Rochester Center for

Economic Research

Backward Induction and Model Deterioration

Francesco Squintani

Working Paper No. 484
July 2001

UNIVERSITY OF

ROCHESTER

Backward Induction and Model Deterioration*

Francesco Squintani
University of Rochester
Department of Economics†

May 2001

Abstract

Unlike formal games, most social interactions are not accompanied by a complete list describing all relevant actions. As a result, the most difficult task faced by the players is often to formulate a model of the interaction. While it is known how players may learn to play in a game they know, the issue of how their model of the game evolves over time is largely unexplored. This paper presents and analyzes a social learning construction that explicitly keeps track of the evolution of models held by players who are able to solve perfect-information extensive-form games (according to their models), and whose models depend on past observation of play. We introduce the possibility of small-probability model deterioration and show that, even when concerning only opponents' unobserved actions, such deterioration may upset the complete-model backward-induction solution, and yield a Pareto-improving long-run distribution of play. We derive necessary and sufficient conditions for the robustness of backward-induction path with respect to model deterioration. These conditions can be interpreted with a forward-induction logic, and are shown to be less demanding than the requirements for asymptotic stability of the backward-induction path under standard evolutionary dynamics. In all games where it may upset the backward-induction path, model deterioration may induce long-run distributions of play that correspond to non subgame perfect Nash equilibria.

JEL CLASSIFICATION: C71, C73, D83

*I am grateful for their comments to the audiences of seminars at Bonn University, Cornell University, Northwestern University, Rochester University, Tel Aviv University, and the University of North Carolina at Chapel Hill. I am indebted for their suggestions to Dieter Balkenborg, Drew Fudenberg, Ehud Kalai, Karl Schlag, Thomas Troger, Juuso Valimaki, and especially Eddie Dekel. Finally, I would like to thank Ahmer Tarar for excellent assistance. The usual caveat applies.

†Harkness Hall, University of Rochester, 14627 NY, sqnt@troi.cc.rochester.edu

1 Introduction

Unlike formal games, most social interactions are not accompanied by a list of written, fixed rules describing all actions that can be taken. As a result, the most difficult task faced by the individuals involved in a strategic interaction is often to formulate a model of the interaction. Once this modeling step is accomplished, solving the model may be relatively easy. The model of players who are repeatedly involved in the same interactions may change over time, and will depend on their past experiences. While the game-theoretical learning literature presents several accounts explaining how players may learn to play in a game they know, the issue of how their model of the game changes over time is largely unexplored. This paper presents and analyzes a social learning construction that explicitly keeps track of the evolution of models held by players who are able to solve games according to the models they formulate, and whose models of the games depend on past observation of play.

In our construction, players from large populations are repeatedly randomly matched to play a generic perfect-information extensive-form game. When matched to play the game, each player formulates a model of the interaction, identified by a subset of the action space. At each decision node the player plays the action corresponding to the unique backward-induction solution of her (possibly incomplete) model. In the lapse of time between two consecutive matches, the player's model may change. It will always include the actions that the player has observed on the path of play at the previous match. But with small probability, it may also *deteriorate,* so as to exclude some feasible actions.[1] While each action can be forgotten from the model only with small probability, we allow the probabilities of forgetting different actions to differ.[2] In order to highlight the effect of

[1] While in economic settings it is not convincing to assume that a player's memory is subject to large-probability imperfections, we believe that infinitesimal-probability forgetfulness should not be dismissed.

[2] When formulating a model of a strategic interaction, some actions may be more salient than others, and thus less likely to be excluded from a player's model. In specific interactions, the relative likelihood of forgetting different actions is determined by extra-game-theoretical issues such as the contextual meaning of actions, or social conventions. We are agnostic with respect to the issue of how this relative likelihood should be formulated in specific interactions, and thus conduct the analysis for any relative forgetfulness probability.

model deterioration, we assume that all players initially hold a complete model of the game, so that initially the play established in each match coincides with the backward-induction path of the complete game. Our results are in terms of the long-run distribution of play, obtained by compounding the long-run model distribution with the play induced in each match.

The first part of our analysis restricts the possibility of model deterioration to opponents' actions that were not taken during the previous period of play, building on the supposition that one is usually less likely to forget one's own possible choices or recently observed actions. We show that the backward-induction path may be upset by small-probability model deterioration, and, in some games, the resulting long-run distribution of play may Pareto-dominate the backward-induction path. We present general results for the case where, because of model deterioration, the long-run distribution of play is different from the backward-induction path, and yet each player formulates a model that yields a solution consistent with the path she observes. When these two conditions are met, we say that the backward-induction path is persistently upset. A necessary condition and a sufficient condition for the backward-induction path to be persistently upset are presented. If the backward-induction path satisfies a forward-induction requirement, then it cannot be upset by model deterioration. In the games where it may persistently upset the backward-induction path, model deterioration may yield long-run paths that are Nash but non subgame perfect, so that population model heterogeneity vanishes in the long-run.[3] In a subclass of these games, the non subgame perfect long-run distribution of play displays a strong forward-induction property.

The second part of the analysis lifts our conservative assumptions on the possibility of model deterioration. When allowing model deterioration about all of the opponents' actions (including observed ones), we determine a single necessary and sufficient condition for the backward-induction path to be upset. When model deterioration concerns all actions on-

[3] In complete-model analysis, the establishment of non-subgame perfect Nash play originates from the ability of players to commit to non-subgame perfect off path choices. Our results show that, by slowly changing the game, model deterioration may mimic this commitment effect.

path (including a player own actions), long-run model distributions where each player's model yields a solution consistent with the observed path may induce distributions of play that do not coincide with any self-confirming equilibrium. Again we establish a necessary and sufficient condition for the backward-induction path to be persistently upset.

Regardless of the restriction on model deterioration, we should also point out that any game can be expanded without altering the backward-induction path, and turned into a game where the backward-induction path can be upset by small-probability model deterioration. Intuitively, the problem of whether the players are able to learn the game that they are playing will be more serious in complex interactions involving many possible choices, rather than in simple interactions with few relevant options available.

The approach introduced in this paper is related to the literature on learning and evolution in games (see Weibull 1992, Samuelson 1997, and Fudenberg and Levine 1998 for comprehensive reviews). *Strictu-sensu* evolutionary game theory analyzes the fitness of genes subject to natural selection forces. This is equivalent to study the learning process of players who are assumed not to have any knowledge of the game beyond their own possible choices. In the contributions on rational learning (cf. for example Kalai and Leher 1993, Fudenberg and Levine 1993b, Nachbar 1997) the main question is how players learn to play in a game they fully know.

This paper situates itself in between these two polar opposite approaches. We assume that players know how to solve game-theoretical models, and we keep track of how their models change over time.[4] This paper focuses on the long-run behavior reached from initial

[4] In a similar vein is the analysis of Stahl (1993), where players are endowed with models ordered so as to represent the complexity of strategic reasoning. If complex reasoning comes at a cost, he shows that evolution selects naive players. In Samuelson (2000) players need to allocate cognitive resources among a ultimatum game, a Rubinstein (1982) bargaining game, and a Machiavellian tournament capturing the opportunity cost of devoting resources to the two bargaining instances. He shows that there is an equilibrium where the players do not distinguish among the two bargaining games, and always play the solution of the Rubinstein game. Jehiel (2001) proposes a model of equilibrium in multi-stage games where players coalesce opponents' different decisional nodes into analogy classes, and formulate a single belief for each class. Under coarse analogies, he shows that initial cooperation may arise in the centipede game, or in the finitely-repeated prisoner's dilemma, and that immediate agreement need not be reached in bargaining games.

full awareness. This assumption is not to be taken as an absolute requirement; rather, it is an obvious benchmark to use in dealing with model deterioration. In order to derive a complete picture of how players learn the rules of the game, one may analyze our model considering any initial model distribution. The results presented in this paper imply that there are games players may never fully learn.

In comparing our framework and analysis with evolutionary game-theory contributions, we should first point out that our approach offers an interpretation of learning dynamics that does not require payoff monotonicity. In evolutionary game theory, people are assumed to imitate those players who hold the highest payoff (see Schlag 1998 or Borgers and Sarin 1999 for a formal argument). However, a player may not always be able to observe the payoff obtained by the other players in the population, whereas she always observes the move made by opponents with whom she is matched. Thus, the diffusion of a strategy in the population may be determined by how often the strategy is used, rather than by its payoff. Consistently with that view, this paper rules out the observation of other players' payoffs, and focuses on the relation between the models formulated by players and their past observation of play.

Secondly, we shall show that our framework may determine long-run distributions of play that do not correspond to a self-confirming equilibrium. This is in sharp contrast with the outcome of natural-selection forces, which in the form of payoff-monotonic regular dynamics (cf. Gale, Binmore and Samuelson 1995, or Cressman and Schlag 1998) select Nash equilibrium, and in the stochastic-stability approach of Noldeke and Samuelson (1993) select self-confirming equilibrium. Our results are also in contrast with the outcome of rational learning, as in the light of the analysis of Fudenberg and Levine (1993b), myopic players who know the game that they are playing coordinate on self-confirming equilibrium, and foresighted players learn Nash equilibrium.

Despite this major difference, our characterization of games whose backward-induction path may be persistently upset by small-probability model deterioration restricted to opponents' unobserved actions turns out to be related to the stability analyses of the backward-induction solution of Noldeke and Samuelson (1993), of Cressman and Schlag (1998) and of Balkenborg and Schlag (2001). Our necessary and our sufficient conditions for the backward-induction solution to be robust with respect to model deterioration are tighter than (respectively) the necessary and the sufficient conditions for the backward-induction solution to be stable in the sense of Noldeke and Samuelson (1993). At the same time we can construct games where the backward-induction solution is persistently upset and yet is stable in the sense of Cressman and Schlag (1998), but we cannot give a definite conclusion with respect to the converse relation. Finally, we can show that the necessary and sufficient condition for the asymptotic stability of the Nash equilibrium component associated with the backward-induction path, identified by Balkenborg and Schlag (2001), is more demanding than both our necessary and our sufficient conditions for the backward-induction path to be robust with respect to model deterioration (restricted to opponents' unobserved actions). In this sense, model deterioration turns out to be less disruptive of the backward-induction path than evolutionary forces.

The paper is presented as follows. The second section informally presents the model and some examples leading to the subsequent analysis. The third section formally presents our dynamic framework. The fourth section restricts the possibility of model deterioration to opponents' unobserved actions, and the fifth section lifts these conservative assumptions. The sixth section compares our characterization results with those of evolutionary game theoretical contributions. The conclusion presents a few possible extensions, and is followed by the Appendix, which lays out the proofs.

2 Leading Examples

In this section we informally present a few examples to introduce the relevance of model deterioration and our characterization results. We restrict the possibility of model's deterioration only to opponents' unobserved actions. Since players initially play the backward-induction path, and in each period all players observe all actions on path, one may expect

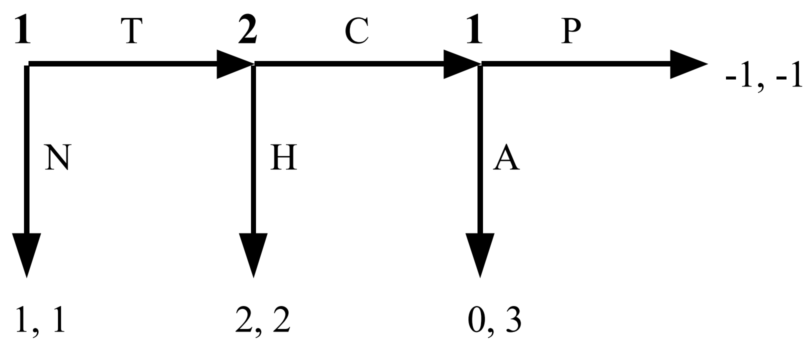

Figure 1: Trust and Punishment Game

that the backward-induction path will not be upset by small-probability model deterioration. In the following example, however, the unique long-run prediction is a non-subgame perfect Nash equilibrium. The backward-induction path is upset because, while model deterioration occurs with small probability, the fraction of players holding models that do not include all actions off path may increase over time, and players' backward-induction choice may crucially depend on whether off-path actions are included in their model or not.

Example 1 Two populations of players are repeatedly randomly matched to play a version of the trust game (depicted in figure 1) which includes the possibility of costly punishment.[5] Each player from population 1 may either trust (T) her opponent, or not (N). If trusted, the second player may honor (H) the first player's trust, or cheat (C). In case her trust is abused, the first player may decide to costly punish (P) the opponent, or to acquiesce (A). The backward-induction path is N. The game has another Nash equilibrium component, which induces the path TH. Notice that the non subgame perfect Nash equilibrium path TH Pareto-dominates the backward-induction path N. If the players in population 1 could credibly commit to play P, they would Pareto-improve upon the backward-induction outcome.

[5]The original trust game has been introduced by Kreps (1990). This expanded version appears in Kohlberg and Mertens (1986) who show that its backward-induction solution fails to satisfy the Never Weak Best Response property together with Admissibility.

We assume that with probability ε, each player from population 2 can forget the possibility that action A is played, in the lapse of time following a match where she has not observed it on the path of play. Also, we say that the players in population 1 may forget the feasibility of action C with probability ε if they did not observe it at the previous match. To simplify calculations, we assume that each player's model always include all other actions.

Initially all players hold a complete model of the game, this induces the path N, and thus both A and C may be forgotten. The backward-induction solution of the incomplete game obtained by deleting action A is THP. Thus players formulating a model that does not include A play H. By playing H, they prevent the possibility of observing action A. Thus the fraction of players whose model does not include A increases over time. Players holding a model that does not include C play T, as the solution of the game where the continuation of action C is deleted is TH. They observe action C only if they face an opponent whose model includes action A. Since the fraction of these players decreases over time, the fraction of players who observe C on path decreases over time. In the long run, there are no players in population 1 whose model includes C and no players in population 2 whose model includes A.

Formally, the populations dynamics are described by letting a_t be the time-t proportion of players whose model does not include A, and c_t be the time-t proportion of players whose model does not include C:

$$\begin{cases} c_{t+1} = c_t a_t + (1 - c_t)\varepsilon \\ a_{t+1} = a_t + (1 - c_t)(1 - a_t)\varepsilon. \end{cases} \tag{1}$$

It is immediate to see that a_t is non-decreasing in t. Thus there must be a value $a \geq a_t$ for all t, such that $a_t \to a$ for $t \to \infty$. Pick any arbitrary t, since $a_{t+1} - c_{t+1} = a_t(1 - c_t)(1 - \varepsilon) \geq 0$, it follows that $1 - c_t \geq 1 - a$. From the second equation in system (1) we thus obtain that $a_{t+1} \geq a_t + (1 - a)(1 - a_t)\varepsilon$. In the limit for $t \to \infty$, this implies that $a \geq a + (1 - a)^2\varepsilon$. As long as $\varepsilon > 0$, this condition is satisfied only if $a = 1$. Since $a_t \to 1$ for $t \to \infty$, it must also be the case that $c_t \to 1$. System (1) asymptotically approaches the state $(c^* = 1, a^* = 1)$.

At the state ($c^* = 1$, $a^* = 1$), the path played in each match is TH: the Pareto-dominant non-subgame perfect Nash equilibrium path. Also note that the state ($c^* = 1$, $a^* = 1$) is the unique stationary state of the system, that it is asymptotically stable, and that it is a global attractor. This underlines that the players will not be able to fully learn the game regardless of the initial model distribution.

In conclusion of this example we would like to point out that if players forget actions P and H instead of A and C, then the long-run distribution of play coincides with the backward-induction path. In general, the long-run distribution of play depends on the relative likelihood that different actions are forgotten by the players.[6] In any game, it is always possible to assign forgetfulness probabilities in such a manner that the backward-induction path is not upset. This paper tackles the (more involved) question of identifying games for which it is possible to assign forgetfulness probabilities that upset the backward-induction path. ◇

By inspecting the strategic structure of the game presented in the above example we can commence an informal presentation of our characterization of games where the backward-induction path may be upset by small-probability model deterioration. First of all, notice that the players in population 1 prefer the outcome TH to the backward-induction path N. Secondly, observe that, while the players in population 2 prefer the path TCA to the path TH, it is also the case that they prefer the path TH to the path TCP, and thus the choice H is rational when holding a model that does not include A. In the fourth section we will show (Corollary 6) that model deterioration may upset the backward-induction path whenever some players may decide to deviate from the backward-induction path with the

[6] In its general analysis, this paper is agnostic with respect to the issue of how the relative likelihood of forgetting actions should be formulated in specific interactions. For this specific example, however, we can follow the guidance of the cognitive studies on framing effects (see Dawes 1988 for a comprehensive presentation). First, it appears that the possibility to punish unfair behavior is very salient and unlikely to be dismissed. Secondly, dishonest, deviant behavior appears to be less salient than behavior conforming to social norms. Applied to this game, these findings suggest that, consistently with our analysis, the likelihood that players' models include actions P and H should be larger than the likelihood that actions A and C are included.

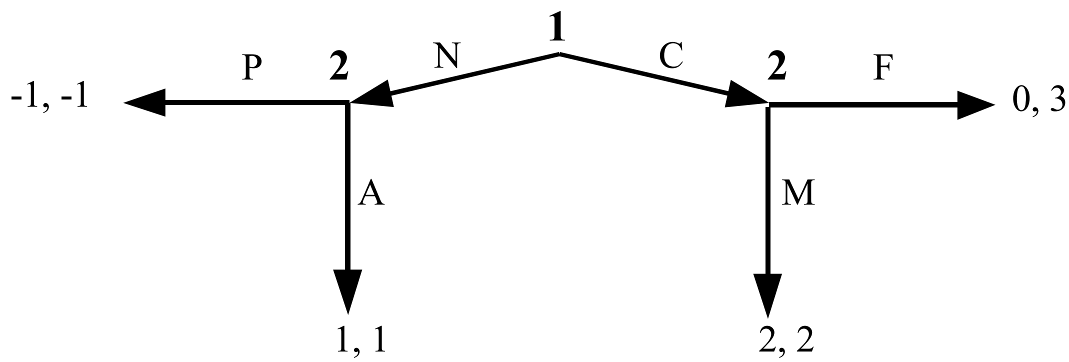

Figure 2: Sequential Contribution and Punishment Game

prospect of reaching a terminal node that gives them a higher payoff, and the path to this terminal node does not require any irrational choice by any player.

At the same time, we will show that the backward-induction path may be upset in games where the terminal nodes that can be reached when some players deviate from the backward-induction path yield these players a higher payoff than some terminal nodes (compatible with players' rationality) that can be reached when not deviating. The backward-induction path is fragile because a player may be lured off path in the illusion of achieving a higher payoff, and then forget actions included in the backward-induction path. When this is the case, the player may believe that her initial backward-induction choice yields a payoff dominated by any payoff allowed by the deviation, so that she chooses not to revert to the backward-induction choice. In order to substantiate this discussion, we present the following example.

Example 2 Two populations of players are repeatedly randomly matched to play the simple version of sequential contribution game presented in Figure 2.[7] Each player from population 1 may choose whether to make a socially-valuable contribution (C), or not (N). If the contribution is not made, the second player may decide either to costly punish her opponent (P), or to acquiesce (A). If the contribution is made, the second player may

[7] The economic theoretical literature studies several sequential models of social contribution (cf. for instance Marx and Matthews 2000 and citations therein).

free ride (F) on her opponent, or match (M) her opponent's contribution. The backward-induction solution of this game is (N, AF) inducing the path NA. The game has another Nash equilibrium component, which induces the path CF.

We assume that each player from population 1 can forget the possibility of action F, or of action A, whenever she does not observe it on the path of play. Again for simplicity, suppose that the probability of forgetting these actions is the same, and equal to ε, and that no other action can be forgotten. The backward-induction solution of the incomplete subgame obtained by deleting action F is CM, and the solution obtained deleting action A is CF. Thus players whose model does not include F or A play C, and observe F on path, but not A. Players holding a complete model play N, so they observe A on path, but not F. Initially all players play the backward-induction path N. They do not observe F on path and may forget it. Once they have forgotten it, they play C and thus may forget A. If they forget A, they will never switch back to N. The proportion of players whose model does not include A is thus strictly increasing and in the long run all players in population 1 play C.

Formally the populations dynamics are described by letting a_t be the time-t proportion of players whose model does not include A, and f_t be the time-t proportion of players whose model does not include F,

$$\begin{cases} f_{t+1} = (1 - f_t - a_t)\varepsilon \\ a_{t+1} = a_t + f_t\varepsilon. \end{cases}$$

It is easy to see that the system asymptotically approaches the state $(f^* = 0,\ a^* = 1)$. At this state, the path played in each match is CF: the non-subgame perfect Nash equilibrium path. Finally, note that the state $(f^* = 1,\ a^* = 1)$ is the unique stationary state of the system, that it is asymptotically stable, and that it is a global attractor. ◇

We conclude this informal presentation by showing a game where small-probability model deterioration cannot upset the backward-induction path. In this game the models of some players may deteriorate so that they are lured off path. But immediately after deviating they will observe their opponents play the actions not included in their incomplete model, and they will immediately switch back to the backward-induction choice. Since model deterioration occurs with very small probability, and the model is immediately restored when the player deviates from the backward-induction path, it follows that the long-run distribution of play is indistinguishable from the backward-induction path. The possibility that players deviate from the backward-induction path lured by the illusion of achieving a higher payoff is not enough to upset the backward-induction path. It is also necessary that after such a deviation, the play can settle on paths that do not reconstruct models whose solutions yield back the initial backward-induction path.

Example 3 In the simple trust game represented in Figure 5, the players in population 1 may only choose whether to trust (T) her opponent or not (N), and the players in population 2 may only choose whether to honor (H) the trust or to cheat (C). The backward-induction solution is N. Suppose that any player in each population i may forget each opponent's action a with probability ε_a^i whenever she does not observe it at the previous match. Regardless of her model, each player in population 2 plays C whenever her decisional node is reached. Any player in population 1 whose model includes action C plays N. The backward-induction choice of any player holding a model that excludes C is to play T. But these players observe action C on path and thus immediately switch back to playing N, because all players in population 2 take C whenever they are called to play.

Letting c_t (respectively h_t) denote the time-t fraction of players in population 1 whose model does not include C (respectively H), the population dynamics are as follows:

$$\begin{cases} c_{t+1} = (1 - c_t - h_t)\varepsilon_C^1 \\ h_{t+1} = h_t + (1 - c_t - h_t)\varepsilon_H^1. \end{cases}$$

For any choice of parameters ε_C^1 and ε_H^1, in the long run, the system will settle on the state (c^*, h^*) such that $c^* \leq \varepsilon_C^1/(1 + \varepsilon_C^1)$. Since ε_C^1 is negligible, the long-run distribution of play is indistinguishable from the backward-induction path. ◇

It is important to notice that the game of Example 1 can be represented as a "backward-induction-irrelevant" expanded version of the game in Example 3. These two games

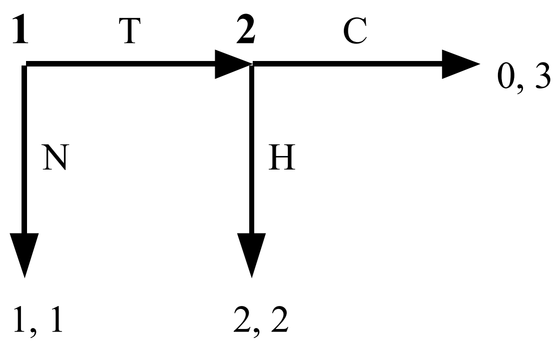

Figure 3: Simple Trust Game

share precisely the same backward-induction path, and yet model deterioration upsets the backward-induction path only in the richer game presented in Example 1. Following this lead, we will also show (Proposition 7) that any game can be expanded without altering the backward-induction path, and turned into a game where the backward-induction path can be upset by small-probability model deterioration. This result comes as no surprise. It is reasonable to expect that the problem of whether or not the players are able to learn the rules of the game that they are playing will be more serious in complex interactions involving many possible choices, rather than in simple interactions with few relevant options available.

3 The Model

There is a finite set I of continuous populations of players. Each population i is represented by a copy of the interval $[0,1]$. At each period t, all players are randomly matched to play the finite generic perfect-information game $\Gamma = (X, Z, A, I, \iota, \mathbf{u})$. Each match is anonymous and it includes one player l_i from each population $i \in I$. Setting $Y = X \cup Z$, the pair $G = (Y, A)$ represents a tree where X is the set of decision nodes, Z is the set of terminal nodes, and the set of arcs $A \subseteq Y \times Y$ is the action set.[8] We denote by $\prec$ the transitive

[8]A tree is an oriented connected, acyclic directed graph with a single root x_0, a leafs Z (see Diestel 1997, first chapter). In this context it must be the case that A is an irreflexive, acyclic order satisfying the following requirements. First, for any $z \in Z$, there is no $y \in Y$ such that $(z, y) \in A$; secondly there is a

closure of A, and (with a minor terminological violation) we say that $x \preceq y$ to mean that either $x \prec y$ or that x coincides with y. A path from node x to node y on the tree $G = (Y, A)$ is the (unique) set of actions $\mathbf{a} = \{a_0, \cdots, a_n\} \subseteq A$ such that $a_0 = (x, y_1)$, $a_1 = (y_1, y_2)$, ..., $a_n = (y_n, y)$, for some finite sequence $\{y_1, y_2, ..., y_n\}$. Each terminal node z uniquely identifies a path from the initial node x_0, and with a standard terminological violation we can also call z a path.

The assignment function $\iota : X \to I$ labels the decision nodes to players. It is extended to actions according to the relation: $\iota(a) = i$ if $a \in A(x)$ and $\iota(x) = i$, and it partitions the sets X and A into the equivalence classes $\{X_i\}_{i \in I}$ and $\{A_i\}_{i \in I}$. The function $\mathbf{u} : Z \to \mathbb{R}^I$ represents the players' payoffs obtained when reaching a terminal node, and each player in the same population has the same utility function. The function is $\mathbf{u}$ extended on $\Delta(Z)$ according to the multilinear expansion formula. In order to avoid trivialities, we will restrict attentions to games such that for any x, $\#A(x) > 1$ and for any $a \in A(x)$, $\iota(a(x)) \neq \iota(x)$. We sometimes express the game $\Gamma = (X, Z, A, I, \iota, \mathbf{u})$ as $\Gamma = (G, I, \iota, \mathbf{u})$, to highlight the tree G. For any node x, we introduce the sets $Y_x = \{x' \in Y : x \preceq x'\}$, $X_x = Y_x \cap X$ and $Z_x = Y_x \cap Z$. The subgame starting at x (which we denote for short as the x-subgame) consists of the game $\Gamma_x = (X_x, Z_x, A|_{Y_x}, I, \iota|_{X_x}, \mathbf{u}|_{Z_x})$.

For any game Γ, the *backward-induction solution* $\mathbf{a}^* \subseteq A$ is defined as follows. Let $Y_0 = Z$, and for any $j \geq 1$, recursively define

$$Y_j = \{x \in Y \setminus (\cup_{k=0}^{j-1} Y_k) : \text{for all } x' \in Y \setminus (\cup_{k=0}^{j-j} Y_k), \quad x \not\prec x'\}.$$

Set $\mathbf{u}^*(z) = \mathbf{u}(z)$ for any $z \in Z$. For any $j \geq 1$, and any $x \in Y_j$, let

$$a_x^* = \arg\max_{a \in A(x)} u_{\iota(x)}^*(a(x)) \qquad \text{and} \quad \mathbf{u}^*(x) = \mathbf{u}^*(a_x^*(x)).$$

The *backward-induction path* z^* identifies the unique terminal node z such that $\{a_0, \cdots, a_n\} \subseteq \mathbf{a}^*$, $a_0^* = (x_0, y_1)$, $a_1 = (y_1, y_2)$, ..., $a_n = (y_n, z)$, for some finite sequence $\{y_1, y_2, ..., y_n\} \subset Y$.

---

unique $x_0 \in Y$ such that $(y, x_0) \notin A$ for any $y \in Y$; finally for any $y \in Y$, $y \neq x_0$, there is a unique $x \in X$ such that $(x, y) \in A$.

Pick any arbitrary period of play. In any match l, at the beginning of the game, each player formulates a (possibly incomplete) model of the game, identified by a (possibly proper) subset of the action space A. For brevity, whenever a player's model does not include an action a, we will say that she is *unaware* of a. As the play develops in the match, it may be the case that some players take actions not included in a player's initial model. Whenever this is the case, this player may find herself on a decision node not specified in her initial model, and will need to formulate a model of the subgame starting at that node. In order to give a well-structured description of players' models at each decision node, we first say that in any match, each player is endowed with a *framework* consisting of a list of actions in the game. A player's framework does not only represent the actions included in her initial model in the match, but also identifies the actions that will be included in her model if the play reaches *any* of the decision nodes, including nodes identified by paths that contain actions not included in her initial model. In order to guarantee that all these models identify well-defined subgames, we assume that each player's framework includes at least one action $a \in A(x)$ for each node $x \in X$.

Definition 1 *For each population i, the set of frameworks is $\mathcal{B}^i = \{B^i \subseteq A :$ for any $x \in X$, $B^i \cap A(x) \neq \varnothing\}$, and a framework assignment is a measurable function $\alpha_i : [0,1] \to \mathcal{B}^i$.*

For further reference, we denote by $B = (B^i)_{i \in I}$ any arbitrary profile of frameworks, by $\mathcal{B} = \times_{i \in I} \mathcal{B}^i$ the set of frameworks profiles, and by α any profile of framework assignments. Also, for any x and B^i, we set $B^i(x) = B^i \cap A(x)$.

Given a player's framework, we can determine her model at any of her decision nodes. Pick an arbitrary player from an arbitrary population i endowed with framework B^i and suppose that the play reaches her decision node x. This player's model of the x-subgame starting at x is determined by all paths that start at x and that include only actions contained in B^i. In order to give a formal definition, note that for any node x the set $Y_x \times Y_x$ identifies all possible arcs connecting the nodes in Y_x, and that $B^i \cap (Y_x \times Y_x)$ thus identifies the set of all actions contained in B^i and connecting nodes in Y_x.

Definition 2 *Take any arbitrary population i and any framework B^i. At any node x such that $\iota(x) = i$, the model of a player l_i such that $\alpha_i(l_i) = B^i$ is denoted by $G(B^i, x) = (Y^i_x, B^i_x)$, and consists of the smallest tree contained in the (not necessarily connected) graph $(Y_x, B^i \cap (Y_x \times Y_x))$.*

At each decisional node x, we assume that each player, given a possibly incomplete model, solves the x-subgame according to the backward-induction solution. Specifically, given any model $G(B^i, x) = (Y^i_x, B^i_x)$, we let $\mathbf{a}^*(B^i, x)$ be the backward-induction solution of game $\Gamma(B^i, x) = \left(G(B^i, x), I, \iota|_{X^i_x}, u|_{Z^i_x}\right)$, we let $\mathbf{u}^*(B^i, x)$ be the associated backward-induction values, and $z^*(B^i, x)$ the associated backward-induction path. We assume that any player l_i endowed with framework B^i takes the (incomplete model) backward-induction action $a^*_x(B^i, x)$ at node x.

In order for the construction to be consistent, it must be the case that, for any population i, the choice $a^*_x(B^i, x)$ is unique at each decision node x. This follows from the fact that the game Γ, and hence any arbitrary game $\Gamma(B^i, x)$, are generic. This fact also implies that our construction determines a unique path of play in game Γ for each match l, and that the path of play depends only on the frameworks held by the players in the match, and is independent of their identity.

Lemma 1 *For any framework profile $B \in \mathcal{B}$, there is a unique path of play $\mathbf{a}(B)$ established in any match l such that $(\alpha_i(l_i))_{i\in I} = B$.*

Proof. Pick any arbitrary population i, framework B^i, and node x such that $\iota(x) = i$. Since the game $\Gamma(B^i, x)$ is generic, there is a unique backward-induction action $a^*_x(B^i, x)$ adopted by all players l_i such that $\alpha_i(l_i) = B^i$. Since i, B^i and x are arbitrary, it follows that for any profile B, there is a unique solution $\mathbf{a}^*(B)$ determined in any match l such that $(\alpha_i(l_i))_{i\in I} = B$. Thus there is a unique path of play $\mathbf{a}(B)$. ■

The model distribution in the populations at any time t is identified by the assignment profile α^t. In order to describe its dynamic transition between time t and time $t+1$, we need to specify the individual framework transition of each player in each population.

In this paper, after any match, each player's framework will always include all actions observed on path. At the same time, in the lapse of time between any two consecutive matches, the player's framework can deteriorate and exclude some actions from consideration. Thus each player's framework transition depends on the observed path, and on random model deterioration. In the remainder of the paper, for brevity, we shall say that "a player *forgets* action a at time t" to mean that during the lapse of time from period t to period $t+1$, the player's framework has deteriorated so as to exclude the possibility that action a can be taken in the game. Conversely, we shall say that "a player *recalls* action a at time t" if her framework at time $t+1$ includes action a and her model at time t does not.

For simplicity, we assume that forgetfulness occurs with probabilities fixed over time and across players of the same population, and independently across actions and opponents. Fix the profile of forgetfulness probabilities $\boldsymbol{\varepsilon} = \{\varepsilon_i(a)\}_{i\in I, a\in A}$: each entry is the probability that a player from population i forgets action a. The case where players cannot forget their own action is formalized as follows.

Assumption 1 *For any population i, and action a, $\varepsilon_i(a) = 0$ if $a \in A_i$.*

Our construction cannot allow actions to be forgotten when that amounts to deleting the entire action set associated with any decisional node, or else the model of the players controlling that node is not well defined. So, for any population i and framework B^i we introduce the set $F(B^i) = \{\hat{B}^i \in \mathcal{B}^i : \hat{B}^i \subseteq B^i\}$, and the probability measure $\{\pi(\hat{B}^i|B^i)\}_{\hat{B}^i\in F(B^i)}$ such that[9]

$$\pi_i(\hat{B}^i|B^i) \propto \prod_{a\in B^i\setminus\hat{B}^i} \varepsilon_i(a) \prod_{a\in\hat{B}^i} (1-\varepsilon_i(a)), \quad \text{for any } \hat{B}^i \in F(B^i).$$

By Lemma 1 the path observed by each player in a match is uniquely pinned down by the framework profile B of the players in the match. This, together with the above construction,

[9] General non-necessarily independent forgetfulness probabilities are defined by arbitrary systems π such that for any i, and any B^i, $\pi(\cdot|B^i) \in \Delta(F(B^i))$. In the proofs we will show that our results also hold when forgetfulness may be correlated.

allows us to say that, for any matching time t, each player's stochastic framework transition in the lapse of time between t and $t+1$ is function of her framework at time t, and of the frameworks of the players that she has met at time t.

The specific form of the framework transition depends on whether it is the case that observed actions can be forgotten or not. We first formalize the case where observed actions cannot be forgotten.

Assumption 2 *For any time t, any match l such that $\alpha^t(l) = B$, and any population i, player l_i assumes framework $\hat{B}^i \cup \mathbf{a}(B)$ at time $t+1$ with probability $\pi(\hat{B}^i|B^i)$.*

The case where observed action can be forgotten is formalized as follows.

Assumption 3 *For any time t, any match l such that $\alpha^t(l) = B$, and any population i, player l_i assumes framework $\hat{B}^i$ at time $t+1$ with probability $\pi(\hat{B}^i|B^i \cup \mathbf{a}(B))$.*

Given the random matching process, the assignment α^t determines a random assignment of framework profiles across matches at time t. This, together with the random individual transition, allows us to determine the random aggregate framework transition $\gamma(\boldsymbol{\varepsilon}) : \alpha^t \mapsto \alpha^{t+1}$.

The analysis of the stochastic process defined by $\gamma(\boldsymbol{\varepsilon})$ is made unmanageable by the structure of the framework assignments space. However, since matching is anonymous, a standard approach of evolutionary game theory allows us to derive a more tractable formulation of the aggregate transition. We partition the populations in finite sets of "types", and then invoke a "Law of Large Numbers" argument, to approximate the stochastic transition $\gamma(\boldsymbol{\varepsilon})$ with a deterministic transition that keeps track of the population type frequencies.[10]

As a consequence of Lemma 1 we have established that players' transitions depend only on their framework and on the framework profile of their match. Thus, the coarsest

[10] Boylan (1993), Proposition 3 shows that for finite large populations there exist anonymous random-matching schemes such that population transition ratios weakly converge to the composition across different types of individual transition probabilities. Alos-Ferrer (1999) extends the analysis to the case of a continuum of players, and shows that there exist random matching processes guaranteeing that the evolution of frequencies is (almost surely) deterministic.

set of types allowing the retention of all the information described by $\gamma(\boldsymbol{\varepsilon})$ is the set of framework profiles $\mathcal{B}$. Given population state α_i, the associated type frequency is denoted by $\rho_i \in \Delta(\mathcal{B}^i)$ such that $\rho_i(B^i) = \nu\{l_i : \alpha_i(l_i) = B^i\}$ for any type B^i (where ν is the Lebesgue measure).[11] Letting any frequency profile be denoted as $\rho = (\rho_i)_{i \in I}$, we express as

$$\xi(\boldsymbol{\varepsilon}) : \times_{i \in I} \Delta(\mathcal{B}^i) \rightarrow \times_{i \in I} \Delta(\mathcal{B}^i),$$

the deterministic transition induced by the transition $\gamma(\boldsymbol{\varepsilon})$.[12] Our construction determines a family $\{\xi(\boldsymbol{\varepsilon})\}$ of dynamic systems parametrized in the profile of forgetfulness probabilities $\boldsymbol{\varepsilon}$. Fixing an initial state ρ^0, each system $\xi(\boldsymbol{\varepsilon})$ determines a solution $\{\rho^t(\boldsymbol{\varepsilon}, \rho^0)\}_{t \geq 0}$, and thus we have obtained a family of solutions $\{\{\rho^t(\boldsymbol{\varepsilon}, \rho^0)\}_{t \geq 0}\}$ parametrized in $\boldsymbol{\varepsilon}$.

This paper's results investigate whether the backward-induction path is upset by small-probability model deterioration. Therefore, we assume that initially all players hold a complete model of the game, and thus that all initial matches establish the backward-induction path.[13] Formally, we study the long-run properties of families of solutions

$$\{\{\rho^t(\boldsymbol{\varepsilon}, \hat{\rho}^0)\}_{t \geq 0}\} \text{ where } \hat{\rho}^0(A) = 1 \text{ for all } i.$$

Since we are interested in the case of small-probability forgetfulness, we will look at the description of the solutions $\{\rho^t(\boldsymbol{\varepsilon}, \hat{\rho}^0)\}_{t \geq 0}$ for $\boldsymbol{\varepsilon}$ close to the vector $\mathbf{0}$. The limit solution may depend on the relative probabilities of forgetting different actions in the limit, thus our results will be given for specific sequences $\{\boldsymbol{\varepsilon}^n\}_{n \geq 0}$ of forgetfulness probability profiles, such that $\boldsymbol{\varepsilon}^n \rightarrow 0$. Since the backward-induction path may in principle be established in matches where some players' model of the game is incomplete, we will need to explicitly

[11] When no confusion can occur, we will simplify our notation, by denoting the frequency of a type unaware of action a in a given population as ρ_a, and the frequency of the complete-model type as ρ_*.

[12] We omit the closed form of ξ as it is not particularly insightful.

[13] It would be inappropriate to derive such a result in terms of stationary or stable states without considering explicitly the initial model distribution. Suppose we proved to be locally stable a state inducing a path distribution that does not coincide with the backward-induction path. If such a state could not ever be reached when players initially hold complete models, the result would not be driven by model deterioration, but by incomplete models initially present in the populations.

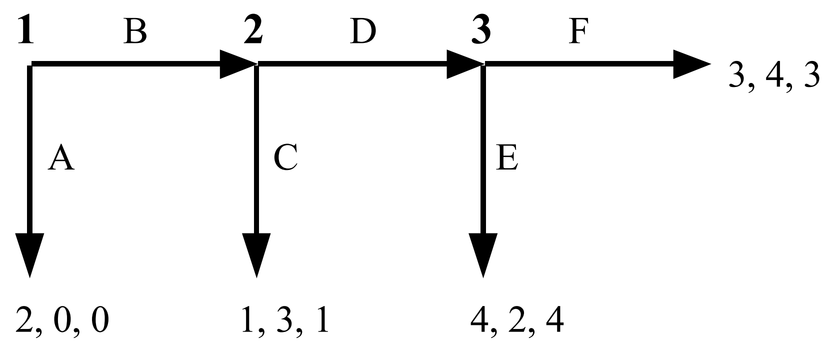

Figure 4: 3-Player 3-Leg Centipede Game

consider the distribution of play induced by the long-run type-frequencies. Formally, for any type frequency ρ, we let $f(\rho) \in \Delta(Z)$ denote the distribution of play induced by the rule

$$f(z) = \sum_{\{B \,:\, z^*(B)=z\}} \rho(B), \text{ for any } z \in Z.$$

Before we formally define the condition under which the backward-induction path is upset by small-probability forgetfulness, we introduce one last qualification. In principle one could say that the backward-induction path is upset whenever it is not the case that the long-run distribution of play coincide with (or at least is close to) the backward-induction path. We argue however that this definition of stability may be considered too restrictive in some games. Consider the following example.

Example 4 *Consider the* 3*-player* 3*-leg version of the centipede game*[14] *presented in figure 4, and say that* $\varepsilon_1(E) > 0$. *For the sake of simplicity, also suppose that for any other action* a, *and population index* i, $\varepsilon_i(a) = 0$. *The backward-induction solution is* (A, C, E), *and the associated path is* A. *Since action* E *is off-path, it may be forgotten by the players in population* 1. *Any player unaware of* E *formulates a model whose backward-induction solution is* (B, D, F), *and thus takes action* B. *Each player in population* 2 *is fully aware*

[14] The centipede game was first introduced by Rosenthal (1982) and is one of the most common examples in the theoretical literature on perfect-information extensive-form games.

of the subgame starting at her decision node, and she thus takes action C. By doing that, she prevents players in population 1 *from recalling action E. The long-run distribution of play concentrates all mass on the path BC, which is different from the backward-induction path.* ◇

In the above example all players in population 1 persistently choose to play B prompted by the mistaken belief that their opponents will play D and F, even though they repeatedly observe that, in reality, players from population 2 play C. If one follows literally the motivation of this work, it is easy to make sense of this phenomenon. Each time that a player in population 1 is matched to play the above game, she first formulates a model of the interaction. She then conjectures that her opponents play the backward-induction solution of her (possibly incomplete) model. If she cannot complete her model, her beliefs on the path of play may be mistaken. As a result, our construction allows for long-run predictions that do not induce any self-confirming equilibrium.

At the same time, we feel that the result that the backward-induction path is upset in Example 4 hinges on extreme features of our construction. The model of players in population 1 does not include action E and thus they infer the incomplete-model backward-induction solution BDF. Upon observing the path BC, they have no chance of learning that action E is feasible, and so they maintain their mistaken beliefs. This is hard to criticize as a description of transient behavior. It is, however, less appealing as a description of *persistent* behavior. It may be the case that in the long run, upon persistently observing players in population 2 playing action C, players in population 1 decide to reject their model's backward-induction solution, and formulate the belief that players in population 2 play C. If that were the case, players in population 1 would play action A, thus reestablishing the backward-induction path in the long run.

The main results of this paper revolve around the theme that small-probability forgetfulness of (opponents' unobserved) actions may upset the backward-induction path. We would like these results to be as robust as possible, and not to depend on extreme or fragile features of our construction. Thus, in order to say that the backward-induction path is upset, we shall not be satisfied by showing that the long-run distribution of play is different from the backward-induction path, but we shall impose the further requirements that in the long run, the beliefs of each player in any population with respect to her opponents' play are not disconfirmed by the path of play that she actually observes.

Our formal definitions first denote as *persistent* any type distribution where all players' beliefs are not disconfirmed by the observed path.

Definition 3 *The type distribution ρ is* persistent *if $z^*(B^i, x_i) \subseteq z^*(B)$ for any profile $B \in Supp(\rho^*)$, for any index i, and node $x_i \prec z^*(B)$.*

Secondly, for any game Γ, and any sequence $\{\boldsymbol{\varepsilon}^n\}_{n\geq 1}$, we introduce the notations

$$\rho^* = \lim_{n\to\infty} \lim_{t\to\infty} \rho^t(\boldsymbol{\varepsilon}^n, \hat{\rho}^0), \quad \text{and } f^* = f(\rho^*),$$

respectively denoting the long-run type frequency, and the induced path distribution.

Definition 4 *The backward-induction path z^* of an arbitrary game Γ is* persistently upset *by the sequence of forgetfulness profiles $\boldsymbol{\varepsilon}^n \to 0$, if $f^*(z^*) \neq 1$ and ρ^* is persistent.*

We conclude this section by showing that any persistent type distribution where all players of any population being aware of all their own actions induces a self-confirming equilibrium. Since we are considering games of perfect information, moreover, self-confirming equilibrium coincides with consistent self-confirming equilibrium.[15] It is important to underline that even though we are restricting attention to generic games, and thus each type or type profile identifies a unique solution and path, it is still possible that a persistent distribution ρ induces a heterogeneous-belief self-confirming equilibrium, because of type heterogeneity in the population.

[15] The formal definitions of self-confirming equilibrium and of consistent self-confirming equilibrium may be found in Fudenberg and Levine (1993a).

Lemma 2 *For any game Γ, any persistent distribution ρ such that $A_i \subseteq B^i$ for any i and any $B^i \in Supp(\rho_i)$ induces a (possibly heterogeneous) consistent self-confirming equilibrium $f(\rho)$.*

Proof. The result that any persistent distribution ρ such that $A_i \subseteq B^i$ for any i and any $B^i \in Supp(\rho_i)$ induces a self-confirming equilibrium $f(\rho)$ is immediately derived by translating our construction into the language of Fudenberg and Levine (1993a) Definition 1. Since Γ is a game of perfect information, $f(\rho)$ must be consistent.

In order to show that $f(\rho)$ may be heterogeneous, consider the game presented in Figure 1 and the type distribution ρ such that $\rho_{1*} = 1/2$, $\rho_{1C} = 1/2$ and $\rho_{2F} = 1$. Note that ρ is persistent and that it induces the distribution of play $f(\rho)$ such that $f(M) = 1/2$, $f(TH) = 1/2$ which is a heterogeneous self-confirming equilibrium. ■

4 Forgetfulness of Opponents' Unobserved Actions

4.1 A Necessary Condition

Throughout this section, we shall assume that players may only forget opponents' unobserved actions. In precise terms, we shall be working under Assumptions 1 and 2. The first part of the section describes and discusses a necessary condition for the backward-induction path to be persistently upset.

Consider a generic perfect-information game Γ. For any index i and any node x, we will define as *backward-inductive maximin* the value $\underline{u}_i(x)$, which is the payoff obtained by i if at any choice following (and including) x the players from population i maximize their utility and their opponents minimize it. Recalling the definition of the sets Y_j from the previous section, the formal definition is as follows.

Definition 5 *For any population i, the* backward-inductive maximin solution $\underline{\boldsymbol{a}}^i$ *is defined as follows. For any $z \in Z$, set $\underline{u}_i(z) = u_i(z)$. For any $j \geq 1$, and any $x \in Y_j$, let*

$$\underline{a}_x^i = \begin{cases} \arg\max_{a \in A(x)} \underline{u}_i(a(x)) & \textit{if } \iota(x) = i \\ \arg\min_{a \in A(x)} \underline{u}_i(a(x)) & \textit{if } \iota(x) \neq i, \end{cases} \qquad \textit{and } u_i(x) = u_i(\underline{a}_x^i(x)).$$

In complete-model analysis, the backward-inductive maximin value of a game may be interpreted as the smallest payoff that a rational player can guarantee herself when participating in the game.

Intuitively, our necessary condition builds on the following two observations. First, since a player cannot forget observed actions and the backward-induction path is initially played by all players, it cannot be the case that a player will deviate from the backward-induction path if all the terminal nodes that follow the deviation make her worse off with respect to the backward-induction path. Since she cannot forget any action on path, in fact, she may decide to deviate only in the hope of improving her payoff.

Secondly, as pointed out by Example 3, establishing that a player might deviate from the backward-induction path is not enough to yield a non subgame perfect long-run distribution of play. It is also necessary that after such a deviation, the observed play does not reconstruct models that yield back the backward-induction path. For the backward-induction path to be persistently upset, there must exist a persistent long-run type distribution different from the backward-induction path. We show that this requires that the game has at least one path (different from the backward-induction path) that awards each player, at each decisional node, a payoff larger than the backward-induction maximin payoff associated with the decisional node (we define any such a path to be *consistent).* Since players cannot forget their own actions, in fact, the worst possible outcome that they can associate with any action when holding an incomplete model coincides with the action's backward-induction maximin value. It would thus be irrational to take an action which is believed to yield a payoff inferior to the backward-inductive maximin value of a different action.

We first formalize the second observation.

Definition 6 *A path z such that $u_{\iota(x)}(z) \geq \underline{u}_{\iota(x)}(x)$ for any $x \prec z$ is defined to be* consistent.

Lemma 3 *If the type distribution ρ^* persistent, then for any type profile $B \in Supp(\rho^*)$, the path $z^*(B)$ is consistent.*

Armed with the above Lemma, we are able to present a necessary condition for the backward-induction path to be persistently upset by small-probability forgetfulness of opponents' unobserved actions.

Definition 7 *Let the triple* (x, a, z) *be called a* deviation *whenever* $x \prec z^*$, $a \in A(x) \backslash \{a^*_x\}$, $z \succeq a(x)$, *and* $u_{\iota(x)}(z) > u^*_{\iota(x)}(x)$.[16]

Proposition 1 *For any game* Γ, *the backward-induction path* z^* *may be persistently upset by some sequence of forgetfulness probabilities* $\boldsymbol{\varepsilon}^n \to 0$ *only if* Γ *admits a deviation* (x, a, z'), *and a consistent path* $z \neq z^*$.

The above result immediately implies that if a game does not admit any deviation, then its backward-induction path cannot be persistently upset by small-probability forgetfulness.[17] This means that the backward-induction path cannot be upset if it satisfies a strong forward-induction requirement: none of the players can force the play off path in order to achieve a higher payoff. This requirement is reminiscent of the forward-induction property proposed by Van Damme (1989), but is considerably stronger. Van Damme (1989), in fact, defines a solution set to be *consistent with forward induction* (in 2-person generic games), if there is no Nash equilibrium in the set such that some player can take a deviation leading with certainty to a subgame in which she is better off under one Nash equilibrium and worse off under all the others. We require that there is simply no deviation that makes a player better off, regardless of whether this deviation can be supported as a Nash equilibrium, or whether it is the only improving deviation.

Proposition 1 also immediately implies the following corollaries.

[16] The concept of deviation may be related to the concept of strict outcome path introduced by Balkenborg (1995). The backward-induction path of a perfect-information generic game is a strict outcome path if and only if the game does not admit deviations.

[17] This result would also hold if we dropped the requirement that the long-run distribution of play is persistent in Definition 4. The proof of Proposition 1 allows us to conclude that if a game does not admit any deviation then the long-run distribution induced by small-probability forgetfulness must be indistinguishable from the backward-induction path.

Corollary 2 *If the backward-induction path includes all decision nodes of a game, then it cannot be persistently upset by small-probability forgetfulness.*

Corollary 3 *If the backward-induction solution is a strict equilibrium, then it cannot be persistently upset by small-probability forgetfulness.*

These two results come as no surprise. The fragility of the backward-induction path originates from the fact that the backward-induction solution depends on off-path play. If there are no nodes off path, it would be unreasonable to think that the backward-induction solution could be upset.[18] At the same time, in this paper forgetfulness is assumed to occur with infinitesimal probability, and it is well known that strict equilibrium cannot be upset by small perturbations.

We conclude this part of the section by presenting a Lemma that will be useful for the comparison of our necessary and our sufficient condition. Keeping in mind the concept of consistent paths, we introduce the concept of x-consistent paths, where x is an arbitrary decision node. Essentially, a path is x-consistent if its restriction to the subgame starting at x is a consistent path. Obviously, each consistent path z is x-consistent for any node $x \prec z$.

Definition 8 *For any node x, the path $z \succ x$ is called x-consistent if $u_{\iota(x')}(z) \geq \underline{u}_{\iota(x')}(x')$ for any node x' such that $x \preceq x' \prec z$.*

We now show that a game has a x-consistent path z (different from the backward-induction path z^*) for some node x on path, if and only if the game has a consistent path z' different from z^* (and possibly also different from z).

[18] Corollary 2 is reminiscent of the results in Reny (1993), which characterize games where common certainty of rationality at the beginning of the game can be maintained throughout the entire path of play, so that the backward-induction path is established. Despite such a similarity, our results are logically independent. We can both construct games where the backward-induction path is persistently upset by model deterioration, and yet common certainty of rationality can be maintained across the entire game, as well as games where the backward-induction path cannot be upset by small-probability forgetfulness, and yet the $\mathcal{WS}^{\infty}$ solution implied by the initial common certainty of rationality (cf. Ben-Porath 1997) is a proper superset of the backward-induction solution. Details are available upon request from the author.

Lemma 4 *Any game* Γ *admits a consistent path* $z \neq z^*$ *if and only if it admits a* x'*-consistent path* $z' \neq z^*$ *for some node* $x' \prec z^*$.

Proposition 1 can thus be equivalently presented as follows.

Corollary 4 *For any game* Γ, *the backward-induction path* z^* *may be persistently upset by some sequence of forgetfulness probabilities* $\boldsymbol{\varepsilon}^n \to 0$ *only if* Γ *admits a deviation* (x, a, z'), *and a* x'*-consistent path* $z \neq z^*$ *for some* $x' \prec z^*$.

4.2 A Sufficient Condition

This part of the section presents a sufficient condition for the backward-induction path to be persistently upset by small-probability forgetfulness. Under this condition, we can assign forgetfulness probabilities such that the long-run distribution of play will be a (non subgame-perfect) Nash Equilibrium.

This sufficient condition is more demanding than the necessary condition presented in Corollary 4. It is not only required that the players controlling node x may be lured off the backward-induction path because of the deviation (x, a, z'), and that there is a x'-consistent path z on which the long-run distribution of play can concentrate (for some node x' on the backward-induction path). The sufficient condition also depends on the relation between the deviation (x, a, z') and the x'-consistent path z. We are able to insure that the backward-induction path can be persistently upset by small-probability forgetfulness in either the case that x' comes after x, or in the case that x' coincides with x and that the choice of the players controlling node x is irrelevant to determining whether path z or z' takes place.[19]

Definition 9 *The paths* z *and* z' *are* i*-independent if for any* x *such that* $\iota(x) = i$, *it is the case that* $[z \succeq a(x) \Leftrightarrow z' \succeq a(x)]$ *for any* $a \in A(x)$.

[19] It may seem that if we dropped the requirement that the long-run distribution of play is persistent in Definition 4, our sufficient condition for the backward-induction path to be upset could be simplified so as to require only that the game admits deviations. Example 3, however, presents a game that admits a deviation and yet, for small-probability model-deterioration, the only long-run distribution of play is indistinguishable from the backward-induction path.

Theorem 5 *Suppose that the game Γ has a deviation (x, a, z') and a x'-consistent path $z \neq z^*$ for some node $x' : x \preceq x' \prec z^*$, such that z and z' are $\iota(x)$-independent if x' coincides with x. Then Γ has a Nash path $z'' \neq z^*$ such that $f^*(z'') = 1$ for some sequence of forgetfulness probabilities $\boldsymbol{\varepsilon}^n \to 0$.*

The proof of this Theorem builds on the following observations. By construction, the x'-consistent path z is established in the x'-subgame if each player forgets the opponents' actions that either do not lead into z or that do not support the backward-induction maximin solution for any of her own actions not leading into z. Some of these actions are off the backward-induction path, and thus can be forgotten by the players. The other ones are on the backward-induction path that starts at x'. In order to make sure that these actions can also be forgotten, suppose that the players controlling node x forget all actions that do not lead into z' in the x-subgame, and deviate from a_x^* to a at x. If x' comes after x, this makes off-path all actions in the x'-subgame (including the backward-induction path of the subgame). If x' coincides with x, since the backward-induction path of the x'-subgame starts with $a_{x'}^*$, the deviation makes off path all actions on the backward-induction path of the subgame. Once the path z is established in the x'-subgame, either some players deviate from the backward-induction path at some node that comes before x', or z is established in the entire game. In either case, a non-subgame perfect Nash path is established.

The difficulty associated with the case where x' comes before x lies in the fact that the players controlling node x' may or may not choose to deviate from the backward-induction path, depending on their model of the $a(x)$-subgame and in the resulting backward-induction solution. If they deviate they may modify upstream players' models and induce them to deviate too, thus starting a chain of possible deviations from the backward-induction path. Given that the long-run distribution of play is determined by the interaction of forgetfulness, recall and incomplete models' backward-induction solutions, it seems unlikely that one would be able to give a concise description of the payoff relations allowing these chains of possible deviations to generate persistent distributions (or even just long-run distributions of play that differ from the backward-induction path).

While these observations provide a guideline for establishing the result, the description of the specific dynamics that upset the backward-induction path and yield a Nash path in the long run is complicated by the interaction of forgetfulness, recall and incomplete model backward-induction. The major difficulty in the proof, in fact, lies in constructing forgetfulness probability profiles that insure that the complexity of the dynamic description is kept under control. Specifically, we set up the system so that it is very likely that the players' models establish the path z in the x'-subgame, before the proportion of players believing that z' is reached becomes non-negligible. Also, to keep under control the deviations in the x-subgame, we assume that players are more likely to forget downstream actions leading to z' rather than upstream ones. This insures that when a player deviates at x, it will be in order to reach z'.

In the case that the paths z' and z originate at the same node x, the requirement of $\iota(x)$-independence is imposed to insure that players from population $\iota(x)$ can forget all opponents' actions that do not lead into z', and at the same time formulate models that establish the path z. In the case that x' comes after x, the additional difficulty is to insure that the play along $a(x)$ does not get stuck on non-persistent paths. This is achieved by setting up the dynamic system so that when the players in population i deviate from a_x^* to a at x, only a negligible proportion of opponents formulate incomplete models of the x-subgame. A single path is thus established in almost all matches where a is taken. Then, by letting players forget all actions not on that path in an ordered manner, it can be insured that all players will correctly anticipate the path that takes place in their match, and the play will eventually settle on a Nash path.[20]

[20] If the play is almost concentrated on a single non-persistent path, we can insure that it will eventually settle on a Nash path by letting the players forget actions in an ordered manner; however, it does not seem possible to generalize this argument to the case where the distribution of play yields non-negligible mass on more than one non-persistent path. Recall occurs much more faster than forgetfulness, in fact, and players could recall actions that induce them to enter subgames where they are unable to correctly anticipate the established path, before forgetfulness takes place.

4.3 Games with Consistent Deviations

The characterization of games whose backward-induction path may be persistently upset presented in Theorem 5 depends on the interaction of several elements of the games. In order to discuss this characterization and elicit its properties in a more transparent manner, we will relate it to our leading examples presented in the second section. This is best accomplished by introducing two subclasses of games that are included in the class of games identified by Theorem 5.

We denote the first class as *games with consistent deviations.* These are essentially games where the terminal node z' that lures players off the backward-induction path at some node x *coincides* with the x'-consistent path z on which the distribution of play concentrates in the long run. This simple class of games includes, for instance, the game of Example 1. The path TH in fact yields players in population 1 a higher payoff than the backward-induction path N, and at the same time it yields players in population 2 a higher payoff than their backward-induction maximin path CF. Thus, TH is a consistent deviation.

Definition 10 *A perfect-information generic extensive-form game is a* game with consistent deviations *if there exists a deviation* (x, a, z) *such that* z *is* x*-consistent.*

Corollary 6 *If* Γ *is a game with consistent deviations, then there is a Nash path* $z'' \neq z^*$ *such that* $f^*(z'') = 1$ *for some sequence of forgetfulness probabilities* $\boldsymbol{\varepsilon}^n \to 0$.

The non subgame perfect long-run distribution of play established in games of consistent deviations display the following forward-induction property. The players from population $\iota(x)$ can "force" the play in the x-subgame onto the persistent path z, which makes them better off with respect to the backward-induction path of the x-subgame. This occurs even though the path z is not subgame perfect, as long as z is a Nash path. The terminal node z is then reached as long as it is not the case that players controlling nodes that come before x on the backward-induction path deviate off path to avoid reaching z. This property may be related to the forward-induction approach of Van Damme (1989), which observes that a player can force the play on her preferred Nash path; however, unlike Van Damme (1989), we do not require that the established path satisfies subgame perfection.

As well as identifying a simple class of games that satisfies the conditions of Theorem 5 and displays forward-induction properties, the concept of games with consistent deviations can be useful to point out that whenever a game is complex enough, the backward-induction path may be persistently upset. Specifically we show that any game can be enriched with actions irrelevant for the backward-induction solution and turned into a game of consistent deviations.[21]

Given a generic game $\Gamma = (X, Z, I, \iota, A, \mathbf{u})$, we construct a generic game Γ' that contains the game Γ, and such that $z^*(\Gamma')$, the backward-induction path of Γ', coincides with $z^*(\Gamma)$, the backward-induction path of Γ.

Definition 11 *The game* Γ' *is a* backward-induction irrelevant expansion *of* Γ *if* $X \subset X'$, $(I, \iota, A, \mathbf{u}|_{Z\cap Z'}) = (I', \iota'|_X, A'|_X, \mathbf{u}'|_{Z\cap Z'})$ *and* $z^*(\Gamma') = z^*(\Gamma)$.

In terms of complete-model analysis, the games Γ' and Γ represent essentially the same strategic interaction, and Γ' is a richer, more complex, model than Γ. Intuitively, a backward-induction-irrelevant expansion can be obtained by attaching trivial or conditionally dominated actions to the original game. Trivial actions are payoff-irrelevant, and since they do not change the backward-induction solution, they are usually interpreted as "inaction". By definition, conditionally dominated actions are never chosen by players holding complete models, and so they are irrelevant for backward induction. Proposition 7 points out that we can construct a game with the payoff structure of Example 1, around the last choice on the backward-induction path of Γ. That choice may thus be upset by small-probability model deterioration, and upstream backward-induction choices may also be upset in a "domino effect". In the statement of Proposition 7, the notation $f^*(\cdot, \Gamma')$ identifies the long-run distribution of play associated with the game Γ'.

[21] Since this result holds when restricting forgetfulness to opponents' unobserved actions, it holds a fortiori when these conservative assumptions are lifted.

Proposition 7 *For any game* Γ*, there exists a* backward-induction-irrelevant expansion Γ' *with a Nash path* $z' \neq z^*$ *such that* $f^*(z', \Gamma') = 1$ *for some sequence of forgetfulness probabilities* $\boldsymbol{\varepsilon}^n \to 0$.

4.4 Games with Fragile Path

In the discussion preceding Example 2, we have argued that the backward-induction path may be persistently upset by small-probability forgetfulness because of its fragility: once a player has deviated from the backward-induction path, her model may deteriorate so that she believes that the backward-induction choice is dominated by the deviation. This part of the section formally introduces the class of games with fragile paths, and shows that it is included in the class of games identified by Theorem 5. The key aspect of the definition is that for some population i, there is a deviation (x_i, a_i, z') such that all terminal nodes that come after action a_i yield the players in population i a higher payoff than the backward-induction maximin value of node x_i. The backward-induction path following x_i is fragile because the model of any player in population i will eventually deteriorate so that any terminal node following the deviation a_i will dominate the original backward-induction action $a^*_{x_i}$. While in games with consistent deviations the backward-induction path is upset because players may force the play off path to achieve a higher payoff, in games with fragile path it is upset because players' backward-induction solutions break down, and thus a smaller payoff may be achieved.

Definition 12 *A perfect-information generic extensive-form game is* a game with fragile path *if there is a population* i *and a deviation* (x_i, a_i, z') *such that* $\iota(x) \neq i$ *for any* $x \succeq a_i(x_i)$*, and* $u_i(z) \geq \underline{u}_i(x_i)$ *for any* $z \succ a_i(x_i)$.

The result that games with fragile path satisfy the condition of Theorem 5 is shown by constructing an appropriate x_i-consistent path that goes through the action a_i. This is always possible because any terminal node that comes after a_i yields player i a payoff larger than the backward-induction maximin value of node x_i.

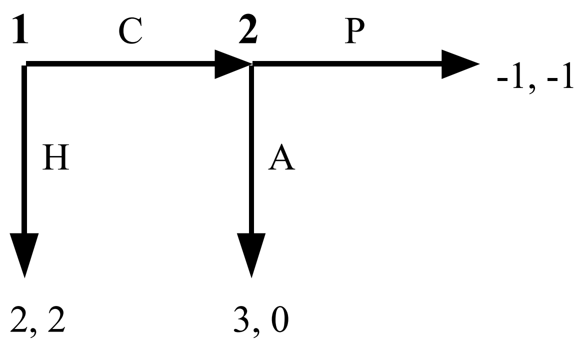

Figure 5: Punishment Game

Lemma 5 *If* Γ *is a game with fragile path, then there is a Nash path* $z'' \neq z^*$ *such that* $f^*(z'') = 1$ *for some sequence of forgetfulness probabilities* $\varepsilon^n \to 0$.

5 Extended Forgetfulness

5.1 Forgetfulness of Observed Actions

The central part of our exploration on the effect of model deterioration on the backward-induction solution has been conducted under the assumption that players can only forget opponents' unobserved actions. For the sake of completeness, we also look into the cases where observed actions or a player's own actions can be forgotten.

The first part of the section considers the case where all opponents' actions can be forgotten. Specifically, we operate under Assumptions 1 and 3. It should be pointed out that in this format forgetfulness is still a very unlikely event, while each player recalls with unit probability any action on path of which she was previously unaware.

With respect to the case where only unobserved actions can be forgotten, this environment is both simpler and less restrictive. As a result, the characterization of games whose backward-induction path can be persistently upset by small-probability model deterioration is both sharper and more disruptive of the backward-induction path. In order to introduce our results, consider the following example.

Example 5 Consider the punishment game presented in Figure 5. Each player in population 1 may either honor (H) or cheat (N) on her opponent. If cheated, a player in population 2 may either costly punish her opponent (P), or acquiesce (A). The backward-induction solution is CA. The players in population 1 forget with probability ε_P^n the action P and play C, and forget with probability ε_A^n the action A and play H, in which case they will never observe A, as they make it off-path. Thus, the system evolves as

$$\begin{cases} \rho_A^{t+1} = \varepsilon_A^n(1-\rho_A^t-\rho_P^t)+\rho_A^t \\ \rho_P^{t+1} = \varepsilon_P^n(1-\rho_A^t-\rho_P^t)+\rho_P^t, \end{cases}$$

and asymptotically reaches a stationary state ρ^n such that $\rho_A^n \in [0,1]$, and $\rho_P^n = 1-\rho_A^n$. The limit stationary distribution is $f^n(H)=\rho_A^n$, $f^n(CA)=\rho_P^n$. By varying the sequence of forgetfulness probabilities, the population play can asymptotically reach any heterogeneous-belief self-confirming equilibrium of the game, and when $\varepsilon_P^n/\varepsilon_A^n \to 0$, the non subgame perfect Nash path H is established. $\diamond$

The key difference with respect to the case where only unobserved actions can be forgotten consists in the fact that the backward-induction path of a game may be persistently upset even if the game does not have any deviation (x,a,z). Still, the presence of a consistent path z is required.

Proposition 8 *The backward-induction path of a game* Γ *is persistently upset by some sequence of forgetfulness probabilities* $\boldsymbol{\varepsilon}^n \to 0$ *if and only if* Γ *has a consistent path* $z \neq z^*$*. In such a case,* Γ *has a Nash path* $z'' \neq z^*$ *such that* $f^*(z'')=1$ *for some sequence* $\{\hat{\boldsymbol{\varepsilon}}^n\}$ *such that* $\hat{\boldsymbol{\varepsilon}}^n \to 0$.

As in the previous section, since a player's own actions cannot be forgotten, a type distribution is persistent only if it induces consistent paths. However, when all of the opponents' actions can be forgotten (including those on path), some players may deviate from the backward-induction path because they forget what terminal nodes it leads to. In order for small-probability forgetfulness to upset the backward-induction path, it is not necessary that there is a deviation (x,a,z'). It is enough to identify a consistent path z, different

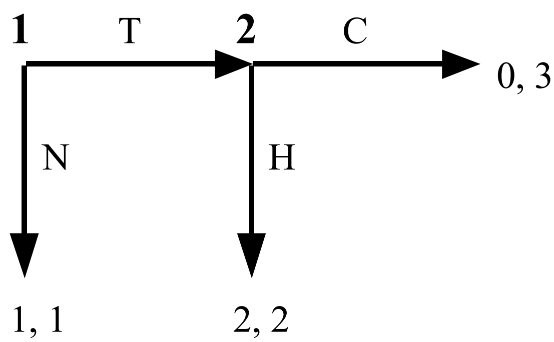

Figure 6: Simple Trust Game

from the backward-induction path, and suppose that the players' models deteriorate so as to assign the backward-induction maximin value to any action that does not lead into z, and so as to exclude all opponents' actions that do not lead into z at each node that comes before z.

5.2 Forgetfulness of Off-Path Actions

In this part of the section we explicitly allow players to also forget their own actions. However, in order to make the analysis comparable with the case of forgetfulness of opponents' unobserved actions, we maintain the assumption that a player cannot forget any action on path. This case is of interest because individuals are more likely to overlook the possibility of taking actions that they have seldom used in the past, rather than actions that they are accustomed to utilize. Formally, we operate under Assumption 2, and remove Assumption 1.

Despite the restriction of forgetfulness to off-path actions, one can construct simple examples where an initially fully-aware population deteriorates into a persistent type distribution inducing a distribution of play that is not even a self-confirming equilibrium.

Example 6 Reconsider the game of Example 3 (reproduced in Figure 6 for the ease of the reader), and assume for simplicity that all players may forget only action C, and that they all forget it with probability ε. Initially, they play the backward-induction solution

(N, C) with path N. As C is off-path, it can be forgotten. Players in population 1 unaware of C play T and players in population 2 unaware of C play H. Therefore, the population dynamics are described by the system:

$$\begin{cases} \rho_{1C}^{t+1} = \varepsilon(1 - \rho_{1C}^t) + \rho_{1C}^t \rho_{2C}^t \\ \rho_{2C}^{t+1} = \varepsilon(1 - \rho_{1C}^t) + \rho_{2C}^t. \end{cases}$$

The system approaches asymptotically the steady state ρ^* such that $\rho_{1C}^* = 1$, $\rho_{2C}^* = 1$, to which corresponds the distribution of play $f^*(TH) = 1$. This path cannot be supported by any self-confirming equilibrium, as the players in population 2, who have forgotten C, are not maximizing their payoff on path. Nevertheless, the profile ρ^* is persistent: all players are unaware of C, and those in population 1 anticipate path TH, while those in population 2 anticipate path H. In concluding this example, note that, by forgetting the possibility of playing H, the players in population 2 improve their payoff as well as that of their opponents. Once again, model deterioration plays a role similar to commitment in Pareto-improving the players' payoffs. $\diamond$

With respect to the characterization of games upset by small-probability forgetfulness, the key difference with respect to the case where only opponents' (unobserved) actions can be forgotten is the fact that the backward-induction path may be persistently upset even in games that do not have any consistent path. Still, the presence of a deviation (x, a, z) is required.

Proposition 9 *The backward-induction path of a game* Γ *is persistently upset by some sequence of forgetfulness probabilities* $\boldsymbol{\varepsilon}^n \to 0$ *if and only if* Γ *has a deviation* (x, a, z).

Intuitively, since actions on path cannot be forgotten, the backward-induction path will not be upset by small-probability forgetfulness if all the terminal nodes that follow a player's deviation make her worse off with respect to the backward-induction path. Since own action can be forgotten, however, the presence of a consistent path is not required. Once identified a deviation (x_i, a_i, z) for some population i, it is just enough to say that all players that do not belong to population i forget all actions that do not lead to z, faster than the players in population i. Eventually, the players in population i will deviate and reach the terminal node z.

6 Discussion

This section compares our construction and results with some enlightening papers in the literature applying evolutionary analysis to extensive-form games. One of the main issues driving this literature is that natural selection is not effective on choices off path. As a result, it is known since Selten (1983) that non-trivial extensive-form games do not admit any evolutionary stable strategy (in the symmetrized representation). Cressman and Schlag (1998), among others, analyze explicit dynamics, and specifically the replicator dynamics. They show that (Lyapounov) stable states are Nash equilibria. Analyzing an interaction equivalent to the game presented in Figure 5, they show that the backward-induction solution will not necessarily be selected by Lyapounov stability (see also Gale Binmore and Samuelson 1995, and Samuelson 1997 for analogous results).[22]

Turning to set-valued concepts, Cressman and Schlag (1998) study sets that are asymptotically stable with respect to trajectories starting in the interior of the mixed strategy profile space. They show that for any perfect-information generic extensive-form game where any path has at most one decision node off the backward-induction path, and this node has at most two choices, the Nash Equilibrium component associated with the backward-induction path is the unique minimal interior asymptotically stable set. In more complex games, however, they show that the Nash Equilibrium component associated with the backward-induction path fails to be interior asymptotically stable (even though the minimal interior asymptotically stable set must contain such a component).

Balkenborg and Schlag (2001) present a sharp characterization of the sets that are asymptotically stable with respect to any trajectory in the mixed strategy profile space,

[22] On a similar account, Roth and Erev (1995) present a different learning model where the backward-induction solution is not necessarily selected in the long run.

under any regular dynamics requiring that the growth rate of any pure best reply be non-negative, and that it be strictly positive unless all strategies played in the corresponding population are best replies.[23] They show that every connected and closed asymptotically stable set of rest points containing a pure strategy profile must be a strict equilibrium set, i.e. a Nash equilibrium component where a player unilaterally deviating from an equilibrium in the component either obtains a strictly smaller payoff or induces another equilibrium in the component.[24] They also prove a converse statement for the replicator dynamics (cf. Balkenborg 1994), for bimatrix games, and for convex strict equilibrium sets. In the environment of generic perfect-information games, they show that the Nash equilibrium component associated with the backward-induction path is asymptotically stable if and only if no player can unilaterally gain by deviating from it.

A learning model that focus on stochastic stability has been proposed by Noldeke and Samuelson (1993). They study finite populations, where each player is endowed with a "characteristic" consisting of a strategy and a conjecture about their opponents' strategies. At each period of play, with some positive probability, she reconciles her conjectures with the actual opponents' population play at nodes that are not precluded by her choice, and she sets her strategy as a best-reply to the conjecture. With small probability she may also randomly mutate her characteristic. In general, stochastically stable states are shown to be self-confirming equilibria. Noldeke and Samuelson (1993) present both a necessary and a sufficient condition for stochastic stability to predict that the backward-induction path is established.

Hart (2000) analyzes an evolutionary stochastic model formalized as a birth-death

[23] Such a class of dynamics is very large, as it includes all payoff-positive and payoff monotonic-dynamics, and hence the replicator dynamics.

[24] The concept of strict equilibrium set can be considered the set-valued analogue of strict equilibrium. Ritzberger and Weibull (1996) show that a pure strategy profile is asymptotically stable if and only if it is a strict equilibrium. Balkenborg (1995) fully describes the existence properties of strict equilibrium set in the environment of repeated games. He shows that strict equilibrium exists for arbitrarily long repetitions if and only if the stage game is of common interest. In such games, there is an upper bound (independent of the number of repetitions) on the number of inefficient outcomes that can occur along the outcome paths induced by strict equilibrium sets.

Markov chain. He assumes that natural selection forces act separately on the different genes that determine players' behavior at different nodes. Specifically, at each moment in time, a single gene (i.e. a single action at a single node) may change. With large probability, it will be subject to natural selection in the form of "better reply" dynamics, and otherwise it randomly mutates. As the mutation rate vanishes, the stationary distribution concentrates on the Nash Equilibrium set. Moreover, in the case that the sizes of the populations increase at the same time as the mutation rate vanishes, the stationary distribution concentrates on a small neighborhood of the backward-induction solution.[25]

Our approach is qualitatively different from evolutionary game theory. For one thing, the joint effect of model deterioration and incomplete-model backward-induction play induces a smaller range of plays than noisy mutation. While a population perturbed by random mutation necessarily induces full-support distribution of play, it is easy to see that, regardless of their model of the game, players who cannot forget their own actions will never play strictly dominated strategies. Secondly, and most importantly, the joint effect of forgetfulness and recall induces dynamic transitions in the distribution of play that are not necessarily monotonic in payoffs. In the analysis of Example 1, for instance, the proportion of players unaware of action A and playing action H is strictly increasing over time, despite the fact that players holding a complete model and thus playing action C achieve a higher payoff.

Given these major differences, it is not surprising that our results are different from those of evolutionary game theory. It is known in fact that evolutionary forces select a Nash Equilibrium (or at least a Self-Confirming Equilibrium in the approach of Noldeke and Samuelson 1993). As illustrated in Example 4, instead, model deterioration can have an even more disruptive effect on standard equilibrium solutions. It is possible to construct games where, starting from a situation of full awareness and backward-induction play, the play deteriorates into a long-run distribution that is not even a self-confirming equilibrium.

[25] On a similar account Hendon, Jacobsen and Sloth (1996) and Jehiel and Samet (2001) present different learning models that favor the backward-induction solution in the long-run.

Despite this major difference, our characterization of games whose backward-induction path may be persistently upset by small-probability forgetfulness of opponents' unobserved actions is comparable to some characterizations of the backward-induction solution's evolutionary stability. In the domain of games where each player moves only once, the necessary condition for the backward-induction solution to be a locally stable outcome provided by Proposition 4 in Noldeke and Samuelson (1993) is logically equivalent to our Corollary 6, which characterizes games with consistent deviations. Our Theorem 5 determines a necessary condition, tighter than Corollary 6, for the backward-induction path to not be persistently upset by small-probability model deterioration.

At the same time, our Proposition 1 identifies a sufficient condition for the backward-induction path to not be persistently upset by small-probability model deterioration. This sufficient condition is tighter than the sufficient condition for the backward-induction solution to be the unique locally stable component provided by Proposition 7 in Noldeke and Samuelson (1993). In the language of this paper, their result in fact states that if a game does not admit deviations, then the backward-induction solution is the unique locally stable component.

With respect to the sufficient conditions for the interior asymptotic stability of the backward-induction solution determined by Cressman and Schlag (1998), our results differ in the sense that we can construct games where the backward-induction solution is the unique minimal interior asymptotically stable set, but the backward-induction path is persistently upset by small-probability model deterioration. Consider the game of Example 2, for which we know that the backward-induction path may be persistently upset. However, this is a game where any path has at most one decision node off the backward-induction path, and this node has at most two choices, and thus the backward-induction solution is the minimal interior asymptotically stable set.

In order to explore the converse question of whether there are games whose backward-induction path is not evolutionarily stable, but may be persistently upset by model deterioration, we have analyzed within this paper's framework the examples provided by Cressman and Schlag (1998) of games where the backward-induction solution is not an interior asymptotically stable set. We have found that in any such example, the backward-induction path can be destabilized by small-probability forgetfulness of unobserved opponents' actions.

In the language of this paper, Balkenborg and Schlag (2001) show that the backward-induction Nash component of a game is asymptotically stable if and only if the game does not admit any deviation. Their characterization turns out to be equivalent to our Proposition 9. This means that model deterioration mimics evolutionary asymptotic stability if and only if it occurs with all actions off-path, and no actions on path. Example 3, however, shows that if model deterioration is further restricted to opponents' actions off path, then it cannot upset the backward-induction path of some games with deviations. In this sense, model deterioration turns out to be less disruptive than evolutionary forces. In fact, both our necessary condition (Proposition 1) and our sufficient condition (Theorem 5) for the backward-induction path to be persistently upset identify a smaller class of games than the one determined by Balkenborg and Schlag (2001).

Finally, our results are clearly quite different from those of Hart (2000), who gives a general account of the evolutionary stability of the backward-induction solution. The key assumption in Hart (2000) is that evolutionary forces act separately on different nodes of the same player's choices. Such an assumption is unrestrictive in games where each player moves only once. In other words, evolutionary forces are more likely to select the backward-induction path in this specific class of games. In our analysis, the opposite is true, since players can only forget their opponents' actions, forgetfulness is more disruptive in games where each player moves only once. Specifically, if the backward-induction path may be upset in an arbitrary extensive-form game, then it may also be upset in the associated (extensive-form) agent-representation game, but the converse is not necessarily true.

7 Conclusion and Possible Extensions

This paper has presented and analyzed a social learning construction that explicitly keeps track of the evolution of models held by players who are able to solve perfect-information extensive-form games according to the models they formulate, and whose models of the games depend on past observation of play. We have introduced the possibility of small-probability model deterioration and have shown that it may upset the complete-model backward-induction path, even when the deterioration is restricted to only the opponents' unobserved actions. Necessary conditions and sufficient condition for the backward-induction path to be upset have been presented. This characterization shows that model deterioration is less disruptive than evolutionary forces for the backward-induction path. In games where the sufficient condition is satisfied, we have shown that model deterioration may yield long-run paths that are Nash but non subgame perfect.

A possible extension of this paper would be a general analysis of the dynamic construction introduced here. The assumption that all players initially hold complete models is made here only to highlight our results on the robustness of the backward-induction path. One may want to address the question of whether players may learn the game starting from a situation of partial awareness. The key issue then becomes the analysis of the basin of attraction of *any* stable states of the dynamics, without restricting attention to those reached from an initial state of full awareness. The characterization of stable sets, and of their attraction sets, would then complete the analysis.

Another possible extension would be a dynamic construction that considers general extensive-form games. The difficulty with this extension lies in the fact that, unlike perfect-information generic games, general games may possess multiple solutions, so that a result analogous to Lemma 1 cannot be provided. In our dynamic construction, the evolution of models depends on past play. Thus, the construction would not be well specified unless a unique solution is selected for each profile of models at each period of play.

A Appendix: Omitted Proofs

In order to simplify notation, in the following proofs for any arbitrary player i, any arbitrary model B^i, any arbitrary node $x \in X_i$, and any node $x' \succ x$, we will denote $u_i^*(B^i, x)(x')$ as $u(B^i, x)(x')$, as well as $u_i^*(B^i, x)(a(x))$ as $u_i^*(B^i, x)(a)$, and $u_i^*(B^i, x)(x)$ as $u(B^i, x)$, also we will denote $a_{x'}^*(B^i, x)$ as $a_{x'}(B^i, x)$ and $a_x^*(B^i, x)$ as $a(B^i, x)$. For any terminal node z, let $\mathbf{a}^z = \{a : a(x) \preceq z \text{ for some } x\}$, and $a_x^z = \mathbf{a}^z \cap A(x)$ for any node $x \prec z$. For any node x, let $\underline{\mathbf{a}}(x)$ be the $\iota(x)$ BI-maximin solution starting at node x. For each set of nodes Y, set $UC(Y) = \{y \in Y : \text{for any } y' \in Y,\ y \succeq y'\}$, and $LC(Y) = \{y \in Y : \text{for any } y' \in Y,\ y \preceq y'\}$

Proof of Lemma 3. Pick any profile $B \in \mathit{Supp}(\rho^*)$, and let z denote $z^*(B)$. Pick any node $x \prec z$, since ρ^* is persistent, it must be the case that $u(B^{\iota(x)}, x) = u_{\iota(x)}(z)$, and that $u(B^{\iota(x)}, x)(a_x^z) = u_{\iota(x)}(z)$. The requirement that $z^*(B) = z$ is satisfied if and only if $a(B^{\iota(x)}, x) = a_x^z$. This requires that $u(B^{\iota(x)}, x)(a_x^z) \geq u(B^{\iota(x)}, x)(a)$ for any $a \in A(x)$.

Since $\rho_{\iota(x)}^*(B^{\iota(x)}) > 0$, it must be the case that $\rho_{\iota(x)}^t(B^{\iota(x)}) > 0$ for any t large enough. By Assumption 1, this implies that for any $x' : \iota(x') = \iota(x)$ it must be the case that $B^{\iota(x)}(x') = A(x')$. By definition of BI maximin, this implies that $u(B^{\iota(x)}, x)(a) \geq \underline{\mathrm{u}}_{\iota(x)}(a)$ for any $a \in A(x)$. In conclusion we obtain that:

$$u_{\iota(x)}(z) = u(B^{\iota(x)}, x)(a_x^z) \geq u(B^{\iota(x)}, x)(a) \geq \underline{\mathrm{u}}_{\iota(x)}(a), \text{ for any } a \in A(x).$$

Since this result holds for any arbitrary B, and $x \prec z(B)$, the claim is proved. ■

Proof of Proposition 1. The proof proceeds in two steps.

Step 1. The BI path z^* may be persistently upset only if Γ admits a deviation (x, a, z'). Pick an arbitrary time $t \geq 1$, and say that $f^{t-1}(z^*) = 1$.

Let $X^* = \{x : x \prec z^*\}$ denote the set of nodes on the BI path. Order the nodes on X^* assigning for each set Y_j the index j to the node $x \in X^* \cap Y_j$.

Take x_1, and denote by i the index $\iota(x_1)$. By construction and Assumption 1, for any $B^i \in \mathit{Supp}(\rho_i^t)$, it must be the case that $B^i(x) = A(x)$. Thus for any $a_i \in A(x_1)$ such that $a_i(x_1) \in Z$, it must be the case that $u(B^i, x_1)(a_i) = u_i(a_i(x_1))$. By construction, $a_{x_1}^*(x_1) \in Z$. Thus $u(B^i, x_1)(a_{x_1}^*) = u_i(a_{x_1}^*(x_1)) = u_i^*(x_1)$. For any B^i, and action $a_i \in A(x_1)$, it must be the case that $u(B^i, x_1)(a_i) \leq \max\{u_i(z') : a_i(x_1) \prec z'\}$. Suppose that for any action $a_i' \in A(x_1) \backslash \{a_{x_1}^*\}$, there are no terminal nodes z', satisfying $a_i(x_1) \prec z'$ and $u_i(z') > u_i^*(x_1)$. Then, it must be the case that $u(B^i, x_1)(a_{x_1}^*) = u_i^*(x_1) > u(B^i, x_1)(a_i)$ for any $B^i \in \mathit{Supp}(\rho_i^t)$.

Pick now an arbitrary $K \geq 2$, denote by l the index $\iota(x_K)$ and say that for any $k : 1 \leq k < K$, it has been shown that, if for any action $a' \in A(x_k) \backslash \{a_{x_k}^*\}$ there are no terminal nodes z' satisfying $a'(x_k) \prec z'$ and $u_{\iota(x_k)}(z') > u_{\iota(x_k)}^*(x_k)$, then $a(B^{\iota(x_k)}, x_k) = a_{x_k}^*$, for any $B^{\iota(x_k)} \in \mathit{Supp}(\rho_{\iota(x_k)}^t)$.

The BI solution $a^*(B^l, x_K)$ of any type B^l at node x_K is constructed by introducing the model $G(B^l, x_K) = (Y^l_{x_K}, B^l_{x_K})$, and then proceeding as follows. Set $\mathbf{u}^*(B^l, x_K)(z') = \mathbf{u}(z')$ for any $z' \in Z \cap Y^l_{x_K}$. For any $j \geq 1$, and for any $x \in Y_j \cap Y^l_{x_K}$, let

$$a^*_x(B^l, x_K) = \arg\max_{a \in B^l_{x_K}(x)} u^*_{\iota(x)}(B^l, x_K)(a(x)); \;\; \mathbf{u}^*(B^l, x_K)(x) = \mathbf{u}^*(B^l, x_K)(a^*_x(B^l, x_K)(x)).$$

Pick an arbitrary $B^l \in Supp(\rho^t_l)$. By Assumption 1, $B^l(x_K) = A(x_K)$. For any $k : 1 \leq k < K$, and any $B^{\iota(x_k)} \in Supp(\rho^t_{\iota(x_k)})$, it must be the case that $B^l(x_k) \subseteq A(x_k) = B^{\iota(x_k)}(x_k)$, and since $f^{t-1}(z^*) = 1$, it follows that $a^*_{x_k} \in B^l$.

Assume that for any $k : 1 \leq k < K$, and any action $a' \in A(x_k) \backslash \{a^*_{x_k}\}$ there are no terminal nodes z' satisfying $a'(x_k) \prec z'$ and $u_{\iota(x_k)}(z') > u^*_{\iota(x_k)}(x_k)$. By the induction hypothesis, for any $k : 1 \leq k < K$, and any $B^{\iota(x_k)} \in Supp(\rho^t_{\iota(x_k)})$, it is the case that $a(B^{\iota(x_k)}, x_k) = a^*_{x_k}$. Since $a^*_{x_k} \in B^l(x_k) \subseteq B^{\iota(x_k)}(x_k)$ for any k, it follows that $a_{x_k}(B^l, x_K) = a^*_{x_k}$ for any k. This establishes that $a(B^l, x_K) = a^*_{x_K}$.

The claim is then proved by induction.

Step 2. The BI path z^* may be persistently upset only if Γ admits a *consistent* path $z \neq z^*$.

Pick a self-confirming type distribution ρ^* inducing path distribution f^* such that $f^*(z^*) < 1$. There must exist a profile $B \in Supp(\rho^*)$ such that $z^*(B) \neq z^*$. By Lemma 3, the path $z^*(B)$ must be consistent. ■

Proof of Lemma 4. The only if part is trivial since $x_0 \prec z^*$ and every consistent path is x_0-consistent by definition.

Pick a x'-consistent path $z' \neq z^*$ for some node $x' \prec z^*$. Say that z' is not consistent and thus the set $\bar{X} = \{x \prec x' | u_{\iota(x)}(z') < \underline{u}_{\iota(x)}(x)\}$ is non-empty. Let $\bar{x}$ denote the unique element of $UC(\bar{X})$. Since $x' \prec z^*$ it follows that $a^{z'}_{\bar{x}}(\bar{x}) \prec z^*$, and since z' is x'-consistent, it must the case that $a^{z'}_{\bar{x}}(\bar{x}) \prec z'$.

Define $\bar{a} = \arg\max_{a \in A(\bar{x})} \underline{u}_{\iota(\bar{x})}(\bar{x})$ and let $\bar{z}$ denote the BI solution of subgame $\Gamma_{\bar{a}(\bar{x})}$. By definition of BI maximin, the path $\bar{z}$ is $\bar{x}$-consistent. By construction $\bar{a} \neq a^{z'}_{\bar{x}}$, thus $\bar{z} \neq z^*$.

In conclusion, we have shown that if the game Γ has x'-*consistent* path $z' \neq z^*$ for some node $x' \prec z^*$, and z' is not x_0-consistent, then Γ must have $\bar{x}$-*consistent* path $\bar{z} \neq z^*$ for some node $\bar{x} \prec x' \prec z^*$. Since the node x' is arbitrary and since the game has a finite set of nodes, it follows that Γ must have a x_0-consistent path $z \neq z^*$. ■

Proof of Theorem 5. For the sake of clarity, the proof is divided in 2 parts.

Part 1. *If there exists a deviation (x_i, a_i, z') for some i, and a x_i-consistent path $z \succ a_i(x_i)$ such that z and z' are is* i-independent*, then for some sequence of forgetfulness probabilities $\boldsymbol{\varepsilon}^n \to 0$, there is a Nash path $z'' \neq z^*(\Gamma)$ such that $f^*(z'') = 1$.*

We shall divide the construction in 5 steps. For notational ease we shall drop the denomination n from the variables.

Step 1. *The forgetfulness probabilities.*

Introduce the sets $X^* = \{x : x_i \preceq x \prec z\}$, $X' = \{x : x_i \prec x \prec z\}$ and $X'' = \{x : x_i \preceq x \prec z'\}$. Set $X^+ = X^* \backslash X''$, and $\bar{x}$ be the unique element of $UC(X^* \backslash X^+)$.

For any population j, introduce the set $X'_j = \{x : x \in X', \text{and } i(x) = j\}$. For population i, also introduce the set $X^*_i = X'_i \cup \{x_i\}$. By hypothesis, for any $x \in X^*_i$ and any $a \in A(x)$, if $z \succeq a(x)$, then $z' \succeq a(x)$, thus for any $x \in X^+$, it must be the case that $\iota(x) \neq i$. Consider the game Γ^z obtained by collapsing the $a_i(x_i)$-subgame into a terminal node $\bar{z}$ such that $\mathbf{u}(\bar{z}) = \mathbf{u}(z)$. Letting $\mathbf{a}^{*z}$ be the BI solution of Γ^z, put $\hat{X} = \{x | \iota(x) \neq i,\ x \prec z^*,\ a^{*z}_x \neq a^*_x\}$, and whenever $\hat{X}$ is non-empty, denote as $\hat{x}$ the unique element of $LC(\hat{X})$.

We enumerate the nodes $x_i \in X^*_i$ with order inverse to $\prec$, so that x^1_i is the node $x \in X^*_i$ such that $x \succeq x'$ for any $x' \in X^*_i$, and for any $k \geq 1$, x^{k+1}_i denotes the node $x \in X^*_i$ such that $x \succeq x'$ for any $x' \in X^*_i \backslash \{x^1_i, ..., x^k_i\}$. Let $K = \# X^*_i$, $x^0_i = z$, and for any $k = 1, ..., K$, we introduce the set $X^k_{-i} = \{x : x \in X'',\ i(x) \neq i \text{ and } x^k_i \prec x \prec x^{k-1}_i\}$.

For any node $x \in X^*$, let $A'(x) = \{a \in A(x) : a(x) \npreceq z\}$ be the set of actions that do not lead into z. Let $A''(x) = \{a \in A(x') : \exists a' \in A'(x),\ a'(x) \prec a(x'),\ \iota(x') \neq \iota(x), \text{ and } a \notin \underline{\mathbf{a}}(a'(x))\}$ be the set of actions alternative to the BI-maximin solution of actions in $A'(x)$. Let $A''_*(x_i) = \{a \in A(x') : a^*_{x_i}(x_i) \prec a(x'),\ \iota(x') \neq \iota(x), \text{ and } a \notin \underline{\mathbf{a}}(a^*_{x_i}(x_i))\}$, and $A''_-(x_i) = A''(x_i) \backslash A''_*(x_i)$. For any population $j \neq i$, introduce the sets $A'_j = \left(\cup_{x \in X'_j} A''(x)\right) \cup \left(\cup_{x \in X' \backslash X'_j} A'(x)\right)$. For population i, introduce the sets $A'_i = \left(\cup_{x \in X'_i} A''(x)\right) \cup A''_-(x_i)$, $A^+_i = \left(\cup_{x \in X^+ \backslash X'_i} A'(x)\right)$, and $\hat{A}_i = \{a \in A(\hat{x}) : a \neq a^{*z}_{\hat{x}}\}$.

Note that for any j, and $x \in X'$ it is the case that $A(x) \backslash A'_j \neq \emptyset$, and that, if $\iota(x) = j$, then $A(x) \cap A'_j = \emptyset$. This allows us to introduce, for any population j, the set of types $\mathbf{B}^j_0 = \{B^j = A \backslash C \text{ for some } C \subseteq A'_j\}$, and the types $B^j_z = A \backslash A'_j$. Note also that $A''_*(x_i) \cap A'_i = \emptyset$, and that for any x it is the case that $A(x) \backslash A''_*(x_i) \neq \emptyset$, and that, if $\iota(x) = i$, then $A(x) \cap A''_*(x_i) = \emptyset$.

For any $x \in X''$, we introduce the sets $A'_{z'}(x) = \{a \in A(x) : a(x) \npreceq z'\}$. We then introduce the set $A^0_i = \emptyset$, the set $A^1_i = A^i_+ \cup A''(\bar{x}) \cup \left(\cup_{x \in X^1_{-i}} A'_{z'}(x)\right)$, and the set $A^k_i = \cup_{x \in X^k_{-i}} A'_{z'}(x)$ for any $k = 2, ..., K$. Since for any $x \in X^+$, $\iota(x) \neq i$, and because of the way that sets $A'(x)$ are constructed, it follows that for any $k = 1, ..., K$, it is the case that $A''_*(x_i) \cap A^k_i = \emptyset$, that $A'_i \cap A^k_i = \emptyset$, and that for any x, $A(x) \backslash A^k_i \neq \emptyset$, and if $\iota(x) = i$, then $A(x) \cap A^k_i = \emptyset$.

Thus, letting $\mathbf{A}^0_i = A'_i$, and for any $k = 1, ..., K$, letting $\mathbf{A}^k_i = A'_i \cup A^i_+ \cup A''(\bar{x}) \cup A^1_i \cup ... \cup A^k_i$, we can introduce the set of types $\mathbf{B}^i_k = \{B^i = A \backslash C \text{ for some } C,\ \mathbf{A}^{k-1}_i \subseteq C \subseteq \mathbf{A}^k_i \cup A''_*(x_i)\}$, and the set of types $\hat{\mathbf{B}}^i_k = \{B^i = A \backslash C \text{ for some } C,\ \mathbf{A}^k_i \subseteq C \subseteq \mathbf{A}^k_i \cup A''_*(x_i)\}$.

Note that $B^i_z \in \mathbf{B}^i_1$, and that for any k, $\hat{\mathbf{B}}^i_k \subset \mathbf{B}^i_{k+1}$.

We also define $B^{iz} = A \backslash \left(\mathbf{A}^K_i \backslash \{a^z_{\bar{x}}\} \cup A''_*(x_i) \right)$, $B^{iz'} = A \backslash \left(\mathbf{A}^K_i \cup A''_*(x_i) \right)$ and $\mathbf{B}^{i*} = \left\{ B^{iz}, B^{iz'} \right\}$. For any $j \neq i$, define the set $\mathbf{B}^{j*} = \{B^j = A \backslash C$ for some C, $A'_j \subseteq C \subseteq A'_j \cup A'(x_i)\}$, and let $B^{j*} = A \backslash \left(A'_j \cup A'(x_i) \right)$.

Construct the forgetfulness probability profiles as follows. For any index $n \geq 1$, any population j, and any action $a \in A'_j$, let $\varepsilon^n_j(a) = 1/n$. For any $k = 1, ..., \#X^*_1 + 1$, let $m_k = \sum_{j \in I} \#(A'_j) + \sum_{l=0}^{k-1} \#(A^l_{z'})$, and set $\varepsilon^n_i(a) = 1/n^{1+m_k}$ for any action $a \in A^k_{-i}$. Also set $\varepsilon^n_i(a) = 1/n^{1+m_1}$ for any action $a \in A^+_i \cup A''_*(x_i)$, $\varepsilon^n_j(a) = 1/n^{1+m_1}$ for any action $a \in A'(x_i)$ and any index $j \neq i$, and $\varepsilon^n_i(a) = 1/n^{1+m_{K+1}}$ for any action $a \in \hat{A}_i$. For any other action a' and population index j, set $\varepsilon^n_j(a') = 0$.

Step 2. *For any $\delta > 0$, there exists a n large enough, and a $T(\delta, n)$ large enough, such that $\rho^T_j(B^j_z) > 1 - \delta$, for any j.*

For any non-empty set X'_j, and any $x \in X'_j$, by construction $a_i(x_i) \preceq x$. Any player from population j that belongs to a match B, where $a(B^i, x_i) \neq a_i$ will not observe any action $a \in A'_j$, thus she can forget a with probability $\varepsilon^n_j(a) = 1/n$. For any action $a' \notin A'_j$, it is the case that $\varepsilon^n_i(a') = o(1/n^m)$, and by construction $\#(A'_j) < m$. It follows that, for any t, as long as $\sum_{B^i \in \mathcal{B}^i} a(B^i, x_i)\rho_i(B^i) \leq o(1/n^m)$, for any $C \subseteq A'_j$,

$$\rho^{t+1}_j(A \setminus C) = \sum_{\emptyset \subseteq C' \subseteq C} \rho^t_j(A \setminus C') \prod_{a \in C \setminus C'} \varepsilon_j(a) \prod_{a \in C'} [1 - \varepsilon_j(a)] + o(1/n^m). \tag{2}$$

Pick any player from population i, holding a model $B^i \in B^i_0$. If X_i is empty, it is trivially the case that $u(B^i, x_i)(a) = u(A, x_i)(a)$, for any $a \in A(x_i)$. Suppose thus that X_i is non-empty. For any $x \in X_i$, because of the manner we have constructed sets $A''(x)$, it follows that $u(B^i, x)(a) \leq u(A, x)(a)$, for any $a \in A(x)$. Since for any $x \in X_i$, it is the case that $a_i \prec x$, it follows that $u(B^i, x_i)(a_i) \leq u(A, x_i)(a_i)$. For the same reason, it follows that, for any action $a'_i \in A(x_i)$, $a_i \neq a'_i$, and any node $x \succeq a'_i(x_i)$, it is the case that $B^i(x) = A(x)$, and thus that $u(B^i, x_i)(a'_i) = u(A, x_i)(a'_i)$. Applying the definition of BI solution, it follows that for any $a \in A(x_i)$,

$$u(B^i, x_i)(a^*_{x_i}) = u(A, x_i)(a^*_{x_i}) = u^*_i(a^*_{x_i}) \geq u^*_i(a) = u(A, x_i)(a) \geq u(B^i, x_i)(a).$$

This shows that if $\rho^t_i(B^i_0) = 1 - o(1/n^m)$, then $\sum_{B^i \in \mathcal{B}^i} a(B^i, x_i)\rho_i(B^i) \leq o(1/n^m)$. Since $\rho^0(A) = 1$, and $A \in B^i_0$, the recursive application of Equation (2) implies that, for any $\delta > 0$, there exists a n large enough, and a $T_0(\delta, n)$ large enough, such that $\rho^{T_0}_j(B^j_z) > 1 - \delta$, for any j.

Step 3. *Take any profile B such that $B^i \in \mathbf{B}^i_k$ for some $k = 1, ..., \#X^*_1$, and $B^j \in \mathbf{B}^{j*}$ for any $j \neq i$. If $x_i \prec z^*(B)$ then $z^*(B) \in \{z^*, z\}$.*

Pick an arbitrary population $j \neq i$, type $B^j \in \mathbf{B}^{j*}$, and node $x \in X'_j$. By construction, $u(B^j, x_j)(a_x^z) = u_j(z) \geq \max_{a \in A(x)} u(B^j, x)(a)$. Thus type B^j will take action a_x^z at node x.

Pick an arbitrary type $B^i \in \mathbf{B}_1^i$. For any $a \in A(x_i) \backslash \{a_i, a_{x_i}^*\}$, by construction, $u(B^i, x_i)(a) = \underline{\mathrm{u}}_i(a(x_i)) < u_i^*(a(x_i))$. Pick an arbitrary node $x \in X_i^{'}$. By hypothesis, it is the case that $z' \succeq a_x^z(x)$. By construction, for any action $a \in A(x) \backslash \{a_x^z\}$, it is the case that $u(B^i, x)(a) = \underline{\mathrm{u}}_i(a(x)) < u_i(z) \leq u_i(z')$. By construction, for any $k > 1$, and any $x \in X_{-i}^k$, it is the case that $B^i(x) = A(x)$. It follows that for any $x \in X_i^*$, $u(B^i, x_i)(a_x^z(x)) > u_i^*(a_x^z(x))$ only if $u(B^i, x_i)(a_x^z(x)) = u_i(z'')$ for some $z'' \succeq x_i^1$. By definition of BI solution, it must then be the case that either $a(B^i, x_i) = a_{x_i}^*$, or that $a(B^i, x) = a_x^z$ for any $x \in X_i^*$.

Consider now any set $B^i \in \mathbf{B}_k^i$ for some $k = 1, ..., \#X_1^*$. By construction, for any $l > k$, and any $x \in X_{-i}^l$, it is the case that $B^i(x) = A(x)$, again it follows that for any $x \in X_i^*$, $x \preceq x_k$, $u(B^i, x_i)(a_x^z(x)) > u_i^*(a_x^z(x))$ only if $u(B^i, x_i)(a_x^z(x)) = u_i(z'')$ for some $z'' \succeq x_i^1$. At the same time, for any $x \in X_i^*$, $x \succ x_k$, it is the case that $u(B^i, x_i)(a_x^z(x)) \geq u_i(z')$. Again, it must then be the case that either $a(B^i, x_i) = a_{x_i}^*$, or that $a(B^i, x) = a_x^z$ for any $x \in X_i^*$. Thus any profile (B_z^{-i}, B^i) such that $B^i \in \mathbf{B}_k^i$, for some $k = 1, ..., \#X_1^*$ establishes path z^* or z.

Step 4. *For any n, it is the case that $\rho_i^t(\mathbf{B}^{i*}) \to 1$ as $t \to \infty$, and that $\rho_j^t(B^{j*}) \to 1$ as $t \to \infty$.*

For any node x, and any action $a \in A''(z') \cap A(x)$, by construction, it must be the case that $a_i(x_i) \preceq x$. Any player from population i of type $B^i \in \mathbf{B}_1^i$ who plays $a(B^i, x_i) = a_{x_i}^*$ will not observe any action $a \in \mathbf{A}_i^1$, and thus she may forget it with probability $\varepsilon_i^n(a) > 0$. For any other action $a' \in B^i$, it is the case that $\varepsilon_i^n(a') \leq o(\sum_{a \in \mathbf{A}_i^1} \varepsilon_i^n(a))$. Since $B_z^i \in \mathbf{B}_1^i$, this implies that either there are types $B^i \in \mathbf{B}_1^i$ such that $a(B^i, x_i) = a_i$, or for any $\delta > 0$, there exists a n large enough, and a $T_1(\delta, n) > T_0(\delta, n)$ large enough, such that $\rho_i^{T_1}(\hat{\mathbf{B}}_1^i) > 1 - \delta$. With analogous reasoning it is concluded that, for any $k > 1$, if for any $B^i \in \mathbf{B}_k^i$, it is the case that $a(B^i, x_i) = a_{x_i}^*$, then for any $\delta > 0$, there exists a n large enough, and a $T_k(\delta, n) > T_k(\delta, n)$ large enough, such that $\rho_i^{T_1}(\hat{\mathbf{B}}_k^i) > 1 - \delta$. By construction, $a(B^i, x_i) = a_i$ for any $B^i \in \hat{\mathbf{B}}_k^i$. This implies that, for any t, any player in population i who is of type $B^i \in \mathbf{B}_k^i$ for some $k = 1, ..., K$ will almost-surely play a_i at some eventual time $\tau \geq t$.

For any $\delta > 0$, there is a T large enough such that for any $t \geq T$, $\rho_{-i}^t(\mathbf{B}^{j*}) > (1 - \delta)^{I-1}$. Pick any arbitrary player of any arbitrary type $B^i \in \mathbf{B}_k^i$ for some k, who plays $a(B^i, x_i) = a_i$ and belongs to the match $B = (B^{-i}, B^i)$ such that $B^j \in \mathbf{B}^{j*}$ for any $j \neq i$. Say first that $x_i \prec z^*(B)$. By Step 4, the player observes the path z. By construction, it must be the case that $u(B^i, x_i)(a_i) \geq u_i(z)$. By construction, for any node x, and any action $a \in A''(x_i) \cap A(x)$, it must be the case that $a(x) \not\prec z$. If it is not the case that $x_i \prec z^*(B)$, then obviously for any node x, and any action $a \in A''(x_i) \cap A(x)$, it must be the case that

$a(x) \not\prec z^*(B)$. Thus, in either case, the player of type B^i may forget any action $a \in A''(x_i)$, with probability $\varepsilon_i^n(a) > 0$. If the player forgets all actions $a \in A''(x_i)$, then at time $t+1$ she will assume a type $\bar{B}^i \in \mathbf{B}^{i*}$ such that

$$u(\bar{B}^i, x_i)(a^*_{x_i}) < \underline{u}_i(a^*_{x_i}(x_i)) < u_i(z) \leq u(\bar{B}^i, x_i)(a_i).$$

The player of type $\bar{B}^i$ cannot ever recall any action $a \in (A \backslash B^i) \cap A''(x_i)$. The player will never play $a^*_{x_i}$ again. If $x_i \prec z^*(\bar{B}^i, B^z_{-i})$, the player will observe path z with probability larger than $(1-\delta)^{I-1}$, and forget all actions $a \in \mathbf{A}_i^K \backslash \{a^z_{\bar{x}}\} \cup A''_*(x_i)$. Obviously, she will forget those actions with positive probability also whenever $x_i \not\prec z^*(\bar{B}^i, B^z_{-i})$. Since the player i and type B^i were arbitrary, we conclude that $\rho_i^t(\mathbf{B}^{i*}) \to 1$, as $t \to \infty$.

Pick an arbitrary index $j \neq i$. First observe that, since for any $B^i \in \mathbf{B}^i_k$, and any k, it is the case that $a(B^i, x_i) \in \{a_i, a^*_{x_i}\}$, any player of type B^j can forget action $a \in A(x_i) \backslash \{a_i, a^*_{x_i}\}$ with positive probability, and will never recall it. The player will forget action $a^*_{x_i}$ with positive probability (at least) when facing a player from population i of type B^i who plays $a(B^i, x_i) = a_i$. The player may observe action $a^*_{x_i}$ on path only if she is facing a player of type B^i who plays $a(B^i, x_i) = a^*_{x_i}$. Since $\rho_i^t(\mathbf{B}^{i*}) \to 1$ as $t \to \infty$, for any fixed n, it follows that $\rho_j^t(B^{j*}) \to 1$ as $t \to \infty$, as well.

Step 5. *There is a Nash path $z'' \neq z^*$ such that $f^t(z'') \to 1$ as $t \to \infty$.*

By Step 2, for any $\delta > 0$, for any $\delta > 0$, there exists a n large enough, and a $T_0(\delta, n)$ large enough, such that for any $t \geq T_0(\delta, n)$, and any $j \neq i$, it is the case that $\rho_j^t(B^j_z) > 1-\delta$, and $\rho_i^{T_0}(B^i_z) > 1-\delta$.

Suppose first that there is a population $j \neq i$, and a node x such that $u^*(x) = u^*(z^*)$, $\iota(x) = j$, and $a(B^{j*}, x) \neq a^*_x$. Thus the set $\hat{X}$ is non-empty, and at any $t \geq T_0$, any player in population i will not observe any action $a \in \mathbf{A}_i^K \cup A''_*(x_i)$, so that $\rho_i^t(B^{iz'}) \to 1$ as $t \to \infty$. Letting z^+ be the BI path of game Γ^z, since $\hat{A}_i \subset \mathbf{A}_i^K$, it follows that for any $x \prec \hat{x} : \iota(x) = i$, it must be the case that $u(B^{iz'}, x)(a^{*z}_{\hat{x}}(\hat{x})) = u_i(z^+)$, and that $a(B^{iz'}, x) = a^{*z}_x$. For any other index j, and node $x \preceq \hat{x} : \iota(x) = j$, by construction it is also the case that $a(B^{j*}, x) = a^{*z}_x$. It immediately follows that for any node $x \preceq \hat{x}$, and any action $a \in A(x)$ there is a path $\bar{z} \succeq a(x)$ such that $u_{\iota(x)}(\bar{z}) \leq u_{\iota(x)}(z^+)$. For any node $x \succeq a^{*z}_{\hat{x}}(\hat{x})$ it is the case that $B^{\iota(x)*}(x) = A(x)$ and that $B^{iz'}(x) = A(x)$, it follows that $a^{*z}_x = a^*_x$. Again immediately follows that for any node $x \succeq a^{*z}_{\hat{x}}(\hat{x})$, and any action $a \in A(x)$ there is a path $\bar{z} \succeq a(x)$ such that $u_{\iota(x)}(\bar{z}) \leq u_{\iota(x)}(z^+)$. These two results together imply that z^+ is a Nash path of Γ. Since $\rho_i^t(B^{iz'}) \to 1$ and $\rho_j^t(B^{j*}) \to 1$ as $t \to \infty$, for any fixed n, it follows that $f^t(z^+) \to 1$ as $t \to \infty$, for any n.

Secondly, say that for all populations $j \neq i$, and nodes $x \prec z$ such that $u^*(x) = u^*(z^*)$, $\iota(x) = j$, it is the case that $a(B^{j*}, x) = a^*_x$. For any x such that $\iota(x) = i$, and $z \succ x$, by hypothesis, it must be the case that $z' \succeq a_x^z(x)$. Thus, it is the case that

$$u(B^{iz'}, x)(a_x^z(x)) = u_i(z') \text{ and } u(B^{iz}, x)(a_x^z(x)) = u(B^{iz}, x)(\bar{x}) \in \{u_i(z'), u_i(z)\}.$$

In the case that $z = z'$, it immediately follows that for any node x such that $\iota(x) = i$, and $z \succ x$, it must be the case that $a(B^{iz'}, x) = a(B^{iz}, x)$, and thus that $f^t(z) \to 1$ as $t \to \infty$, for any n. By hypothesis, for any $x : x_i \preceq x \prec z$, and any $a \in A(x)$, there is a path $\bar{z} \succeq a(x)$ such that $u_{\iota(x)}(\bar{z}) \leq u_{\iota(x)}(z)$. At the same time, for any node $x \prec x_i$, it is the case that $u(B_z^{\iota(x)}, x)(a_x^*(x)) = u_{\iota(x)}(z)$, and that $u(B_z^{\iota(x)}, x)(a(x)) = u^*_{\iota(x)}(a(x))$ for any $a \in A(x) \backslash \{a_x^*\}$. It is thus immediate to see that for any $x \prec z$ and $a \in A(x)$, there is a path $\bar{z} \succeq a(x)$ such that $u_{\iota(x)}(\bar{z}) \leq u_{\iota(x)}(z)$. Thus z is a Nash path of Γ.

Suppose that $z \neq z'$. By construction, for any node $x \in X_i^*$, it is still the case that $a(B^{iz'}, x) = a(B^{iz}, x)$. So suppose that there is a node x, $\iota(x) = i$ and $x \prec x_i$, such that $a(B^{iz'}, x) \neq a(B^{iz}, x)$. Pick any arbitrary such a node, since it must be the case that

$$u(B^{iz'}, x)(a(x)) = u(B^{iz}, x)(a(x)), \quad \text{for any action } a \in A(x) \backslash \{a_x^z\},$$
$$\text{and } u(B^{iz}, x)(a_x^z(x)) = u_i(z) < u_i(z') = u(B^{iz'}, x)(a_x^z(x)),$$

it must be the case that $a(B^{iz'}, x) = a_x^z$, and $a(B^{iz}, x) \neq a_i$. But then, any player of type $B^{iz'}$ observes path z, and thus becomes of type B^{iz}. Again, it must be the case that $f^t(z) \to 1$ as $t \to \infty$, for any n. To see that also in this case z must be a Nash path of Γ, notice that $A_+^i \cap B^{iz} = \emptyset$. This implies that, again, for any $x \prec z : \iota(x) = i$, it must be the case that $u(B^{iz}, x)(a_x^z(x)) = u_i(z)$.

Part 2. *If there exists a deviation (x_i, a_i, z') for some i, and a x-consistent path $z \neq z^*$ for some $\tilde{x} : x_i \prec \tilde{x} \prec z^*$, then for some sequence of forgetfulness probabilities $\varepsilon^n \to 0$, there is a Nash path $z'' \neq z^*(\Gamma)$ such that $f^*(z'') = 1$.*

We shall divide the construction in 2 steps.

Step 1. *The forgetfulness probabilities.*

Introduce the sets $X' = \{x : \tilde{x} \prec x \prec z\}$, $X'' = \{x : x_i \preceq x \prec z'\}$, $X^+ = \{x : x \prec x_i\}$. For any population j, introduce the set $X'_j = \{x : x \in X'$, and $i(x) = j\}$, $X''_j = \{x : x \in X''$, and $i(x) = j\}$ and $X_j^+ = \{x : x \in X^+$, and $i(x) = j\}$.

As in part 1, we enumerate the nodes $x \in X''_i$ from 1 to $\bar{K}$ with order inverse to $\prec$, and introduce the set X_{-i}^k for each k. Now however we also enumerate the nodes $x \in X'_{\iota(\tilde{x})}$ from 1 to $\bar{\kappa}$ with order inverse to $\prec$, and introduce the set $X^{\kappa}_{-\iota(\tilde{x})}$ for each κ. We let $K = \max\{\bar{K}, \bar{\kappa}\}$. Also, we enumerate the nodes $x \in X'' \cup X^+$ from 1 to L with order inverse to $\prec$, and introduce the notation x^l.

For any $x \in X' \cup \{\tilde{x}\}$ the sets $A'(x)$ and $A''(x)$ are as defined in part 1, so are the sets $A'_{z'}(x)$ for any $x \in X''$, the sets $A''_*(\tilde{x})$, and $A''_-(\tilde{x})$, the sets A'_j, the set of types $\mathbf{B}_0^j$, and the types B_z^j for any population j.

We then introduce the set $A_i^0 = \emptyset$, and the sets $A_i^k = \cup_{x \in X_{-i}^k} A'_{z'}(x)$ for any $k = 1, ..., \bar{K}$, as well as the sets $A_0^{\iota(\tilde{x})} = \emptyset$, and the sets $A_\kappa^{\iota(\tilde{x})} = \cup_{x \in X_{-\iota(\tilde{x})}^\kappa} A'(x)$ for any $\kappa = 1, ..., \bar{\kappa}$. We also put $\mathbf{A}_i^0 = A'_i$, $\mathbf{A}_i^k = A'_i \cup A_i^1 \cup ... \cup A_i^k$ for any $k = 1, ..., \bar{K}$, and $\mathbf{A}_0^{\iota(\tilde{x})} = A'_i$, $\mathbf{A}_\kappa^{\iota(\tilde{x})} = A'_i \cup A_1^{\iota(\tilde{x})} \cup ... \cup A_\kappa^{\iota(\tilde{x})}$ for any $\kappa = 1, ..., \bar{\kappa}$.

We introduce the set of types $\mathbf{B}_\kappa^{\iota(\tilde{x})} = \{B^{\iota(\tilde{x})} : \mathbf{A}_\kappa^{\iota(\tilde{x})} \cap B^{\iota(\tilde{x})} = \emptyset$, and $\mathbf{A}_{\kappa+2}^{\iota(\tilde{x})} \subseteq B^{\iota(\tilde{x})}\}$, for any κ, and for any $j \neq \iota(\tilde{x})$, define the set $\mathbf{B}^j = \{B^j : A'_j \cap B^j = \emptyset\}$. We introduce also $\hat{\mathbf{B}}_k^i = \{B^i : \mathbf{A}_k^i \cap B^i = \emptyset$, and $\mathbf{A}_{k+2}^i \subseteq B^i\}$ for any k, $\mathbf{B}^{*j} = \{B^j = A \backslash C$ for some C, $A'_j \subseteq C \subseteq A'_j \cup A'(x_i)\}$ for any $j \neq i$, and $\mathbf{B}^{*i} = \hat{\mathbf{B}}_{\bar{K}}^i$.

For any $j \neq i$, let $B_0^{j**} = A \backslash \left(A'_j \cup A'(x_i)\right)$ and $B_0^{i**} = A \backslash \left(\mathbf{A}_K^i\right)$. Let $B_0^{\iota(\tilde{x})*} = B_0^{\iota(\tilde{x})**} \backslash \mathbf{A}_{\bar{\kappa}}^{\iota(\tilde{x})}$, and for any other j, let $B_0^{j*} = B_0^{j**} \backslash \{a_{\tilde{x}}^*\}$. Also let $z_0 = z^*\left(B_0^*\right)$. For any $l = 1, ..., L$, and any j, let $B_l^{j*} = (B_{l-1}^{j*} \backslash A(x^{l-1})) \cup \mathbf{a}^{z_{l-1}}$ for any $j \neq \iota(x^{l-1})$, and $B_l^{j*} = B_{l-1}^{j*} \cup \mathbf{a}^{z_{l-1}}$, and $z_l = z^*(B_l^*)$, where $B_l^* = (B_l^{j*})_{j \in I}$.

Construct the forgetfulness probability profiles as follows. For any index $n \geq 1$, any population j, and any action $a \in A'_j$, let $\varepsilon_j^n(a) = 1/n$. For any $k = 1, ..., K$, let $m_k = \sum_{j \in I} \#(A'_j) + \sum_{q=0}^{k-1} \#(A_{z'}^q \cup A_q^{\iota(\tilde{x})})$, and set $\varepsilon_i^n(a) = 1/n^{1+m_k}$ for any action $a \in A_i^k$, and $\varepsilon_{\iota(\tilde{x})}^n(a) = 1/n^{1+m_k}$ for any action $a \in A_k^{\iota(\tilde{x})}$. Also set $\varepsilon_{\iota(\tilde{x})}^n(a) = 1/n^{1+m_1}$ for any action $a \in A''_*(\tilde{x})$, $\varepsilon_j^n(a) = 1/n^{1+m_1}$ for any action $a \in A'(\tilde{x})$ and any index $j \neq \iota(\tilde{x})$.

For any $l = 1, ..., L$, let $m'_l = m_K + \sum_{q=0}^{l-1} \#(A(x^q))$, and set $\varepsilon_j^n(a) = 1/n^{1+m'_l}$ if $a \in A(x^q)$ and $j \notin \{i, \iota(x^q)\}$, or if $j = i \neq \iota(x^q)$ and $a \in A(x^q) \backslash \mathbf{A}_i^K$. For any other action a' and population index j, set $\varepsilon_j^n(a') = 0$.

Step 2. *Dynamics*

The proof of Step 2 in part 1 yields that for any $\delta > 0$, there exists a n large enough, and a $T(\delta, n)$ large enough, such that $\rho_j^T(B_z^j) > 1 - \delta$, for any j.

The proof of Step 3 in part 1 yields that for any profile B such that $B^{\iota(\tilde{x})} \in \mathbf{B}_k^{\iota(\tilde{x})}$ for some k, and $B^j \in \mathbf{B}^j$ for any $j \neq \iota(\tilde{x})$, if $\tilde{x} \prec z^*(B)$ then $z^*(B) \in \{z^*, z\}$. Also the proof of Step 4 shows for any n and t large enough $\rho_i^t(\mathbf{B}_{\bar{K}}^{\iota(\tilde{x})}) \approx 1$ and $\rho_j^t(\mathbf{B}^j) \approx 1$.

The proof of Step 3 in part 1 yields that for any profile B such that $B^i \in \hat{\mathbf{B}}_k^i$ for some k, and $B^j \in \mathbf{B}^{*j}$ for any $j \neq i$, if $a_i(x_i) \prec z^*(B)$ then $z^*(B) = z_0$. Also the proof of Step 4 shows that for any $\delta > 0$ and any n, there is a T_0 large enough such that $\rho^{T_0}(B_0^*) > 1 - \delta$.

Since $a_i(x_i) \preceq z^*(B_0^*)$, it follows that $f^{T_0}(z_0) > 1 - \delta$. Since only observed actions can be forgotten, there must exist times $T_L > T_{L-1} > ... > T_1 > T_0$ large enough such that for any l, $\rho^{T_l}(B_l^*) > 1 - \delta$, and that thus $f^{T_l}(z_l) > 1 - \delta$.

For any l, and any x such that $x^l \prec x \prec z_l$, by construction, it must be the case that $a_x(B_l^{j*}, x_0) = a_x^{z'}$ for any player j. Note also that by construction the set $\{l : z_l$ is a Nash path$\}$ is non-empty. It follows that letting $\bar{l} = \min\{l : z_l$ is a Nash path$\}$, then for any $l \geq \bar{l}$, it must be the case that $z_l = z_{\bar{l}}$. If it is the case that $a_{x_i}^*(x_i) \not\preceq z_{\bar{l}}$, we have shown that there is a Nash path z'' such that $f^*(z'') = 1$. The same is the case whenever $x_i \not\prec z_{\bar{l}}$. Thus suppose that $a^*_{x_i}(x_i) \preceq z_{\bar{l}}$. Pick any $x : \tilde{x} \preceq x \prec z$, by construction, for any j and any l, it must be the case that $a(B_l^{j*}, x) = a^z_x$. Thus it follows that $f^*(z) = 1$, and by hypothesis $z \neq z^*$. ■

Proof of Proposition 7. We say that the game Γ' is obtained from Γ by the addition of the *trivial* action $a' \in A'$ if there is $z \in Z$, and $x' \in X'$, such that the trees G and G' satisfy

$$(X, Z \setminus \{z\}, A) = (X' \setminus \{x'\}, Z' \setminus \{a'(x')\}, A' \setminus \{a'\}),$$

that the population indexes satisfy $(I, \iota) = (I', \iota'|_{X' \setminus \{x'\}})$, and the utility functions are such that

$$\mathbf{u}|_{Z\setminus\{z\}} = \mathbf{u}'|_{Z'\setminus\{a'(x')\}} \text{ and } \mathbf{u}(z) = \mathbf{u}'(a'(x')).$$

Similarly, we say that the game Γ' is obtained from Γ by the addition of the *conditionally strictly dominated* action $a' \in A'$, $a' \in A(x)$ if the trees G and G' satisfy

$$(X, Z, A) = (X', Z' \setminus \{a'(x)\}, A' \setminus \{a'\}),$$

the population indexes (I, ι) coincide with (I', ι'), the utility function $\mathbf{u}$ coincides with $\mathbf{u}'|_{Z\setminus\{a'(x)\}}$, and $u'_{\iota'(x)}(a'(x)) < u_{\iota(x)}(a(x))$, for some $a \in A(x)$.

It is straightforward to see that, if there exist a finite sequence $\{\Gamma_k\}_{k=0}^K$ such that for any k, the game Γ_k can be obtained from Γ_{k-1} by the addition of a trivial or conditionally dominated action, and $\Gamma_0 = \Gamma$, $\Gamma_K = \Gamma'$, then the game Γ' is a *BI-irrelevant expansion of* Γ.

Take any game Γ, consider the BI set ordering $\{Y_j\}$. Consider node x such that $\mathbf{u}^*(x) = \mathbf{u}^*(x_0)$ and $x \in Y_1$ (x is the last node on BI path). Relabel a^*_x as N and $\iota(x)$ as population 1. Consider the sequence of actions $\{T, H, C, A, P\}$ and add them sequentially to the game Γ, in the following manner. Introduce $T = (x, z^T)$, $u_1(z_T) > u_1(z^*)$ and $u_2(z_T) > u_2(z^*)$, thus T is strictly dominated addition. Introduce the node $x'_2 : \iota(x'_2) = 2$, relabel T as (x_1, x'_2), and introduce $C = (x'_2, z_C)$, such that $\mathbf{u}(z_C) = \mathbf{u}(z_T)$, so C is a trivial addition. Introduce $H = (x'_2, z_X)$, such that $u_1(z_H) > u_1(z^*)$ and $u_2(z^*) < u_2(z_H) < u_2(z_C)$, so that H is a strictly dominated addition. Introduce the node $x'_1 : \iota(x'_1) = 1$ and relabel C as (x'_2, x'_1) introduce $A = (x'_1, z_A)$, such that $\mathbf{u}(z_A) = \mathbf{u}(z_C)$, thus A is a trivial addition. Finally, introduce $P = (x'_1, z_P)$, such that $u_1(z_P) < u_1(z_A)$, so that P is strictly dominated addition.

It is immediate to show that Γ' is a game with consistent deviations. Assigning $\varepsilon_2(A) > 0$, $\varepsilon_1(C) > 0$ and $\varepsilon_i(a)$ for any other pair (i, a), the claim is proved. ■

Proof of Lemma 5. We shall show that any game with fragile path satisfies the Hypothesis of Theorem 5.

Let $x^0 = x_i$, and $a^0 = a_i$. For any $k \geq 1$, such that $a_{k-1}(x^{k-1}) \in X$, let $x^k = a^k(x^{k-1})$. Let $\underline{\mathbf{a}}^k_{\iota(x^k)}$ be the BI maximin solution for player $\iota(x_k)$ of the x^k-subgame. Set $a^k = A(x^k) \cap \underline{\mathbf{a}}^k_{\iota(x^k)}$. The construction is repeated iteratively until an index $K+1$ is reached such that $a^K(x^K) \in Z$, let $z = a^K(x^K)$.

By hypothesis, $u_i(z) \geq \underline{u}_i(a(x))$, for any $a \in A(x_i)$. For any $k = 1, ..., K$, by definition of BI maximin solution, it must be the case that $\underline{u}_{\iota(x^k)}(a^k(x^k)) \geq \underline{u}_{\iota(x^k)}(a(x^k))$ for any $a \in A(x^k)$, and that $\underline{u}_j(a^k(x^k)) \geq \underline{u}_j(x^k)$ for any j.

It immediately follows that $u_{\iota(x_K)}(z) \geq \underline{u}_{\iota(x_K)}(a(x^K))$ for any $a \in A(x^K)$. For any $m = 1, ..., K-1$, the above implies that

$$u_{\iota(x^{K-m})}(z) \geq \underline{u}_{\iota(x^{K-1})}(x^K) = \underline{u}_{\iota(x^{K-m})}(a^{K-1}(x^{K-1})),$$

and that for any k running from $K-1$ to $K-m$,

$$\underline{u}_{\iota(x^{K-m})}(a^k(x^k)) \geq \underline{u}_{\iota(x^{K-m})}(x^k), \quad \text{and } u_{\iota(x^{K-m})}(x^k) = \underline{u}_{\iota(x^{K-m})}(a^{k-1}(x^{k-1}));$$

it thus follows that

$$u_{\iota(x^{K-m})}(z) \geq \underline{u}_{\iota(x^{K-m})}(a^{K-m}(x^{K-m})) \geq \underline{u}_{\iota(x^{K-m})}(a(x^{K-m})) \text{ for any } \ a \in A(x^{K-m})$$

where the last inequality follows from the definition of BI maximin solution. ■

Proof of Proposition 8. The necessity part of the first claim follows from Lemma 3.

In order to prove sufficiency and the second claim, by Lemma 4 it is enough to show that if Γ has a x'-*consistent* path $z \neq z^*$ for some node $x' \prec z^*$, then Γ has a Nash path $z'' \neq z^*$ such that $f^*(z'') = 1$ for some sequence $\{\boldsymbol{\varepsilon}^n\}$ such that $\boldsymbol{\varepsilon}^n \to 0$.

We introduce the set $X' = \{x : x' \prec x \prec z\}$ and $X^* = \{x : x' \preceq x \prec z\}$. For any j, introduce the set $X'_j = \{x \in X' : \iota(x) = j\}$, and $X^*_j = \{x \in X^* : \iota(x) = j\}$. Enumerate the nodes $x_j \in X'_j$ with order inverse to $\prec$ with index k_j running from 1 to K_j, let $x^0_j = z$ and $x^{K_j+1}_j = x'$, and introduce $X^{k_j}_j = \{x \in X^* \backslash X^*_j : x^{k_j+1}_j \prec x \prec x^{k_j}_j\}$ for any $k_j = 0, ..., K_j$. For any node $x \in X'$, the sets $A'(x)$ and $A''(x)$ are as defined in the step 1 of part 1 in the proof of Theorem 5. For any population j, let $A'_j = \cup_{x \in X'_j} A''(x)$ and $B^j_z = A \backslash A'_j$; also introduce $A^{k_j}_j = \cup_{x \in X^{k_j}} A'(x)$, $\mathbf{B}^j_{k_j} = \left\{ A \backslash C \text{ for some } C : A'_j \cup \left(\cup_{l=0}^{k_j-1} A^l_j \right) \subseteq A'_j \cup \left(\cup_{l=0}^{k_j} A^l_j \right) \right\}$, and $B^{*j} = A \backslash \left(A'_j \cup \left(\cup_{l=0}^{K_j} A^l_j \right) \right)$. Construct the forgetfulness probability profiles as follows. For any index $n \geq 1$, any population j, and any action $a \in A'_j$, let $\varepsilon^n_j(a) = 1/n$; also any $k_j = 0, 1, ..., K_j$, let $m_{k_j} = \sum_{i \in I} \#(A'_i) + \sum_{l=0}^{k_j-1} \#(A^l_j)$, and set $\varepsilon^n_j(a) = 1/n^{1+m_{k_j}}$ for any action $a \in A^{k_j}_j$. For any other action a' and population index i, set $\varepsilon^n_i(a') = 0$.

The proof of Step 2 in part 1 of the proof of Theorem 5 yields that for any $\delta > 0$, there exists a n large enough, and a $T(\delta, n)$ large enough, such that $\rho_j^T(B_z^j) > 1-\delta$, for any j. The proof of Step 3 in part 1 of the proof of Theorem 5 yields that for any profile B such that for any j, $B^j \in \mathbf{B}_{k_j}^j$ for some k_j, if $x' \prec z^*(B)$ then $z^*(B) \in \{z^*, z\}$. Assumption 3 then implies that for any $\delta > 0$ and any j, there must exist times $T^* > T_{K_j} > T_{K_j - 1} > ... > T_1 > T_0$ large enough such that for any k_j, $\rho^{T_{k_j}}(\mathbf{B}_{k_j}^j) > 1 - \delta$, and that for any $t \geq T^*$, $\rho_j^t(\mathbf{B}^{*j}) > 1 - \delta$. Since it is the case that if $x' \prec z^*(B^*)$, then $z^*(B^*) = z$, it follows that there is a Nash path $z'' \neq z^*$ such that $f^*(z'') = 1$. ■

Proof of Proposition 9. Necessity follows from Step 1 in the proof of Proposition 2.

In order to prove sufficiency, introduce the set $X' = \{x : x' \preceq x \prec z\}$. For any node $x \in X'$, the sets $A'(x)$ and $A''(x)$ are as defined in part 1 of the proof of Theorem 5. For any j, introduce the set $X'_j = \{x \in X' : \iota(x) = j\}$. Enumerate the nodes in X'_i with order inverse to $\prec$ with index k running from 1 to K, set $x_i^0 = z$, and for any $k = 1, ..., K$, we introduce the set $X_i^k = \{x : x \in X' \backslash X'_i,\ i(x) \neq i \text{ and } x_i^k \prec x \prec x_i^{k-1}\}$. For any population $j \neq i$, let $A'_j = \cup_{x \in X'_j} A'(x)$ and $B_z^j = A \backslash A'_j$, also let $B_z^i = A$. Introduce $A_i^k = \cup_{x \in X_i^k} A'(x)$, $\mathbf{B}_k^i = \left\{ A \backslash C \text{ for some } C : \cup_{l=0}^{k-1} A_i^l \subseteq \cup_{l=0}^{k} A_i^l \right\}$, and $B^{*i} = A \backslash \left(\cup_{l=0}^{K} A_i^l \right)$. Construct the forgetfulness probability profiles as follows. For any index $n \geq 1$, any population j, and any action $a \in A'_j$, let $\varepsilon_j^n(a) = 1/n$; also for any $k = 0, 1, ..., K$, let $m_k = \sum_{j \in I} \#(A'_j) + \sum_{l=0}^{k-1} \#(A_i^l)$, and set $\varepsilon_i^n(a) = 1/n^{1+m_k}$ for any action $a \in A_i^k$. For any other action a' and population index i, set $\varepsilon_i^n(a') = 0$.

The proof of Step 2 in part 1 of the proof of Theorem 5 yields that for any $\delta > 0$, there exists a n large enough, and a $T(\delta, n)$ large enough, such that $\rho_j^T(B_z^j) > 1 - \delta$. The proof of Step 3 in part 1 of the proof of Theorem 5 yields that for any profile (B_z^{-i}, B^i) such that $B^i \in \mathbf{B}_k^i$ for some k, if $x' \prec z^*(B)$ then $z^*(B) \in \{z^*, z\}$. It follows that for any $\delta > 0$, there must exist times $T^* > T_{K_j} > T_{K_j - 1} > ... > T_1 > T_0$ large enough such that for any k, $\rho_i^{T_k}(\mathbf{B}_k^i) > 1 - \delta$, and that for any $t \geq T^*$, $\rho_i^t(\mathbf{B}^{*i}) > 1 - \delta$. Since it is the case that if $x' \prec z^*(B_z^i, B^{i*})$, then $z^*(B_z^i, B^{i*}) = z$, it follows that there is a persistent path $z'' \neq z^*$ such that $f^*(z'') = 1$. ■

References

[1] Alos-Ferrer C. [1999]: "Dynamical Systems with a Continuum of Randomly Matched Agents", *Journal of Economic Theory* **86**: 245-267

[2] Aumann R. [1995] "Backward Induction and Common Knowledge of Rationality" *Econometrica* 6-19

[3] Balkenborg D. [1994]: "Strictness and Evolutionary Stability" mimeo *Hebrew University of Jerusalem*

[4] Balkenborg D. [1995]: "Evolutionary Stability, Strictness and Repeated Games with Common Interest" mimeo *University of Bonn*

[5] Balkenborg D. and K. Schlag [2001]: "A Note on the Evolutionary Selection of Nash Equilibrium Components" mimeo *University of Exeter*

[6] Ben-Porath E. [1997]: "Rationality, Nash Equilibrium and Backwards Induction in Perfect-Information Games" *Review of Economic Studies* **64:** 23-46

[7] Boylan R. [1992]: "Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals" *Journal of Economic Theory* **57**: 473-504

[8] Borgers T. [1994]: "Weak Dominance and Approximate Common Knowledge" *Journal of Economic Theory* **64:** 265-76

[9] Cressman, R.; Schlag, K. H. [1998]: "The Dynamic (In)Stability of Backwards Induction" *Journal of Economic Theory* **83:** 260-85

[10] Dawes R. [1988]: *Rational Choice in an Uncertain World* Harcourt New York

[11] Diestel R. [1997]: *Graph Theory* Springer, New York

[12] Dudley R. M. [1989]: *Real Analysis and Probability* Chapman and Hall, New York

[13] Durrett R. [1996]: *Probability Theory and Examples* Belmont CA, Wadsworth Pu.

[14] Fudenberg D. and D.K. Levine [1993a]: "Self-Confirming Equilibrium" *Econometrica* **61**: 523-545

[15] Fudenberg D. and D.K. Levine [1993b]: "Steady State Learning and Nash Equilibrium" *Econometrica* **61**: 547-573

[16] Fudenberg D. and D.K. Levine [1998]: *The Theory of Learning in Games,* MIT Press

[17] Fudenberg D. and Tirole J. [1991]: *Game Theory* Cambridge MA, MIT press

[18] Gale J. K. Binmore and L. Samuelson [1995]: "Learning to Be Imperfect: The Ultimatum Game" *Games and Economic Behavior* **8:** 56-90

[19] Hale J. [1969]: *Ordinary Differential Equations*, Wiley, NY

[20] Hart S. [2000]: "Evolutionary Dynamics and Backward Induction", mimeo, Jerusalem University

[21] Hendon E, H. Jacobsen and B. Sloth [1996]: "Fictitious Play in Extensive Form Games" *Games and Economic Behavior* **15:** 177-202

[22] Jehiel P. [2001]: "Analogy-Based Expectation Equilibrium" mimeo University College of London

[23] Jehiel P. and D. Samet [2001]: "Learning to Play Games in Extensive Form by Valuation" mimeo *University College of London*

[24] Kalai E. and E. Lehrer [1993]: "Rational Learning Leads to Nash Equilibrium" *Econometrica* **61:** 1019-45

[25] Kohlberg E. and J.F. Mertens [1986]: "On the Strategic Stability of Equilibria" *Econometrica* **54**: 1003-1038

[26] Kreps D. [1990] "Corporate Culture and Economic Theory", in *Perspectives on Positive Political Economy*, J. Alt and K. Shepsle eds. Cambridge University Press

[27] Marx L. and Matthews S. [2000]: "Dynamic Voluntary Contribution to a Public Project" *Review of Economic Studies* **67:** 327-58

[28] Nachbar J. [1997] "Prediction, Optimization, and Learning in Repeated Games" *Econometrica* **77:** 275-309

[29] Noldeke G. and L. Samuelson (1993): "An Evolutionary Analysis of Backward and Forward Induction" *Games and Economic Behavior* **5:** 425-54

[30] Reny [1993]: "Common Beliefs and the Theory of Games with Perfect Information" *Journal of Economic Theory* 257-274

[31] Ritzberger K. and J. Weibull [1996]: "Evolutionary Selection in Normal Form Games" *Econometrica* **63:** 1371-1399

[32] Rosenthal R. [1982]: "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox" *Journal of Economic Theory* **25:** 83-96

[33] Roth A. and I. Erev [1995]: "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term" *Games and Economic Behavior* **8:** 164-212

[34] Rubinstein A. [1982]: "Perfect Equilibrium in a Bargaining Model" *Econometrica* **50:** 97-110

[35] Schlag K. [1998]: "Why Imitate, and If So, How?" *Journal of Economic Theory* **78**: 130-156

[36] Samuelson L. [1997]: *Evolutionary Games and Equilibrium Selection* MIT Press

[37] Samuelson [2000]: "Analogies, Adaptation, and Anomalies" *Journal of Economic Theory* forthcoming

[38] Selten R. [1975]: "A Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games" *International Journal of Game Theory* **4**: 25-55

[39] Selten R. [1983]: "Evolutionary Stability in Extensive Two-Person Games" *Mathematical Social Sciences* **5:** 269-363

[40] Stahl D. [1993]: "The Evolution of Smart-n Players" *Games and Economic Behavior* **10:** 218-54

[41] Van Damme [1989]: "Stable Equilibrium and Forward Induction" *Journal of Economic Theory* **48:** 476-509

[42] Weibull J. W. [1995]: *Evolutionary Game Theory* Cambridge MA, MIT press