Advances in Theoretical Economics

Volume 4, *Issue* 1 2004 *Article* 2

Backward Induction and Model Deterioration

Francesco Squintani*

*University College London, f.squintani@ucl.ac.uk

Backward Induction and Model Deterioration*

Francesco Squintani

Abstract

The issue of how players' model of a game may evolves over time is largely unexplored. We formalize this issue for games with perfect information, and show that small-probability model deterioration may upset the complete-model backward induction solution, possibly yielding a Pareto-improving long run distribution of play. We derive necessary and sufficient conditions for the robustness of backward induction. These conditions can be interpreted with a forward-induction logic, and are shown to be closely related to the requirements for asymptotic stability of the backward induction path under standard evolutionary dynamics.

KEYWORDS: Evolutionary Game Theory, Backward Induction, Subgame Perfect Equilibrium

*Department of Economics, Gower Street, London WC1E 6BT, United Kingdom. I am grateful for their comments to the seminar participants at 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 Drew Fudenberg, Ehud Kalai, Georg Noldeke, Larry Samuelson, Juuso Valimaki, and especially Eddie Dekel. The usual caveat applies.

1 Introduction

Unlike formal games, most social interactions are not accompanied by a complete list of written, fixed rules describing all actions that can be taken. While many actions may be relevant, some are more salient, and others can be more easily overlooked, depending on the context of the interaction. As a result, the most difficult task faced by the players is often to formulate a model of the interaction, a list of all actions relevant in the game. Once this modeling step is accomplished, solving the model may be relatively easy. The players' model may change over time, and depend both on their past experiences and on the context of the interactions they are involved in. The issue of how the players' model of a 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 continuous populations live for two periods. In the second period, they are randomly matched to play a perfect-information extensive-form game. In the first period, they observe their parents' play. Matching is anonymous and independent across generations. After being matched, each player formulates a model of the game, identified by a subset of the action space and based on the play observed in her first life period as well as on the information transmitted by her parent. Players do not know or learn each others' models while playing the game. A player whose parent is unaware of a given action, will also be unaware of such action, unless she observes it in her parent's match. At each decision node the player plays the action corresponding to the unique backward induction solution of her (possibly incomplete) model.

Information transmission across generations is imperfect: with small probability, models may deteriorate across generations, so as to exclude some feasible actions. To account for the fact that some actions can be more easily overlooked than others, we allow for different probabilities of forgetting different actions. In order to highlight the effect of 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 aggregate long run distribution of play, obtained by compounding the long run model distribution with the play induced in each match. The core of our analysis restricts the possibility of model deterioration to opponents' actions not played in the parents' match, building on the supposition that one is usually less likely to forget one's own possible choices or recently observed actions. We extend the analysis in an extension section to allow for players to forget also their own actions, as long as they were not played in their parents' matches. We also discuss preliminary results of a variation of our model where players tremble when making their choices.

First, we show that while all players initially play the backward induction path, and in each period all players observe all actions on path, the backward induction path may be upset by small-probability model deterioration, and the resulting long run distribution of play may be Pareto-superior. As long as the off-path actions that support the complete-game backward induction path can be more easily forgotten than the actions that upset it, model deterioration generates more and more players who deviate from the backward induction path. In some games these deviations prevent the players from regaining awareness of the actions that their predecessors forgot. As a result, while model deterioration occurs with small probability, the fraction of players who do not play the backward induction solution increases over time, and eventually overcomes the population. When this is the case, we say that the complete-model backward induction path is *upset* by model deterioration.

We characterize games where the backward induction path may be upset. It is necessary that the backward induction path admits "profitable deviations" (i.e. by deviating from the backward induction path, a player enters a subgame with a terminal node that makes her better off). This condition is also sufficient when players can forget their own actions. When they can only forget opponents' actions, our sufficient condition further requires that there be a non-subgame perfect Nash equilibrium in any subgame originating on the backward induction path at the deviation node or after such a node. This sufficient condition shows that the backward induction path may be upset by model deterioration in any game that is complex enough. However, if the backward induction path satisfies a strong forward-induction requirement, in the spirit of Van Damme (1989), then it cannot be upset by model deterioration.[1]

This paper is related in scope 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

[1] Among evolutionary stability concepts that display forward-induction properties, Swinkels (1992) proposes the concept of equilibrium evolutionary stable sets, and shows that they are robust to iterated removal of weakly dominated strategies and satisfy the Never Weak Best Response Property (see Kohlberg and Mertens 1986). The forward-induction property displayed by our solution concept is closer to the property analyzed in Balkenborg (1994).

fitness of genes subject to natural selection forces. This is equivalent to studying the learning process of non-strategic players. The literature on rational learning (see for example Kalai and Leher 1993, Fudenberg and Levine 1993a and 1993b, Nachbar 1997) focuses on how players learn opponents' strategies. We assume that players know how to fully solve game-theoretical models, and hence play an equilibrium given their model of the game. We keep track of how their models change over time. Our results imply that there are games that players may never fully learn.

In comparing our framework and analysis with evolutionary game-theory contributions, we should first point out that our learning dynamics are not payoff-monotonic. In evolutionary game theory, natural selection is assumed to favor strategies yielding the highest payoff. One rationale for such learning dynamics is that people imitate those players who achieve the highest payoff (see Schlag 1998 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 the opponents with whom she is matched. Consistent with that view, this paper focuses on the relation between the players' models and their past observation of play.

Despite this major difference, our characterization of games where the backward induction path is robust with respect to model deterioration unexpectedly turns out to be in close logical relation to standard evolutionary stability analyses. The backward induction solution is not necessarily selected by Lyapounov stability under the replicator dynamics;[2] Balkenborg and Schlag (2001) fully characterize sets that are asymptotically stable with respect to any trajectory in the mixed strategy profile space, under any evolutionary dynamics that satisfies mild requirements.[3] In the language of this paper, they show that the backward induction Nash component of a game is asymptotically stable if and only if the game does not admit any profitable deviation.[4] This

[2] Hart (2000), Hendon, Jacobsen and Sloth (1996) and Jehiel and Samet (2001) instead, present different learning models that favor the backward induction solution in the long run.

[3] Specifically, they require that the dynamics be regular and 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. Such a class of dynamics is very large, as it includes all regular payoff-positive and payoff monotonic-dynamics, and hence the replicator dynamics.

[4] Studying sets that are asymptotically stable with respect to trajectories starting in the *interior* of the mixed strategy profile space, instead, Cressman and Schlag (1998) show that the Nash Equilibrium component associated with the backward induction path is the unique minimal interior asymptotically stable set in 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.

characterization is logically equivalent to our characterization of games where model deterioration may upset the backward induction path, for the case when players can forget their own actions. When players can only forget opponents' actions, we show that the backward induction path of any game without profitable deviations cannot be upset by model deterioration, but we also show games with profitable deviations where the backward induction path cannot be upset.

In the stochastic learning model by Noldeke and Samuelson (1993), each player is endowed with a "characteristic" consisting of a strategy and a conjecture on the opponents' choices. In each period, she may reconcile her conjectures with the opponents' population play off path and choose a best-response. But with small probability she may also randomly mutate her characteristic. Our sufficient condition for the robustness of the backward induction path with respect to model deterioration is logically equivalent to their sufficient condition (Proposition 7) for the backward-induction solution to be the unique locally stable component. Their necessary condition for the backward-induction solution to be a locally stable outcome (Proposition 4) is that by deviating from the backward induction path, any player cannot enter a subgame where there is a (non-subgame perfect) Nash equilibrium that makes her better off. Our necessary conditions for the backward-induction path to be robust to model deterioration are tighter. When players can forget their own actions, it is necessary that any player cannot enter a subgame where there is an outcome that makes them better off. When players can only forget opponents' actions, it is necessary that, by deviating from the backward-induction path, any player cannot enter a subgame where there is a non-subgame perfect Nash equilibrium, together with a possibly different outcome that makes the player better off with respect to the backward-induction solution.[5]

The paper is presented as follows. The second section describes a simple example leading to the subsequent analysis. The third section formally presents our dynamic framework, and the fourth section contains our characterization results. The fifth section discusses a few possible extensions. The sixth section concludes, and it is followed by the Appendix, which lays out the proofs.

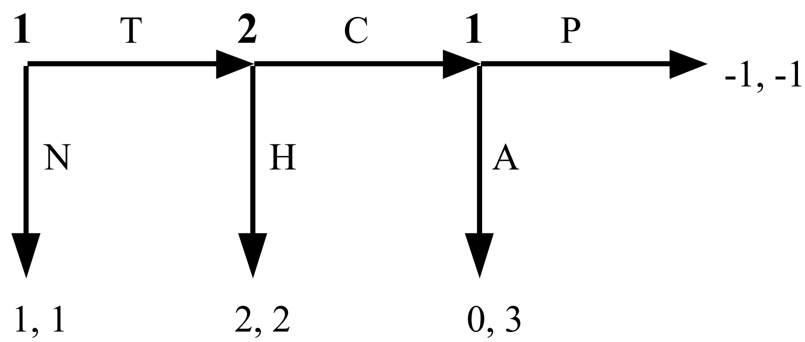

Figure 1: Trust and Punishment Game

2 A Simple Example

Example 1 Two populations of players of size 1 are randomly matched to play a version of the trust game (depicted in figure 1) which includes the possibility of costly punishment.[6] 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 subgame-perfect equilibrium is (NA, C) with backward induction path N. The game has another Nash equilibrium component, which induces the path TH, and Pareto-dominates the backward induction path N.

We assume that with probability ε, players from population 2 can forget the possibility that action A is played, and players from population 1 can forget action C, if they have not observed them in their parents' match.[7] To simplify the calculations, we assume that each player's model always include

[5]Unlike our characterization, their characterization is limited to the domain of games where each player moves only once along each path. Hence the comparison is staged in that domain only.

[6]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.

[7]If a player in population 1 is unaware of C, the possibility of being cheated, when choosing whether to trust (T) her opponent or not (N) at her first decision node, then she also does not consider the choice between punishing the opponent (P) or acquiescing (A) in the case that she is cheated. This choice is meaningful only if the opponent plays C, a possibility of which she is unaware. This does not mean that she is unaware of her own actions P and A. If her opponent subsequently cheats her, the player will quickly realize that she may respond by playing either P or A at her second decision node. This issue is further clarified in the exposition of our general model at the beginning of section 3.2.

all other actions.

We first present our results informally. In the initial match, all players hold a complete model of the game. Players from population 1 play the backward induction path N, and thus both A and C may be forgotten. Population-2 players unaware of A play H, because the backward induction solution of their incomplete model of the game is THP. Their play prevents their offsprings from observing A. Because the offspring of a parent unaware of A will also be unaware of it, unless she observes A, the fraction of players unaware of A increases over time. Since TH is the solution of the game where the continuation of action C is deleted, population-1 players unaware of C play T. Their opponents play C only if they are aware of A.[8] Since the fraction of these players decreases over time, the fraction of offsprings 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, so that the Nash equilibrium path TH is established.

Formally, the populations dynamics are described by letting a_t be the time-t proportion of population-2 players unaware of A, and c_t be the proportion of population-1 players unaware of C:

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

In the first equation, the proportion a_{t+1} of players unaware of A at time $t+1$ is determined as follows. The fraction a_t of parents unaware of A play H hence preventing their offspring from regaining awareness of A. A fraction $1 - c_t$ of the $1 - a_t$ parents aware of A is matched with fully aware opponents who play N: their offspring forget A with probability ε. The remaining offspring of fully-aware parents maintain a complete model. In the second equation, a fraction $1 - c_t$ of parents are fully aware and play N: their offsprings do not observe C and forget it with probability ε. Of all the c_t parents unaware of C, a fraction a_t is matched with opponents unaware of A who play H. All their offsprings remain unaware of C. The remaining $(1 - a_t)$ offspring regain awareness of C and become fully aware.

[8] Our model assumes that each fully-aware player in population 2 plays C after her opponent plays T. In the context of this example, this can happen because the player makes sense of the opponent's choice T by concluding that the opponent is unaware of C. In our general model (see section 3.2) we assume that when a player find herself at a decision node that contradicts her model's backward-induction solution, she expands her model by conjecturing that the opponent is aware of actions that she is not aware of. We assume that this model revision does not change her choice. This assumption is reasonable because the revision does not generate any new information about the payoff consequences of her choices.

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 $1 - c_t \geq 1 - a$, and hence 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. ◇

We conclude this section by underlying a key feature of the above example, and of subsequent analysis: *the possibility that the backward-induction path is upset by model deterioration crucially depends on the relative likelihood that different actions are forgotten by the players.* In the above example, we study the polar case where players only forget the actions A and C that belong to the backward-induction solution. It is immediate to see that in the polar opposite case where players forget actions P and H instead of A and C, the long run distribution of play coincides with the backward induction path. Because players unaware of P (respectively, of H) play the backward-induction solution strategies C and respectively NA, the actions P and H are not observed on path, and in the long run all players play the backward-induction solution and are unaware of the opponents' possible deviations. Formally, letting p_t be the time-t proportion of population-2 players unaware of P, and h_t be the proportion of population-1 players unaware of H, we obtain

$$\begin{cases} p_{t+1} = p_t + (1-p_t)\varepsilon, \\ h_{t+1} = h_t + (1-h_t)\varepsilon. \end{cases}$$

The system converges to the steady state state ($p^* = 1$, $h^* = 1$) and the backward-induction path N is established.

The relative likelihood that different actions are forgotten by the players does not depend on the characteristics of the game form, but on the contextual meaning of different actions in a social interaction. In its general game-theoretical analysis, this paper is agnostic with respect to the issue of how the relative likelihood of forgetting actions should be formulated in specific social interactions.[9] It is easy to show that in any game form, it is always possible to

[9]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

assign model-deterioration probabilities in such a manner that the backward induction path is not upset. This paper tackles the more interesting question of identifying game forms where model deterioration may upset the backward induction path for at least some forgetfulness probabilities across actions.

3 The Learning Model

3.1 Preliminaries

There is a finite set I of populations of players. In each population i, there is a continuum of players of size 1. At each period t, all players are randomly matched to play the finite perfect-information "generic" game $\Gamma = (X, Z, A, I, \iota, \mathbf{u})$ without any chance move. Each match is anonymous and it includes one player from each population; matches are independent over time. 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. For any node x, $A(x)$ is the set of arcs exiting x. For any arc $a \in A(x)$, we denote as $a(x)$ the node successor of x reached through arc a. We denote by $\prec$ the transitive closure of A, by $x \prec y$ we mean that x precedes y in the game tree, and by $x \preceq y$ that either $x \prec y$ or that x coincides with y. The path $\mathbf{a}(x, y)$ from node x to node y on the tree $G = (Y, A)$ is the (unique) set of actions $\{a_0, \cdots, a_n\}$ such that $a_0 = (x, y_1)$, $a_1 = (y_1, y_2)$, ..., $a_n = (y_n, y)$, for some set $\{y_1, y_2, ..., y_n\}$. Each terminal node z uniquely identifies a path from the initial node x_0, and 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 in the standard fashion, and partitions the sets X and A into $\{X_i\}_{i \in I}$ and $\{A_i\}_{i \in I}$. The function $\mathbf{u} : Z \to \mathbb{R}^I$ represents the players' payoffs, and each player in the same population has the same utility function. We assume that there are no ties in payoffs, $u_i(z) \neq u_i(z')$ for any distinct paths z, z' and any i; the set of games satisfying this property is generic in the set of finite perfect-information games. To avoid trivialities, we focus on games where for any x, $\#A(x) > 1$ and for any $a \in A(x)$, $\iota(a(x)) \neq \iota(x)$. 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 consists of the game $\Gamma_x = (X_x, Z_x, A|_{Y_x}, I, \iota|_{X_x}, \mathbf{u}|_{Z_x})$.

that the possibility to punish unfair behavior is very salient and unlikely to be dismissed. Second, dishonest and 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.

For any game Γ, any profile $\mathbf{a} \in \times_{x\in X} A(x)$ is a pure-strategy profile (note that $\mathbf{a}$ identifies a proper subset of A), the strategy associated to any index i is $\mathbf{a}_i = (a_x)_{x\in X_i}$ which identifies a proper subset of A_i, and the opponents' strategy profile is $\mathbf{a}_{-i} = (a_x)_{x\in X\backslash X_i}$, which identifies a proper subset of A_{-i}. The path (or outcome) induced by $\mathbf{a}$ is the unique terminal node z such that $a_0 = (x_0, y_1)$, $a_1 = (y_1, y_2)$, ..., $a_n = (y_n, z)$ for some set $\{y_1, y_2, ..., y_n\}$ and some set of actions $\{a_0, \cdots, a_n\} \subset \mathbf{a}$. The definition of outcomes and payoffs are extended as customary when introducing the behavioral strategies $\boldsymbol{\alpha}_i \in \times_{x\in X_i} \Delta(A(x))$, for any index i, and the behavioral strategy profiles $\boldsymbol{\alpha}_{-i} \in \times_{x\in X\backslash X_i} \Delta(A(x))$ and $\boldsymbol{\alpha} \in \times_{x\in X} \Delta(A(x))$. Where for any finite set Q, the notation $\Delta(Q)$ denotes the set of distributions over Q.

The *backward induction solution* $\mathbf{a}^* \in \times_{x\in X} A(x)$ is defined as follows. Let $Y_0 = Z$, and for any $j \geq 1$, recursively define

$$Y_j = \{x \in Y \backslash (\cup_{k=0}^{j-1} Y_k) : \text{for all } a \in A(x), \quad a(x) \in (\cup_{k=0}^{j-1} Y_k)\}.$$

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 associated *backward induction path* is denoted by z^*.

3.2 The Individual Game Models

In any period of play and any match, 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. As the play develops in the match, it may be that some players play actions not included in a player's initial model. Whenever this is the case, this player may find herself at 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 assume 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 $x \in X$. Furthermore we assume that players are aware of all their own actions.

Definition 1 *For each population i, the set of admissible frameworks is* $\mathcal{B}^i = \{B^i \subseteq A : A_i \subset B^i$ *and for any* $x \in X$, $B^i \cap A(x) \neq \varnothing\}$.

For future 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.

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 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 i and B^i. At any of her own decision nodes x, the* model *of any player in population i with framework B^i is denoted by (Y_x^i, B_x^i), and consists of the largest 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 subgame Γ_x according to the unique backward induction solution. Specifically, given any model (Y_x^i, B_x^i), we let $\mathbf{a}^*(B^i, x)$ be the backward induction solution of the game $\Gamma(B^i, x) = \left((Y_x^i, B_x^i), I, \iota|_{X_x^i}, \mathbf{u}|_{Z_x^i}\right)$, where $X_x^i = Y_x^i \cap X$ and $Z_x^i = Y_x^i \cap Z$, and we let $\mathbf{u}^*(B^i, x)$ be the associated backward induction values.

Assumption 1 *Take any i and B^i. Any player in population i endowed with framework B^i plays the (unique) backward induction action $a_x^*(B^i, x)$ at any of her decision nodes x.*

It is important to stress that when players make their choices, they have no knowledge of the models of the other players in their matches. This assumption is crucial for our analysis. If players knew everybody's model in their match, then they would share the same model of the game. Because each player is fully aware of her actions, all her opponents would also know her actions, and all the players' models would be complete. While players may still forget actions in the lapse of time between their parents' match and the moment they are called to play, such forgetfulness would have no bite. All players would regain awareness of forgotten actions by observing their matched opponents' models.

We also underline that each player plays the backward induction solution of her (possibly incomplete) model even if she finds her at a decision node that contradicts her model's backward-induction solution. This is crucial for our characterization results. We shall now argue that it is reasonable in this context. Suppose that the opponents' choices leading to a node x are inconsistent with the solution of a player's incomplete model at x. To make sense of these opponents' choices, the player may either conjecture that the opponents do not know some actions that she is aware of, or revise her model by conjecturing that they know some opponents' actions of which she is not aware.[10] We assume that she does the latter: she expands her model at x until the backward-induction solution of the expanded model is consistent with the play reaching node x.[11] This assumption is not innocuous; if the player were to draw any inferences on the opponents' models, she may want to change her backward-induction solution.

As is shown below, any incomplete model at x can be expanded in such a way that the backward-induction solution of the expanded model is consistent with the play reaching x, and still it yields the same choice at x as the original model. Indeed, one key observation is that making sense of the opponents' choices before x places no restrictions on –and hence provides no information about– the player's own payoffs induced by the actions that are added to the model. Because the player cannot acquire any additional information about the payoff consequences of her choices at x through her model revision process, it is reasonable to assume that after expanding her model, she makes the same choice at x that she would make with her original model.

Formally, the incomplete model of the subgame Γ_x at node x, identified by the game $\Gamma(B^i, x)$, may be expanded into the *rationalized model* $\hat{\Gamma}(B^i, x) = ((\hat{Y}^i_x, \hat{B}^i_x), I, \hat{\imath}, \hat{\mathbf{u}})$ of the whole game Γ, as follows. Let $(\tilde{Y}^i_x, \tilde{B}^i_x)$ be the largest tree contained in $(Y_x, B^i \cup \mathbf{a}(x_0, x))$, and $\mathbf{a}^*(\hat{\Gamma}(B^i, x))$ the backward induction solution of game $\hat{\Gamma}(B^i, x)$. Construct $\hat{\Gamma}(B^i, x)$ so that $(\tilde{Y}^i_x, \tilde{B}^i_x) \subseteq (\hat{Y}^i_x, \hat{B}^i_x)$, $\hat{\imath}|_{\tilde{X}^i_x} = i|_{\tilde{X}^i_x}$, $\hat{\mathbf{u}}|_{\tilde{Z}^i_x} = \mathbf{u}|_{\tilde{Z}^i_x}$, $\mathbf{a}(x_0, x) \subseteq \mathbf{a}^*(\hat{\Gamma}(B^i, x))$ and $\mathbf{a}^*(B^i, x) \subseteq \mathbf{a}^*(\hat{\Gamma}(B^i, x))$. It is immediate to see that a rationalized model $\hat{\Gamma}(B^i, x)$ exists for any B^i and

[10] The issue is related to the literature on common certainty of rationality. There, the game is common knowledge and, by finding herself off path, a player may only infer that rationality is not commonly certain. Whether or not she will play according to the backward induction solution depends on such inferences and on her higher order beliefs (see, for example, Battigalli 1996, and Battigalli and Siniscalchi 2002). This paper assumes that that it is common knowledge that players are rational, but that they may have different models of the game.

[11] Alternatively, our assumption can be motivated in an extension of our model with trembles. In such a setting, when a player observes opponents' actions that are inconsistent with her solution of the game, she can rationalize them them as opponents' mistakes. Full fledged analysis of a model with trembles is beyond the scope of this paper, but we discuss this model further in Section 5.

x. By construction, it yields the same backward induction solution as $\Gamma(B^i, x)$ at x, and the node x is reached on its backward-induction path.

3.3 Dynamics

The framework distribution in the populations at any time t is identified by ρ^t, where for each i, $\rho_i^t \in \Delta(\mathcal{B}^i)$. In order to describe its dynamic transition between time t and time $t+1$, we need to specify the framework transition of each player in each population. Each player's framework always include all actions observed on path in her parent's match. Of those actions that were not observed, the player is unaware of all the actions that her parent is unaware of. In addition, the player may randomly forget some unobserved actions, so that her model randomly deteriorates.

For simplicity, we assume that forgetfulness occurs with probabilities fixed over time and across players of the same population. For any population i and framework B^i we introduce the probability distribution $\pi_i(\cdot|B^i)$ with support $\{\hat{B}^i \in \mathcal{B}^i : \hat{B}^i \subseteq B^i\}$. The offspring of each player with framework B^i may forget actions according to the probability distribution $\pi_i(\cdot|B^i)$. Because all frameworks B^i in the collection $\mathcal{B}^i$ include all actions of the players in population i and at least one action at each node, players cannot forget their own actions, and cannot forget all actions at any node. Restricting the support to sets $\hat{B}^i \subseteq B^i$ represents the assumption that players cannot become aware of actions their parents are unaware of, unless they observe them on path.

Because for any population i and framework B^i, the choice $a_x^*(B^i, x)$ is unique at each decision node x, the path observed by each player in her parent's match is uniquely pinned down by the frameworks of the players in the match. Hence, for any matching time t, a player's stochastic framework transition between t and $t+1$ is only a function of her parent's framework at time t, and of the frameworks of the players in her parent's match. This leads us to the following assumption which also incorporates the restrictions that observed actions cannot be forgotten. While possibly deteriorating according to the probability system $\boldsymbol{\pi}$, each player's model always includes the actions observed in her parent's match.

Assumption 2 *Take any time t and any match where the players' frameworks are B. For any i, the offspring of the player in population i is endowed with framework $\hat{B}^i \in \mathcal{B}^i$ at time $t+1$ with probability*

$$\sum_{\tilde{B}^i \,:\, \tilde{B}^i \subseteq B^i,\ \hat{B}^i = \tilde{B}^i \cup \mathbf{a}(B)} \pi_i(\tilde{B}^i|B^i).$$

This assumption, together with a standard "Law of Large Number" argument (see Alos-Ferrer, 1999), allows us to derive the framework distribution transition function $\xi(\boldsymbol{\pi}) : \times_{i\in I}\Delta(\mathcal{B}^i) \to \times_{i\in I}\Delta(\mathcal{B}^i)$, such that

$$\xi_i(\boldsymbol{\pi})(\hat{B}^i) = \sum_{(B,\tilde{B}^i)\,:\,\tilde{B}^i\subseteq B^i,\,\hat{B}^i=\tilde{B}^i\cup\mathbf{a}(B)} \rho(B)\,\pi_i\left(\tilde{B}^i|B^i\right), \qquad \text{for any } i \text{ and } \hat{B}^i.$$

Our construction determines a family $\{\xi(\boldsymbol{\pi})\}$ of dynamic systems parametrized in the profile of forgetfulness probability systems $\boldsymbol{\pi}$. Fixing an initial state ρ^0, each system $\xi(\boldsymbol{\pi})$ determines a solution $\{\rho^t(\boldsymbol{\pi},\rho^0)\}_{t\geq 0}$, and thus we have obtained a family of solutions $\{\{\rho^t(\boldsymbol{\pi},\rho^0)\}_{t\geq 0}\}$ parametrized in $\boldsymbol{\pi}$.

For any type distribution ρ, we let $f\in\Delta(Z)$ be the aggregate play induced by ρ:

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

This paper investigates whether complete game models (and hence the induced backward induction path) can be upset by small-probability model deterioration.[12] Formally, we study the families of solutions $\{\{\rho^t(\boldsymbol{\pi},\rho^0)\}_{t\geq 0}\}$ where $\rho_i^0(A) = 1$ for all i, for $\boldsymbol{\pi}$ close to the family $\boldsymbol{\pi}^* = (\pi_i^*)_{i\in I}$ such that

$$\pi_i^*\left(B^i|B^i\right) = 1 \text{ for all } B^i \in \mathcal{B}^i.$$

As already pointed out, the limit solution may depend on the relative probabilities of forgetting different actions in the limit. Therefore our results are given for specific sequences $\{\boldsymbol{\pi}^n\}_{n\geq 0}$ of model-deterioration probability profiles. We say that the backward induction path is fragile, if it differs from the long run play for some sequence $\{\boldsymbol{\pi}^n\}_{n\geq 0}$ approaching $\boldsymbol{\pi}^*$ in the limit.

Definition 3 *The backward induction path z^* may be* upset *by model deterioration if there is a sequence of model-deterioration profiles $\boldsymbol{\pi}^n \to \boldsymbol{\pi}^*$ such that*

$$\lim_{n\to\infty}\inf \lim_{t\to\infty}\inf f\left(\rho^t(\boldsymbol{\pi}^n,\rho^0)\right)(z^*) < 1.$$

[12] It would be inappropriate to study the persistence properties of outcomes allowing for initially incomplete models. Whenever assuming that players are initially unaware of some relevant actions, and do not have a chance to observe and learn them, the analysis would be more appropriate to describe the effect of model deterioration on a smaller game where these actions are simply impossible. A fortiori, it would be inappropriate to study the stability of the backward induction path in terms of (Lyapounov or asymptotically) stable states without explicitly considering the initial model distribution.

4 The Characterization

This section characterizes necessary and sufficient conditions for the backward induction path of a game Γ to be upset by small-probability model deterioration. Our necessary condition builds on the following simple observations. Players cannot forget observed actions and the backward induction path is initially played by all players. Because players cannot forget actions on the backward induction path, they may decide to deviate only in the hope of improving their payoff. A player will never 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. To give a formal account of this intuition, we introduce the concept of *profitable deviation.*

Definition 4 *Let* (x, z) *be called a* (profitable) deviation *whenever* $x \prec z^*$, $x \prec z$, $a^*_x(x) \not\preceq z$ *and* $u_{\iota(x)}(z) > u^*_{\iota(x)}(x)$.[13]

Theorem 1 *For any game* Γ, *the backward induction path* z^* *may be upset by small-probability model deterioration only if* Γ *admits a profitable deviation* (x, z).

The above result implies that if the backward induction path satisfies a forward-induction requirement similar in spirit to (but stronger than) the one introduced by Van Damme (1989), then it cannot be upset by model deterioration. Van Damme (1989) defines a subset S of the Nash equilibrium set *consistent with forward induction* in 2-person generic games if there is no equilibrium in S such that some player i can take a deviation leading with certainty to a subgame in which she is better off under one equilibrium of the subgame and worse off under all the others. We modify this forward induction requirement by stipulating that a subset S of the Nash equilibrium set is *consistent with strong forward induction* if there is no equilibrium in S such that some player i can take a deviation leading with certainty to a subgame in which she is better off given any strategy profile of the subgame. Evidently, this requirement is stronger than the one introduced by Van Damme (1989). Theorem 1 states the backward induction path cannot be upset by model deterioration if the associated Nash equilibrium component is consistent with strong forward induction.

[13]The concept of profitable deviation is closely related to the concept of strict outcome path introduced by Balkenborg (1995). It is immediate to show that 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.

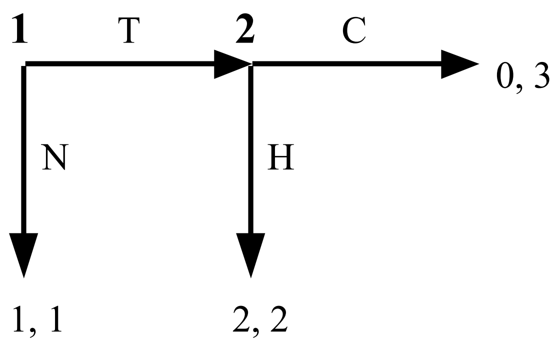

Figure 2: Trust Game

While profitable deviations are necessary for the backward-induction path to be upset by small-probability model deterioration, the next example shows that they are not sufficient. When a forgetful player deviates from the backward induction path in the beliefs of improving her payoffs, the opponent plays the forgotten backward-induction solution actions and reminds her that she would be better off staying on the backward induction path. Thus, the fraction of forgetful players cannot accumulate and will remain negligible over time.

Example 2 The trust game represented in Figure 2 has a unique Nash Equilibrium (N, C), with path N. If players in population 1 forget C, then they deviate and play T, but regardless of her model, each player in population 2 plays C if her decisional node is reached. As a result the unaware player 1 is immediately reminded of action C and she switches back to N. Formally, we let c_t (respectively h_t) denote the time-t fraction of players that are unaware of C (respectively H), and obtain:

$$\begin{cases} c_{t+1} = (1 - c_t - h_t)\pi_C \\ h_{t+1} = h_t + (1 - c_t - h_t)\pi_H, \end{cases}$$

where we use the shorthand notations π_C and π_H, to denote the probability that a player from population 1 forgets action C and H respectively. For any choice of the parameters $\pi_C^1 \geq 0$ and $\pi_H^1 \geq 0$, in the long run, the system settles on the state (c^*, h^*) such that $c^* \leq \pi_C^1/(1 + \pi_C^1)$. Since π_C^1 is negligible, the long run distribution of play is indistinguishable from the backward induction path. $\diamond$

The key feature of Example 2 that makes the backward-induction path robust to model deterioration is that, once any forgetful player deviates from the backward induction path, she is immediately reminded of the forgotten actions, so that unawareness cannot accumulate over time. This stands in sharp contrast to our motivating Example 1, where the backward induction path N is upset because players in population 1 who are unaware of C deviate from N in the hope of reaching path TH, and players in population 2 who forget action A *indeed play* H, so that their forgetful opponents remain unaware of C. As a result, the proportion of matches inducing the path TH increases over time and the backward induction path N is upset by small-probability model deterioration.

The salient features of the path TH are that (i) it induces a profitable deviation, (ii) *it is induced by the non-subgame perfect Nash equilibrium* (TP, H). Hence, the forgetful players' models of the game are consistent with the opponents' play, and they do not regain awareness of the actions they have forgotten. Proposition 2 shows that model deterioration may upset the backward-induction path in all games with a path z that has the same characteristics as TH. We call a Nash path of a game Γ any path that is associated with a Nash equilibrium component of Γ.[14]

Proposition 2 *If a game* Γ *has a deviation* (x, z) *such that* z *is a Nash path of the subgame* Γ_x*, then the backward induction path of* Γ *may be upset by small-probability model deterioration.*

For an intuitive explanation, let a be the action at x such that $a(x) \preceq z$, and consider any node x' that lies on the path between x and z. By construction x' is off the backward induction path z^*, and all actions in the subgame $\Gamma_{x'}$ can be forgotten, as long as players play a_x^* at x. Let a' be the action at x' such that $a'(x') \preceq z$, and consider any alternative action a''. Once a player forgets all opponents' actions in the subgame $\Gamma_{a''(x')}$ that do not belong to the Nash equilibrium supporting z, and all opponents' actions that do not lead the play into z at any nodes between x' and z, she will choose a' whenever the node x' is reached. This precludes the player from ever recollecting such forgotten actions, as they are off the path established by action a'. Hence the frequency of these forgetful types cannot decrease in time, and it is strictly increasing as long as node x' is not reached with probability one. Since the node x' is arbitrary, this shows that the path z is established in the long run, if the action a is played at node x.

[14] In generic games of perfect information, any Nash equilibrium component identifies a unique (degenerate) path of play z, see Kreps and Wilson (1982), and Kohlberg and Mertens (1986).

Suppose that a player at node x forgets all the opponents actions in $\Gamma_{a(x)}$ that do not lead the play into z. Since (x, z) is a profitable deviation, she chooses a at x. In the long run, these forgetful types most likely observe actions on the path from x to z, and will not be reminded of forgotten actions. As a result, these forgetful types eventually overcome the population. Since they deviate the play from z^* to z whenever they are called to play at node x, the backward induction path z^* is upset.

While the argument is less transparent, the characterization of Proposition 2 can be extended also for the case in which the profitable deviation (x', z') that lures players off path does not induce a Nash path in $\Gamma_{x'}$, as long as there is a Nash path z different from z^* in any subgame Γ_x such that x either coincides or comes after x' on the backward induction path. The result is driven by the observation that, when players deviate from the backward induction solution at node x', they make it possible to forget all actions on the backward induction path that starts at x. Once players forget all opponents' action that do not induce the Nash path z in the subgame Γ_x they will believe that z is reached in that subgame. Hence they either do not enter the subgame or they establish the path z. In either case, they deviate from the backward-induction path z^*.

Theorem 3 *If a game Γ has a deviation (x', z'), and for some $x : x' \preceq x \prec z^*$ the subgame Γ_x has a Nash path $z \neq z^*$, then the backward induction path of Γ may be upset by small-probability model deterioration.*

Obviously, the sufficient condition of Theorem 3 is stronger than the necessary condition of Theorem 1. While the latter only requires that a game admits a profitable deviation (x', z'), Theorem 3 further requires that there is a non-subgame perfect Nash equilibrium in a subgame Γ_x where x comes after x' on the backward induction path, or coincides with x'. This naturally begs the question of whether the backward induction path is upset by small-probability model deterioration when there are not any non-subgame perfect Nash equilibrium in any such subgames, but still the game Γ has non-subgame perfect Nash equilibria. As it turns out, the backward induction path may or may not be upset depending on further characteristics of the game. The issue is best illustrated by the following two examples.

Example 3 Consider the game presented in figure 3. The backward induction solution is (BE, C) with path BC. The Nash equilibrium (AF, D) induces the Nash path A, there are no further Nash paths. The players in population 1 have no profitable deviations from the backward induction path, whereas the players in population 2 may deviate and play D if they expect their opponents to play

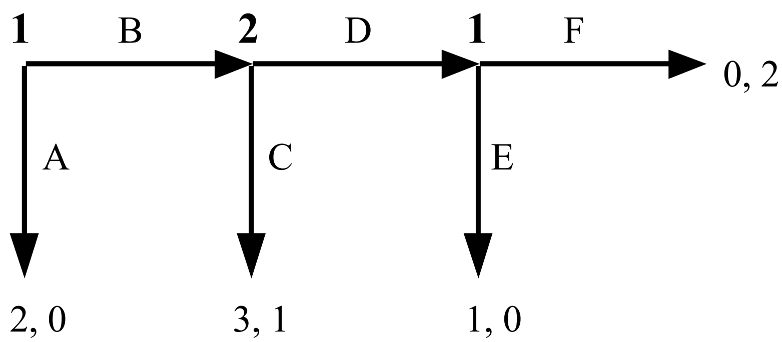

Figure 3: The Backward induction path is persistently upset.

F. The profile (D, F) is not a Nash equilibrium in the subgame starting with the choice of players in population 2. Despite this, the backward induction path may be upset. Suppose that actions C and E can be forgotten with probability $\varepsilon > 0$, and that no other action can be forgotten.

Initially, players in population 2 play C; their offspring do not observe E, and may forget it. Players unaware of E play D, their opponents' offsprings do not observe C and may forget it. In such a case, they play A, thus preventing their offsprings and their opponents' forgetful offsprings from recollecting forgotten actions. Hence the fraction of players playing A grows over time and eventually takes over the population.

Formally, let c_t be the time-t proportion of players unaware of C, and e_t be the time-t proportion of players unaware of E:

$$\begin{cases} c_{t+1} = c_t + (1 - c_t)e_t\varepsilon, \\ e_{t+1} = e_t c_t + (1 - e_t)\varepsilon. \end{cases}$$

This system asymptotically approaches the state $(c^* = 1,\ e^* = 1)$, the long run path of play is A.

Example 4 The backward induction solution of the game in figure 4 is (B, E, F) with path BE. The players in population 1 have no profitable deviations, and those in population 2 may only profitably deviate by playing D in expectation that their opponents play G. The profile (D, G) is not a Nash equilibrium in the subgame starting with the choice of players in population 2, but the Nash equilibria (A, C, F) and (A, C, G) induce the path A.

The backward induction path of this game cannot be upset by model deterioration. Intuitively, each type from population 2 chooses to play either D or E. Hence any player from population 1 either observes D or E on path. Since she does not expect her opponent to ever play C, she plays B. If a player from population 2 is unaware of F and plays D, her offspring observes F. As in Example 2, the long run path of play is indistinguishable from the backward induction path BE.

Formally, we suppose that action C cannot be forgotten, so as to give the best fighting chance to the Nash path A. Without loss of generality, we further simplify the analysis by assuming that the probability of forgetting either action E or action D is independent of the players' model. Letting d_t (respectively e_t, f_t, g_t) be the time-t fraction of players unaware of D (respectively E, F, G), we obtain:

$$\begin{cases} d_{t+1} = (1 - f_t)(\pi_D + \pi_{DE}); & e_{t+1} = f_t(\pi_E + \pi_{DE}); \\ f_{t+1} = (1 - f_t - g_t)\pi_F; & g_{t+1} = g_t + (1 - f_t - g_t)\pi_G, \end{cases}$$

where we use the shorthand notations π_F and π_G to denote the probability that a player from population 2 forgets action F or G respectively, and the notations π_D, π_E and π_{DE} to denote the probability that a player from population 1 forgets either action D or E only, or both actions D and E at the same time.

For any forgetfulness profile π, in the long run, the system settles on a steady state (e^*, f^*) where $f^* \leq \pi_F/(1 + \pi_F)$, and $e^* \leq (\pi_E + \pi_{DE})\pi_F/(1 + \pi_F)$. Since π_E, π_{DE} and π_F are all negligible, the long run distribution of play approximates the backward induction path BE.

This example is extreme because C is conditionally dominated by both E and D: if a player in population 2 is called to play, she prefers D and E over C regardless of her opponents' choice (see Shimoji and Watson (1998), page 169, for the formal definition of conditional dominance). Nevertheless this game is suggestive of why it may be that the backward induction path is not persistently upset in games outside the class identified by Theorem 3. Moreover, it is not difficult to expand this game, so that C is not dominated by either D or E, and so that the same result still obtains. ◇

We conclude this section by studying a simple version of centipede game (see Rosenthal, 1982), a well known perfect-information game that is not covered in our characterization. Despite the often questioned plausibility of its backward-induction solution, we find that model deterioration cannot upset the backward-induction path.

Example 5 Figure 5 represents a 3-leg version of the centipede game: the backward-induction solution (AE, C) yields path A even though the players may achieve a higher payoff by proceeding in the game until the last decision node. We will show that the backward-induction path cannot be upset. To

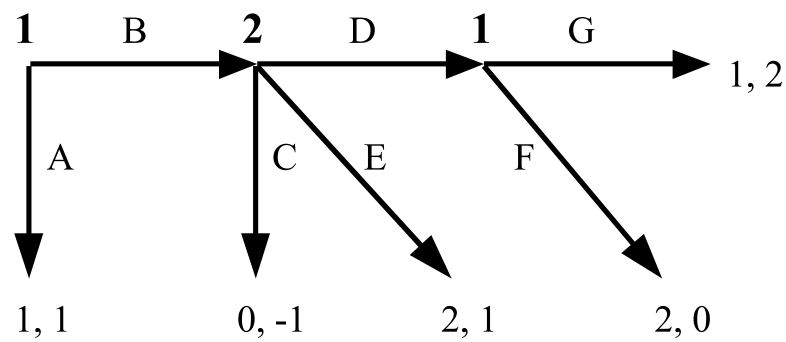

Figure 4: The Backward induction path is not persistently upset.

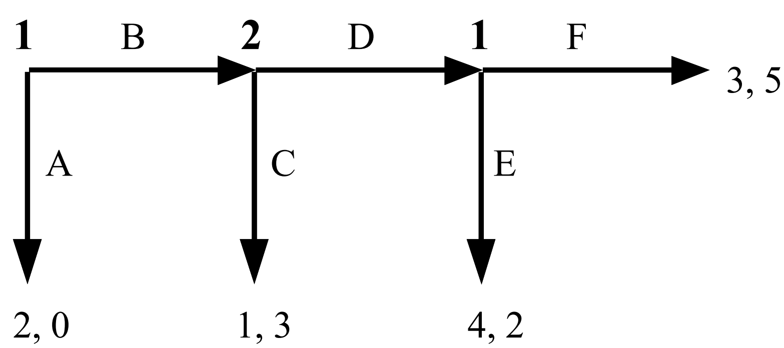

Figure 5: A simple version of centipede game.

give model deterioration the best fighting chance in upsetting the backward-induction path, we assume that actions D and F cannot be forgotten.

Players in population 2 may forget E because their fully-aware parents play C. If their decision node is reached, such forgetful players play D. Their offspring are then reminded of E, because all players in population 1 play E, regardless of their model of the game. Players in population 1 may forget C because their fully-aware parents play A. Such forgetful players play B: if matched with fully-aware opponents, their offsprings observe C and become fully aware. In sum, it cannot be the case that the path of play differs from A and that at the same time model deterioration accumulates in the populations. Hence the backward-induction path cannot be upset by small model deterioration.

Formally, letting c_t be the time-t proportion of players unaware of C, and e_t be the time-t proportion of players unaware of E, we obtain:

$$\begin{cases} c_{t+1} = c_t e_t + (1 - c_t)\pi_C, \\ e_{t+1} = e_t(1 - c_t) + (1 - e_t)\pi_E; \end{cases}$$

where π_C denotes forgetfulness of action C and π_E forgetfulness of E. For any forgetfulness profile π, in the long run, the system settles on a steady state (c^*, e^*) where $c^* \leq \frac{1}{2(\pi_C+1)}\left[\pi_C(1-\pi_E) - \sqrt{4\pi_E\pi_C + \pi_C^2 + 2\pi_E\pi_C^2 + \pi_E^2\pi_C^2}\right]$. Since π_C is negligible, the long run distribution of play approximates the backward induction path A. ◇

5 Extensions

Our construction and results can be extended in several directions. We shall now discuss some of the most interesting ones.

Players' forgetfulness of their own actions. In order to allow players to forget also their own actions within our dynamic construction, it is just sufficient to lift the restriction that admissible framework always contain all players' own actions in Definition 1. For each population i, we let the set of admissible frameworks be $\mathcal{B}^i = \{B^i \subseteq A : \text{for any } x \in X,\ B^i \cap A(x) \neq \varnothing\}$, and again consider any forgetfulness systems $\boldsymbol{\pi}$ with support $\{\hat{B}^i \in \mathcal{B}^i : \hat{B}^i \subseteq B^i\}$ for any probability distribution $\pi_i(\cdot|B^i)$, any B^i and any i. We maintain the assumption that a player cannot forget any action observed in her parent's match.

When players can forget their own actions, we can provide a *complete* characterization of the games where the backward-induction path may be upset by small-probability model deterioration. The condition identified in Theorem 1 is necessary and sufficient: the backward-induction path may be upset if and only if the game admits profitable deviations. This characterization is logically equivalent to the characterization by Balkenborg and Schlag (2001) of games where the backward induction Nash component of a game is asymptotically stable in a general class of evolutionary dynamics.

Theorem 4 *Suppose that players can forget all actions not observed in their parents' matches, including their own. For any game Γ, the backward induction path z^* may be upset by small-probability model deterioration if and only if Γ admits a profitable deviation (x, z).*

Because actions observed in the parents' match cannot be forgotten, the backward-induction path z^* will not be upset by small forgetfulness if all the terminal nodes that follow a player's deviation make her worse off with respect to the backward-induction path. Suppose that there is a profitable deviation (x, z) for players in an arbitrary population i. Say that the forgetfulness profile $\boldsymbol{\pi}$ is such that all players that do not belong to population i forget all actions in subgame Γ_x (including their own) that do not lead to z, much faster than the players in population i. These actions are initially off path, because players in population i play the backward induction action a_x^* at x. Hence unawareness of these actions accumulates over time. Eventually the players in population i forget these actions too, and deviate at node x. Because the opponents have become largely unaware, the path reaches the terminal node z. Because z makes them better off with respect to the original backward-induction outcome z^*, they have no reason to switch back.

A variation of our model with trembles. While the above variation of our model is easily solvable, introducing trembles in the players' decisions leads to cumbersome calculations. A model with trembles could be formulated as follows. Given a game Γ, an arbitrary tremble profile is $\boldsymbol{\delta} = (\delta_a)_{a \in A}$, where $\delta_a \in (0, 1)$ for any action a. In order to keep the focus of the analysis on model deterioration, it is worth assuming that trembles are smaller than model deterioration probabilities. Formally, this is done by restricting attention to perturbation sequences $\{\boldsymbol{\pi}^n, \boldsymbol{\delta}^n\}_{n \geq 1}$ such that $\delta_a^n < \pi_i^n(\hat{B}^i | B^i)$ for any $a \in A$, for any i, $B^i \in \mathcal{B}^i$, and any $\hat{B}^i \in \mathcal{B}^i : \hat{B}^i \subseteq B^i$, i.e. any $\hat{B}^i$ in the support of $\pi_i^n(\cdot | B^i)$. Any player in population i endowed with framework B^i plays at each of her decision nodes x the (unique) backward induction action $a_x^*(B^i, x)$ with probability $1 - \sum_{a \in A(x) \setminus \{a_x^*(B^i, x)\}} \delta_a^n$, and any other action $a \in A(x)$ with probability δ_a^n. The remainder of the construction is unchanged.

Our preliminary analysis, available upon request, reconsiders the introductory Example 1, and derives the dynamic system that describes the evolution of framework distributions over time. While we do not derive analytical conclusions, our numerical simulations strongly suggest that along some sequences $\{\boldsymbol{\pi}^n, \boldsymbol{\delta}^n\}_{n\geq 1}$, the system converges in the long run to a fixed point such that all players in population 2 are unaware of A and all players in population 2 are unaware of C. Hence the path TH is established in the long run, and backward-induction path N is upset by model deterioration.

Intuitively, small trembles counteract the accumulation of model incompleteness in the populations because they give a chance to players in population 2 unaware of A to play C instead of H. Because their opponents play A if receiving C, these players' offspring will regain awareness of A and become fully aware. This in principle can undo the forces of model deterioration, because fully-aware players in population 2 play C, hence reminding unaware players in population 1 that they should play the backward-induction path N instead of T. But if the trembles and the probabilities that players forget the backward induction actions C and A are sufficiently small relative to the probabilities that they forget H and P, the effect of trembles is small relative to the accumulation of unawareness of A. In the long run, almost all players in population 2 are unaware of A.[15]

Technical generalizations. A possible technical extension of this paper would be a general analysis of our dynamic construction. The assumption that all players initially hold complete models is made here to highlight our results on the robustness of complete model solutions. 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. Because this general analysis is quite complex, we postpone it to further research.

Another possible extension would be a dynamic construction that considers general extensive-form games. The difficulty with this extension lies in the

[15] While we cannot provide formal proofs, we do not see any reason why this logic should not extend to the constructions in the proof of our Proposition 2 and Theorem 3, which identify sufficient conditions for model deterioration to upset the backward-induction path. Instead we have derived a formal argument (available upon request) of a result that extends Theorem 1. Also in the presence of trembles, the backward induction path of any game without profitable deviations cannot be upset by small-probability model deterioration.

fact that, unlike perfect-information generic games, general games may possess multiple solutions, so that in our dynamic construction the evolution of models depends on past play. Thus, the dynamic construction would not be well specified unless a unique solution is selected for each profile of models at each period of play.

6 Conclusion

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 and sufficient condition for the backward induction path to be upset have been presented. When players can only forget opponents' actions, this characterization shows that model deterioration may upset the complete-model backward induction path in a smaller class of games than the class of games where the backward induction Nash component is asymptotically stable under standard evolutionary dynamics. Still, it also shows that model deterioration may upset the backward induction path in all games that are complex enough.

Appendix: Omitted Proofs

Proof of Theorem 1. Pick an arbitrary time $t \geq 1$, and suppose that $f^{t-1}(z^*) = 1$: the "complete-model" BI path z^* is established in all matches at time $t-1$. We will show that this implies that z^* is played in any match B such that $B^i \in Supp(\rho_i^t)$ for all i, and that hence $f^t(z^*) = 1$. This result proves our claim by induction over t because the initial condition $\rho^0(A) = 1$ implies that $f^0(z^*) = 1$.

We order the set of nodes on the BI path $X^* = \{x : x \prec z^*\}$ by assigning for each set Y_j the index $j - \min\{l : X^* \cap Y_l \neq \varnothing\} + 1$ to the node $x \in X^* \cap Y_j$.

Consider x_1, the last decision node on the BI path z^*. By construction, each player in population $\iota(x_1)$ is aware of her own actions: $A(x_1) \subset B^{\iota(x_1)}$ for any $B^{\iota(x_1)} \in \mathcal{B}^{\iota(x_1)}$. Because x_1 is the last node on the BI path, $a^*_{x_1}(x_1) = z^*$, and thus $u^*_{\iota(x_1)}(B^{\iota(x_1)}, x_1)(a^*_{x_1}(x_1)) = u_{\iota(x_1)}(z^*)$. Because game Γ has no profitable deviations, it must be the case that $u_{\iota(x_1)}(z') < u_{\iota(x_1)}(z^*)$ for all terminal nodes z' such that $a(x_1) \preceq z'$ for some $a \in A(x_1)$, $a \neq a^*_{x_1}$. It follows that $u^*_{\iota(x_1)}(B^{\iota(x_1)}, x_1)(a^*_{x_1}(x_1)) = u^*_{\iota(x_1)}(x_1) > u^*_{\iota(x_1)}(B^{\iota(x_1)}, x_1)(a(x_1))$ for any $a \in A(x_1)$, $a \neq a^*_{x_1}$; and hence the BI choice at node x_1 of any player in population $\iota(x_1)$ with any possible framework $B^{\iota(x_1)} \in \mathcal{B}^{\iota(x_1)}$ is $a^*_{x_1}(B^{\iota(x_1)}, x_1) = a^*_{x_1}$, the complete-model BI action at node x_1.

For any arbitrary node x_K, with $K \geq 2$, and framework $B^{\iota(x_K)} \in Supp(\rho^t_{\iota(x_K)})$, consider the BI solution $\mathbf{a}^*(B^{\iota(x_K)}, x_K)$. Pick any node x_k on the BI path z^* with index $k : 1 \leq k < K$. Because $f^{t-1}(z^*) = 1$, i.e. the "complete-model" BI path z^* is observed in all matches at time $t-1$, it must be the case that $a^*_{x_k} \in B^{\iota(x_K)}$. Because game Γ has no profitable deviations, it must be the case that $u_{\iota(x_k)}(z') < u_{\iota(x_k)}(z^*)$ for all terminal nodes z' such that $a(x_k) \preceq z'$ for some $a \in A(x_k)$, $a \neq a^*_{x_k}$. Proceeding by backward induction along the nodes x_k with $k : 1 \leq k < K$, the BI solution $\mathbf{a}^*(B^{\iota(x_K)}, x_K)$ must be such that $u^*_{\iota(x_k)}(B^{\iota(x_K)}, x_K)(x_k) = u_{\iota(x_k)}(z^*)$ and hence $a^*_{x_k}(B^{\iota(x_K)}, x_K) = a^*_{x_k}$ for all k.

Because the game Γ has no profitable deviations, and every player in population $\iota(x_K)$ is aware of her own actions, i.e. $A(x_K) \subset B^{\iota(x_K)}$, it must be the case that $a^*_{x_K}(B^{\iota(x_K)}, x_K) = a^*_{x_K}$: every player with framework $B^{\iota(x_K)} \in Supp(\rho^t_{\iota(x_K)})$ plays the "complete-model" BI action $a^*_{x_K}$ at node x^*_K. Because K and $B^{\iota(x_K)} \in Supp(\rho^t_{\iota(x_K)})$ are arbitrary, this concludes that the BI path z^* is established in all matches B such that $B^i \in Supp(\rho^t_i)$ for any $i \in I$, and hence that $f^t(z^*) = 1$. ■

The proof of the remaining results is simplified by the following Lemma, its proof is available upon request.

Lemma 1 *Any Nash equilibrium component of a perfect-information game Γ without chance moves and with no ties in the payoffs contains a pure-strategy Nash equilibrium* $\mathbf{a}$.

Proof of Proposition 2. Let $\iota(x)$ be i to simplify notation. Let $\mathbf{a}'$ be a pure-strategy Nash equilibrium of subgame Γ_x that induces the path z. Let a be the Nash equilibrium action at node x (i.e. $a \in A(x) \cap \mathbf{a}$) and $\mathbf{a}$ be the restriction of $\mathbf{a}'$ on the subgame $\Gamma_{a(x)}$.

Step 1. *Construction of the forgetfulness probability profiles* $\{\boldsymbol{\pi}^n\}_{n\geq 1}$.

For any player j, let $\bar{B}^j = A \backslash (A|_{Y_{a(x)}})_{-j} \cup \mathbf{a}_{-j}$. Each player j of type $\bar{B}^j$ is unaware of any opponent's action that does not support the Nash equilibrium path z in subgame $\Gamma_{a(x)}$. Clearly, all actions on the path z belong to $\bar{B}^j$. Moreover, it is also the case that $z^*(\bar{B}) = z$, the path z is established if all players in the match are of type $\bar{B}^j$ for all population j. This is because (i) the profile $\mathbf{a}$ is a Nash equilibrium of the subgame $\Gamma_{a(x)}$ that induces z, and (ii) any player in population i who believes that the terminal node z is reached if playing a will deviate from the BI action a_x^* to play a at node x, since $u_i(z) > u_i(z^*)$ by definition of deviation (x, z).

For any index $n \geq 1$, and any framework $B^i \in \mathcal{B}^i$, $\bar{B}^i \subsetneq B^i$, we let $\pi_i\left(\bar{B}^i | B^i\right) = 1/n^I$ and $\pi_i\left(B^i | B^i\right) = 1 - 1/n^I$; for any other type B^i, we let $\pi_i\left(B^i | B^i\right) = 1$. For any population $j \neq i$ and any framework $B^j \in \mathcal{B}^j$, $\bar{B}^j \subsetneq B^j$, we let $\pi_j\left(\bar{B}^j | B^j\right) = 1/n$ and $\pi_j\left(B^j | B^j\right) = 1 - 1/n$, for any other type B^j, we let $\pi_j\left(B^j | B^j\right) = 1$. Intuitively, we let players in population i forget the actions in $(A|_{Y_{a(x)}})_{-i} \backslash \mathbf{a}_{-i}$ infinitely slower than the players in population $j \neq i$ forget the actions in $(A|_{Y_{a(x)}})_{-j} \backslash \mathbf{a}_{-j}$. This will give the time to players in population $j \neq i$ to accumulate unawareness in the subgame $\Gamma_{a(x)}$. When the players in population i deviate from the BI path at node x_i and play a they will most likely be matched with unaware opponents so that the path $z = z^*\left(\bar{B}\right)$ is established.

Step 2. *For any $\delta_1 > 0$, there exists a n large enough, and a $T_1(\delta_1, n)$ large enough, such that $\rho_j^{T_1}(\bar{B}^j) > 1 - \delta_1$ for any $j \neq i$.*

Any player in population $j \neq i$ in a match B, where $a_x^*(B^i, x) \neq a$ will not observe any action in $(A|_{Y_{a(x)}})_{-j} \backslash \mathbf{a}_{-j}$. It follows that, for any t, as long as

$$\sum_{B^i \in \mathcal{B}^i \,:\, a_x^*(B^i, x) \neq a} \rho_i(B^i) \leq o(1/n^{I-1}),$$

the type transition in population j is approximated by the equation

$$\rho_j^{t+1}(\bar{B}^j) = \rho_j^t(\bar{B}^j) + \sum_{B^j \,:\, \bar{B}^j \subsetneq B^j} \rho_j^t(B^j)/n + o(1/n^{I-1}). \tag{2}$$

Given the forgetfulness probability system π_j and because $\rho_j^0(A) = 1$, for any t, $\rho_j^t(B^j) = 0$ unless $\bar{B}^j \subseteq B^j$. Because $\rho_i^0(A) = 1$, $a_x^*(A, x) = a_x^* \neq a$ and

$$\rho_i^{t+1}(A) = \rho_j^t(A) + o(1/n^{I-1}),$$

the recursive application of Equation (2) implies the result.

Step 3. *For any $\delta_2 > 0$, there exists a n large enough, and a $T_2(\delta_2, n) > T_1(\delta_1, n)$ large enough, such that $\rho_i^t(\bar{B}^i) > 1 - \delta_2$* for any $t \geq T_2$.

Pick any $t > T_1(\delta_1, n)$, because $\rho_i^0(A) = 1$, for any t, $\rho_i^t(B^i) = 0$ unless $\bar{B}^i \subseteq B^i$.

For any player in population i of any type B^i such that $\rho_i^t(B^i) > 0$, any action $a' \in A(x)$, $a' \neq a$, and any node $x' \succeq a'(x)$, it is the case that $A(x') \subset B^i$ and hence that

$$u_x(B^i, x)(a_x^*) = u_i^*(a_x^*) \geq u_i^*(a') = u_i(B^i, x_i)(a').$$

Thus any player of type B^i such that $\rho_i^t(B^i) > 0$ either plays a_i or a_x^* at node x. If she plays a_x^*, then all actions that do not belong to $\bar{B}^i$ are off-path by construction. If she plays a, with probability $(1 - \delta_1)$ she is matched with the profile of types $\bar{B}^j$. In such a case, as we have previously shown, $z^* \left(\bar{B}\right) = z$ and all actions that do not belong to $\bar{B}^i$ are off path by construction. Because δ_1 can be taken arbitrarily small, the type transition in population j is approximated by the equation

$$\rho_i^{t+1}(\bar{B}^i) = \rho_i^t(\bar{B}^i) + \sum_{B^i \,:\, \bar{B}^i \subsetneq B^i} \rho_i^t(B^i)/n^I.$$

The result is again implied by the recursive application of this Equation.

The proof is then concluded because $z^* \left(\bar{B}\right) = z \neq z^*$, and so $f^t\left(z^*\right) < \delta_2$, for any $t \geq T_2$, the BI path z^* is upset by small probability model deterioration. ■

Proof of Theorem 3. The proof extends the construction of the proof of Proposition 2 in two Lemmata.

Lemma 2 *If Γ has a deviation (x, z'), and the subgame Γ_x has a Nash path $z \neq z^*$, then z^* may be upset by small probability model deterioration.*

Proof. As in the proof of Proposition 2, let $\iota\left(x\right) = i$. The case where $u_i(z) > u_i\left(z^*\right)$ has already been covered in Proposition 2.

Let **a** be a pure-strategy Nash equilibrium of subgame Γ_x that induces the path z, and a' be the action in $A\left(x\right)$ such that $a'\left(x\right) \prec z'$. Let x^+ be the node $x \preceq x^+ \prec z$ and $x \prec x^+ \prec z'$, where there are distinct actions in $A\left(x^+\right)$, called a_+ and a'_+, such that $a'_+\left(x^+\right) \prec z'$, $a_+\left(x^+\right) \prec z$.

For index i, let

$$A_i^- = \left\{a \in A_{-i} \cap A\left(x'\right) : x^+ \preceq x' \prec z', a\left(x'\right) \npreceq z'\right\}$$

be the set of opponents actions that deviate from the path z', and

$$A_i^+ = \left\{a \in A_{-i} \cap A\left(x'\right) : x^+ \preceq x' \prec z', a\left(x'\right) \preceq z'\right\}$$

be the set of opponents actions on the path z'.

Case 1. Suppose that $\iota(x^+) = i$.

For any $j \neq i$, let the type $\bar{B}^j = A \backslash (A|_{Y_{a(x)}})_{-j} \cup \mathbf{a}_{-j}$. For index i, we introduce the framework $\bar{B}^i = A \backslash (A|_{Y_{a(x)}})_{-i} \cup \mathbf{a}_{-i} \backslash A_i^- \cup A_i^+$: a player of type $\bar{B}^i$ believes that the opponents will not deviate from the path z' if she plays action a'_+ at node x^+. For any $j \in I$ let the assignments π_j be the same as the ones defined in the proof of Proposition 2.

Step 2 in the proof of Proposition 2 extends without modifications: for any $\delta_1 > 0$, there are n, $T_1(\delta_1, n)$ such that $\rho_j^{T_1}(\bar{B}^j) > 1 - \delta_1$ for any $j \neq i$. Pick any $t > T_1(\delta_1, n)$. Any player in population i of type B^i such that $\rho_i^t(B^i) > 0$ either plays a' or a_x^* (or a_+ in the case that $x^+ = x$) at node x. If the player plays a_x^*, then her offspring may forget all actions in subgame $\Gamma_{a'(x)}$ and in the subgame $\Gamma_{a_+(x)}$ and become of type $\hat{B}^i = \bar{B}^i \cup \mathbf{a}(x, z^*)$, note that she cannot forget the actions on the BI path z^*. Because (x, z') is a profitable deviation, any player of type $\hat{B}^i$ plays a': she is lured off the BI path z^* by the prospect of obtaining a payoff at least as large as $u_i(z')$. With probability $(1 - \delta_1)$ she is matched with the type profile $\bar{B}^{-i}$ and the node x^+ is on the path $z^*(\hat{B}^i, \bar{B}^{-i})$. Her offspring does not observe the BI path z^* and, with positive probability, she may forgets actions in the set $\mathbf{a}(x, z^*)$, including all the actions that do not belong to the Nash equilibrium $\mathbf{a}_{-i}$ in the subgame $\Gamma_{a_x^*(x)}$. This player believes that the payoff achieved when reaching path z (or path z') is larger than the payoff she achieves by playing the BI action a_x^* at node x. Because she deviates from the BI path z^* at node x, her offspring cannot regain awareness of any actions in the subgame $\Gamma_{a_x^*(x)}$. In particular, she cannot regain awareness of the actions that do not punish her for playing a_x^* at node x, including those on the BI path z^*.

Since δ_1 can be taken arbitrarily small, we conclude that for any $\delta_2 > 0$, there exists a n large enough, and a $T_2(\delta_2, n) > T_1(\delta_1, n)$ large enough, such that such that

$$\sum_{B^i \in \mathcal{B}^i \,:\, a_x^*(B^i, x) = a_x^*} \rho_i^t\left(B^i\right) < \delta_2 \text{ for any } t \geq T_2,$$

and hence $f^t(z^*) < \delta_2$ for any $t \geq T_2$, analogously to step 3 in the proof of Proposition 2.

Case 2. Suppose that $\iota(x^+) = k \neq i$.

For any $j \neq i$, let the type $\bar{B}^j$ and the assignments π_j be defined as in the proof of Proposition 2. For player i, let $\bar{B}^i = A \backslash (A|_{Y_x})_{-\underline{i}} \cup \mathbf{a}_{-i} \backslash A_i^- \cup A_i^+$ as in case 1 and $\bar{B}_2^i = \bar{B}^i \cup \{a_+\} \backslash \{a'_+\}$. A player of type $\bar{B}^i$ believes that the opponents will not deviate from the path z' if she plays action a at node x. A player of type $\bar{B}^i_2$ instead believes that the path z is established if she plays action a at node x. For any index $n \geq 1$, and any framework $B^i \in \mathcal{B}^i$, $\bar{B}^i \subsetneq B^i$ or $\bar{B}^i_2 \subsetneq B^i$, we let $\pi_i\left(\bar{B}^i|B^i\right) = 1/n^I$, $\pi_i\left(\bar{B}^i_2|B^i\right) = 1/n^{2I}$ and $\pi_i(\hat{B}^i|B^i) = 0$ for any type $\hat{B}^i \notin \{\bar{B}^i, \bar{B}^i_2, B^i\}$. For any other type B^i, we let $\pi_i\left(B^i|B^i\right) = 1$. For n approaching zero, we let players in population i forget actions in $\bar{B}^i \backslash \bar{B}^i_2$ infinitely faster than actions in $\bar{B}^i_2 \backslash \bar{B}^i$.

Again, step 2 in the proof of Proposition 2 extends without modifications. Pick any $t > T_1(\delta_1, n)$. Any player in population i of type B^i such that $\rho^t_i(B^i) > 0$ either plays a' or a^* at node x. Consider the offspring of any player that plays a^*_x at node x. For any n small enough, the probability that she forgets actions in $\bar{B}^i_2 \backslash \bar{B}^i$ (leading to path z) is infinitesimal with respect to the probability that she forgets actions in $\bar{B}^i \backslash \bar{B}^i_2$ (leading to path z') and becomes of type $\hat{B}^i = \bar{B}^i \cup \mathbf{a}(x, z^*)$. Because (x, z') is a profitable deviation, any player of type $\hat{B}^i$ plays action a' at node x. With probability $(1 - \delta_1)$ she is matched with the profile of types $\bar{B}^{-i}$ so that the path $z^*(\hat{B}^i, \bar{B}^{-i}) = z$ is established. With positive probability, her offspring forgets the action a'_+ (leading to path z') as well as all actions in $\Gamma_{a^*_x(x)}$ other than those in the Nash equilibrium profile $\mathbf{a}_{-i}$, thus becoming of type $\bar{B}^i_2$. In such a case, she plays a' at node x because she believes to obtain a payoff at least as large as $u_i(z)$, which is larger than the payoff that she believes to obtain by playing action a^*_x. Because this player plays a' at x, her offspring cannot regain awareness of any actions in the subgame $\Gamma_{a^*_x(x)}$. With probability $(1 - \delta_1)$ she is matched with the profile of types $\bar{B}^{-i}$ so that the path $z^*(\hat{B}^i, \bar{B}^{-i}) = z$ is established and she cannot regain awareness of any forgotten action in $\Gamma_{a'(x)}$. Since δ_1 can be taken arbitrarily small, for any $\delta_2 > 0$, there exists a n large enough, and a $T_2(\delta_2, n) > T_1(\delta_1, n)$ large enough, such that $f^t(z^*) < \delta_2$, for any $t \geq T_2$, analogously to case 1. ■

Lemma 3 *If Γ has a deviation (x', z'), and for some $x : x' \prec x \prec z^*$, the subgame Γ_x has a Nash path $z \neq z^*$, then z^* may be persistently upset.*

Proof. The cases such that there is a deviation (x, z''), or the subgame $\Gamma_{x'}$ has a Nash path $z'' \neq z^*$ are covered in Lemma 2. Let $\mathbf{a}$ be a pure-strategy Nash equilibrium of subgame Γ_x that induces the path z, a be the action in $A(x)$ such that $a(x) \prec z$, and a' be the action in $A(x')$ such that $a'(x') \prec z'$. As in the proof of Lemma 2, let $\iota(x') = i$, and introduce

$$A^-_i = \{a \in A_{-i} \cap A(x'') : x' \preceq x'' \prec z', a(x'') \not\preceq z'\}$$

be the set of opponents actions that deviate from the path z'.

Case 1. Suppose that $\iota(x) = i$.

For any $j \neq i$, let the type $\bar{B}^j = A \backslash (A|_{Y_x})_{-j} \cup \mathbf{a}_{-j}$. For index i, we introduce the framework $\bar{B}^i = A \backslash (A|_{Y_x})_{-i} \cup \mathbf{a}_{-i} \backslash A_i^-$: note that the path z' does not reach the subgame Γ_x, hence there are no actions in $\mathbf{a}_{-i}$ that prevent reaching z'. For any $j \in I$ let the assignments π_j be the same as the ones defined in the proof of Proposition 2.

The second step in the proof of Proposition 2 extends because all actions in $\Gamma_{a(x)}$ are off path in any match where the population-i player plays a_x^*. Pick any $t > T_1(\delta_1, n)$, any type B^i such that $\rho_i^t(B^i) > 0$ either plays a' or $a_{x'}^*$ at node x'. If she plays $a_{x'}^*$, her offspring may forget all actions in subgame $\Gamma_{a'(x')}$ and become of type $\hat{B}^i = \bar{B}^i \cup \mathbf{a}(x, z^*)$. Because (x', z') is a profitable deviation, this offspring then plays a'. With positive probability, her offspring forgets actions in A_i^-, including those on the path z^*; she thus become of a type $\bar{B}^i \cup \mathbf{a}(\hat{B}^i, B^{-i})$ depending on the type profile B^{-i} she is matched with. Regardless of B^{-i}, the path $z^*(\hat{B}^i, B^{-i})$ does not reach the subgame $\Gamma_{a(x)}$. So this player does not play action the BI action a_x^* at node x, as she believes that the payoff of path z is larger than the payoff obtained when playing a_x^*. As she deviates from the BI path z^* at node x, her offspring cannot regain awareness of any actions in the subgame $\Gamma_{a(x)}$. Since δ_1 can be taken arbitrarily small, for any δ_2, there exists a n large enough, and a $T_2(\delta_2, n) > T_1(\delta_1, n)$ large enough, such that

$$\sum_{B^i \in \mathcal{B}^i \,:\, a_x^*(B^i, x) = a_x^*} \rho_i^t\left(B^i\right) < \delta_2 \text{ for any } t \geq T_2.$$

In any match B, such that the player from population i does not play a_x^* at node x, either the path $z^*(B)$ does not reach the node x or it deviates from the BI path z^* at node x. In either case, $f^t(z^*) < \delta_2$, for any $t \geq T_2$.

Case 2. Suppose that $\iota(x) \equiv k \neq i$.

For any index j, the type $\bar{B}^j$ and forgetfulness profiles π_j are defined as in case 1. Step 2 in the proof of Proposition 2 extends with the following modification: for any $\delta_1 > 0$, there exists a n large enough, and a $T_1(\delta_1, n)$ large enough, such that $\rho_j^{T_1}(\bar{B}^j) > 1 - \delta_1$ for any $j \notin \{i, k\}$ and $\sum_{B^i \in \mathcal{B}^i \,:\, B^i \subseteq \bar{B}^i} \rho_i^{T_1}(\bar{B}^i) > 1 - \delta_1$. The index i has been changed with k and players in population i (who may forget actions in the subgame $\Gamma_{a'(x')}$) may be unaware of a larger set of actions than $\bar{B}^i$.

Pick any $t > T_1(\delta_1, n)$, any player in population i of type B^i such that $\rho_i^t(B^i) > 0$ either plays a' or $a_{x'}^*$ at node x'. If the player plays $a_{x'}^*$, her offspring may forget all actions in subgame $\Gamma_{a'(x')}$ and become of type $\bar{B}_2^i$: note that because $i \neq k$, there are no actions on the path z^* in $A \backslash \bar{B}_2^i$. Because (x', z') is a profitable deviation, a player of type $\bar{B}_2^i$ plays a', deviating from the BI path

z^* at the node x' that precedes x. The offspring of the player in population k matched with such a player of type $\bar{B}^i_2$ may forget with positive probability all the actions on Γ_x, including those on the path z^*, thus becoming of type $\bar{B}^k$. With probability $(1-\delta_1)$ this offspring in population k is then matched with a type profile B^{-k}, such that $B^i \subseteq \bar{B}^i$ and $B^j = \bar{B}^j$ for all $j \neq k, i$. In such a match the path $z^*\left(\bar{B}^k, B^{-k}\right) = z$ is established. Again, since δ_1 can be taken arbitrarily small, for any $\delta_2 > 0$ there exists a n large enough, and a $T_2(\delta_2, n) > T_1(\delta_1, n)$ large enough, such that $\sum_{B^k \in \mathcal{B}^k \,:\, a^*_x(B^k,x)=a^*_x} \rho^t_k\left(B^k\right) < \delta_2$, and hence $f^t(z^*) < \delta_2$ for any $t \geq T_2$, analogously to case 1. ■ ■

Proof of Theorem 4. Necessity follows from the proof of Theorem 1.

In order to prove sufficiency, let $\iota(x)$ be i to simplify notation, let a be the action at node x that leads into the terminal node x, and $\mathbf{a}$ be the set of actions in the subgame $\Gamma_{a(x)}$ that lead into the terminal node z, i.e. $a \in A(x)$ such that $a(x) \preceq z$ and $a' \in \mathbf{a}$ if there is a node x' such that $a(x) \preceq x'$ and $a'(x') \preceq z$. For any population j, let $\bar{B}^j = A \backslash (A|_{Y_{a(x)}}) \cup \mathbf{a}$. Each player j of type $\bar{B}^j$ is aware only of the actions that lead into the terminal node in subgame $\Gamma_{a(x)}$.

Note that $z^*\left(\bar{B}\right) = z$, the path z is established if all players in the match are of type $\bar{B}^j$ for all population j. This is because (i) all players at nodes other than x think they have no choice but to play the action leading in z, (ii) because (x, z) is a profitable deviation, the players in population i prefer to play a at x and lead the path into z, rather than playing the backward-induction action a^*_x.

For any index $n \geq 1$, and any framework $B^i \in \mathcal{B}^i$, $\bar{B}^i \subsetneq B^i$, we let $\pi_i\left(\bar{B}^i|B^i\right) = 1/n^I$ and $\pi_i\left(B^i|B^i\right) = 1 - 1/n^I$; for any other type B^i, we let $\pi_i\left(B^i|B^i\right) = 1$. For any population $j \neq i$ and any framework $B^j \in \mathcal{B}^j$, $\bar{B}^j \subsetneq B^j$, we let $\pi_j\left(\bar{B}^j|B^j\right) = 1/n$ and $\pi_j\left(B^j|B^j\right) = 1 - 1/n$, for any other type B^j, we let $\pi_j\left(B^j|B^j\right) = 1$.

Intuitively, we let players in population i forget the actions in $(A|_{Y_{a(x)}}) \backslash \mathbf{a}$ infinitely slower than the players in population $j \neq i$ forget the actions in $(A|_{Y_{a(x)}}) \backslash \mathbf{a}$. This will give the time to players in population $j \neq i$ to accumulate unawareness in the subgame $\Gamma_{a(x)}$. When the players in population i deviate from the BI path at node x_i and play a they will most likely be matched with unaware opponents so that the path $z = z^*\left(\bar{B}\right)$ is established.

Indeed, steps 2 and 3 in the proof of Proposition 2 extend without modification. For any $\delta_1 > 0$, there exists a n large enough, and a $T_1(\delta_1, n)$ large enough, such that $\rho^{T_1}_j(\bar{B}^j) > 1 - \delta_1$ for any $j \neq i$. For any $\delta_2 > 0$, there exists a n large enough, and a $T_2(\delta_2, n) > T_1(\delta_1, n)$ large enough, such that $\rho^t_i(\bar{B}^i) > 1 - \delta_2$ for any $t \geq T_2$. It follows that the BI path z^* is upset by small probability model deterioration because $z^* \left(\bar{B}\right) = z \neq z^*$, and so $f^t \left(z^*\right) < \delta_2$, for any $t \geq T_2$. ■

References

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

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

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

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

[5] Battigalli P. [1996]: "Strategic Rationality Orderings and the Best Rationalization Principle", *Games and Economic Behavior* **13:** 178-200

[6] Battigalli P. and M. Siniscalchi [2002]: "Strong Belief and Forward Induction Reasoning", *Journal of Economic Theory* **106:** 356-91

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

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

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

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

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

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

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

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

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

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

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

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

[19] Kreps D. and R. Wilson [1982]: "Sequential Equilibria", *Econometrica,* **50:** 863- 894

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

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

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

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

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

[25] Shimoji M. and J. Watson [1998]: "Conditional Dominance, Rationalizability, and Game Forms," *Journal of Economic Theory* **83:** 161-195.

[26] Swinkels J. [1992]: "Evolutionary Stability with Equilibrium Entrants" *Journal of Economic Theory* ***57**: 306-332*

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

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