in below-minmax long run payoffs, the individuals’ beliefs following off-path messages are critical. Step 2 describes these beliefs \mathbf{b} and establishes the consistency property, namely that these beliefs are obtained as the limit from a sequence of completely mixed trembles on the strategy profile. Finally, Step 3 establishes sequential rationality of the equilibrium, i.e., deviations are not profitable after any history for δ close to one.

Step 1. Construction of the equilibrium profile $\mathbf{f} = (\boldsymbol{\alpha}, \boldsymbol{\mu})$. Fix N as in Proposition 1. Before proceeding with the formal definition, the action strategy profile $\boldsymbol{\alpha}$ is first described heuristically in terms of its continuation phases. These are depicted in a schematic way in Figure 1. The determination of which phase the players are in depends on the messages to be specified shortly. The “equilibrium” phase consists of an infinite repetition of an N -period cycle. In each such cycle, the first $N - 1$ periods are classified as “war periods.” These war periods are followed by a single “peace period.” The cycle then repeats *ad infinitum*.¹⁰ A “punishment phase” for i consists of T rounds (with T finite) of a similar type of cycle as in the equilibrium phase. As in the equilibrium phase, each cycle lasts N periods with the first $N - 1$ periods classified as “war periods.” The difference here is that the prescribed behavior in the “peace period” that follows is different than in the peace periods of the equilibrium phase. We now describe exactly “who does what” in each phase.

According to Figure 1, in every war period regardless of phase, the profile WW is chosen. The two countries start off in the *equilibrium phase* during which the cooperative action pair CC is chosen in peace periods. They remain in this phase unless a “deviation” (yet to be described) occurs. If individual \mathcal{F}^t deviates from the prescriptions of the equilibrium, then the *\mathcal{F} -punishment phase* starts (or re-starts, as appropriate) at which time the pair CD (recall that \mathcal{F} is the row player) is chosen in the peace periods. Similarly, any deviation by any individual \mathcal{G}^t starts (or re-starts, as appropriate) the *\mathcal{G} -punishment phase* in which DC is chosen in the peace periods. Double deviations are ignored. Each of these phases, if no subsequent deviations occur, lasts until the next T *peace periods* have passed, after which time the two countries return to the equilibrium phase.

¹⁰Since the first period is $t = 0$, the first peace period is $t = N - 1$. All periods with $t = \ell N - 1$ with ℓ a positive integer greater than 1 are also peace periods. All other periods are war periods.

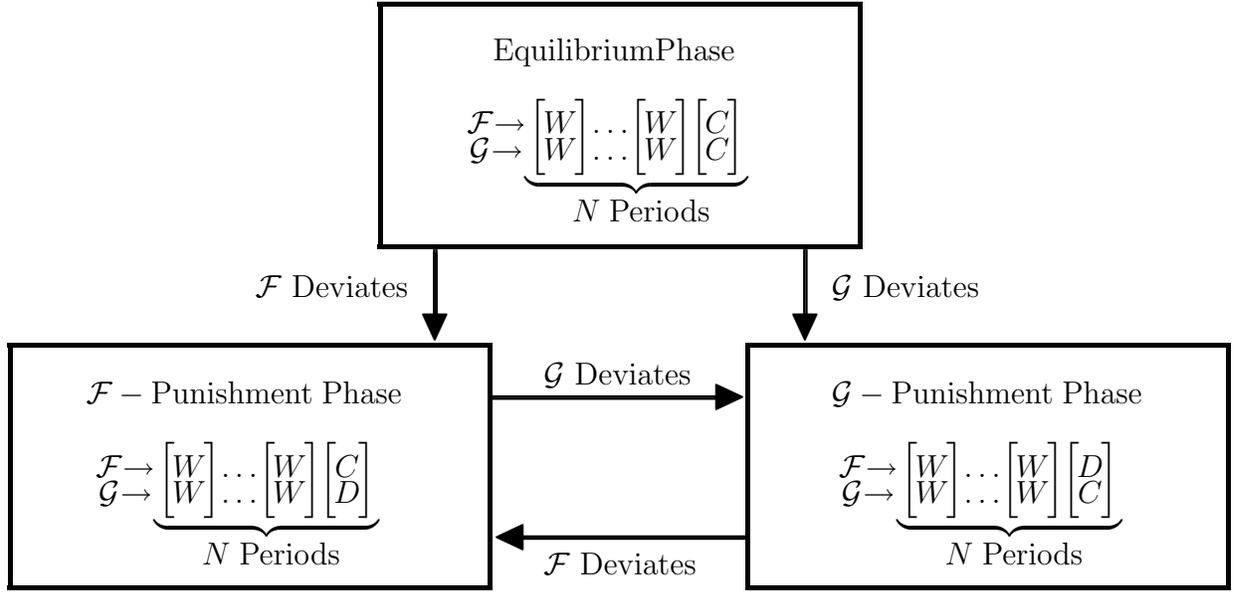


Figure 1

The formal definitions of action and message strategies are as follows. We start by defining the message space. For all periods $t \geq TN - 1$, the message space is the same for each dynasty $i = \mathcal{F}, \mathcal{G}$ and is given by

$$M_i = \{m^*, m^{\mathcal{F}, \tau}, m^{\mathcal{G}, \tau}\}_{\tau=1, \dots, T}$$

At the beginning of each period, messages in M_i determine the place in each phase. The message m^* is the equilibrium phase message. The message $m^{\mathcal{F}, \tau}$ indicates an \mathcal{F} -punishment phase message in which there are τ cycles remaining in the punishment phase. Intuitively, the index $\tau = T, T - 1, \dots, 1$ is used to “count down” the remaining peace periods that must pass before the \mathcal{F} punishment phase ends.¹¹ A symmetric description applies to $m^{\mathcal{G}, \tau}$. For the first $NT - 1$ periods these message spaces have superfluous messages because not enough periods have elapsed since the game started at $t = 0$. Consequently, one would have to “shrink” the message spaces for these periods in order to have only meaningful messages. We omit the details.

¹¹Notice that the messages do not track the *place* in each cycle, i.e., q periods left in the τ th cycle. This is because calendar time is common knowledge, and the calendar clock can always be used to distinguish war periods from peace periods. The assumption of common knowledge of calendar time can easily be dropped, however, an added notational burden to the message space would be then required — one that adds very little insight to the analysis.

In all that follows, we define the strategy for \mathcal{F} with the understanding that \mathcal{G} 's strategy is completely symmetric. The prescriptions of $\alpha_{\mathcal{F}}$ and $\mu_{\mathcal{F}}$ will not vary with the signal s^t in any way.

The action strategy $\alpha_{\mathcal{F}}$ is defined as follows. In all war periods $t > 0$, and for any message $m_{\mathcal{F}}^t$ and signal realization s^t , $\alpha_{\mathcal{F}}^t(m_{\mathcal{F}}^t, s^t) = W$. In words, the prescribed action in all war periods in all phases is the same. Namely, $\alpha_{\mathcal{F}}$ prescribes W in every war period, *regardless* of the phase. Hence, the only difference between on-path and off-path phases is in the prescribed peace-time behavior.

Here and throughout the rest, we treat the initial null message $m_{\mathcal{F}}^0$ as equal to m^* . Next, for any peace period $t \geq 0$ let

$$\alpha_{\mathcal{F}}^t(m_{\mathcal{F}}^t, s^t) = \begin{cases} C & \text{if } m_{\mathcal{F}}^t = m^* \text{ or } m^{\mathcal{F},\tau}, \text{ for any } \tau \\ D & \text{if } m_{\mathcal{F}}^t = m^{\mathcal{G},\tau}, \text{ for any } \tau \end{cases}$$

Hence, $\alpha_{\mathcal{F}}$ prescribes that in peace periods, the individual chooses C in both the equilibrium phase and the \mathcal{F} -punishment phase. He chooses D in the \mathcal{G} -punishment phase.

We now define the message strategy $\mu_{\mathcal{F}}$ for dynasty \mathcal{F} as follows. Intuitively, the messages “truthfully” indicate to the next player how and when to switch between phases. First, let $t > 0$ be any war period. Then for any message-signal pair, $(m_{\mathcal{F}}^t, s^{t+1})$,

$$\mu_{\mathcal{F}}^t(m_{\mathcal{F}}^t, s^{t+1}, a^t) = \begin{cases} m_{\mathcal{F}}^{t+1} = m_{\mathcal{F}}^t & \text{if } a^t = WW \text{ or if bilateral deviation from } WW \\ m^{\mathcal{F},T} & \text{if } \mathcal{F}^t \text{ unilaterally deviates from } WW \\ m^{\mathcal{G},T} & \text{if } \mathcal{G}^t \text{ unilaterally deviates from } WW \end{cases}$$

In war periods, \mathcal{F}^t passes on the existing message if WW is observed (or if both individuals deviate from W). Otherwise, \mathcal{F}^t 's message indicates the beginning of a punishment phase for a deviator.

Next suppose that $t \geq 0$ is a peace period. For any realized signal s^{t+1} , the message strategy is described in each of the three phases as follows. First, in the equilibrium phase we have

$$\mu_{\mathcal{F}}^t(m^*, s^{t+1}, a^t) = \begin{cases} m^{\mathcal{F},T} & \text{if } a^t = DC \text{ or } WC \\ m^{\mathcal{G},T} & \text{if } a^t = CD \text{ or } CW \\ m^* & \text{otherwise} \end{cases}$$

In words, an individual who received m^* sends message $m^{\mathcal{F},T}$ (resp., $m^{\mathcal{G},T}$) to his successor if a unilateral action deviation by \mathcal{F}^t (resp., \mathcal{G}^t) is observed during a peace period. Next, in the \mathcal{G} -punishment phase,

$$\mu_{\mathcal{F}}^t(m^{\mathcal{G},\tau}, s^{t+1}, a^t) = \begin{cases} m^{\mathcal{G},\tau-1} & \text{if } a^t = DC \text{ and } \tau > 1 \\ m^{\mathcal{G},T} & \text{if } a^t = DD, DW, CD, \text{ or } CW \\ m^{\mathcal{F},T} & \text{if } a^t = WC, \\ m^* & \text{otherwise} \end{cases}$$

During a \mathcal{G} -punishment phase the individual \mathcal{F}^t who received $m^{\mathcal{G},\tau}$ sends message $m^{\mathcal{G},\tau-1}$ to his successor if no action deviation is observed in peace period t . If $\tau = 1$ then the \mathcal{G} -punishment phase is ending, and so he sends m^* . Action deviations simply (re)start a punishment phase. Finally, in the \mathcal{F} -punishment phase, \mathcal{F}^t 's strategy is symmetric to that in the \mathcal{G} -phase.

$$\mu_{\mathcal{F}}^t(m^{\mathcal{F},\tau}, s^{t+1}, a^t) = \begin{cases} m^{\mathcal{F},\tau-1} & \text{if } a^t = CD \text{ and } \tau > 1 \\ m^{\mathcal{F},T} & \text{if } a^t = DD, WD, DC, \text{ or } WC \\ m^{\mathcal{G},T} & \text{if } a^t = CW \\ m^* & \text{otherwise} \end{cases}$$

Step 2. Construction of beliefs \mathbf{b} . Recall that we defined the beginning-of-period belief $\mathbf{b}_{\mathcal{F}}^t(m_{\mathcal{F}}^t, s^t)$ of \mathcal{F}^t as a conditional probability distribution over the entire history of messages and actions (the beliefs for \mathcal{G}^t have the same domain). As a point of notation, since the beliefs we construct below will not depend on the signal realization s^t in any way, we suppress it in the notation and express the above belief by $\mathbf{b}_{\mathcal{F}}^t(m_{\mathcal{F}}^t)$. The only beliefs that matter after any history are \mathcal{F}^t 's belief about the *message received by \mathcal{G}^t given that \mathcal{F}^t receives message $m_{\mathcal{F}}^t$ at the beginning of the period.* This marginal distribution over \mathcal{G}^t 's received messages is denoted by $\bar{\mathbf{b}}_{\mathcal{F}}^t(m_{\mathcal{F}}^t)$. Abusing notation, we write $\bar{\mathbf{b}}_{\mathcal{F}}^t(m_{\mathcal{F}}^t) = m$ to mean that “ \mathcal{F}^t believes with probability one that his rival received message m .”

Beginning-of-period beliefs are given by: $\bar{\mathbf{b}}_i^t(m^*) = m^*$ in the equilibrium phase; $\bar{\mathbf{b}}_{\mathcal{F}}^t(m^{\mathcal{G},\tau}) = m^{\mathcal{G},\tau}$ in the \mathcal{G} -punishment phase; and $\bar{\mathbf{b}}_{\mathcal{F}}^t(m^{\mathcal{F},\tau}) = m^*$ in the \mathcal{F} -punishment phase. In other words, the individual \mathcal{F}^t believes his rival receives the same message in either the equilibrium or \mathcal{G} -punishment phase. However, in the \mathcal{F} -punishment phase, \mathcal{F}^t believes his rival receives

m^* instead of $m^{\mathcal{F},\tau}$. To understand why this can be the case, we refer the reader to the Consistency argument below.

Similarly, let $\bar{\mathbf{b}}_{\mathcal{F}}^t(m_{\mathcal{F}}^t, a^t)$ denote the updated belief at the end of the period of \mathcal{F}^t about the message his rival *received* at the beginning of the period, given that \mathcal{F}^t received message $m_{\mathcal{F}}^t$ and observed action profile a^t . If t is peace period, then $\bar{\mathbf{b}}_{\mathcal{F}}^t(m^{\mathcal{F},\tau}, a_{\mathcal{F}}^t, a_{\mathcal{G}}^t) = m^*$ for any action $a_{\mathcal{F}}^t$ and any action $a_{\mathcal{G}}^t = C, W$. In all other cases, i.e., for all other pairs $(m_{\mathcal{F}}^t, a^t)$, $\bar{\mathbf{b}}_{\mathcal{F}}^t(m_{\mathcal{F}}^t, a^t) = m_{\mathcal{F}}^t$. If t is war period, then $\bar{\mathbf{b}}_{\mathcal{F}}^t(m_{\mathcal{F}}^t, a^t) = m_{\mathcal{F}}^t$ if $m_{\mathcal{F}}^t = m^*, m^{\mathcal{G},\tau}$. However, if $m_{\mathcal{F}}^t = m^{\mathcal{F},\tau}$ then $\bar{\mathbf{b}}_{\mathcal{F}}^t(m^{\mathcal{F},\tau}, a^t) = m^*$.

Consistency. Consistency of all these beliefs is established as follows. Recall, first, that consistency entails that beliefs are obtained in the limit of a sequence of completely mixed strategies defined by “trembles” on the prescribed strategy profile \mathbf{f} . We first describe the trembles formally and then give an intuitive account of their meaning.

Formally, the trembles are all expressed in terms of a common variable ε to be shrunk to zero as follows. The trembles at the two different action and message stages remain of course distinct. The difference between them is identified by the power to which ε is raised in each case. Whenever the equilibrium prescribes an action different from D , then D is actually played with probability ε^2 . Whenever the equilibrium prescribes an action different from W in period t , W is actually played with probability $\varepsilon^{\frac{t+3}{t+2}}$. Whenever the equilibrium prescribes an action different from C in period t , C is actually played with probability $\varepsilon^{\frac{t+3}{t+2}}$. The probability that any \mathcal{F}^t (resp \mathcal{G}^t) sends any message $m^{\mathcal{F},\tau}$ (resp $m^{\mathcal{G},\tau}$) when instead m^* is prescribed by the equilibrium strategies is ε . The probability of all other messages being sent instead of any other equilibrium ones is ε^4 .

Intuitively, the most common interpretation, one that we adopt here, is that trembles occur due to small errors at either the action stage or the message stage. Hence, consider an individual \mathcal{F}^t who receives an off-the-equilibrium-path message m , at the beginning of the period. He must weigh several (infinitesimally unlikely) possibilities. He could have received m when play is in fact in the equilibrium phase and individual \mathcal{F}^{t-1} mistakenly sent message m instead of m^* . Alternatively, he could have received m because some previous individual mistakenly took the wrong action during the equilibrium phase. Given the trembles described above, if a \mathcal{G} -punishment phase message $m^{\mathcal{G},\tau}$ is observed, then individual \mathcal{F}^t 's beliefs is the limit of a sequence where a message error is of order strictly less than that of an action error. In the limit, \mathcal{F}^t therefore believes (with probability one) that message $m^{\mathcal{G},\tau}$ correctly

reveals that play is in the \mathcal{G} -punishment phase. However, if $m^{\mathcal{F},\tau}$ is observed, then \mathcal{F}^t 's beliefs is the limit of a sequence where a message error is of order strictly larger than that of an action error. Specifically, he believes that his predecessor erred in sending $m^{\mathcal{F},\tau}$ and that the equilibrium phase is, in fact, the true phase.¹²

A similar reasoning applies to end-of-period beliefs. If at the end of the period and given his received message, individual \mathcal{F}^t observes an unexpected action by his rival then in all circumstances except one, he attributes it to an action error by \mathcal{G}^t , rather than a message error by a predecessor. The one exception is where he observes the unexpected action D by his rival in any peace period of the \mathcal{F} -phase. In that case, and recalling his beginning-of-period beliefs in that phase, he learns that $m^{\mathcal{F},\tau}$ was not, in fact, a mistaken message as he had earlier believed. For if $m^{\mathcal{F},\tau}$ was, in fact, a mistake, then \mathcal{F}^t should have observed action C played by \mathcal{G}^t . Since instead \mathcal{G}^t played D , and the trembles we use imply that it is more likely that action D is an equilibrium response by \mathcal{G}^t rather than a further deviation, he should not believe that his predecessor's message was wrong.

Notice that in the construction, beliefs are always updated from trembles of different orders. Certain off-path actions are infinitesimals of lower or higher order than certain messages, depending on the conditioning event. For this reason, the beliefs of all individuals \mathcal{F}^t and \mathcal{G}^t are *degenerate* in the sense that they always assign probability one to a particular past history of messages and actions, both on and off the equilibrium path. Since the beliefs are degenerate in this sense while, at the same time, the per period signals have full support, it is easy to see that the belief system \mathbf{b} does not depend on signal realizations in any way.

Step 3. Incentives. From the constructions in Steps 1 and 2, observe that neither beliefs nor strategies depend on signal realizations (i.e., the evidence) in any way. Hence, in what follows, we omit the signals from the notation and consider the incentives of an individual \mathcal{F}^t . The argument for \mathcal{G}^t is completely symmetric.

We examine all incentive constraints supposing that t is a peace period. Once equilibrium incentives in peace periods have been established, incentives in war periods can be easily checked since (a) punishment only occurs during peace periods, and (b) because the

¹²One can see why such beliefs are plausible. They belong to a population of individuals who have inherent difficulty in accepting blame/fault/responsibility for their own group's role in a long standing conflict. It's a good deal easier to believe that fault lies with the other group. Clearly, one could envision alternative error structures for different societies. Each would have different implications for equilibrium actions and payoffs sustained on path.

equilibrium prescribes WW during war periods regardless of the message, no end-of-period learning/updating can take place.

Let $V_{\mathcal{F}}^t(m_{\mathcal{F}}^t)$ denote the beginning-of-period continuation payoff to individual \mathcal{F}^t in period t after *receiving* message $m_{\mathcal{F}}^t$ from his predecessor, given his belief and the fact that both dynasties follow their prescribed action and message strategies thereafter. Similarly, we let $U_{\mathcal{F}}^t(m)$ denote the continuation payoff to \mathcal{F}^t if both he and his rival \mathcal{G}^t *send* the same message m to their successors. Note that at the end of any peace period an individual who is supposed to send message m believes with certainty that this rival sends the same message m .

Action incentives. Since t is a peace period, then the construction of (\mathbf{f}, \mathbf{b}) in Steps 1 and 2 implies

$$\begin{aligned} V_{\mathcal{F}}^t(m^*) &= (1 - \delta)\pi_{\mathcal{F}}(CC) + \delta U_{\mathcal{F}}^t(m^*), \\ V_{\mathcal{F}}^t(m^{\mathcal{G},\tau}) &= (1 - \delta)\pi_{\mathcal{F}}(DC) + \delta U_{\mathcal{F}}^t(m^{\mathcal{G},\tau-1}), \text{ where } m^{\mathcal{G},0} \equiv m^* \\ V_{\mathcal{F}}^t(m^{\mathcal{F},\tau}) &= V_{\mathcal{F}}^t(m^*) = (1 - \delta)\pi_{\mathcal{F}}(CC) + \delta U_{\mathcal{F}}^t(m^*) \end{aligned}$$

Notice that the payoff to individual \mathcal{F}^t after receiving the message $m^{\mathcal{F},\tau}$ is as it would be in the equilibrium phase. This is because he believes \mathcal{G}^t received m^* from his predecessor \mathcal{G}^{t-1} . Remember, however, that in the \mathcal{F} -punishment phase his assessment is incorrect. The actual profile in peace period t is CD rather than CC . The “after-the-fact” learning discussed at the end of Step 2 occurs for decision makers in peace periods. Hence, the end-of-period beliefs of \mathcal{F}^t imply $U_{\mathcal{F}}^t(m^*) = V_{\mathcal{F}}^{t+1}(m^*)$ and $U_{\mathcal{F}}^t(m^{\mathcal{G},\tau}) = V_{\mathcal{F}}^{t+1}(m^{\mathcal{G},\tau})$ for the equilibrium and \mathcal{G} -punishment phases, respectively. However, in the \mathcal{F} -phase, $U_{\mathcal{F}}^t(m^{\mathcal{F},\tau}) \neq V_{\mathcal{F}}^{t+1}(m^{\mathcal{F},\tau}) = V_{\mathcal{F}}^{t+1}(m^*)$. This is because \mathcal{F}^t learns the truth after observing action profile CD in period t , but he cannot communicate this to his successor \mathcal{F}^{t+1} who instead interprets \mathcal{F}^t 's message as an error.

It is not hard to check that for any $\tau = 1, \dots, T$,

$$U_{\mathcal{F}}^t(m^{\mathcal{G},\tau}) > U_{\mathcal{F}}^t(m^*) > U_{\mathcal{F}}^t(m^{\mathcal{F},\tau}),$$

and that $U_{\mathcal{F}}^{t-1}(m^{\mathcal{G},\tau})$ is increasing in τ while $U_{\mathcal{F}}^{t-1}(m^{\mathcal{F},\tau})$ is decreasing in τ .¹³ Using this string of inequalities and given that δ is close to one, it is straightforward to check that it is not in \mathcal{F}^t 's interest to deviate from the prescribed behavior in each of the phases as follows. If, in the equilibrium phase, \mathcal{F}^t deviates by choosing W or D , this initiates the \mathcal{F} -punishment phase, in which he receives (through his successors' actions) a continuation of $U_{\mathcal{F}}^t(m^{\mathcal{F},T})$. While he may gain in the current period from the deviation, in the next T peace periods, play switches from CC in the equilibrium phase to CD in the \mathcal{F} -punishment phase. If T is large enough, and if δ is close enough to one, this decline in \mathcal{F} 's long run continuation payoff outweighs the current one-shot deviation gain.

Next consider the \mathcal{G} -punishment phase. Here, if \mathcal{F}^t deviates by choosing W , this also initiates the \mathcal{F} -punishment phase. If he deviates by choosing C , his rival \mathcal{G}^t interprets this as an equilibrium path action and therefore ignores any message to the contrary. Given this belief by \mathcal{G}^t , the best possible continuation for \mathcal{F}^t following his choice of C is $U_{\mathcal{F}}^t(m^*)$. But this also leaves \mathcal{F}^t worse off than in the \mathcal{G} -punishment phase. Finally, during the \mathcal{F} -punishment phase, individual \mathcal{F}^t mistakenly believes he is in the equilibrium phase. Hence his action incentives are exactly as in that phase.

Incentives to Send Message m^ .* If he sends his prescribed message m^* (either because play is in the equilibrium phase, or because he knows that play is at the end of a punishment phase) his continuation payoff is $U_{\mathcal{F}}^t(m^*)$ which always gives him $\pi_{\mathcal{F}}(CC)$ during peace periods. If, instead, he sends a message $m^{\mathcal{F},\tau}$ then his payoff will also be $U_{\mathcal{F}}^t(m^*)$. This is because his successors will interpret $m^{\mathcal{F},\tau}$ as a message error. Alternatively, if he sends $m^{\mathcal{G},\tau}$ his payoff gives him $\pi_{\mathcal{F}}(DC)$ in the next peace period, but leaves him with $\pi_{\mathcal{F}}(DD)$ in every peace period thereafter (we ignore war periods since they always yield $\pi_{\mathcal{F}}(WW)$ regardless of the message). The reason is as follows. A message $m^{\mathcal{G},\tau}$ leads future individuals in dynasty \mathcal{F} to believe that the \mathcal{G} -punishment phase is occurring. This, in turn, leads to a choice of D by a successor in \mathcal{F} in the next peace period, $t + N$. But at that point, individual \mathcal{G}^{t+N} interprets his rival's choice of D as a deviation since \mathcal{G}^{t+N} receives m^* from *his* predecessor in \mathcal{G} . Hence, expectations are permanently mismatched, and play in these circumstances locks into DD

¹³The precise calculations of all for these payoffs are tedious though not hard to carry out given the war-and-peace cycle. For instance, in any peace period t ,

$$U_{\mathcal{F}}^t(m^{\mathcal{F},\tau}) = (1 - \delta^N)\pi_{\mathcal{F}}(WW) + \delta^N(1 - \delta)\pi_{\mathcal{F}}(CD) + \delta^{N+1}U_{\mathcal{F}}^{t+N}(m^{\mathcal{F},\tau-1}).$$

in every future peace period. For a high enough δ , \mathcal{F}^t clearly prefers message m^* to this alternative.

Incentives to Send Message $m^{\mathcal{G},\tau}$. By sending his prescribed message, say $m^{\mathcal{G},\tau}$, the continuation payoff for \mathcal{F}^t is $U_{\mathcal{F}}^t(m^{\mathcal{G},\tau})$. By alternatively sending messages m^* or $m^{\mathcal{F},\tau'}$ for any τ' , his payoff will be $U_{\mathcal{F}}^t(m^*)$ since these both induce CC in future peace periods. Since $U_{\mathcal{F}}^t(m^{\mathcal{G},\tau}) > U_{\mathcal{F}}^t(m^*)$ his best response is clearly to send $m^{\mathcal{G},\tau}$. Finally, by sending a different \mathcal{G} -phase message, say $m^{\mathcal{G},\tau'}$ with $\tau' \neq \tau$, individual \mathcal{F}^t cannot gain. If, for instance, $\tau' < \tau$, then the \mathcal{G} -phase ends “too soon” which is bad for \mathcal{F}^t since he benefits from the profile DC chosen during the peace periods of the \mathcal{G} -phase. If, on the other hand, $\tau' > \tau$, then this leads to mismatched expectations between future rivals in \mathcal{F} and \mathcal{G} , leading eventually to play that locks into DD in all future peace periods.

Incentives to Send Message $m^{\mathcal{F},\tau}$. When \mathcal{F}^t is supposed to send $m^{\mathcal{F},\tau}$, he correctly infers that the true phase is \mathcal{F} -punishment. However, his options are now limited because his successor \mathcal{F}^{t+1} interprets message $m^{\mathcal{F},\tau}$ as an error. By sending the prescribed message $m^{\mathcal{F},\tau}$, he then receives $\pi_{\mathcal{F}}(CD)$ for the next τ peace periods, but play will then eventually return the countries to the equilibrium phase. To verify that he will indeed send $m^{\mathcal{F},\tau}$ if δ is close enough to one, observe that a deviation to m^* , $m^{\mathcal{F},\tau'}$ for $\tau' < \tau$, or $m^{\mathcal{G},\tau''}$ would cause mismatched expectations, as outlined before, eventually leading to DD in *all* future peace periods. Finally, one can show that \mathcal{F}^t is indifferent between $m^{\mathcal{F},\tau}$ and $m^{\mathcal{F},\tau'}$ for $\tau' > \tau$. With this we conclude the argument. ■

4. Social Memory, Messages, and Evidence

The logic of the result highlights the critical roles of private, intra-dynastic messages and the error structure on those messages and actions. In the next few subsections, we re-examine these roles in more detail and ask whether and to what extent the particular features of our equilibrium are essential to sustaining below-minmax payoffs.

4.1. All Out War and Private Communication

Social memory hinges on information transmission through time. How important is the private communication between one generation of individuals and the next in sustaining frequent all out wars? Consider the dynastic game with messages taken out entirely. Then, for every t , the long-run continuation payoffs in (2) cannot be below the minmax payoff for either \mathcal{F} or \mathcal{G} . Therefore the all out war equilibrium of Proposition 1 is no longer viable.

This follows from the fact that, with messages taken out the dynastic game is easily seen to be equivalent to a repeated game with “imperfect public monitoring” (Fudenberg, Levine, and Maskin, 1994) in which individuals are forbidden from using “private strategies.”¹⁴ Every “perfect public equilibrium” in Fudenberg, Levine, and Maskin (1994) yields long-run payoffs that cannot be below a player’s minmax payoff. Hence, in our case not below the minmax payoffs $\pi_{\mathcal{F}}(DW)$ and $\pi_{\mathcal{G}}(WD)$, respectively.

4.2. Messages Versus Evidence

Proposition 1 is false if the evidence s^t is perfectly informative about the past history of action profiles. In this case social memory must be correct. In particular, if we set $\gamma = 1$, the long-run continuation payoffs in (2) cannot be below the minmax for either \mathcal{F} or \mathcal{G} . Therefore the all out war equilibrium of Proposition 1 is no longer viable.

To see why this is true we can proceed by induction in the following way. Individuals \mathcal{F}^0 and \mathcal{G}^0 of course observe nothing about past play and do not receive any messages at all. So, in equilibrium they forecast correctly the action chosen by the opposing individual.¹⁵ Then, during period 0, an action pair is played, and finally messages $m_{\mathcal{F}}^1$ and $m_{\mathcal{G}}^1$ are sent to individuals \mathcal{F}^1 and \mathcal{G}^1 .

At the beginning of period 1, individuals \mathcal{F}^1 and \mathcal{G}^1 observe $m_{\mathcal{F}}^1$ and $m_{\mathcal{G}}^1$ and the action pair that took place at 0 (since $\gamma = 1$). Now we can ask whether the beginning-of-period beliefs of individual \mathcal{F}^1 about what message individual \mathcal{G}^1 has received (and hence about what he will play) can possibly depend on $m_{\mathcal{F}}^1$. The answer must be no. This is because individual \mathcal{F}^1 observes exactly what individual \mathcal{G}^0 observed when he selected the message $m_{\mathcal{G}}^1$ he sent to \mathcal{G}^1 . Hence, in equilibrium, individual \mathcal{F}^1 can forecast exactly the message $m_{\mathcal{G}}^1$ simply on the basis of what he observes about the past. Therefore his beginning-of-period beliefs about $m_{\mathcal{G}}^1$ do not depend on $m_{\mathcal{F}}^1$. A symmetric argument can be used to see that the beginning-of-period beliefs of individual \mathcal{G}^1 about $m_{\mathcal{F}}^1$ cannot depend on the message $m_{\mathcal{G}}^1$ he receives.

Proceeding by induction forward in time, this line of argument shows that the beginning-of-period beliefs of any individuals \mathcal{F}^t and \mathcal{G}^t about what the opposing individual is about

¹⁴The latter is because individuals \mathcal{F}^t and \mathcal{G}^t observe s^t at the beginning of t , but do not observe directly the actions of any previous individuals, not even those of their own predecessors in the same country.

¹⁵The argument is more delicate but essentially the same if mixed strategies are allowed. In this case of course each individual forecasts correctly the probabilities with which each action will be taken by the opposing individual.

to choose cannot depend on the messages $m_{\mathcal{F}}^t$ and $m_{\mathcal{G}}^t$ that they respectively receive.

From the insensitivity of beliefs to messages we have just shown we can then deduce that when the evidence s^t is perfectly informative, the model must behave in the same way as the model with private messages taken out entirely.¹⁶ This, as we saw above, implies that the long-run continuation payoffs in (2) can never be below the minmax for either country.

On the other hand, things change dramatically if $\gamma < 1$. We remarked already that the beliefs of individuals \mathcal{F}^t and \mathcal{G}^t in the equilibrium in Proposition 1 do not depend on the realization of the observed evidence s^t . This is neither necessary nor sufficient to generate “incorrect social memory.” To see what this means formally, start with an individual’s beginning-of-period belief $\mathbf{b}_i^t(m_i^t, s^t)$, and take the appropriate marginal distributions to derive belief $\mathbf{q}_i^t(m_i^t, s^t)$ over *past* histories of actions. Social memory can thus be identified with belief \mathbf{q}_i^t which, intuitively, coincides with the wider notion of social memory discussed earlier (see, for instance, Crumley (2002)). By this definition, \mathbf{q}_i^t could be incorrect even though it varies with the signal s^t . Alternatively, \mathbf{q}_i^t could be full and correct even though it does not vary with the signal. This is because the messages from one generation to the next can in principle convey all available information about the past history of action profiles.¹⁷

Intuitively, equilibrium beliefs can be independent of the evidence from the past as in the equilibrium in Proposition 1 because the messages are sent after individuals \mathcal{F}^t and \mathcal{G}^t observe the actual play in period t , and hence, in principle, can convey the true record of play. Physical evidence, on the other hand, cannot possibly contain all information about what happened in the past. It is then possible that the individuals place sufficiently more trust in the messages than in the physical evidence and effectively ignore the latter. In a sense, it is precisely because messages can in principle noiselessly encode the actual history that they may entirely override the evidence in the equilibrium beliefs. We view this as an appealing and realistic attribute of the model.

¹⁶The argument here becomes a little more involved than our intuitive description suggests. This is because upon receiving different messages, an individual may take different actions even though his beliefs are the same. However, this difficulty can be circumvented noting that, since his beliefs are the same, he must be indifferent between the different actions he takes.

¹⁷It is in fact possible to construct “truthful” equilibria in which all information about the past history of play is correctly conveyed from one set of individuals to the next. The perturbations that must be used to construct the individuals beliefs in such equilibria are far from “natural.” In essence full truthfulness requires that a single message mistake be infinitely less likely than an *unbounded* number of action mistakes combined.

4.3. Imperfect Current Monitoring

In light of our previous discussion of messages versus evidence, we consider the following modification of the model. In each period individuals \mathcal{F}^t and \mathcal{G}^t , after choosing their action at t , no longer observe the opponent's action at t while they still observe the imperfect public signal p^t of the action pair taken at t , as before. All other details of the model are unchanged. Refer to this modified model as one of *imperfect current monitoring*.

Under imperfect current monitoring, physical evidence is the main engine behind equilibrium behavior. Crucially, this is reflected in the impossibility of social memory that is systematically wrong. In fact there is a well defined sense in which under imperfect current monitoring social memory must be *effectively correct*.

Given any $\bar{t} \leq t - 1$, it will be convenient to denote by $\mathbf{q}_{\mathcal{F}}^t(\bar{t}, m_{\mathcal{F}}^t, s^t)$ and $\mathbf{q}_{\mathcal{G}}^t(\bar{t}, m_{\mathcal{G}}^t, s^t)$ the probability distributions over action profiles in periods $\bar{t}, \dots, t - 1$ only implied by (the marginals of) $\mathbf{q}_{\mathcal{F}}^t(m_{\mathcal{F}}^t, s^t)$ and $\mathbf{q}_{\mathcal{G}}^t(m_{\mathcal{G}}^t, s^t)$. Let $h_+^{\bar{t}} \in H^{\bar{t}} \times \left(\times_{\tau=0}^{\bar{t}-1} (M_{\mathcal{F}}^{\tau} \times M_{\mathcal{G}}^{\tau}) \right)$. That is, $h_+^{\bar{t}}$ describes a full history of actions and messages up to and including the action stage of $\bar{t} - 1$.¹⁸

We wish to consider histories that follow $h_+^{\bar{t}}$, assuming no further deviations after the action stage of $\bar{t} - 1$. Let $\mathbf{e}^t(h_+^{\bar{t}}, s^t)$ be the actual equilibrium distribution over action profiles in periods $\bar{t}, \dots, t - 1$, following $h_+^{\bar{t}}$, given the evidence s^t .

The following result shows in essence that with imperfect current monitoring after any history, on or off path, the two countries' social memory can only differ from the actual equilibrium distribution over action profiles in a payoff irrelevant way.

In the following Proposition, two equilibria are said to be *equivalent* if they have the same payoff continuations to each player after every history.

Proposition 2. *Correct Social Memory:* Fix any precision level $\gamma \in [0, 1)$. Any pure strategy sequential equilibrium $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$ under imperfect current monitoring is equivalent to an equilibrium $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ with the following properties: (i) neither the beliefs nor the strategies of any individual depend on his message received (ii) given any \bar{t} and any history $h_+^{\bar{t}}$ as above, for every $t > \bar{t}$ and s^t then $\mathbf{q}_{\mathcal{F}}^t(\bar{t}, s^t) = \mathbf{q}_{\mathcal{G}}^t(\bar{t}, s^t) = \mathbf{e}^t(h_+^{\bar{t}}, s^t)$.¹⁹

A formal proof of Proposition 2 is in the Appendix. According to its statement, under imperfect current monitoring, any equilibrium is equivalent to one in which the social memory

¹⁸So far we have used the notation h^t to indicate a history of *action profiles* from 0 to $t - 1$. We are now using a subscript “+” to indicate a history of *both* action and message profiles.

¹⁹In light of (i), we drop message m in the arguments of beliefs $\mathbf{q}_{\mathcal{F}}^t, \mathbf{q}_{\mathcal{G}}^t$.

of both countries about previous action profiles chosen after \bar{t} must be correct. Our claim is surprising in the following sense. Going from the dynastic game we considered before to the one with imperfect current monitoring we *degrade the individuals' information* about the current action profile. Yet, as a result we find that the equilibrium has the same payoffs after any history as one in which strategies and beliefs do not vary with messages. Moreover, the individuals' assessments of the history of play qualitatively *improves* in the sense that it cannot be systematically wrong.

This contrasts with our original model where social memory can fail to predict even the actions of the individuals who follow their equilibrium strategies. To give an example, consider the equilibrium described in Proposition 1. Let $h_+^{\bar{t}}$ be defined as before such that there is one and only one deviation. At the end of period $\bar{t} - 2$, individual $\mathcal{F}^{\bar{t}-2}$ sends some message $m^{G,\tau}$ instead of m^* as prescribed. Suppose that all the players in periods $\bar{t} - 1, \bar{t}, \dots$ play according to the equilibrium of Proposition 1. According to the equilibrium, the action profile will eventually lock into DD in future peace periods. However, every player in period $t > \bar{t} - 1$ assigns probability zero to the event that DD has ever been played before t . In other words, the individuals living after period $\bar{t} - 1$ are *systematically wrong* about the action profiles played in the peace periods after $\bar{t} - 1$.

It is worth explicitly making the following observation about imperfect current monitoring. Using the same logic as in Section 4.1, one can show the following. Fix any precision level $\gamma \in [0, 1)$ and consider the model with imperfect current monitoring. Consider any pure strategy equilibrium. Then the long-run payoff to both countries cannot be below the minmax payoff. Therefore the all out war equilibrium of Proposition 1 is no longer viable. In this sense, degrading the information that individuals have about actions currently taken makes both countries better off.

It is legitimate to ask what happens if we allow for a current signal that is imperfect but with a higher level of precision than those concerning past actions; an intermediate case between our baseline model of Section 2 and the imperfect current monitoring case considered in this section. In this case, it is not too hard to show, using examples, that equilibrium behavior can indeed depend on messages in a non-trivial way that allows for equilibrium payoffs that are not available in the standard repeated game with infinitely-lived players. Whether in the limit as δ approaches 1 payoffs below the minmax can be sustained is an open question at this point.

5. Relation to the Literature

Earlier, we compared this model to other models of war. Unlike most of the literature, this paper analyzes conflicts so destructive as to deny participants the usual modicum of self protection (inherent in the minmax outcome). We examine the role of social memory in generating these types of conflicts.

As far as we know, ours is the only dynastic game model that deals specifically with issues of war and social memory. There are other dynastic game models of communication: Anderlini and Lagunoff (2005), Kobayashi (2007), Lagunoff and Matsui (2004), and AGL. However, they have a different focus. Their interest is in characterizing the broadest possible equilibrium set rather than in either war or social memory. Indeed, because physical evidence is absent from those models, they cannot address the questions we pose here on whether/how social memory incorporates tangible evidence alongside private messages.²⁰

There are, however, two other papers that bear mentioning because they deal with “memory” and conflict, albeit in very different settings. Glaeser (2005) models the political economy of group hatred. In his model, “entrepreneurial” politicians can supply hate-creating stories as a signal of an out-group’s threat to the rest of the uninformed citizenry. Verifying these stories is costly, and so a politician’s partially revealing (i.e., mixed) strategy leads the citizens to put more weight in their beliefs on the truthfulness of these stories. Dessí (2008) studies the role of collective memory when individuals’ investment decisions exhibit spillovers. She models the problem of an informed principal (the older generation) who selectively informs two younger agents about the value of a noisy signal of the past. Though her model is not about war per se, the negative spillovers create conflicts among the young, and therefore must be endogenously controlled by the principal who chooses to selectively withhold some information.

The present model differs from these and, in our view, contributes to the literature in one critical respect. In our case a society’s history has no direct effect on current payoffs. In game theoretic parlance, there are no payoff types. While an individual’s private information affects his beliefs about past history, it does not reflect any fundamentals. This distinction is not inconsequential. For instance, in Glaeser’s model, there must be some states of the world, arising with positive probability, in which the out-group really *is* dangerous in order for citizens to believe this even when the out-group is, in fact, not dangerous. Our results

²⁰One other critical difference is that the results of the prior papers do not apply to the case of two dynasties.

hold without the presence of payoff types: destructive forces can be unleashed even when there is no objective danger to the current generation in *any* state of the world.

This is noteworthy, because even in the larger literature on endogenous “memory” that deals non-war issues, the signals are either tangible, costly, or both. Clearly, the literature on herding falls into this category. For this massive literature we refer the reader to surveys on herding by Bikhchandani, Hirshleifer, and Welch (1998) and Gale (1996), and the work of Ahn and Suominen (2001), Banerjee and Fudenberg (2004), Jackson and Kalai (1999) and Moscarini, Ottaviani, and Smith (1998).

Among the intergenerational models more closely aligned to our framework, Bisin and Verdier (2001) propose a model of intergenerational transmission of cultural traits. They examine a learning process in which imperfectly altruistic parents choose whether to pass on certain cultural traits to their kids. A society with heterogeneous traits is shown to emerge if parental and societal transmission mechanisms are substitutes. Tabellini (2008) studies a similar mechanism for parental transmission of values in a Prisoner’s Dilemma setting. Benabou and Tirole (2006) propose a signalling model of economic ideology. They analyze the relationship between the beliefs about the causes of wealth and poverty and redistributive policies. In their model, people want to use their private productivity as a signal to motivate their children to work harder when taxes and redistribution are low. In this political regime, there are greater incentives to send a signal indicating that economic success depends on effort. At the same time, when people become convinced that economic success depends on effort, they vote for a regime of low taxes and redistribution. Because of the complementarities, there are two equilibria: an optimistic “American” equilibrium with laissez-faire public policy, and a pessimistic “European” equilibrium with a more extensive welfare state.

We emphasize that, just as in Glaeser (2005)’s and Dessí (2008)’s frameworks, individual memory in these models is derived from tangible or costly signals. The signal need not be an asset that has value but can, nevertheless, effectively encode an individual’s past behavior (Johnson, Levine, and Pesendorfer, 2001, Kandori, 1992b). Corbae, Temzelides, and Wright (2003), Kocherlakota (1998), Kocherlakota and Wallace (1998) and Wallace (2001) show how fiat money, for example, plays this role.

Finally, we note that the present results are roughly consistent with a number of unusual findings in lab experiments designed to capture intergenerational environments. Chaudhuri, Schotter, and Sopher (2009) and Schotter and Sopher (2003, 2007) show that word-of-mouth

learning is a stronger force for perpetuating conventions, good or bad, than simply having access to the historical record. In related experiments, Duffy and Feltovich (2005) report that word-of-mouth communication frequently makes things worse than observing history alone. Obviously, the link between theory and the experiments is imperfect. Nevertheless, their findings largely accord with our results, provided that subjects are never absolutely certain that the historical evidence they receive is exact. They also help in understanding the demonstrable effect that private versus public communication has for sustaining bad outcomes found by Chaudhuri, Schotter, and Sopher (2009).

6. Conclusion

This paper studies the role of social memory in creating and sustaining conflicts. Social memory is embodied in a society's vicarious beliefs about the past. These beliefs are shaped by both intergenerational communication and the imperfect physical evidence from the past. To formalize it entails a detailed model of the intergenerational communication within dynastic societies.

We show that there exist equilibria in a canonical Game of Conflict in which "all out war" occurs with arbitrarily high frequency. In these equilibria physical evidence is ignored and, in fact, beliefs of one or both parties can be incorrect after certain events.

Significantly, these equilibria can occur despite the fact that there are no objective states of the world in which the conflict is desirable from anyone's point of view. These outcomes could not be attained in a standard infinitely repeated game. Because messages can, in principle, convey more information than any imperfectly informative physical evidence, there are equilibria in which the current generation focuses only on the messages. Ironically, social memory can be incorrect precisely because it relies on sources that can be more informative than hard evidence.

A few issues bear mentioning here. The first concerns robustness with respect to the demographic structure. We initially assumed a "canonical demographic" whereby there is full replacement within each dynasty every period. However, using a technique known at least since Ellison (1994), one can show that our results are fairly robust to the demographic structure. Specifically, it turns out that any equilibrium of the model with canonical demographics has a corresponding equilibrium of a model in which replacement within a dynasty can occur at any time provided that there is a uniform upper bound L on individuals' lifetimes. The idea is

to construct L interleaved “copies” of the equilibrium with canonical demographics.²¹ Here, the strategies and beliefs of individuals alive in periods $0, L, 2L, 3L$ and so on, “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 information concerning period 0 , individuals alive at $2L$ will only consider information concerning periods 0 and L and so on, forward without bound.²² 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. It is fairly straightforward in this generalized model to show that for any equilibrium of the model with canonical demographics and discount factor δ , the construction outlined here yields an equilibrium for the L -bounded demographics case for a discount factor of $\delta^{\frac{1}{L}}$.

Second, we have focused our attention entirely on “bad” equilibria with frequent all out wars. The implication is that wrong social memory is a bad thing. But there is a flip-side to this which highlights the possible “good” consequences of wrong social memory. Precisely because very bad payoffs can be sustained on path, these payoffs can be used as “punishments” off it. We have examples that show that the cooperative outcome CC can be sustained in equilibrium for a lower δ in the dynastic game than in the standard repeated game if γ is low. The construction is not as straightforward as it might look at first sight since one cannot simply “plug in” — say — the equilibrium of Proposition 1 as a punishment phase of another equilibrium in which CC is sustained. The reason is that the equilibrium of Proposition 1 is viable for a high δ in the first place.

Whether cooperation can *in general* be sustained more easily in the dynastic game is an open question at this point. The question of how the possibility of inaccurate social memory might lead to the emergence of better equilibria is clearly both interesting and potentially important. We leave this issue for future research.

²¹The same would be true if we considered the case where the L -bounded demographics are stochastic in the sense that no individual lives more than L periods with probability one. The same technique would also allow us to use the model with canonical demographics to handle demographics with overlapping generations of individuals within a dynasty. The details are well beyond the scope of this paper.

²²Clearly, in the model with the L -bounded demographics, the messages sent must be constructed so that the relevant information is passed along (unused) via the intervening individuals.

References

- AHN, I., AND M. SUOMINEN (2001): “Word-Of-Mouth Communication and Community Enforcement,” *International Economic Review*, 42, 399–415.
- ANDERLINI, L., D. GERARDI, AND R. LAGUNOFF (2008): “A ‘Super’ Folk Theorem in Dynastic Repeated Games,” *Economic Theory*, 37, 357–394.
- ANDERLINI, L., AND R. LAGUNOFF (2005): “Communication in Dynastic Repeated Games: ‘Whitewashes’ and ‘Coverups,’” *Economic Theory*, 26, 265–299.
- BANERJEE, A., AND D. FUDENBERG (2004): “Word-of-mouth learning,” *Games and Economic Behavior*, 46(1), 1–22.
- BAUMEISTER, R., AND S. HASTINGS (1997): “Distortions of Collective Memory: How Groups Flatter and Deceive Themselves,” in *Collective Memory of Political Events*, ed. by J. W. Pennebaker D. Paez, and B. Rimè. Mahwah, NJ: Lawrence Erlbaum Associates.
- BENABOU, R., AND J. TIROLE (2006): “Belief in a Just World and Redistributive Politics,” *Quarterly Journal of Economics*, 121(2), 699–746.
- BHASKAR, V. (1998): “Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems,” *Review of Economic Studies*, 65, 135–149.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1998): “Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades,” *Journal of Economic Perspectives*, 12(3), 151–170.
- BISIN, A., AND T. VERDIER (2001): “The Economics of Cultural Transmission and the Dynamics of Preferences,” *Journal of Economic Theory*, 97, 298–319.
- CATTELL, M. G., AND J. J. CLIMO (2002): *Social Memory and History: Anthropological Perspectives*. Walnut Creek, CA: Altamira Press.
- CHASSANG, S., AND G. P. I MIGUEL (2008): “Conflict and Deterrence under Strategic Risk,” mimeo, Princeton University.
- CHAUDHURI, A., A. SCHOTTER, AND B. SOPHER (2009): “Talking Ourselves to Efficiency: Coordination in Inter-Generational Minimum Effort Games with Private, Almost Common and Common Knowledge of Advice,” *Economic Journal*, 119, 91–122.
- CONNERTON, P. (1989): *How Societies Remember*. Cambridge: Cambridge University Press.
- CORBAE, D., T. TEMZELIDES, AND R. WRIGHT (2003): “Directed Matching and Monetary Exchange,” *Econometrica*, 71, 731–756.
- CRUMLEY, C. (2002): “Exploring Venues of Social Memory,” in *Social Memory and History: Anthropological Perspectives*, ed. by M. G. Cattell and J. J. Climo. Walnut Creek, CA: Altamira Press.
- DESSÍ, R. (2008): “Collective Memory, Cultural Transmission and Investments,” *American Economic Review*, 98, 534–560.

- DUFFY, J., AND N. FELTOVICH (2005): "Words, Deeds, and Lies: Strategic Behavior in Games with Multiple Signals," *Review of Economic Studies*, 73, 669–688.
- ELLISON, G. (1994): "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61, 567–588.
- FEARON, J. (1995): "Rationalist Explanations for War," *International Organization*, 49, 379–414.
- FENTRESS, J., AND C. WICKHAM (1992): *Social Memory*. Oxford, UK: Blackwell.
- FOLEY, R. (2005): *German Strategy and the Path to Verdun: Erich von Falkenhayn and the Development of Attrition, 1870-1916*, Cambridge Military Histories. Cambridge UK: Cambridge University Press.
- FUDENBERG, D., D. K. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1039.
- FUDENBERG, D., AND E. S. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533–556.
- FUDENBERG, D., AND J. TIROLE (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236–260.
- GALE, D. (1996): "What Have We Learned from Social Learning?," *European Economic Review, Papers and Proceedings of the Tenth Annual Congress of the European Economic Association*, 40, 617–628.
- GLAESER, E. (2005): "The Political Economy of Hatred," *Quarterly Journal of Economics*, 120, 45–86.
- JACKSON, M., AND M. MORELLI (2008): "Strategic Militarization, Deterrence and Wars," mimeo, Stanford University.
- JACKSON, M. O., AND E. KALAI (1999): "Reputation Versus Social Learning," *Journal of Economic Theory*, 88(1), 40–59.
- 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.
- KOBAYASHI, H. (2007): "Folk Theorems for Infinitely Repeated Games Played by Organizations with Short-Lived Members," *International Economic Review*, 48, 517–549.
- 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.
- MOSCARINI, G., M. OTTAVIANI, AND L. SMITH (1998): "Social Learning in a Changing World," *Economic Theory*, 11, 657–665.
- OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. Cambridge, MA: M.I.T. Press.
- PENNEBAKER, J., D. PAEZ, AND B. RIMÈ (eds.) (1997): *Collective Memory of Political Events*. Mahwah, NJ: Lawrence Erlbaum Associates.
- ROGERS, Y. (1997): "A Brief Introduction to Distributed Cognition," Discussion paper, Interact Lab, School of Cognitive and Computing Sciences, University of Sussex.
- SALANT, D. (1991): "A Repeated Game with Finitely Lived Overlapping Generations of Players," *Games and Economic Behavior*, 3, 244–259.
- SHELLING, T. (1960): *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- SCHOTTER, A., AND B. SOPHER (2003): "Social Learning and Coordination Conventions in Inter-Generational Games: An Experimental Study," *Journal of Political Economy*, 111, 498–529.
- SCHOTTER, A., AND B. SOPHER (2007): "Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach," *Games and Economic Behavior*, 58, 365–393.
- SCHWARZ, M., AND K. SONIN (2007): "A Theory of Brinkmanship, Conflicts, and Commitments," *Journal of Law, Economics, and Organization*, Online Access.
- SMITH, L. (1992): "Folk Theorems in Overlapping Generations Games," *Games and Economic Behavior*, 4, 426–449.
- SUTTON, J. (2005): "<http://www.phil.mq.edu.au/staff/jsutton/socialmemory.htm>," On-Line Reading List, Department of Philosophy, Macquarie University, Sidney.
- TABELLINI, G. (2008): "The Scope of Cooperation: Values and Incentives," *Quarterly Journal of Economics*, forthcoming.
- WALLACE, N. (2001): "Whither Monetary Economics?," *International Economic Review*, 42, 847–869.
- YARED, P. (2009): "A Dynamic Theory of Concessions and War," mimeo, Columbia University.

Appendix

A.1. Proof of Proposition 2

Let $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$ be a pure strategy equilibrium under imperfect current monitoring. We establish the Proposition from the point of view of an individual in dynasty \mathcal{F} concerning the actions of individuals in dynasty \mathcal{G} . The flip-side of the argument is, *mutatis mutandis*, identical hence omitted.

Step 1. We first establish that in the equilibrium $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$, the beliefs of \mathcal{F}^t concerning the history of actions and messages taken by dynasty \mathcal{G} do not depend on $m_{\mathcal{F}}^t$.

Recall that the per-period signals defining s^t have *full support*. That is, for any t , any realized s^t has positive probability given any possible true history of action profiles h^t .

Given the equilibrium strategies of all individuals, and evidence s^t we can recurse forward to compute all equilibrium actions and messages of all individuals up to an including period $t - 1$ in the obvious way. Fix $s^t = (p^0, \dots, p^{t-1})$. We begin with \mathcal{F}^0 and \mathcal{G}^0 , who of course observe nothing about the past. So their action strategies determine the equilibrium action profile a^{0*} . Given p^0 and the equilibrium message strategies we can now compute the equilibrium messages $m_{\mathcal{F}}^{1*}$ and $m_{\mathcal{G}}^{1*}$. Given these messages and p^0 we can then use the equilibrium action strategies of individuals \mathcal{F}^1 and \mathcal{G}^1 to compute the equilibrium action profile a^{1*} . Recursing forward in this way, we can compute all equilibrium action and message profiles up to and including a^{t*} and $m^{t+1*} = (m_{\mathcal{F}}^{t+1*}, m_{\mathcal{G}}^{t+1*})$.

For a given any s^t , for every $\tau = 0, \dots, t$ let $\mathbf{a}(\tau, s^t) = (\mathbf{a}_{\mathcal{F}}(\tau, s^t), \mathbf{a}_{\mathcal{G}}(\tau, s^t))$ and $\mathbf{m}(\tau, s^t) = (\mathbf{m}_{\mathcal{F}}(\tau, s^t), \mathbf{m}_{\mathcal{G}}(\tau, s^t))$ be the equilibrium action and message profiles $a^{\tau*}$ and $m^{\tau*}$ we have just computed, but with the dependence on s^t made explicit, so that we can now vary it.

Consider individual \mathcal{F}^t at the beginning of t . The first case we consider is that he observes evidence s^t and receives the corresponding equilibrium message $\mathbf{m}_{\mathcal{F}}(t, s^t)$. It then follows from completely standard arguments that his beliefs about the *entire* history of action and message profiles must assign probability one to the sequences $\{\mathbf{a}(\tau, s^t)\}_{\tau=0}^{t-1}$ and $\{\mathbf{m}(\tau, s^t)\}_{\tau=1}^t$. At the end of period t , since we are in the imperfect current monitoring case, it must also be the case that \mathcal{F}^t assigns probability one to individual \mathcal{G}^t having played $\mathbf{a}_{\mathcal{G}}(t, s^t)$ and having sent $\mathbf{m}_{\mathcal{G}}(t+1, s^t)$ to his successor \mathcal{G}^{t+1} .

Now consider again individual \mathcal{F}^t at the beginning of t , but consider the complementary case in which he observes evidence s^t and receives an off-path message $m_{\mathcal{F}}^t \neq \mathbf{m}_{\mathcal{F}}(t, s^t)$. Clearly, in this case he must conclude that some deviation from equilibrium has occurred. In fact, since we are in the imperfect current monitoring case, it *must* be that one or more individuals in dynasty \mathcal{F} has deviated before t . Any message $m_{\mathcal{F}}^t \neq \mathbf{m}_{\mathcal{F}}(t, s^t)$ could not possibly be observed otherwise.

Can the beginning-of-period equilibrium beliefs of \mathcal{F}^t assign positive probability to one or more individuals in dynasty \mathcal{G} having deviated before t ? The answer must be “no.” A routine check reveals that in this case, since the per-period signals have full support, \mathcal{F}^t must assign probability zero to this event after observing s^t and $m_{\mathcal{F}}^t \neq \mathbf{m}_{\mathcal{F}}(t, s^t)$. Intuitively this is because a deviation by dynasty \mathcal{F} is *necessary* to reach this “information set,” and it is also the case that a *single* deviation (for instance at the message stage of $t - 1$) is *sufficient* to reach it. So, ascribing a deviation to dynasty \mathcal{G} after observing s^t and $m_{\mathcal{F}}^t \neq \mathbf{m}_{\mathcal{F}}(t, s^t)$ involves more deviations than the minimum necessary to actually reach the given information set. By standard arguments, this can never be the case in a Sequential Equilibrium. Hence, at the beginning of period t the

beliefs of \mathcal{F}^t about dynasty \mathcal{G} must be as in the first case, in which \mathcal{F}^t observes s^t and the equilibrium $m_{\mathcal{F}}^t = \mathbf{m}(t, s^t)$. At the end of period t , since we are in the imperfect current monitoring case, it must also be the case that \mathcal{F}^t assigns probability one to individual \mathcal{G}^t having played $\mathbf{a}_{\mathcal{G}}(t, s^t)$ and having sent $\mathbf{m}_{\mathcal{G}}(t+1, s^t)$ to his successor \mathcal{G}^{t+1} .

Comparing the first and second case we just considered, it is apparent that the beliefs of \mathcal{F}^t on path concerning the history of actions and messages taken by dynasty \mathcal{G} do not depend on $m_{\mathcal{F}}^t$.

Step 2. In this Step, we construct a new equilibrium $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ from the original one, and then argue that (a) it is equivalent to the original equilibrium in the sense that it delivers the same beginning and end of period payoffs after every history of play; and (b) in this new equilibrium, neither the message or action strategies nor the beliefs of any individual depend on the message he receives.

To get $\tilde{\mathbf{f}}$ from $\hat{\mathbf{f}}$ we set the behavior of all individuals after any off-path message to be the same as if they had instead observed the equilibrium message. As a consequence, under $\tilde{\mathbf{f}}$ the behavior of all individuals (at the action and the message stage) does not depend on the message they receive.

To get $\tilde{\mathbf{b}}$ from $\hat{\mathbf{b}}$ we set the beliefs of all individuals after any off-path message to be the same as if they had instead observed the equilibrium message. As a consequence, under $\tilde{\mathbf{b}}$ the beliefs of all individuals do not depend on the message they receive.

Clearly, by construction of $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$, neither the message or action strategies nor the beliefs of any individual depend on the message he receives.

It remains to argue that $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ is in fact an equilibrium and that it is equivalent to $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$ in the sense of (a) above. It is convenient to argue the latter first. As before, we focus on dynasty \mathcal{F} . The details for dynasty \mathcal{G} are symmetric.

Given the argument in Step 1, it must be that going from $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$ to $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ the beliefs of any individual \mathcal{F}^t about the history of actions and messages for dynasty \mathcal{G} are the *same* in the two equilibria.

Since following the receipt of an on-path message the strategies are the same across the two equilibria, there is clearly nothing to prove in this case.

Consider then an on-path and an off-path message for the same s^t in the original equilibrium $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$. That is consider \mathcal{F}^t and two information sets, an on-path one given by s^t and $m_{\mathcal{F}}^t = \mathbf{m}_{\mathcal{F}}(t, s^t)$, and the other given by the same s^t and any off-path message $m_{\mathcal{F}}^t \neq \mathbf{m}_{\mathcal{F}}(t, s^t)$. At these two information sets, and at all the ones following them after s^{t+1} is realized, the beliefs of \mathcal{F}^t about dynasty \mathcal{G} must also be the *same*. But this implies that the continuation expected payoff of \mathcal{F}^t must be the same at these two information sets and at those following them after s^{t+1} is realized. Because of the way we constructed $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ from $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$, this is clearly enough to show that the two equilibria are equivalent.

To see that $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ is in fact an equilibrium, we need to verify two things. The first is that the strategies specified by $\tilde{\mathbf{f}}$ are sequentially rational given $\tilde{\mathbf{b}}$, and the second is that the beliefs $\tilde{\mathbf{b}}$ are admissible for a Sequential Equilibrium. Sequential rationality is a simple consequence of the fact that $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$ is an equilibrium and of how we constructed $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ from it. If sequential rationality failed at any information set in $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$, then it would have to fail (on path) at a corresponding information set in $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$. Lastly, to check that the beliefs $\tilde{\mathbf{b}}$ have the requisite properties is a routine exercise. It is enough to assume that message deviations are sufficiently “more likely” than deviations at the action stage. We omit the details.

Step 3. The last step establishes part (ii). Given the equilibrium $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$, the beginning of period equilibrium beliefs of all individuals can be characterized as follows.

Consider individual \mathcal{F}^t at the beginning of period t , observing $s^t = (p^0, p^1, \dots, p^{t-1})$. Since the actions of all individuals do not depend on messages, and the per-period signals have full support, his beliefs about the previous action profiles chosen can be computed recursively as follows. In period 0 of course no one observes anything and so the action profile is directly given by the equilibrium action strategies of individuals \mathcal{F}^0 and \mathcal{G}^0 . In period 1, we can compute the equilibrium action strategies of \mathcal{F}^1 and \mathcal{G}^1 , as a function of $s^1 = p^0$, and so on forward in time. In other words, the beginning of period beliefs of any individual \mathcal{F}^t or \mathcal{G}^t , given $s^t = (p^0, p^1, \dots, p^{t-1})$ can be written as period-by-period functions of the per-period signals p^τ with $\tau = 0, \dots, p^{t-1}$. So, we can write the social memory of country \mathcal{F} at t , as $\mathbf{q}_{\mathcal{F}}^t(s^t) = (q(0, s^t), q(1, s^t), \dots, q(t-1, s^t))$, where each $q(\tau, s^t)$ (with $\tau = 0, \dots, t-1$) denotes the action profile in period τ . Note that the function q does *not* have a country subscript or superscript since it is the same for \mathcal{F} and \mathcal{G} . Hence $\mathbf{q}_{\mathcal{G}}^t(s^t) \equiv \mathbf{q}_{\mathcal{F}}^t(s^t)$.

Now fix any \bar{t} , $h_{+}^{\bar{t}}$, $t > \bar{t}$ and s^t . Let $\mathbf{e}(h_{+}^{\bar{t}}, s^t)$ be as in the statement of the Proposition.

By construction, in the equilibrium $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ the actions of all individuals depend only on the evidence they observe. Hence, conditional on a given s^t , given that the per-period signals have full support, both individuals at time t will behave as if they were in equilibrium, even if the initial history $h^{\bar{t}}$ is in fact off-path.

It then follows immediately that we can write $\mathbf{e}(h_{+}^{\bar{t}}, s^t)$ as period-by-period functions of the per-period signals p^τ with $\tau = 0, \dots, p^{t-1}$. So, we can write this as $\mathbf{e}^t(h_{+}^{\bar{t}}, s^t) = (e(\bar{t}, s^t), \dots, e(t-1, s^t))$, where each $e(\tau, s^t)$ (with $\tau = \bar{t}, \dots, t-1$) denotes the action profile in period τ .

Since both the functions $e^\tau(\cdot, \cdot)$ and $q^\tau(\cdot, \cdot)$ are derived directly from the equilibrium strategies, they must be the same. Hence the claim follows immediately. ■