

Social Memory and Evidence from the Past*

LUCA ANDERLINI
(Georgetown University)

DINO GERARDI
(Yale University)

ROGER LAGUNOFF
(Georgetown University)

January 2007

Abstract. Examples of repeated destructive behavior abound throughout the history of human societies. This paper examines the role of social memory — a society’s vicarious beliefs about the past — in creating and perpetuating destructive conflicts. We examine whether such behavior is consistent with the theory of rational strategic behavior.

We analyze an infinite-horizon model in which two countries face off each period in an extended Prisoner’s Dilemma game in which an additional possibility of mutually destructive “all out war” yields catastrophic consequence for both sides. Each country is inhabited by a dynastic sequence of individuals who care about future individuals in the same country, and can communicate with the next generation of their countrymen using *private messages*. The two countries’ actions in each period also produce physical evidence; a sequence of informative but imperfect public signals that can be observed by all current and future individuals.

We find that, provided the future is sufficiently important for all individuals, *regardless of the precision of physical evidence from the past* there is an equilibrium of the model in which the two countries’ *social memory is systematically wrong*, and in which the two countries engage in all out war with arbitrarily high frequency.

Surprisingly, we find that *degrading* the quality of information that individuals have about current decisions may “improve” social memory so that it can *no longer be systematically wrong*. This in turn ensures that arbitrarily frequent all out wars cannot take place.

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

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

ADDRESS FOR CORRESPONDENCE: Luca Anderlini, Georgetown University, 37th and O Streets NW, Washington DC 20057, USA. 1a2@georgetown.edu

*All three authors are grateful to the National Science Foundation for financial support (Grant SES-0617789). We are also grateful to Helena Poon and numerous seminar participants for helpful comments. All errors are our own.

1. Introduction

Between February 21 and December 19 of 1916, the French and German armies clashed near the French city of Verdun. The French suffered 378,000 casualties of whom 120,000 dead. The Germans suffered 337,000 casualties of whom 100,000 dead. The final positions of the French and German armies on December 19 were no more than a few kilometers away from what they were on February 21, with no major population center having gone from the control of one side to the other.

This and many other bloody pages of European history seem even more paradoxical and pointless viewed from the later perspective of March 25, 1957, when France and (West) Germany, together with four other European countries (Italy, Belgium, The Netherlands and Luxembourg) signed the Treaty of Rome, establishing what was later to become the European Union.

Further conflict between major European powers now seems unimaginable. However, for many centuries the end of one conflict in Europe clearly did not make future ones any less likely. In some cases one conflict sowed the seeds for one to follow. How could the same societies that endured the battle of Verdun and many other mass slaughters of World War I clash again, on an even grander scale, barely 20 years later? How could the memory of one conflict not stop the next one from igniting? Generations of scholars have struggled with these questions.

In this paper, we carry out a parallel exercise to the detailed and specialized understanding that historians provide. Using the tools of game theory as a supplement to the historian's perspective, we examine the role of memory in creating and perpetuating destructive behavior on a large scale.

Imagine that we strip from the analysis the political layers within each society that may give positive incentives for conflict, for instance because one social stratum stands to gain from what is an overall destructive conflict with another society. Consider then a situation in which all out conflict is catastrophically bad for all involved in a way that everyone understands. Suppose also that we endow the actors with unlimited capacity for strategic thinking and data analysis. Is it then the case that all these "optimistic hypotheticals" rule out histories with the equivalent of the battle of Verdun and World War II starting 23 years later? In this paper, we establish that the answer is in fact "no." The basic mechanism that sustains disastrous outcomes under our optimistic hypotheticals revolves around the possibility that

a society's memory may be far from accurate, even in the presence of compelling physical evidence from the past.

It is no surprise that memory matters. Yet, in many instances, the *way* in which a society remembers its past influences its decisions in the present, sometimes in profoundly disturbing ways. Consider just two examples:

“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).

“You know the Spanish crusade against Muslims and the expulsions from Al Andalus [in 1492] are not so long ago.” — Al Qaeda spokesman (after claiming responsibility for the 2004 Madrid bombing).¹

In these instances, horrific outcomes were seemingly shaped by oral or written accounts passed on through generations. Social scientists in anthropology, psychology, and distributed cognition study these phenomena which they refer to as *social memory*. A small sample of this large literature includes Cattell and Climo (2002), Connerton (1989), Fentress and Wickham (1992), Pennebaker, Paez, and Rimè (1997), Rogers (1997) and Sutton (2005). The literature identifies a number of specific features that define social memory. According to Carole Crumley,

“Social memory is 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 [...].” — Crumley (2002)

In our model, the practical ties and relationships among the generations are forged by intergenerational communication within “dynasties.” These are placeholders for individual decision makers who are finitely lived but care about what happens to future generations who inhabit the same dynasty. Families, tribes, ethnic groups, and nation states are all ongoing entities with these features.

¹As quoted in Reston (2005, Prologue).

Within each dynasty, each decision maker directly observes the current outcome then chooses what and how much of this information to pass on to his successor in the dynasty. 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 — the written or oral historiography of the dynasty — received from his predecessor. Significantly, we assume that at least some part of the communication within a dynasty is not publicly available to other dynasties. The other source is the physical evidence — history’s “footprint” — in the form of a sequence of informative but imperfect public signals of events. It is impossible to reconstruct perfectly past events once time has confounded the evidence.

Our results demonstrate the possibility that participants, even in the face of imperfect but mounting evidence pointing to the truth, rely exclusively on the intra-dynastic messages to gauge the best course of action. Weighing the various possible intentions and mistakes that may generate their predecessors’ behavior, it is possible that the beliefs of the current participants may stray very far from what is actually taking place. In other words, social memory within and across societies can be *systematically incorrect*. Once this is granted, catastrophic and perpetual all out conflicts may actually materialize in a variety of ways provided everyone making decisions today is sufficiently concerned about the future.

To illustrate the mechanism by which this occurs, consider a world in which there are two dynasties, in the form of two countries — \mathcal{F} (France) and \mathcal{G} (Germany). These countries are engaged in a repeated *Game of Conflict*. In a Game of Conflict, there is at least one “destructive” action profile which takes both countries below their “individually rational” (minmax) payoff levels in the stage game. We refer to this profile as “all out war.” Notice that if both countries had perfect memory of the past then all out war could not occur with high frequency since, with perfect memory, long-run equilibrium payoffs must lie above the stage game minmax.

We show that all out wars can start on the basis of a seemingly irrelevant “trigger.” In our construction they start at pre-determined dates, but it is immediate to see that they could also be brought about by some, possibly random, exogenous payoff irrelevant event. At this point both sides plunge into destructive behavior because they believe it to be in the long run interest of their country. But how can this be the case? Both sides realize that at some, possibly distant, point in the future the slaughter will be over, and peace will prevail again. The question in the mind of both sides is how good will peace be for them once it prevails. Both sides think that peace-time will be a lot worse for them if they fail to wage war today.

Since the future matters a lot for everyone this is sufficient to sustain a state of all out war. This is so even if both sides can see that all out war will happen again, even very frequently.

Under our optimistic hypotheticals, the first difficulty comes once we look carefully at what happens in the Game of Conflict during peaceful times. Peaceful times come in three, very different, 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 the two sides are reversed. The idea is that if all sides have followed the (war-mongering) precepts of the trigger then peaceful times will be good for both, but otherwise they will be bad for the side that did not and good for the side that did. But now suppose that country \mathcal{F} did not follow the precepts and therefore peace just consists of domination of \mathcal{F} by \mathcal{G} . If we are in a world with very frequent all out wars, soon a trigger will dictate that both sides switch from (asymmetric) peace to all out war. But why should \mathcal{F} engage in such costly slaughter? After all, country \mathcal{F} does not have anything more to lose during peace times: these are already good for \mathcal{G} and bad for \mathcal{F} . This is where the possibility of systematically incorrect social memory becomes critical. We demonstrate how country \mathcal{F} , relying on its own historiography (passed down from one generation to the next) rather than on available evidence, suffers from the illusion that peace times are *not* asymmetrically biased against it. If such illusion can be maintained for long enough in the face of mounting evidence against it, then the circle can be closed, and a world with very frequent destructive all out wars is compatible with the optimistic hypotheticals posited by the theory of rational equilibrium behavior.

To summarize, we show that there exist equilibria in the 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) In every history in which there is a deviation from equilibrium, social memory is incorrect (from the point of view of at least one side); in fact, there are certain possible histories after which social memory is *maximally incorrect* in the following sense. Using a standard hypothesis testing procedure, and given any level of confidence, the null hypothesis that a certain event took place is rejected by a statistician who observes the accumulated evidence. However, all participants’ beliefs assign probability 1 to the very same event.

Our model behaves in a fundamentally different way from a standard repeated game with infinitely-lived players. Crucially, 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 point, we examine a variation of the dynastic 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 therefore could not arise.

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

2. Relation to the Literature

Our framework can be compared to other models of endogenous “memory.” For example, 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. Their paper belongs to a strand in the literature in which individuals’ memory is derived from tangible or costly assets. The asset need not itself have 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.

It is worth mentioning separately two papers in this strand that examine issues more closely related to ours. Dessi (2005) 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. Negative or positive spillovers in the younger agents’ investments determine whether or not the principal withholds some information. 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.

Our model differs from these in one critical respect. In our model 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 show that the presence of payoff types is unnecessary: destructive forces can be unleashed even when there is no objective danger to the current generation in *any* state of the world.

The present model is also closely related to Anderlini and Lagunoff (2005), Kobayashi (2003), Lagunoff and Matsui (2004), and to Anderlini, Gerardi, and Lagunoff (2006); the latter also examines private communication in a dynastic game. But while these focus on characterizing the broadest possible equilibrium set, the present paper concerns the creation and consequences of social memory in such games. Crucially, these results do not apply to the case of two dynasties and physical evidence from the past alongside private messages.

Our results also provide one coherent explanation for a number of unusual findings in lab experiments designed to capture intergenerational environments. Chaudhuri, Schotter, and Sopher (2001) and Schotter and Sopher (2001, 2003) 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. These findings accord with our results, provided that subjects are never absolutely certain that the historical evidence they receive is exact. Our results also help to explain the demonstrable effect that private versus public communication has for sustaining bad outcomes found by Chaudhuri, Schotter, and Sopher (2001).

A large strand of literature examines endogenous memory created by communication of payoff-relevant signals. For this, we refer the reader to surveys on herding by Bikhchandani, Hirshleifer, and Welch (1998) and Gale (1996), and recent work by Ahn and Suominen (2001), Banerjee and Fudenberg (2004), Jackson and Kalai (1999) and Moscarini, Ottaviani, and Smith (1998).

The role of messages in games with infinitely-lived players and private monitoring is the focus of Ben-Porath and Kahneman (1996), Compte (1998) and Kandori and Matsushima (1998). Finally, a sizeable literature has been concerned with overlapping generation games (Bhaskar, 1998, Kandori, 1992a, Salant, 1991, Smith, 1992, among others) in which participants have full memory and preferences do not encompass a dynastic component.

3. Model

3.1. A Game of Conflict

In each period two countries — France (\mathcal{F}) and Germany (\mathcal{G}) — face off in a stationary environment. The class of stage games to which our analysis applies is broad.² However, for the sake of concreteness we work with the numerical example below throughout the paper.

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

If we restrict attention to the strategies C and D for both countries, the game in (1) is a version of the Prisoners' dilemma, with efficient payoffs reached when both sides “cooperate” (play C), and a unique inefficient dominant strategy Nash equilibrium when both countries “defect” (play D). The third strategy (W) available to each country is interpreted as “war.” Both countries choosing W is an “all out war.”

War is a dominated strategy for both players (as is C). If \mathcal{F} and \mathcal{G} were to face off once only, neither could possibly choose W since D yields a higher payoff for every possible strategy of the opponent. All out war pushes both sides below their individually rational (minmax) payoffs; using the numbers in (1) these are both -8 . War is a good way to punish the other side, but it hurts the country that wages war. A country that chooses W gets a higher payoff if the other cooperates. Its payoff decreases if the other side defects (“defends itself”), and decreases further if the other side responds with war to war. Moreover, from a

²For example, our results trivially extend to any 3×3 stage game with the same ordinal payoffs as in (1). Our results also extend to a broader class of games, which we term “games of conflict.” This is composed of all finite 2-player games in which each player has three strategies that correspond to C , D and W in (1) in the following sense. If we restrict attention to the 9 combinations of these strategies (and ignore any others) the payoffs have the same ordering as in (1).

“utilitarian” point of view war is inefficient; the sum of payoffs if one or both sides choose W is at most -13 , while it is at least 0 if neither country chooses W .

It is critical to observe that if \mathcal{F} and \mathcal{G} were to face off indefinitely and were modeled as standard long-run players with perfect recall, then we could be sure that “very frequent” all out wars are ruled out. 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 are directly related to the possibility of a systematically incorrect social memory when the two countries face off repeatedly in a *dynastic game*. This is what we turn to next.

3.2. Dynastic Game

Each of our 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. The individuals who inhabit \mathcal{F} and \mathcal{G} in period t are denoted by \mathcal{F}^t and \mathcal{G}^t respectively.³ The two countries’ per-period payoffs at t are denoted by $\pi_{\mathcal{F}}^t$ and $\pi_{\mathcal{G}}^t$ respectively, and are as in (1).

In their general form the demographics of the model satisfy only one constraint. There is a *uniform upper bound* L on the length of life of all individuals in all countries.⁴ Other than this, the individuals in each country may be replaced at the same time or at different times, the length of their lives in each or both countries may be constant or changing through time. We refer to these as L -bounded demographics.

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

³Notice that at this point there is a redundancy in our notation. If the same person is alive, say, in \mathcal{F} at times t and $t + 1$, then \mathcal{F}^t and \mathcal{F}^{t+1} identify the same individual. This will become convenient in Subsection 3.4 below.

⁴The demographics we explicitly consider are *deterministic*. All our results extend easily to the case of stochastic L -bounded demographics in which no individual lives for more than L periods with probability one. This generalization involves no conceptual difficulties, but is very costly in notational terms.

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

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

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 from time $t = 0$, then the model would be equivalent to the standard repeated game world of, say, Fudenberg and Maskin (1986);⁵ incorrect social memory would be ruled out by assumption, and no frequent all out wars could take place.

When an individual (\mathcal{F}^t or \mathcal{G}^t , or both simultaneously) enters the game, he observes two things: imperfect publicly available evidence of the past history of actions s^t , and a *private* message about the previous history of actions from his predecessor, who was alive at $t - 1$ in his own country and whom he replaces. (The order in which these are observed, or if they are observed simultaneously or not, does not matter.) Let these messages be $m_{\mathcal{F}}^t$ and $m_{\mathcal{G}}^t$ respectively, depending on which individual(s) is replaced at t . All individuals alive at t (regardless of when they entered the game) observe the opponent's action at t when play takes place, and afterwards the realization of s^{t+1} . The evidence s^t is the *footprint* that the history of actions taken up to $t - 1$ leaves for the future. We think of this as physical evidence from the past; the “fossil record” of the past history of behavior.

The evidence s^t observed by individuals \mathcal{F}^t and \mathcal{G}^t at the beginning of t takes the form of a past history $h^t \in H^t$ of action profiles a^t played in each period $0, 1, \dots, t - 1$, and by convention $s^0 = \emptyset$. Of course the observed evidence may or may not correspond to the history that actually took place. The precision of s^t is parameterized by $\gamma \in [0, 1]$; 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.

The exact parametric form of the distribution of observed evidence is unimportant, but to fix ideas we assume it to be as follows. In each period, conditional on a particular action profile a^t taking place at t , a signal p^t is realized. The value of p^t is equal to 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

⁵This 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).

$(1 - \gamma)/9$. The evidence s^1 observed at the beginning of time $t = 1$ (and at the end of period $t = 0$) is simply the signal p^0 . The evidence observed at the beginning of $t = 2$ (and at the end of period $t = 1$) is $s^2 = (p^0, p^1)$, and so forth so that $s^t = (p^0, \dots, p^{t-1})$. 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 the (imprecise) historical footprint is set, it remains as given through time.

A private message ($m_{\mathcal{F}}^t$ or $m_{\mathcal{G}}^t$) also takes the form of a complete past history of action pairs h^t from 0 to $t - 1$.⁶ Individuals \mathcal{F}^t and \mathcal{G}^t first observe the evidence s^t , their private messages if they are just entering the game, and then select actions. Both then observe the action selected by the opposing individual at time t and subsequently s^{t+1} . Once period t is over, if they are to be replaced at $t + 1$, individuals \mathcal{F}^t and \mathcal{G}^t send private messages $m_{\mathcal{F}}^{t+1}$ and $m_{\mathcal{G}}^{t+1}$ to individuals \mathcal{F}^{t+1} and \mathcal{G}^{t+1} .

To summarize, each country is inhabited by a representative decision maker each period who cares about his own and future countrymen's payoffs. The demographic pattern of entry and exit of individual decision makers within a country is arbitrary (subject to a bound on lifetimes). 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.

3.3. Equilibrium

In every period individuals \mathcal{F}^t and \mathcal{G}^t choose actions as a function of what they have observed during their lives. When they are about to be replaced they also choose messages to send to individuals \mathcal{F}^{t+1} and \mathcal{G}^{t+1} . We refer to the first component of their behavior as the *t-action strategies*, denoted $\alpha_{\mathcal{F}}^t$ and $\alpha_{\mathcal{G}}^t$ respectively, and to the second as the *t-message strategies*, denoted $\mu_{\mathcal{F}}^t$ and $\mu_{\mathcal{G}}^t$ respectively. Taken together, the *t-action strategy* and the *t-message strategy* of an individual make up the individual's *full t-strategy*. The full *t-strategies* of individuals \mathcal{F}^t and \mathcal{G}^t are denoted by $\mathbf{f}_{\mathcal{F}}^t = (\alpha_{\mathcal{F}}^t, \mu_{\mathcal{F}}^t)$ and $\mathbf{f}_{\mathcal{G}}^t = (\alpha_{\mathcal{G}}^t, \mu_{\mathcal{G}}^t)$ respectively. The

⁶There 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.

array that specifies the strategies of all individuals in both countries in all periods is denoted by \mathbf{f} ; a *strategy profile*.

The beliefs of individuals \mathcal{F}^t and \mathcal{G}^t about the previous history of actions and messages may change during their lives, as the actions taken by the individual in the opposing country and new evidence from the past is observed. We distinguish between *beginning-of-period- t beliefs* and *end-of-period- t beliefs*. For individuals \mathcal{F}^t and \mathcal{G}^t , $\mathbf{b}_{\mathcal{F}}^t$ and $\mathbf{b}_{\mathcal{G}}^t$ respectively specify both the beginning and end of period t beliefs. The array of beliefs of all individuals in all periods is denoted \mathbf{b} and is called a *system of beliefs*.

Our focus is *social memory* which is embodied in the individuals' beliefs about the *action profiles* played in the past. Of course, it is the beliefs at the beginning of every period that determine the actions chosen by the current individuals. Therefore, it is convenient to establish a piece of notation for the beginning-of-period- t beliefs of individuals \mathcal{F}^t and \mathcal{G}^t about the past history of action profiles implied by $\mathbf{b}_{\mathcal{F}}^t$ and $\mathbf{b}_{\mathcal{G}}^t$. These are denoted $\mathbf{q}_{\mathcal{F}}^t$ and $\mathbf{q}_{\mathcal{G}}^t$, and correspond to *social memory* in the two countries at the beginning of period t .⁷

Since our main results revolve on the possibility of systematically incorrect social memory, it is necessary to proceed with caution as to what is required of an equilibrium system of beliefs. The most stringent conditions on equilibrium beliefs generally accepted in game theory are those imposed by the notion of *Sequential Equilibrium* (Kreps and Wilson, 1982). This requires the equilibrium beliefs to be recoverable as the limit of beliefs entirely determined by Bayes' rule using a "fully mixed" perturbation of the equilibrium strategies. This is more than a technical condition. Crucially, it requires all individuals to have a *complete* and *common* theory of possible mistakes of all other individuals. Of course, in a Sequential Equilibrium, the equilibrium strategies of all individuals must be (sequentially) rational given the equilibrium beliefs.

From now on, we refer to (\mathbf{f}, \mathbf{b}) as an equilibrium of the model if and only if it constitutes a Sequential Equilibrium of the dynastic game.

3.4. Canonical Demographics

Consider a special case of the model in which there is *full replacement every period*. That is, consider the case in which the demographics are that all individuals live *one period*, after

⁷Of course, individuals who have been alive for some time have trivial beliefs about the history of action profiles played during their lives since they directly observe them.

which they are replaced. This simpler demographic structure is sufficient to generate all the phenomena we are interested in: frequent all out wars, incorrect social memory and so on.

As it turns out, there is relatively little loss of generality in restricting attention to the “canonical” demographics with full replacement every period. This is because, using a technique known at least since Ellison (1994), any equilibrium of the model with canonical demographics has a corresponding equilibrium in the model with L -bounded demographics with the same “qualitative features.” In essence, the equilibrium in the L -bounded demographics case is constructed as L interleaved “copies” of the equilibrium with canonical demographics.⁸

In the L -bounded demographics case, 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.⁹ 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 to verify that if we are given an equilibrium of the model with canonical demographics and discount factor δ , then the construction we have outlined yields an equilibrium for the L -bounded demographics case for a discount factor of $\delta^{\frac{1}{L}}$.

Since the model with full replacement every period is both analytically and notationally much simpler, from now on we *restrict attention to this case*.¹⁰

Before finally moving on, we remark that the equilibria of the model with L -bounded demographics we have just sketched out may seem contrived. However, the point is that there are equilibria in the case of L -bounded demographics with the properties we focus on below. It is clear that they do not all have the form we have sketched out here; more “natural”

⁸The same would be true if we considered the stochastic L -bounded demographics mentioned in footnote 4 above. 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.

⁹Clearly, 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.

¹⁰Hence, from now on our notation no longer has any redundancy. This was built in before precisely so as to avoid having to change it at this point. See footnote 3 above.

equilibria with the same qualitative properties exist in virtually all cases.

4. Frequent All Out Wars

As we anticipated several times, it is possible that in equilibrium the two countries \mathcal{F} and \mathcal{G} are caught in a perpetual cycle of frequent all out wars.

Proposition 1. *Arbitrarily Frequent All Out Wars: Fix any positive integer N , and assume that δ is sufficiently close to 1.*

Then there exists a pair $(\mathbf{f}^, \mathbf{b}^*)$ in which the action profile (W, W) is played $N - 1$ times out of any N consecutive periods, which is an equilibrium of the model for any $\gamma \in [0, 1)$.*

We do not give a full formal proof of Proposition 1. This would inevitably be notationally very intensive and hardly offer more clarity than the extended sketch we give here. The equilibrium is depicted in a highly schematic way in the diagram below, and the argument proceeds in several steps.

Step 1. We first argue that it is possible to construct an equilibrium of the model in which the beliefs of every individual (beginning- and end-of-period) do not depend on the observed evidence from the past.


To see that this is possible, consider a variation of the model with the physical evidence from the past taken out entirely. Suppose that in this model we can construct an equilibrium in which 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.

Now consider adding imperfectly informative evidence from the past (so that $\gamma \in [0, 1)$) to the model. Is the equilibrium we just described still viable? Clearly, we only need to argue that the postulated beliefs can still obtain in equilibrium. (The strategies of all individuals must be sequentially rational given these beliefs, simply because they constitute an equilibrium of the model without physical evidence from the past.) To see that this is the case consider the same completely mixed strategies (as required for a Sequential Equilibrium) that were used in the model without evidence, and construct completely mixed strategies in the model with imperfectly informative evidence, using the same mixtures for *every possible* realization of the physical evidence from the past. Since the beliefs were degenerate in the equilibrium of the model without evidence from the past, and the per period signals have full support, it is now easy to verify that using Bayes' rule and shrinking the perturbations to

zero yields the same beliefs in the model with an imperfect imperfectly informative evidence from the past, regardless of its realization.

Step 2. In the construction outlined below, we work with message spaces that are smaller than the sets of finite histories H^t . This is without loss of generality since it is well known that enlarging the message spaces can never shrink the set of equilibrium outcomes. The reason for this is exactly the same as in many other cheap-talk models (see Farrell and Rabin (1996) for a survey), and we do not expand further on it here.

Step 3. The strategies that sustain the outcome described in Proposition 1 can be specified in terms of “phases” and how the time periods are partitioned.¹¹ These are depicted in a highly schematic way in the diagram below.



Fix N as in the statement of Proposition 1. The first $N - 1$ periods are “war periods,” and the N^{th} period is a “peace period.” The next $N - 1$ periods are again war periods, and these are again followed by a single peace period. This cycle of length N is repeated without bound.¹²

¹¹Clearly, Proposition 1 identifies a *class* of equilibria, parameterized by N . Since they all share the same structure we refer to them as a single equilibrium. Moreover, from now on, when we refer to the equilibrium in Proposition 1, we mean the actual equilibrium described in this extensive sketch

¹²Since 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.

The action pair (W, W) is chosen in every war period, *regardless* of which phase we are in. On the other hand, the messages sent during war periods do change according to the phase. So, in the description of phases that follow, we only describe which actions are taken during the peace periods. During the war periods, provided that the all out war pair (W, W) is chosen, both individuals simply repeat to their successors the message they receive from their predecessors.

The two countries start off in the *equilibrium phase*, and remain in this phase unless a deviation occurs. During the equilibrium phase the cooperative action pair (C, C) is chosen in all peace periods. During the equilibrium phase both individuals \mathcal{F}^t and \mathcal{G}^t send message m^* to their successors provided that no deviations are observed at t .

If individual \mathcal{F}^t at any point (in a war or in a peace period, in any phase) deviates from the prescriptions of the equilibrium, then the \mathcal{F} *punishment phase* starts (or re-starts, as appropriate). Similarly, any deviation by any individual \mathcal{G}^t starts (or re-starts, as appropriate) the \mathcal{G} *punishment phase*. Double deviations are ignored. Each of these phases, if no subsequent deviations occur, lasts until the next T , with T “large,” *peace periods* have passed, after which time the two countries return to the equilibrium phase. During the \mathcal{F} punishment phase the pair (C, D) is chosen in the peace periods, while during the \mathcal{G} punishment phase the pair (D, C) is chosen in the peace periods. During the \mathcal{F} punishment phase, both individuals \mathcal{F}^t and \mathcal{G}^t send message $m^{\mathcal{F},\tau}$ to their successors. 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.¹³ During the \mathcal{G} punishment phase both individuals \mathcal{F}^t and \mathcal{G}^t send message $m^{\mathcal{G},\tau}$ to their successors, counting down in the same way.

Step 4. The individuals’ beliefs in equilibrium are independent of the evidence s^t . By Step 1 above, so long as the beliefs are degenerate this is a feature of equilibrium.

The individuals’ beliefs are “correct” in the equilibrium phase. In other words, upon receiving m^* both individuals \mathcal{F}^t and \mathcal{G}^t believe with probability one that the other has also received message m^* . Given these beliefs and the strategies of all future individuals in both countries, given that δ is close to one, it is in the interest of both individuals not to deviate from choosing W in a war period and C in a peace period. This is because any deviation will cause the payoff to their country to decrease (from 2 to -1) in a large number (T) of future

¹³If $t < TN - 1$ the smaller values of the count-down index clearly have no meaning. The message sets we have in mind are suitably smaller when $t < TN - 1$, so that only meaningful messages are included.

peace periods.

Step 5. Now consider the \mathcal{F} punishment phase. In this phase the beliefs of individual \mathcal{G}^t are also “correct.” After receiving $m^{\mathcal{F},\tau}$ he believes with probability one that individual \mathcal{F}^t has received the same message $m^{\mathcal{F},\tau}$. It is not hard to see how this might be the case in a Sequential Equilibrium and we do not pursue this further here (see footnote 14 for further details). Given these beliefs an argument similar to the one we used for the equilibrium phase is sufficient to see that no individual \mathcal{G}^t has an incentive to deviate during the \mathcal{F} punishment phase. A symmetric argument establishes that no individual \mathcal{F}^t has an incentive to deviate during the \mathcal{G} punishment phase.

Step 6. Consider an individual \mathcal{F}^t during the \mathcal{F} punishment phase, or an individual \mathcal{G}^t during the \mathcal{G} punishment phase. We concentrate on the former case since the latter is completely symmetric. The first difficulty arises when we consider the beginning-of-period beliefs of individual \mathcal{F}^t during the \mathcal{F} punishment phase. Imagine that we are at the beginning of such phase. The continuation equilibrium is bad for country \mathcal{F} for many periods to come. In fact its continuation payoff is sufficiently worse than the continuation payoff during the equilibrium phase so as to deter deviations during the equilibrium phase.

Why should then individual \mathcal{F}^t go along and submit his country to the punishment prescribed by the equilibrium strategies? The answer is that, on receiving $m^{\mathcal{F},\tau}$ he *incorrectly believes* that individual \mathcal{G}^t has received the message m^* . But how can these beliefs be consistent, as required by equilibrium? The answer lies in the fact that upon receiving $m^{\mathcal{F},\tau}$ (or any other off-the-equilibrium-path message) an individual must weigh several (infinitesimally unlikely) possibilities. He could have received $m^{\mathcal{F},\tau}$ because play is in fact in the equilibrium phase and individual \mathcal{F}^{t-1} mistakenly sent message $m^{\mathcal{F},\tau}$ instead of m^* . Alternatively, he could have received $m^{\mathcal{F},\tau}$ because some previous individual in country \mathcal{F} mistakenly took the wrong action during the equilibrium phase. If the ratio of the second infinitesimal divided by the first is zero, then individual \mathcal{F}^t will in fact have the beliefs we have described. He will choose C , mistakenly thinking that individual \mathcal{G}^t will do the same.¹⁴

¹⁴ For instance, the fact that the beliefs we are postulating satisfy consistency as required for a Sequential Equilibrium can be verified in the following way. The “trembles” used to construct the fully mixed strategies, all expressed in terms of a common variable ε to be shrunk to zero are as follows. Whenever the equilibrium prescribes an action different from D , 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

Step 7. Finally, consider the end-of-period beliefs of an individual \mathcal{F}^t during the \mathcal{F} punishment phase (the argument is symmetric for an individual \mathcal{G}^t during the \mathcal{G} punishment phase). At the beginning of period t , when he receives $m^{\mathcal{F},\tau}$, he mistakenly expects individual \mathcal{G}^t to choose C . But at the end of period t he discovers that the equilibrium is in the \mathcal{F} punishment phase since he *observes* that individual \mathcal{G}^t has in fact chosen D . However, whatever attempt he makes at communicating to his successor \mathcal{F}^{t+1} that play is in the \mathcal{F} punishment phase is doomed to fail. Given the strategies of all successive individuals, it is straightforward to check that sending any message other than $m^{\mathcal{F},\tau-1}$ (m^* if $\tau = 1$) will eventually cause the equilibrium to lock into (D, D) in *every* future peace period. On the other hand sending $m^{\mathcal{F},\tau-1}$ (m^* if $\tau = 1$) will eventually return the countries to the equilibrium phase. This is preferable from the point of view of \mathcal{F}^t because δ is close to 1 and in the equilibrium phase the action profile (C, C) is chosen during every peace period.

The sketch of the proof of Proposition 1 is now complete. Before moving on, we proceed with two remarks.

Remark 1. *All Out Wars and Private Communication:* The possibility of private communication between one generation of individuals and the next is critical for the possibility of frequent all out wars. To see this, consider the dynastic game described above but with messages taken out entirely. Then, for every t , the long-run continuation payoffs in (2) cannot be below -8 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.”¹⁵ 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 -8 .

Remark 2. *All Out Wars and Perfect Signals:* Proposition 1 is false if the evidence s^t is perfectly informative about the past history of action profiles.

In particular, if we set $\gamma = 1$, the long-run continuation payoffs in (2) cannot be below -8 for either \mathcal{F} or \mathcal{G} , and for every t . Therefore the all out war equilibrium of Proposition 1 is no longer viable.

equilibrium strategies is ε . The probability of all other messages being sent instead of any other equilibrium ones is ε^4 .

¹⁵The 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.

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.¹⁶ 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 (since $\gamma = 1$) the action pair that took place at 0. 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 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.¹⁷ This, as we saw above, implies that the long-run continuation payoffs in (2) can never be below -8 for either country.

5. Social Memory

In the equilibrium of Proposition 1, after some histories, social memory is incorrect in an obvious intuitive sense (see Step 6 of the sketch of proof above). On the other hand, the individuals have access to imperfectly informative evidence from the past that can be arbitrarily precise. Our focus here is on how this can occur in the first place, and to examine

¹⁶The 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.

¹⁷The 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.

“how wrong” social memory can really be.

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. Social memory could be incorrect even though beliefs are sensitive to the evidence from the past. And social memory could be full and correct even though physical evidence from the past is disregarded entirely because the messages from one generation to the next can in principle convey all available information about the past history of action profiles.¹⁸

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 be truthful that they may override the evidence entirely in the equilibrium beliefs. We view this as an appealing and realistic attribute of the model. Witnesses of the Holocaust have more information than could possibly be conveyed by the physical remains of Auschwitz.

In the equilibrium of Proposition 1 messages override the evidence from the past and create social memory that is incorrect in a strong sense of the word. Our next task is to make this precise.

5.1. *Statistical Testing*

Because the individuals observe imperfect but informative evidence of the past history of action profiles, there is a natural candidate to “measure” how incorrect social memory really is. This way of proceeding seems particularly appealing if we want to argue that social memory can be “maximally incorrect,” as turns out to be the case in the equilibrium in Proposition 1.

In short, we can compare the individuals’ social memory at t — the beginning-of-period- t

¹⁸It 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.

beliefs about previous action profiles, denoted by $(\mathbf{q}_{\mathcal{F}}^t, \mathbf{q}_{\mathcal{G}}^t)$ — with the inferences that an outside observer (a “statistician”) — who knows nothing except for the realized evidence s^t and its conditional distributions — would draw about the past history of play. Since the evidence is informative about the past it could, in the limit as time goes by, make the statistician “fully convinced” that certain statements are false (or true) about the past. If the beliefs of one or more individuals were the reverse of what the statistician came to believe, there would be good reason to say that social history is *maximally incorrect* in this case.

To flesh out this idea we need to be precise about what tests the statistician performs to make inferences, and above all what it means that, in the limit, she becomes fully convinced that some event did not or did occur.

The null hypothesis is a subset H_0^t of the set H^t of histories of action profiles of length t . For instance, this could be of the form “the action pair (C, C) has been taken n or more times in the first t periods,” or “the action pair (D, D) never took place before t ,” and so on. Given the evidence s^t the statistician then asks whether the null hypothesis can be rejected or not at a given significance level using a standard likelihood ratio test.¹⁹

Denote by $\Lambda(s^t, H_0^t)$ the actual value of the likelihood ratio for the null hypothesis H_0^t , given the evidence s^t . The test rejects the null H_0^t at the significance level $\alpha \in (0, 1)$ if $\Lambda(s^t, H_0^t)$ is between zero and a cut-off level $\lambda(\alpha, H_0^t)$ that depends on α and H_0^t .

Further, given a particular history of action profiles, say h^{t*} , not contained in the null H_0^t we can ask the following question. If h^{t*} were in fact the true history, what is the probability that the observed evidence is such that the likelihood ratio test will reject H_0^t at the α confidence level?²⁰ If this probability is η or more we say that, given h^{t*} the null hypothesis H_0^t is rejected at the α confidence level with probability at least η .

If, in the limit as t grows large, α can be made arbitrarily small and η arbitrarily large,²¹ we can legitimately say that, given h^{t*} the statistician would eventually become *fully convinced* that the event “the true history is an element of H_0^t ” did not occur.

¹⁹See Section A.1 in the Appendix for further details concerning the statistical procedures we refer to here. Since these are completely standard, they are omitted from the main body of the paper for the sake of brevity. See for instance the textbook by Silvey (1975).

²⁰Given h^{t*} and α , this probability is known as the *power* of the test.

²¹In other words if we fix any $\alpha > 0$ and $\eta < 1$, then for t sufficiently large given h^{t*} the null H_0^t is rejected at the α confidence level with probability η . See Section A.1 in the Appendix for details.

5.2. Maximally Incorrect Social Memory

Grotesquely incorrect accounts of the past are regrettably easy to find.

“They have created a myth today that they call the massacre of Jews and they consider it a principle above God, religions and the prophets” — Iranian President Mahmoud Ahmadinejad on live state television in December 2005.²²

Intuitively, social memory is maximally incorrect if for some null H_0^t we find that the individuals believe H_0^t to be true with probability one, while the statistician will eventually become fully convinced that H_0^t is in fact false. The individuals’ beliefs are that H_0^t is true despite mounting evidence that it is false. A little extra notation is needed to proceed with the actual claim.

Consider the equilibrium in Proposition 1. Fix any time \bar{t} and assume that up to and including period $\bar{t} - 1$ no deviations from equilibrium occur. Consider then any deviation by individual $\mathcal{F}^{\bar{t}}$ at the action stage of period \bar{t} , followed by a deviation by the same individual at the message stage of period \bar{t} , sending some $m^{\mathcal{G},\tau}$ instead of $m^{\mathcal{F},T}$ as prescribed.²³

Now consider any period $t > \bar{t}$, and assume no further deviations occur after \bar{t} . Let \hat{h}^t be the actual history of play of length t (following the deviations at \bar{t}) prescribed by the equilibrium of Proposition 1.

Proposition 2. *Maximally Incorrect Social Memory:* Fix any precision level $\gamma \in (0, 1)$ for the physical evidence of past history and consider the equilibrium of Proposition 1.

Let \hat{h}^t be as above, and for every t let \hat{H}_0^t represent the null hypothesis “the action pair (D, D) has never been chosen before t .”

Then, the equilibrium of Proposition 1 has the following features.

(i) Following \hat{h}^t , the social memory $\mathbf{q}_{\mathcal{F}}^t$ and $\mathbf{q}_{\mathcal{G}}^t$ of all individuals \mathcal{F}^t and \mathcal{G}^t is that the null hypothesis \hat{H}_0^t is true with probability one.

(ii) Following \hat{h}^t , the statistician observing the evidence will eventually (as t becomes large) become fully convinced that \hat{H}_0^t is in fact false.

Thus, after certain histories of play, social memory in both countries is maximally incorrect. In spite of mounting evidence to the contrary, both individuals \mathcal{F}^t and \mathcal{G}^t at the beginning

²²As quoted in English translation by the B.B.C. (2005).

²³We could consider the symmetric deviations, by $\mathcal{G}^{\bar{t}}$ at the action and message stages, reaching exactly the same conclusions as in Proposition 2 below.

of t believe that play is in the punishment phase for the opposing country, while (as is clear from Step 7 of the sketch of the proof of Proposition 1) it is in fact locked into a perpetual cycle of $N - 1$ periods in which all out war (W, W) takes place followed by one period in which (D, D) is chosen.

Maximally incorrect social memory does not require that the individuals be unanimous in their mistaken beliefs. After certain histories of play (involving one deviation, by one individual, at the action stage), social memory in one country alone (say \mathcal{F}) is maximally incorrect.²⁴ Despite mounting evidence to the contrary, individual \mathcal{F}^t at the beginning of period t believes that play is in the equilibrium phase while it is (and has been for a long time) in fact in the \mathcal{F} punishment phase. This takes place while social memory in country \mathcal{G} is correct; individual \mathcal{G}^t believes with probability one that play is in the \mathcal{F} punishment phase.

5.3. Evidence-Driven and Evidence-Determined Equilibria

We mentioned above that one of the critical ingredients that yields equilibrium beliefs driven by messages rather than by evidence from the past is the fact that messages are sent after individuals \mathcal{F}^t and \mathcal{G}^t observe the actual play in period t . Therefore they can, in principle, convey more information than the evidence from the past, no matter how precise the latter is provided it is not perfect: eyewitness' accounts can override imperfect physical evidence.

To put this claim to the test, 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*.

In this modified dynastic game with imperfect current monitoring, social memory is determined by physical evidence alone in a well defined sense. We will show that in this case social memory cannot be systematically wrong. In turn this implies that frequent all out war equilibria like the one in Proposition 1 are no longer viable.

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 in equilibrium their assessment

²⁴Strictly speaking, we need δ to approach 1 and the length of the punishment phases to become unboundedly large to ensure that social memory is maximally incorrect in this case.

of the history of play qualitatively *improves* in the sense that it cannot be systematically wrong as before. We need to be more specific about what is meant by social memory being determined by physical evidence alone.

Consider either the original dynastic game or the one with imperfect current monitoring. Suppose now that all individuals *ignore the messages* they receive, and only base their beliefs about the previous *actions* taken and *messages* sent by the *opposing country* on the evidence s^t . In particular, the individuals' beginning-of-period beliefs about what action the opposing individual is going to choose must also depend only on the evidence s^t . Hence, in this case the individuals' beginning-of-period expected payoffs from taking any given action must also depend only on s^t . We call an equilibrium in which the individuals' beliefs have the property we have just described an *evidence-driven* one.

As we anticipated, degrading the information that each individual has about the actions in the *current period* has a dramatic effect on social memory. Our first characterization of this effect is embodied in the next proposition.²⁵

Proposition 3. *Evidence-Driven Equilibria: Every pure strategy equilibrium under imperfect current monitoring is evidence-driven.*

In an evidence-driven equilibrium expected continuation payoffs only depend on s^t . However, an individual's choice of action may depend on the message he receives, provided his expected continuation payoff does not change. Hence, to sharpen the characterization of the extent to which the essential behavior of the model under imperfect current monitoring is effectively determined by physical evidence only, we refer to a notion of equivalence among equilibria.

Consider the dynastic game with either perfect or imperfect current monitoring. We say that two equilibria are *equivalent* if and only if they yield the same beginning-of-period and end-of-period expected payoffs, after any possible history of play, in terms of actions, messages and observed evidence.

An equilibrium is *evidence-determined* if and only if the beliefs (about any part of the previous history of actions and/or messages) and both the action and message strategies of all individuals do not depend on the messages received. Our second characterization of the effects of degraded information is the following.

²⁵Here, and in the remainder of the paper, we focus exclusively on pure strategy equilibria. The full-blown analysis of the mixed strategy case is well beyond the scope of this paper.

Proposition 4. *Evidence-Determined Equilibria: Every pure strategy equilibrium under imperfect current monitoring is equivalent to an equilibrium that is evidence-determined.*

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* after any history of action profiles, whether deviations from equilibrium have occurred or not.

Building on the notation we established in Subsection 3.3 above, let $\mathbf{q}_{\mathcal{F}}^t(s^t)$ and $\mathbf{q}_{\mathcal{G}}^t(s^t)$ be the social memory of the two countries at the beginning of t , with the dependence on s^t made explicit. Recall that both $\mathbf{q}_{\mathcal{F}}^t(s^t)$ and $\mathbf{q}_{\mathcal{G}}^t(s^t)$ are probability distributions over the set H^t of previously chosen action profiles. Given any $\bar{t} < t - 1$, it will be convenient to denote by $\mathbf{q}_{\mathcal{F}}^t(\bar{t}, s^t)$ and $\mathbf{q}_{\mathcal{G}}^t(\bar{t}, s^t)$ the probability distributions over action profiles in periods $\bar{t}, \dots, t - 1$ only implied by (the marginals of) $\mathbf{q}_{\mathcal{F}}^t(s^t)$ and $\mathbf{q}_{\mathcal{G}}^t(s^t)$.

Fix any time \bar{t} and any history $h_+^{\bar{t}}$ of *actions and messages* up to and including the action stage of $\bar{t} - 1$.²⁶ As in Subsection 5.2, we then wish to consider histories that follow $h_+^{\bar{t}}$, assuming no further deviations after the action stage of $\bar{t} - 1$. Given any $t \geq \bar{t}$, let $\mathbf{e}^t(h_+^{\bar{t}}, s^t)$ be the actual equilibrium distribution over action profiles in periods \bar{t}, \dots, t , following $h_+^{\bar{t}}$, given the evidence s^t .

In essence, in the dynastic game 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. Any equilibrium under imperfect current monitoring must be equivalent to an evidence-determined equilibrium in the payoff sense we have specified before. In any evidence-determined equilibrium social memory must be correct on or off path.

Proposition 5. *Correct Social Memory: Fix any precision level $\gamma \in [0, 1]$ and any evidence-determined equilibrium with either perfect or imperfect current monitoring. Consider any \bar{t} and any history $h_+^{\bar{t}}$ as above.*

Then, for every $t > \bar{t}$ and s^t it must be that $\mathbf{q}_{\mathcal{F}}^t(\bar{t}, s^t) = \mathbf{q}_{\mathcal{G}}^t(\bar{t}, s^t) = \mathbf{e}_{\mathcal{F}}^t(h_+^{\bar{t}}, s^t) = \mathbf{e}_{\mathcal{G}}^t(h_+^{\bar{t}}, s^t)$. In other words, in any evidence-determined equilibrium, the social memory of both countries about previous action profiles chosen after \bar{t} must be correct.

²⁶So 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.

Clearly, Proposition 5 undoes the possibility of any discrepancy between the “findings” of a statistician observer and social memory. In any evidence-determined equilibrium, social memory correctly identifies the process that generates the actual data observed by the statistician (possibly aside from a finite number of periods \bar{t}). As t becomes large, after any $h_+^{\bar{t}}$, any statistical test of social memory as a null hypothesis will be no more likely to reject it than if the null were in fact true.

When we apply the definition of evidence-determined (or evidence-driven) to the model with either perfect or imperfect current monitoring, all equilibria with frequent all out wars are ruled out. Using the same logic as in Remark 1 it is straightforward to see that the following is true.

Remark 3. *Above Min-Max Payoffs:* Fix any precision level $\gamma \in [0, 1)$ and consider the model with either perfect or imperfect current monitoring. Consider any evidence-determined equilibrium. Then the long-run the payoff to both countries cannot be below -8 . Therefore the all out war equilibrium of Proposition 1 is no longer viable.

Thus, there is a sense in which degrading the information that individuals have about actions currently taken makes both countries better off.

6. Concluding Remarks

History is littered with instances of destructive conflicts. Some of them appear so utterly incomprehensible that, in our view, existing theories have trouble explaining them.

A standard signalling model, for example, could possibly be used to justify conflict as a rational response to a partially revealing strategy. The drawback of this explanation is that it requires 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 cases (World War I being a case in point), the conflict appears so universally destructive that it would be hard to imagine rational agents believing this.

Another justification could come from the theory of repeated games. In this scenario, the repeated game model might explain a given conflict as one of many possible equilibrium outcomes of the Folk Theorem. The problem is that even the Folk Theorem puts limits on what can happen — only payoff outcomes above each player’s minmax can be possibly sustained. Yet, the long-lasting devastation in some conflicts is so total that it seems hard to conceive that such outcomes would lie above the minmax.

In this paper, we wish to maintain the “optimistic hypotheticals” of these potential explanations. That is, we want to maintain the assumption of fully rational individuals who have common knowledge about the destructive nature of the conflict. Among other things, this rules out explanations that rely on exogenous states or preferences in which the conflict is actually desirable.

Consequently, we study 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 are maximally 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 maximally incorrect precisely because it relies on sources that can be more informative than hard evidence.

Two further issues bear mentioning here. First, a natural question concerns the issue of social memory on path. Our use of equilibrium theory leaves little scope to address the issue. On path, social memory cannot be systematically wrong. Yet, the phenomena concerning social memory we have highlighted here take place off-path in the equilibrium of Proposition 1. Hence, it is important that the equilibrium we construct in Proposition 1 passes a critical robustness test for Sequential Equilibrium. The strategies used in this equilibrium remain sequentially rational even if we consider the trembles that generate the equilibrium beliefs when they are arbitrarily small, but before they have completely shrunk to zero.²⁷ Intuitively, the robustness property we have just claimed tells us that the equilibrium on which we have focused throughout the paper survives the possibility that play does in fact stray off the

²⁷In jargon, we are claiming that the equilibrium in Proposition 1 is “trembling-hand perfect in the agent normal form” of the game.

equilibrium path with positive (albeit arbitrarily small) probability. On the other hand, systematically incorrect social memory may be possible on path for reasons other than what we focus on in this paper. It might emerge, for example, from subjective beliefs that do not come from a common prior. At this stage we leave this issue for future research.

Second, we have focused our attention entirely on “bad” equilibria with frequent all out wars. The implication is that systematically 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 (C, C) 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 (C, C) 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 issue of how the possibility of systematically inaccurate social memory might lead to the emergence of better equilibria is clearly both interesting and potentially important. This is another issue we leave for future research.

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 (2006): “A ‘Super’ Folk Theorem for Dynastic Repeated Games,” Georgetown University, *mimeo*, <http://www.georgetown.edu/faculty/la2/folktheorem.pdf>.
- 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.
- B.B.C. (2005): “http://news.bbc.co.uk/2/hi/middle_east/4527142.stm,” British Broadcasting Corporation Web-Site, London.
- BEN-PORATH, E., AND M. KAHNEMAN (1996): “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 70, 281–297.
- 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.
- CHAUDHURI, A., A. SCHOTTER, AND B. SOPHER (2001): “Talking Ourselves to Efficiency: Coordination in an Inter-generational Minimum Game with Private, Almost Common and Common Knowledge of Advice,” Working Paper 01-11, C.V. Starr Center for Applied Economics, New York University.
- COMPTE, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66, 597–626.
- 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.

- DESSI, R. (2005): “Memory, Cultural Transmission and Investments,” University of Toulouse, IDEI Working Paper No. 316, <http://idei.fr/doc/wp/2005/memory.pdf>.
- 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.
- FARRELL, J., AND M. RABIN (1996): “Cheap Talk,” *The Journal of Economic Perspectives*, 10, 103–118.
- FENTRESS, J., AND C. WICKHAM (1992): *Social Memory*. Oxford, UK: Blackwell.
- FUDENBERG, D., D. K. LEVINE, AND E. S. 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): *Game Theory*. Cambridge: MIT Press.
- GALE, D. (1996): “What Have We Learned From Social Learning?,” *European Economic Review*, 40(3-5), 617–628.
- GLAESER, E. (2005): “The Political Economy of Hatred,” *Quarterly Journal of Economics*, 120, 45–86.
- 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.

- KANDORI, M., AND H. MATSUSHIMA (1998): "Private Observation, Communication and Collusion," *Econometrica*, 66, 627–652.
- KOBAYASHI, H. (2003): "Folk Theorem for Infinitely Repeated Games Played by Organizations with Short-Lived Members," Osaka Prefecture University, mimeo.
- KOCHERLAKOTA, N. (1998): "Money is Memory," *Journal of Economic Theory*, 81, 232–251.
- KOCHERLAKOTA, N., AND N. WALLACE (1998): "Incomplete Record Keeping and Optimal Payout Arrangements," *Journal of Economic Theory*, 81, 272–289.
- KREPS, D. M., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.
- LAGUNOFF, R., AND A. MATSUI (2004): "Organizations and Overlapping Generations Games: Memory, Communication, and Altruism," *Review of Economic Design*, 8, 383–411.
- MOSCARINI, G., M. OTTAVIANI, AND L. SMITH (1998): "Social Learning in a Changing World," *Economic Theory*, 11(3), 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.
- RESTON, J. J. (2005): *Dogs of God: Columbus, the Inquisition, and the Defeat of the Moors*. New York, NY: Random House.
- 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.
- SCHOTTER, A., AND B. SOPHER (2001): "Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach," Working Paper 01-04, C.V. Starr Center for Applied Economics, New York University.

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.

SILVEY, S. D. (1975): *Statistical Inference*. New-York, NY: Chapman and Hall.

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.

WALLACE, N. (2001): “Whither Monetary Economics?,” *International Economic Review*, 42, 847–869.

Appendix

A.1. The Likelihood Ratio Test

Purely for the sake of completeness, we make explicit the formal structure of the statistical testing procedures used in the text. As we remarked already, these are completely standard (see for instance Silvey, 1975).

Remark A.1. *Likelihood Ratio Test:* For any fixed $h^t \in H^t$, let $L(s^t, h^t)$ be the probability that the realized evidence is s^t given that the actual history is h^t . Given s^t , as h^t varies $L(s^t, h^t)$ is known as the likelihood function of h^t .

Fix a set $H_0^t \subset H^t$ of histories of length t and let

$$\Lambda(s^t, H_0^t) = \frac{\max_{h^t \in H_0^t} L(s^t, h^t)}{\max_{h^t \in H^t} L(s^t, h^t)} \quad (\text{A.1})$$

The quantity $\Lambda(s^t, H_0^t)$ is known as the likelihood ratio. For any given $h^t \in H^t$, $\Lambda(s^t, H_0^t)$ can be viewed as the realization of a random variable (as s^t takes different random values).²⁸ Given any $E \subseteq \mathbb{R}_+$ and any h^t , denote by $\Pr_{h^t}[\Lambda(s^t, H_0^t) \in E]$ the probability that this random variable takes any value in E given that h^t is the true history of past action profiles.

Let $\alpha \in (0, 1)$ be given. It is common to say that the likelihood ratio test rejects the null hypothesis H_0^t at the α significance level if $\Lambda(s^t, H_0^t) \leq \lambda(\alpha, H_0^t)$, where $\lambda(\alpha, H_0^t) = \min_{h^t \in H_0^t} \lambda(\alpha, h^t, H_0^t)$ and each $\lambda(\alpha, h^t, H_0^t)$ is defined as the maximum $\lambda \geq 0$ such that

$$\Pr_{h^t} \{ \Lambda(s^t, H_0^t) \in [0, \lambda] \} \leq \alpha \quad (\text{A.2})$$

²⁸Throughout, we use the same notation for a random variable and its realizations since this will not cause any ambiguity.

Definition A.1. *Rejection at α Level With Probability η :* Let a null hypothesis H_0^t be given. Let also a particular history h^{t*} be given.

We say that given h^{t*} the null hypothesis H_0^t is rejected at the α confidence level with probability η if

$$\Pr_{h^{t*}} \{ \Lambda(s^t, H_0^t) \in [0, \lambda(\alpha, H_0^t)] \} \geq \eta \quad (\text{A.3})$$

A.2. Proof of Proposition 2

To facilitate the exposition, we break the argument into 7 steps.

Step 1. From the strategies in Step 3 of the sketch of proof of Proposition 1, it is straightforward to check that any history of action profiles \hat{h}^t as in the statement of the Proposition must be as follows.

Before \bar{t} , the action pair (W, W) is chosen in all war periods, and the pair (C, C) is chosen in all peace periods. If \bar{t} is a war period, then at \bar{t} , individual $\mathcal{F}^{\bar{t}}$ chooses either C or D and individual \mathcal{G} chooses W . If \bar{t} is a peace period, then at \bar{t} , individual $\mathcal{F}^{\bar{t}}$ chooses either W or D , and individual $\mathcal{G}^{\bar{t}}$ chooses C . After \bar{t} the action pair (W, W) is chosen in every war period, and the action pair (D, D) is chosen in every peace period.

Step 2. From the strategies in Step 3 of the sketch of proof of Proposition 1, it is straightforward to check that at the beginning of each period, following any \hat{h}^t , individual \mathcal{F}^t receives message $m^{\mathcal{G}, T}$ and individual \mathcal{G}^t receives message $m^{\mathcal{F}, T}$.

From Steps 4, 5 and 6 of the sketch of proof of Proposition 1, it is then immediate to conclude that at the beginning of each period, following any \hat{h}^t , individual \mathcal{F}^t believes with probability one that play has been in the equilibrium phase from period 0 to period $t - 2$, and that individual \mathcal{G}^{t-1} has deviated in period $t - 1$ and hence started the \mathcal{G} punishment phase. Symmetrically, at the beginning of each period, following any \hat{h}^t , \mathcal{G}^t believes with probability one that play has been in the equilibrium phase from period 0 to period $t - 2$, and that individual \mathcal{F}^{t-1} has deviated in period $t - 1$ and hence started the \mathcal{F} punishment phase.

Clearly, at the beginning of each period, following any \hat{h}^t , both \mathcal{F}^t and \mathcal{G}^t believe with probability one that the action pair (D, D) has never been played before. Thus, we have already shown (i) of the statement of Proposition 2 to be true.

Step 3. Given the structure of the per-period signals we have described in the text, it is immediate that, for any t , s^t and $h^t \in H^t$ it must be that

$$s^t = \arg \max_{h^t \in H^t} L(s^t, h^t) \quad (\text{A.4})$$

and hence that for any t and any $h^t \in H^t$ it must be that

$$\left[\frac{1 + 8\gamma}{9} \right]^t = \max_{h^t \in H^t} L(s^t, h^t) \quad (\text{A.5})$$

Step 4. Let \hat{H}_0^t be the null consisting of the set of histories of action profiles length t such that $a^\tau \neq (D, D)$ for every $\tau = 0, \dots, t - 1$. This is the null of the statement of Proposition 2.

For any s^t , let $\mathbf{D}(s^t)$ the number of occurrences of the action profile (D, D) in the string of action profiles that make up s^t . Given the structure of the per-period signals we have described in the text, it is immediate

that, for any t , s^t and $\gamma \in [0, 1)$ we have that

$$\left[\frac{1+8\gamma}{9} \right]^{t-\mathbf{D}(s^t)} \left[\frac{1-\gamma}{9} \right]^{\mathbf{D}(s^t)} = \max_{h^t \in \hat{H}_0^t} L(s^t, h^t) \quad (\text{A.6})$$

Step 5. From Steps 3 and 4 above, it is immediate that for any t , s^t and $\gamma \in [0, 1)$

$$\left[\frac{1-\gamma}{1+8\gamma} \right]^{\mathbf{D}(s^t)} = \Lambda(s^t, \hat{H}_0^t) = \frac{\max_{h^t \in \hat{H}_0^t} L(s^t, h^t)}{\max_{h^t \in H^t} L(s^t, h^t)} \quad (\text{A.7})$$

Step 6. From (A.7), given the structure of the per-period signals we have described in the text, it is clear that the probability distribution of $\Lambda(s^t, \hat{H}_0^t)$ is the same under any true history $h^t \in \hat{H}_0^t$. In particular, under \hat{H}_0^t the probability that $\mathbf{D}(s^t)$ equals any given value between 0 and t is given by a binomial with a “success probability” equal to $(1-\gamma)/9$. From the (weak) Law of Large Numbers we then know that, given any sequence of histories of action profiles $\{h^t \in \hat{H}_0^t\}_{t=1}^\infty$, we must have that

$$\lim_{t \rightarrow \infty} \Pr_{h^t} \left\{ \mathbf{D}(s^t) < t \left[\frac{1-\gamma}{9} + \frac{\gamma}{9N} \right] \right\} = 1 \quad (\text{A.8})$$

As we have seen above, after $\bar{t} + 1$, the histories \hat{h}^t that follow the deviations at \bar{t} display the profile (D, D) in every peace period, and hence every N periods. A completely analogous argument to the one used to establish (A.8) can then be used to establish that, given any sequence of histories $\{\hat{h}^t\}_{t=1}^\infty$ as in the statement of the Proposition, we must have that

$$\lim_{t \rightarrow \infty} \Pr_{\hat{h}^t} \left\{ \mathbf{D}(s^t) > t \left[\frac{1-\gamma}{9} + \frac{\gamma}{9N} \right] \right\} = 1 \quad (\text{A.9})$$

Step 7. Let $\lambda(\alpha, \hat{H}_0^t)$ be a threshold rejection value as in Remark A.1. Then, from (A.7) and (A.8) it follows that given any $\alpha \in (0, 1)$, we can find a \tilde{t} such that

$$t \geq \tilde{t} \quad \Rightarrow \quad \lambda(\alpha, \hat{H}_0^t) > \left[\frac{1-\gamma}{1+8\gamma} \right]^t \left[\frac{1-\gamma}{9} + \frac{\gamma}{9N} \right] \quad (\text{A.10})$$

Finally, using (A.7), (A.10) and (A.9) it is clear that, given any $\alpha \in (0, 1)$ and any $\eta \in [0, 1)$, we can find a t^* such that

$$t \geq t^* \quad \Rightarrow \quad \Pr_{\hat{h}^t} \left\{ \Lambda(s^t, \hat{H}_0^t) \in [0, \lambda(\alpha, \hat{H}_0^t)] \right\} \geq \eta \quad (\text{A.11})$$

and hence the proof is now complete. ■

A.3. Proof of Proposition 3

We argue that the Proposition is true for the beliefs of the individuals 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.

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 and 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 (pure) action strategies determine the equilibrium action profile a^{0*} . Given p^0 and the (pure) 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 (pure) 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 concerning the history of actions and messages taken by dynasty \mathcal{G} do not depend on $m_{\mathcal{F}}^t$. Hence the proof is now complete. ■

A.4. Proof of Proposition 4

Consider the evidence-driven equilibrium described in detail in the proof of Proposition 3 (Section A.3). Denote the strategy profile and system of beliefs constituting this equilibrium by $\hat{\mathbf{f}}$ and $\hat{\mathbf{b}}$ respectively, with all their components also identified by a “hat” whenever necessary. We construct a new equilibrium $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ from the original one (with all components identified by “tilde” whenever necessary), and then argue that it is evidence-determined and that it is equivalent to $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$.

We describe in detail how the construction proceeds for dynasty \mathcal{F} . The construction for dynasty \mathcal{G} is symmetric and we omit the details.

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, $(\tilde{\mathbf{f}}, \tilde{\mathbf{b}})$ is evidence-determined. By construction 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 $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$. It is convenient to argue the latter first. As before, we focus on dynasty \mathcal{F} . The details for dynasty \mathcal{G} are symmetric.

Since $(\hat{\mathbf{f}}, \hat{\mathbf{b}})$ is in fact evidence-driven, 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 $m_{\mathcal{F}}^t = \hat{\mathbf{m}}_{\mathcal{F}}(t, s^t)$, and the other given by the same s^t and any off-path message $m_{\mathcal{F}}^t \neq \hat{\mathbf{m}}_{\mathcal{F}}(t, s^t)$. 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. ■

A.5. Proof of Proposition 5

Note that the statement is trivial if $\gamma = 1$. Hence, we proceed assuming that $\gamma \in [0, 1)$. Let an evidence-determined equilibrium (\mathbf{f}, \mathbf{b}) of either the original dynastic game or the one with imperfect current monitoring be given. In either case 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 (possibly mixed) 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 (possibly mixed) 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 $\tau = 0, \dots, t-1$. 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 $\tau = 0, \dots, t-1$) is a probability distribution over action profiles 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 assumption, in an evidence-determined equilibrium 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 $\tau = 0, \dots, t-1$. So, we can write as $\mathbf{e}^t(h_{+}^{\bar{t}}, s^t) = (e(\bar{t}+1, s^t), \dots, e(t-1, s^t))$, where each $e(\tau, s^t)$ (with $\tau = \bar{t}+1, \dots, t-1$) is a probability distribution over action profiles 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. ■