

Squintani, Francesco

Working Paper

Games with Small Forgetfulness

Discussion Paper, No. 1273

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Squintani, Francesco (1999) : Games with Small Forgetfulness, Discussion Paper, No. 1273, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221629>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Games with Small Forgetfulness*

Francesco Squintani
Northwestern University[†]

August 1999

Abstract

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 introduces small forgetfulness and shows that it may destabilize standard full-memory solutions. Players are repeatedly matched to play a game. After any match, they forget with infinitesimal probability the feasibility of any opponents' unobserved action, and they are reminded of all actions that they observe. During each period, they play an equilibrium consistent with their perception of the game. We show that the unique backward induction path drifts into a non-Nash, self-confirming equilibrium, in a class of extensive-form games that are fully characterized. Such a long-run prediction is always Pareto-undominated, and may Pareto dominate the original backward induction path. In one-shot simultaneous-move games, forgetfulness yields a refinement stronger than trembling hand perfection. Our results imply that there are games that players may never fully learn.

JEL CLASSIFICATION: C71, C73, D83

*I am indebted to Jeffrey Ely, Drew Fudenberg, Ehud Kalai, George Mailath, Steven Matthews, Larry Samuelson, and especially Eddie Dekel, and Juuso Valimaki, for guiding me in developing the ideas contained in this paper, and in organizing the exposition. I would also like to thank all the participants of a Math Center seminar at Northwestern University, of the 1999 North American Summer Meeting of the Econometric Society at the University of Wisconsin at Madison, of the 1999 Summer in Tel Aviv, and of the 1999 Stony Brook International Conference on Game Theory. The usual caveat applies.

[†]Department of Economics, 2005 Sheridan Rd. Evanston, IL, 60208-2009. E-mail: squinzio@nwu.edu
Webpage: <http://pubweb.nwu.edu/~fsq395/>

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. Individuals who are repeatedly involved in the same social interaction may hold different perceptions of available actions at different times. Their mental image of the game typically depends on their past experiences. Our model allows any player to forget with infinitesimal probability some possible opponents' actions, particularly if these actions have not been observed in previous periods. While large-probability forgetfulness is not a convincing assumption in economic settings, infinitesimal forgetfulness should not be dismissed. This paper shows that perfect-memory solution concepts may be destabilized by infinitesimal forgetfulness.

In our model, players from different large populations are repeatedly matched to play a game. During the first period, each player is fully aware of all possible actions in the game; however, with infinitesimal probability, the player may later forget some of them. During each period, players coordinate on an equilibrium consistent with their possibly partial model of the game, and they recall any observed action previously forgotten. Equilibrium play given awareness, together with the evolution of awareness resulting from forgetfulness, define a dynamic transition on the players' perception of the game. The results concern the long-run aggregate distributions of play.

We first allow players to forget only 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. Even under these conservative assumptions, full-awareness solutions may be *destabilized*. In extensive-form games, we prove that perfect equilibria may drift to a *non-Nash*, and not even unitary, self-confirming equilibrium. Such a long-run prediction is always Pareto-undominated by the original backward induction path, and may Pareto-dominate it. We characterize the class of generic perfect-information games in which the unique backward induction path is destabilized. When opponents' observed actions are also allowed to be forgotten, the backward induction

path is destabilized in games with a non-credible threat (formally defined in section 4.3). In this paper, players are always fully aware of their own feasible options. Some issues related to own actions' forgetfulness are presented in the Conclusion. The interested reader may request our extended version, Squintani (1999).

Turning to normal-form games, we first remark that such a representation is not appropriate to deal with forgetfulness in games that display a dynamic nature. Thus we restrict attention to one-shot simultaneous move games, and we prove that forgetfulness of opponents' actions which are unobserved at equilibrium may destabilize entire Nash equilibrium components. That result would hold even if all forgetfulness were only temporary, so that all players always recall any forgotten action after a short time regardless of whether they observe it or not. To relate our work to other normal form solution concepts, we characterize the set of Nash equilibria not destabilized by forgetfulness. Weak dominance, trembling hand perfection, and properness are shown not to be stronger than our refinement. However, the assumption that observed actions may not be forgotten renders all pure-strategy equilibria stable. Once that assumption is relaxed, our stability concept is *stronger* than trembling hand perfection.

This paper introduces a new question in the literature about learning and evolution in games (see Weibull 1992, Samuelson 1997, and Fudenberg and Levine 1998). In many treatments, players are assumed not to have any knowledge of the game beyond their own possible choices, in other ones, the main question is how players learn to play in a game they fully know. Instead, by explicitly analyzing how players forget, and recall what they and their opponents can do in a social interaction, we address the issue of how players learn the game itself. Our work is motivated by the supposition that, when learning how to deal with a complex economic interaction, the most difficult task is often to establish all the relevant possibilities, and once the modeling step is accomplished, coordinating on a model's equilibrium becomes relatively easy.

This paper focuses on the long-run behavior reached from initial full awareness. This assumption is not to be taken as a definite; rather it is an obvious benchmark to use in

dealing with forgetfulness. To derive a complete picture of how players learn the rules of the game, one may analyze our model considering any initial awareness state. The results we achieve imply that there are games players may never fully learn.

Some conclusions presented in this paper may be related to the results obtained by Fudenberg and Levine (1993b) on learning with experimentation. In that paper, players play a best reply to last-period observed play. However, with a small probability, they experiment with different strategies. While players initially coordinate on a (possibly non-Nash) self-confirming equilibrium, if they are patient enough, the population play eventually reaches a Nash equilibrium. In contrast, our players coordinate on a Nash equilibrium in the first period. Allowing for small forgetfulness, we prove that the population play deteriorates to a (possibly non-Nash) self-confirming equilibrium. It would be interesting to meld these two analyses to ascertain whether experimentation leads to Nash equilibrium when small forgetfulness is possible.

Finally, our analysis offers an interpretation of learning dynamics that does require payoff monotonicity. In evolutionary game theory, people are assumed to imitate those players who hold 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 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 awareness of actions.

The paper is presented as follows. The second section informally presents the model and some introductory examples. The third section concerns normal form games, and the fourth section extends the presentation to the perfect-information games. The conclusion presents a few possible extensions, and is followed by the Appendix, which lays out the proofs.

2 Motivational Examples

In this section we informally explain the main features of the model, and present some examples to introduce the relevance of forgetfulness.

Players from large populations are repeatedly, randomly and anonymously matched to play a game. Each match is formed by a player from each population. All players in the same population have the same action set and utility function. However, each player may hold a different perception of the game, identified by a subset of the action space. At each period, before being matched to play, each individual gathers information in order to formulate a conjecture regarding her opponents' strategies. That information is meaningful only when it is consistent with her interpretational model of the interaction. To represent that scenario, we say that, while she may not infer her opponents' strategies, each player correctly assesses the aggregate distribution of strategies in the populations, as long as they refer to actions of which the player is aware. When that is not the case, she completes her conjecture by attaching positive probability to actions that are rationalizable given her description of the game.¹ Given her conjecture, and her perception of the game, each player takes a sequentially rational strategy. This construction yields an equilibrium concept consistent with incomplete perception of the game in the population.²

After any period of play t , each player may forget with infinitesimal probability some of the actions, but she will be reminded of all actions that she has observed on the path of play. The path observed by each player, and thus her individual perception transition, depend on the strategy of the players with whom she is matched. The aggregate transition of the awareness distribution is obtained compounding the probability of a match between

¹It would be unfair to present results that depend on players holding completely unjustifiable beliefs, when their opponents play actions they are unaware of.

²Instead of assuming players to hold correct beliefs on the aggregate distribution of equilibrium strategies, we could propose that they formulate beliefs based on their past observations of play. While intuitively appealing, that modeling approach is much less tractable than the one we present in this paper, and yields almost the same results. In fact, the only difference with respect to the characterization presented in this paper concerns normal-form games full-support mixed strategies equilibria. Also, since we want to isolate the effect of forgetfulness on standard analysis, we think it meaningful to maintain as close as possible to the spirit of standard equilibrium concepts.




Figure 1: Extended Trust Games

different players, with the probability of individual transition in that match. Given the time $t + 1$ awareness distribution, we can calculate time $t + 1$ equilibrium. We assume that all players initially hold a full description of the game. Our results are in terms of the long-run distribution of play, obtained compounding the awareness distribution with each player's strategy.

In this informal presentation, we restrict forgetfulness to opponents' unobserved actions, and consider first extensive-form games. Since our model requires that players play a Perfect Bayesian Equilibrium (given the possibly incomplete perception of the game), and at each period forgetfulness is infinitesimal, one expects the long run prediction to be close to perfect equilibrium. In the following example, however, the unique long-run prediction is a non-Nash, non-unitary self-confirming equilibrium.³ Infinitesimal forgetfulness of opponents' unobserved actions has such a radical effect on extensive form games equilibria because off-path unawareness can accumulate over time, and upset the backward induction path.

Example 1 (Extended Trust Game) Two populations of players are repeatedly anonymously randomly matched to play a version of the trust game by Kreps (1990), which we depict in figure 1. Each player from population 1 needs to choose whether to 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 enter a costly fight (F) with

³In our extended version (Squintani 1999) we show that any *stationary* distribution of play is a heterogeneous self-confirming equilibrium.

the opponent, or accept (A) her abuse. The backward induction path is N : the players in population 1 will not trust their opponents.⁴

We assume that each player from population 2 can forget the possibility of action A with infinitesimal probability π_A whenever she does not observe it on the path of play. Since the backward induction path is N , action A is off-path and can be forgotten. Players unaware of A will honor the trust, as they believe to be punished if they do not. The ratio a_t of players unaware of A will increase over time as long as it is less than $1/2$. In fact, each player from population 1 will expect to face with probability a_t an opponent who plays H and with probability $1 - a_t$ one that plays C . For $a_t < 1/2$, she prefers not to enter the agreement, but by doing so, she renders A off-path, allowing a_t to increase with increment $\Delta a_t = (1 - a_t)\pi_A$. However, when $a_t > 1/2$, the player prefers to play TA . Players unaware of A cannot be reminded of it, as they play H which makes A off-path. None of the aware ones can forget A , as their observed path is TCA . The unique full-awareness Perfect Bayesian Equilibrium is destabilized, even though, at each period, players coordinate on the only PBE conditional on their awareness. Moreover, for π_A small enough, the only aggregate stable state is approximately $(TA; H/2, C/2)$ which is not a full awareness PBE, nor a full awareness Nash equilibrium, and not even a unitary beliefs self-confirming equilibrium (it is an heterogeneous beliefs self-confirming equilibrium).

The aggregate stable path Pareto-dominates the backward-induction solution, and predicts that players from population 1 will trust their opponents, and their trust will be honored with nonnegligible probability. $\diamond$

By repeating the analysis for *any* values of initial awareness $a_0 \in [0, 1)$, one can easily verify that the players will never be able to fully learn the game in Example 1. For $a_0 \in [1/2, 1)$, in fact, the unique steady state with infinitesimal forgetfulness is $a^* = 1/2$, whereas for $a_0 \in [0, 1/2)$, the steady state a^* coincides with a_0 . In both instances, at steady

⁴While the reputation literature points out that in a repeated game the first player will punish her opponent in order to establish a reputation for toughness (see Kreps and Wilson 1982, Milgrom and Roberts 1982, Kreps 1990), that explanation does not apply to random matching games.

state, there is a nonnegligible portion of the population that is not aware of at least part of the game.

The result depends on the structure of Example 1. The action F can be interpreted as a non-credible threat “nested” within the credible threat C . The second player poses a credible threat to the first one: if she plays T , the second player will play C . The first player may then respond to that threat by announcing: “If you play C , I will play F ”. But that threat is non-credible: if asked to choose between F and A , the first player will play A . Because of that, if required to pick H or C , the second player will take C , so that her threat is credible. Nevertheless, any player who forgets the feasibility of action A believes the first player’s threat: for him, it is credible. She thus takes H , and her threat to play C is non-credible. So when enough players in the population 2 have forgotten A , the players in population 1 rationally switch from action N to action A .

The assumption that players in population 1 always know how many players in population 2 have forgotten action A is *not necessary* to let forgetfulness destabilize the backward induction path of Example 1. Assume in fact that players’ beliefs are adaptive. In particular, when initially playing the backward induction solution, players who never play T always maintain the belief that all players in population 2 play C in response to T . Say, however, that the players in population 1 may forget action C with probability π_c if they do not observe it on their path of play. Players who forgot C at time $t - 1$ will play T at time t . At time $t + 1$ they will switch back to N if and only if they faced an aware opponent at time t . So, the proportion c of players unaware of C , and the proportion a of players unaware of A follow this transition.

$$\begin{cases} c_{t+1} = c_t a_t + (1 - c_t) \pi_c \\ a_{t+1} = a_t + (1 - c_t)(1 - a_t) \pi_a \end{cases}$$

The system reaches the steady state ($c^* = 1, a^* = 1$), which corresponds to the path TH : a Nash equilibrium outcome. Consistently with Fudenberg and Levine (1993b), if players form beliefs in an adaptive manner, and strategies are slightly perturbed at each period of play, a Nash equilibrium is reached at steady state. Their perturbations are motivated by

experimentation, ours are implied by forgetfulness.

Since Example 1 relies on forgetfulness of unobserved actions, one may believe that in normal-form games no Nash equilibrium is destabilized, as all the NE actions are observed. While this is true for pure strategy equilibria, when a mixed strategy is played, the players will not observe all the actions in the support, only the realized one. As they will be able to forget some unobserved actions in the support of the strategy, their choice may change. Each opponent's equilibrium belief changes to account for these new equilibrium actions, and this change may be enough to offset the original Nash equilibrium. Unlike the previous example, destabilization does not require unawareness to accumulate over time. Destabilization will occur also when forgetfulness is temporary, i.e. players recall any forgotten action after just one period of unawareness, even if they do not observe them. The result is informally illustrated in the following example.

	D	E
A	3,1	0,2
B	0,1	3,2
C	2,0	2,0

Example 2 Each player in population 1 is randomly matched with a player in population 2 to play the above game. The game has 2 Nash Equilibrium components: the singleton (B, E) and the weakly dominated component $(C, \sigma_D \in [1/3, 2/3])$. Let the population initially play any Nash equilibrium from the second component. After the first repetition of play, a fraction σ_D of players in population 1 observes action D , and forgets action E with probability π_E . A fraction $1 - \sigma_D$ observes E , and forgets D with probability π_D . At the second repetition of play, $\pi_D(1 - \sigma_D)$ players are unaware of D , and play B , whereas $\pi_E\sigma_D$ are unaware of E and play A . In equilibrium, each player in population 2 knows that if she is matched with a player not fully aware, she will not receive action C . As that happens with probability $\pi_D(1 - \sigma_D) + \pi_E\sigma_D > 0$, her only best response is E . Fully-aware players in population 1 anticipate such behavior and respond playing B . The “bad” equilibrium

has been upset in only one period of play. With infinitesimal forgetfulness, $\pi \rightarrow 0$, the aggregate play reaches instantaneously the “good” equilibrium (B, E) . $\diamond$

Remark 1 In the paper, whenever we refer to normal-form games, we mean one-shot simultaneous-move games. For them, normal form is equivalent to extensive form representation. In general, instead, the normal-form representation is inappropriate to deal with forgetfulness. Consider Example 1 in normal-form, and let the initial state Nash equilibrium (NA, C) . Suppose each player in population 1 may⁵ forget C . Now she is indifferent between TF and TA , and she prefers them to NF and NA . If she plays TF , her opponent responds with H , so that forgetfulness yields (TF, H) . Repeating the analysis in an extensive-form representation, however, clarifies that it is inappropriate to say that the unaware player is indifferent between the strategies TF and TA . In fact, she is called to choose between F and A only if her opponent has chosen C . Even if she is unaware of action C at the beginning of the game, the player cannot be unaware of action C at the moment she chooses between F and A . The choice between F and A is contingent on the occurrence of C , it is thus meaningless to talk about a strategy TF , by which the player commits to play F , if she is not even aware of C . $\diamond$

3 Normal Form Games

In the interest of simplicity, we first present the model for normal-form games. In the next section, we will extend the presentation to perfect-information extensive-form games. Alternatively, we could present only the model for general extensive-form games, but since that is rather involved, we relegate it to our extended version (Squintani 1999).

⁵Not to burden the exposition, we make our point using Example 1, and allowing for observed actions’ forgetfulness. We could present the same argument with forgetfulness of unobserved actions only, modifying Example 1 so as to have H on backward-induction path, instead of C .

3.1 The Model

We first define the equilibrium behavior at each period of play, then deal with the dynamics implied by forgetfulness and learning.

Consider the normal-form game $G = (I, A, u)$. Each $i \in I$ denotes a continuous population of players: $A = \times_I A_i$ is the finite action space and $u_i : A \rightarrow \mathbb{R}, i \in I$ the utility functions. Each player $l \in [0, 1]$, in any population i holds a model of the game, denoted by R^i , that consists of a subset of the opponents' action space assigning at least one action to each player. Formally, $R^i \in \mathcal{R}^i := \times_{j \neq i} (2^{A_j} \setminus \emptyset)$, let R_j^i be the j -th component of the set R^i . Let $\alpha_i : [0, 1] \rightarrow \mathcal{R}^i$, be the assignment of models, and let $\alpha = (\alpha_i)_I$. A strategy is a profile $\sigma = (\sigma_i)_I$ of Lebesgue-measurable functions $\sigma_i : [0, 1] \rightarrow \Delta(A_i)$.

At each period t , each player enters a match that includes one player from each population. The matches are formed randomly, anonymously, and independently over time. We assume the players play an equilibrium called ABE (short for Awareness Bayesian Equilibrium) described as follows. Players do not know the strategy of the players with whom they are matched, they formulate a conjecture that coincides with the average population play, whenever referring to actions the player is aware of. The conjecture is then completed with a belief that is rationalizable given the player's possibly partial model of the game.⁶ For any model R^i , denote by $E[\sigma_j | R^i]$ the average distribution of strategies in population j restricted to the actions contained in R_j^i .

Definition 1 *Given the game G and the assignment α , the strategy profile σ is an Awareness Bayesian Equilibrium whenever, for any player l , in population i , holding model R^i ,*

⁶A simple example may help with the intuition. Consider a player in population 1, who is unaware of action A . Players in population 2 play A or B with probability $1/2$ and they do not play C . Action D is strongly dominated by B . We require player 1 to correctly assess that one half of her possible opponents play B ; but, being unaware of A , she cannot believe the remaining ones to play A . It would not make much sense to require her to believe them to play D either, as she knows that to be dominated by B . As long as a mixture of B and C is rationalizable, instead, she may believe the players in population 2 play it. Thus her beliefs would be restricted so that half probability is given to B , and the remaining half to a mixture of B and C that is rationalizable when action A is ruled out.

1. (Rationality) for any $a_i \in A_i$, $\sigma_{il}(a_i) > 0$ only if

$$a_i \in \arg \max_{\hat{a}_i \in A_i} \sum_{a_{-i} \in A^{-i}} \left[\prod_{j \neq i} \beta_j^l(a_j) \right] u_i(\hat{a}_i, a_{-i})$$

2. (Conjectures)⁷ for any $j \neq i$,

$$\beta_j^l = E[\sigma_j | R_j^i] + \left[1 - \sum_{a_j \in R_j^i} E[\sigma_j | R_j^i](a_j) \right] \hat{\beta}_j[R^i]$$

where $\hat{\beta}_j[R^i]$ is any rationalizable⁸ belief for the game $G|_{R^i} = (I, R^i, u|_{R^i})$.

As the action set A is finite, an ABE always exists by standard argument.

Each player's individual awareness transition will depend on her current model, on the action profile she observes, and on random forgetfulness. For simplicity, we assume forgetfulness to occur with probabilities fixed over time and independently across actions and opponents.⁹ Fix the matrix of forgetfulness probabilities $\pi = \{\pi_{a_j}^i\}_{i,a_j}$: each entry is the probability for player i to forget action a_j , where $j \neq i$. One cannot allow actions to be forgotten when that amounts to deleting some player's action set, or else that player's mental model is not well defined. So, for any mental model R^i define the *forgettable actions* set $\hat{F}(R^i) := \{a_j : R_j^i \setminus \{a_j\} \neq \emptyset\}$, and the *forgotten actions* random set $F(R^i) \subseteq \hat{F}(R^i)$ such that $a_j \in F(R^i)$ with probability $\pi_{a_j}^i$, independently across i and j .

To avoid burdening the formalization of individual awareness transition, we specialize it to the case in which forgetfulness is permanent and observed actions cannot be forgotten.

Assumption 1 *Each player with R^i , who observes profile a and forgets $F(R^i)$ will hold the model $\hat{R}^i = (R^i \setminus F(R^i)) \cup \{a\}$.*

⁷This functional form is chosen to guarantee ABE existence. Say we assumed conjectures to consist of conditional population distribution whenever possible, and rationalizable beliefs otherwise. When composing beliefs with our rationality assumption, we would obtain a Best Reply correspondence that need not be upper-hemi-continuous.

⁸Cf. Bernheim 1984.

⁹General non-independent forgetfulness may be captured defining a forgetfulness probability system π s.t. $\forall i, \forall R^i$ let $\pi(\cdot | R^i) \in \Delta(2^{R^i})$ and $\pi(\hat{R}^i | R^i) = 0$ if $\hat{R}^i \setminus \hat{F}(R^i) \neq \emptyset$. In Squintani (1999) we show our results to be invariant to general forgetfulness unless (in some normal form games) forgetfulness of different actions is perfectly correlated.

To allow forgetfulness of *observed* actions, let $\hat{R}^i = (R^i \cup \{a\}) \setminus F(R^i)$. Were forgetfulness to be only *temporary*, Assumption 1 would be modified by substituting A_{-i} for R^i in all the assertions, so that players would recall all forgotten actions after one period of play.

Given the time t assignment α^t , Definition 1 allows to derive the set of ABE, $E^*(G, \alpha^t)$. In case of multiple ABE, each player's awareness transition depends on which ABE is considered. We require *inertia*: at each period t , the players must coordinate on one of the ABE that are closest to that played at period t . Our metric is the average across players of the (Euclidean) distance between each player's strategies at time $t + 1$ and at time t . We thus ensure that our destabilization results are due to forgetfulness only: whenever there are multiple equilibria, a non-inertial population can drift away from the starting point equilibrium for no reason whatsoever. Formally, given G, α^{t+1} , let the set of *inertial* ABE be:

$$\arg \min_{\sigma \in E^*(G, \alpha^{t+1})} \sum_I E ||\sigma_i - \sigma_i^t||$$

Our main result is that standard solution concept are upset by infinitesimal forgetfulness, thus we concentrate our analysis on the long-run behavior of a population of individuals that hold initially a complete model of the game, and that coordinate on a equilibrium strategy.¹⁰ Therefore we stipulate that at time $t = 0$, all players l in population i hold the fully-aware model A_{-i} , and coordinate on the same equilibrium strategy $\sigma^\circ \in \times_I \Delta(A_i)$. Formally, we assume that:

$$\forall l, \quad \alpha_{il}^0 = A_{-i}, \text{ and } \sigma_{il}^0 = \sigma_i^\circ.$$

Aggregating individual awareness transitions in the population, one obtains a random population transition description $\gamma : (\alpha^t, \sigma^t) \mapsto \alpha^{t+1}$. Given the aggregate transition γ , and the initial assignment α^0 , and equilibrium σ^0 , the random population assignment α^t and equilibrium σ^t may be defined recursively for any $t \geq 0$. That formulation, however,

¹⁰It would be inappropriate to derive such a result in terms of stationary or stable states without considering explicitly the initial awareness. Suppose we proved stationary a non-Nash state with some unawareness. If such state could not ever be reached from initial full awareness, the result would not be driven by infinitesimal forgetfulness, but by unawareness originally present in the population.

requires us to keep track of a rather complicated stochastic process on the space of all the assignments and all the distributional strategies. In order to define the awareness transition in a more tractable manner, we follow a standard approach of evolutionary game theory. We partition the population in a finite set of “types”, and then invoke a “Law of Large Numbers” argument,¹¹ to approximate the stochastic transition γ with a deterministic transition defined on the frequencies of types.

Specifically, we restrict attention to inertial ABE where players choose the same strategy, if they hold the same mental model, and if they took the same strategy at the previous period.¹² Formally, given σ^t and α^{t+1} , an inertial ABE σ^{t+1} is defined *simple* if:

$$\forall i, \sigma_{il}^{t+1} = \sigma_{il'}^{t+1} \text{ whenever } \alpha_{il}^{t+1} = \alpha_{il'}^{t+1} \text{ and } \sigma_{il}^t = \sigma_{il'}^t.$$

For each population i , the restriction to simple ABE allows to assigns each player, at each period of play, a type $T^i = (R^i, \hat{\sigma}_i)$, where $\hat{\sigma}_i \in \Delta(A_i)$ denotes the strategy played at the last period of play. Let $\mathcal{T}^i$ denote the set of types in population i , let λ_i be the distribution in the population, and define the product measure λ on $\times_i \mathcal{T}^i$. The strategy σ_l of any player l of type T^i will be denoted as σ_{T^i} , and her conjectures as $\beta[T^i]$. Since the distribution λ includes the distribution of strategies in the population, we can subsume all the dynamics in the distribution λ , write the set of simple inertial ABE as $\phi(G, \lambda)$, and define the set of associated aggregate distributions of play as $\varphi(\lambda, \phi(G, \lambda))$.

In order to describe the population transition, since we are dealing with continuous populations, and for any t , the support of λ^t is finite, we can invoke a Law of Large Numbers argument. We identify each population transition ratio λ_i^t , with the composition of the individual transition probabilities, across different types T^i , and across different observed profiles a , distributed according to $\varphi(\lambda^t, \phi(G, \lambda^t))$. Given the game G , and the

¹¹Boylan (1993), Proposition 3 shows that for finite large populations there exists 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.

¹²If anything, this restriction strengthens our destabilization results.

forgetfulness probabilities π , we obtain the transition correspondence $\xi(G, \pi, \cdot)$, that assigns a set of finite support distributions λ^{t+1} given a distribution λ^t . We omit the closed form of ξ as it is not particularly insightful.

Given the initial equilibrium σ° , and the transition $\xi(G, \pi, \cdot)$, we derive a set of solutions $\{\lambda^t\}_{t \geq 0}$. For each solution $\{\lambda^t\}_{t \geq 0}$, we can determine the average distribution of play:

$$\lim_{T \rightarrow \infty} \sum_{t=0}^T \frac{\varphi(\lambda^t, \phi(G, \lambda^t))}{T}$$

Holding the game G , and the initial equilibrium σ° fixed, we pick a sequence of matrices $\pi^n \rightarrow 0$, so that the model generates a sequence of sets f^n of average distributions of play that depend on π^n . We denote as f^* the set of the limit points from any selection of f^n . With a minor notational violation, we will often identify the sets f^n and f^* with their elements, if the sets are singleton.

Definition 2 *Given the game G , and the sequence of matrices $\pi^n \rightarrow 0$, the equilibrium profile σ° is F -stable if $f^* = \sigma^\circ$.*

To ensure almost full memory, the analysis is conducted for sequences of matrices $\pi^n \rightarrow 0$. What gives some degrees of freedom is the sequence along which $\pi^n \rightarrow 0$. Whereas all actions are infinitesimally forgotten, the relative probability to forget different actions may be very large.

In showing our results, it will often be the case that the game admits a unique sequence of simple inertial ABE as a function of π , which shall be denoted by $\{\sigma^t\}_{t \geq 0}$. Moreover, often the players' equilibrium strategies will not depend on their previous period's strategy. So we introduce the distribution of models in each population i , denoted by $\rho_i \in \Delta(\mathcal{R}^i)$, with $\rho = (\rho_1, \dots, \rho_I)$, and denote strategies as σ_{R^i} . When no confusion can occur, we will further simplify our notation, denoting the type unaware of an arbitrary action a by $[a]$, the type's proportion in the population by ρ_a , and her strategy by $\sigma_{[a]}$, and denoting the fully aware type by $[\star]$.

3.2 Forgetfulness of Opponents' Unobserved Actions

In the normal-form refinement literature (cf. Selten (1975), Myerson (1978) etc.) each Nash equilibrium is selected only if it survives a given stability check. A Nash equilibrium is trembling hand perfect, for instance, if there exists a sequence of full support perturbations along which the equilibrium strategy is still played by all players. Analogously, a Nash equilibrium is F -stable (for a certain sequence of matrices π^n) if and only if it survives a “forgetfulness” stability check: it is still played as an equilibrium after players forget actions with probability π^n . Thus F -stability is a refinement. We provide a new substantive reason for refining Nash equilibrium. Traditionally, players are assumed to fully understand the game, but to make small mistakes while playing it. We require them to be able to correctly play any game once they have written down its representation. However, we presuppose that when facing a complex social interaction, they may be unable to fully represent all its relevant alternatives.

We begin the exposition by pointing out some simple general properties of F -stability.

Proposition 1 *For any sequence of matrices $\pi^n > 0, \pi^n \rightarrow 0$, if σ° is a pure-strategy, or a full-support isolated¹³ mixed-strategy equilibrium, then it is F -stable.*

Because forgetfulness of observed actions is not allowed, if σ° is a pure strategy equilibrium, no player may ever forget an opponent's equilibrium actions. Inertia then requires all players to stick to the pure strategies initially played, so that σ° is F -stable. Full-support mixed-strategy equilibria are shown to be F -stable as follows. Say at time 1 some unaware types take a strategy different than σ° . Inertia requires fully-aware players to slightly adjust their strategy so that the aggregate distribution of play is σ° . In this case, the ratio of unaware players always remains very small, so that fully-aware players may always adjust to yield σ° .

¹³A Nash equilibrium σ^* is isolated if $\exists \varepsilon$ s.t. $\forall \sigma : \|\sigma - \sigma^*\| < \varepsilon$, σ is not a Nash equilibrium. There are (somewhat uncommon) games with full-support, non-isolated Nash equilibria.

Since there exist pure-strategy weakly dominated equilibria, F -stability is not stronger than weak dominance, and hence trembling-hand perfection or properness. However, we can show that F -stability is also *not weaker* than properness, and hence trembling-hand perfection or weak dominance.

	A	B	C	V	W
D	1,0	0,2	1,3	2,1	2,2
E	2,3	0,2	2,0	1,2	1,1

Example 3 Extending the analysis of Example 2, one sees that all the Nash equilibria in the component $(\sigma_D^\circ \in [1/3, 2/3], \sigma_B^\circ = 1)$ are F -destabilized for any sequence of strictly positive π^n matrices. In fact, $(1 - \sigma_D^\circ)\pi_D$ players in population 2 will forget D and play A , and $\sigma_D^\circ\pi_E$ will forget E and play C . As V and W are strictly dominated by mixtures of (A, B) and (B, C) respectively, aware players in population 2 will never play them. Thus aware players in population 1 must play E at ABE , and aware players in population 2 respond with A , so that $f^*(A, E) = 1$. However, all the equilibria in the NE $(\sigma_D^\circ \in [1/3, 2/3], \sigma_B^\circ = 1)$ are undominated because V renders D undominated. In 2-player games, undomination implies trembling hand perfection, so all NE $(\sigma_D^\circ \in [1/3, 2/3], \sigma_B^\circ = 1)$ are also THP. Properness is stronger than THP and, allowing only trembles ordered with equilibrium utility, it follows closer F -stability, as it mimics the requirement of players to play an ABE at each period of play. However, the Nash equilibrium $(\sigma_D^\circ = 1/2, \sigma_B^\circ = 1)$ is proper, and F -destabilized. In fact all non-Nash actions of player 2 yield the same¹⁴ utility against $\sigma_D^\circ = 1/2$. Pick trembles assigning $\varepsilon/4$ probability to each action A, C, V , and W : then $U_1(D, \sigma_\varepsilon^\circ) = 2(1 - \varepsilon) + 3\varepsilon/2 = U_1(E, \sigma_\varepsilon^\circ)$, so player 1 may randomize against the perturbation, and the equilibrium is proper. $\diamond$

Example 3 suggests that, unlike some apparently more appealing refinements as weak dominance, trembling hand perfection and properness, F -stability is invariant to deletion of

¹⁴The example makes use of a tie in the upsetting actions. It would be nice to find out whether properness implies F -stability in games with generic upsetting strategies.

strictly dominated strategies. That intuition is confirmed by Proposition 2 below: whatever model of the game a player may hold, since she is always aware of all her strategies, she will never play a strictly dominated strategy. So the addition of a strictly dominated strategy does not change the evolution of the population in games of normal-form. Moreover, any normal-form solution concept meant to be stable with respect to forgetfulness, should not be invariant to deletion of strictly dominated strategies.

Proposition 2 *Let $F^*(G)$ be the set of F -stable NE of a game G . If G' is derived from G by deleting strictly dominated strategies, then $F^*(G) = F^*(G')$.*

We conclude our analysis of F -stability by giving a complete characterization for 2×2 games. First, we say that a game with action space $\{A, B\} \times \{C, D\}$ is a (normal-form) *entry game* if¹⁵ $u_2(B, C) = u_2(B, D)$, and $[(u_1(A, C) \sqcap u_1(B, C), u_1(A, D) > u_1(B, D) \text{ or } (u_1(A, C) > u_1(B, C), u_1(A, D) \sqcap u_1(B, D))]$. The name is inspired by the following game.

	C: fight	D: accept
A: enter	0,1	3,2
B: out	2,0	2,0

Theorem 3 *For any 2×2 game G , and any sequence of matrices $\pi^n > 0, \pi^n \rightarrow 0$, the Nash equilibrium σ° is F -destabilized if and only if G is an entry-game¹⁶ and $\sigma_C^\circ \in (0, 1)$.*

Proof. Consider an entry game, without loss of generality pick $u_2(B, D) = u_2(B, C)$, $u_1(A, C) \sqcap u_1(B, C)$ and $u_1(A, D) > u_1(B, D)$. Consider the case $u_2(A, D) \neq u_2(A, C)$, without loss of generality, say $u_2(A, D) > u_2(A, C)$. We want to show that any NE $(\sigma_C^\circ \in (0, 1), \sigma_B^\circ = 1)$ is F -destabilized for any $\pi^n \rightarrow 0, \pi^n > 0, \forall n$. In fact, say any such NE is played at stage 0. At time 1, $\rho_C^1 = (1 - \sigma_C^\circ)\pi_C > 0$ and $\sigma_{[C]}(A) = 1$: a positive fraction ρ_C^1

¹⁵The characterization is obviously invariant to relabeling of strategies, and of players. So that, for instance, it includes also $u_2(A, C) = u_2(A, D)$, and $u_1(A, C) < u_1(B, C)$, $u_1(A, D) \geq u_1(B, D)$ or $u_1(A, C) \geq u_1(B, C)$, $u_1(A, D) < u_1(B, D)$.

¹⁶Actually, not to burden the statement, we left out one special case: in entry-games where $u_2(\sigma)$ is independent of σ , all non-full-support equilibria are F -stable. See the proof in Appendix, for details.

of players in population 1 will forget C and therefore play A . As D is the only best response when A is played with positive probability, all aware players in population 2 must respond D : $\sigma_*(D) = 1$. As $\rho_*^2 = (1 - \pi_A)(1 - \pi_B)$, for π small enough, the unique best-response is A , so that all the aware players in population 1 play A . So, for π small, the population play jumps very close to (A, D) . As $u_2(A, D) > u_2(A, C)$ and $u_1(A, D) > u_1(B, D)$, a π small enough can be found so that forgetfulness at periods further than 1 is irrelevant. Since $f^* = (A, D)$, σ° is F -destabilized. The case for $u_2(A, D) = u_2(A, C)$ and the only if part of the proof are in the Appendix. ■

The above analysis characterized F -stable NE; a natural question is whether a destabilized population will settle on a (F -stable) Nash equilibrium. We show a slightly different result: any *stationary* state of the transition ξ (with infinitesimal forgetfulness) is a NE.¹⁷ The result complements the analysis of F -stability as a refinement: when a Nash equilibrium is destabilized by infinitesimal forgetfulness, if the population is ever to rest, it will reach a F -stable Nash equilibrium. Given a game G , let $S_n^* = \{\lambda_n | \lambda_n = \xi(G, \pi^n, \lambda_n)\}$ denote the set of stationary points of the transition ξ , given the forgetfulness matrix π^n .

Proposition 4 *For any sequence of matrices $\pi^n > 0, \pi^n \rightarrow 0$, if λ^* is a limit point of a sequence from $\{S_n^*\}_{n>0}$, then $f^* = \varphi(\lambda^*, \phi(G, \lambda^*))$ is a Nash equilibrium of G .*

Players are fully aware of their own actions and will remember any action they observe. So at a stationary state, when forgetfulness is very small, their beliefs on path must be almost right, and by continuity, they play a self-confirming equilibrium. As forgetfulness generates different mental models in the population, the players' beliefs may be heterogeneous: since G is a normal form game, f^* is a correlated equilibrium. As the awareness distribution is independent across populations, f^* is a Nash equilibrium.

On the other hand, the population play need not reach a stable state: it allows for cyclical behavior, even with initial full awareness and infinitesimal forgetfulness.

¹⁷We think that the result may be extended to include time averages over stable sets.

	D	E
A	3,2	0,0
B	2,1	2,1
C	0,0	3,2

Example 4 Assume $\pi_D^n = \pi_E^n =: \pi^n, \forall n$ and for simplicity, say that $\forall a_1 \in A_1, \pi_{a_1}^n = 0$. We will show that a population with initial full awareness will jump from one extreme of the NE component $(B, \sigma_D^\circ \in [1/3, 2/3])$ to the opposite one, hitting $(B, \sigma_D^\circ = 2/3)$ in even periods, and $(B, \sigma_D^\circ = 1/3)$ in odd ones. Start with $(\sigma_D^\circ = 2/3, \sigma_B^\circ = 1)$. At time 1, $\pi/3$ players in population 1 forget D and $2\pi/3$ forget E so that $\rho_E^1 = 2\pi/3$ and $\rho_D^1 = \pi/3$ and the ABE is s.t. $\sigma_{[E]}(A) = 1, \sigma_{[D]}(C) = 1$. As we want to find the inertial ABE, we want aware players in population 2 to randomize, so it must be that $2[\rho_E^1 + \sigma_*(A)] = f_A^1 = f_C^1 = 2[\rho_E^1 + \sigma_*(C)]$. As $\rho_E > \rho_D$ we need $\sigma_*(C) > 0$. Consider aware players in population 1: their best reply is A if $f_E^2 \leq 1/3$, it is B if $f_E^2 \in [1/3, 2/3]$ and it is C if $f_E^2 \geq 2/3$. Thus, their inertial ABE strategy is: $\sigma_*(C) = (\rho_E - \rho_D)$, $\sigma_*(A) = 0$, $\sigma_*(E) = 2/3$. At time 1, $f^1 = ([2\pi/3]A, [1 - 4\pi/3]B, [2\pi/3]C; [1/3]D, [2/3]E)$. Because of that at time 2, $1/3$ of the players in population 1 will not observe E and $2/3$ will not observe D : now $\rho_E^2 = \rho_E^1/3 + \pi(1 - \rho_E^1)/3 = 2\pi/9 + \pi/3 - o(\pi)$ and, $\rho_D^2 = 2\rho_D^1/3 + 2\pi(1 - \rho_D^1)/3 = 2\pi/9 + 2\pi/3 - o(\pi)$. Again we want aware players in population 2 to randomize, but now $\rho_E^2 < \rho_D^2$, so we need $\sigma_*(A) > 0$. The inertial ABE aware strategy is: $\sigma_*(A) = (\rho_D - \rho_E)$, $\sigma_*(C) = 0$, $\sigma_*(E) = 1/3$, so that the inertial ABE displays $f^2 = ([2/3]D, [1/3]E)$. By induction, for any t , if $f^{t-1} = ([2/3]D, [1/3]E)$, then $\rho_E^t = \rho_D^t + \pi/3 - o(\pi)$, so that $f^t = ([1/3]D, [2/3]E)$, and then $\rho_E^{t+1} = \rho_D^{t+1} - \pi/3 - o(\pi)$, so that $f^{t+1} = ([2/3]D, [1/3]E)$. So that, for $\pi \rightarrow 0$, one obtains that the sequence f^t oscillates hitting $([1/3]E, [2/3]D, B)$ in even periods and $([2/3]E, [1/3]D, B)$ in odd ones. Nevertheless, the time-average distribution of play is $f^* = ([1/2]E, [1/2]D, B)$, which is a Nash equilibrium. $\diamond$

3.3 Forgetfulness of Observed Actions

In this subsection we extend the analysis by allowing also forgetfulness of opponents' observed actions. Repeating the proof, it can be shown that Proposition 4 continues to hold. Nevertheless, the characterization of F -stability is different. Since we are allowing a larger set of actions to be forgotten, we will upset a larger set of Nash equilibria. In particular, we will get rid of the “bad” pure-strategy equilibria that survived F -destabilization in the previous section. In fact, pure-strategy equilibria were unaffected by forgetfulness only because observed actions could not be forgotten, and inertia thus required all players to stick to the pure strategies initially played. If observed actions may be forgotten, on the other hand, some players may deviate from the initial equilibrium, and upset it. As a consequence of that, we will be able to establish the main result of this section, Theorem 5, and show that F -stability is stronger than Trembling Hand Perfection if forgetfulness of observed action is allowed.

Unlike Theorem 1, the results in Theorem 5 do not hold for any sequence of forgetfulness matrices. Example 5 shows that there are some very peculiar games, in which some non-THP equilibria are F -stable along some forgetfulness sequence.

	C	D	E
A	1,1	0,0	2,1
B	1,1	1,0	2,1

Example 5 Consider the non-THP equilibrium (A, C) . At any time t , regardless of what they forget, inertia requires all players in population 2 to play C . At time 1 a ratio $\rho_{CE}^1 = \pi_C \pi_E$ of players in population 1 forget C and E , and play B . However all players who forget only C or E still play A by inertia, as well as do all players who forget D . At time 2, all players of type $[CE]$ will be reminded of C , but by inertia they will still play B . Players who forget D will always remain unaware of it, and so may not forget both C and E anymore. So, if $\pi_E^n = \pi_D^n = \pi_C^n, \forall n$, then (A, C) is F -stable. In fact, for any t ,

let $\lambda_{<CE>}^t$ denote the ratio of players that have been unaware of C and E at any time $T \square t$, and define $\lambda_{<D>}^t$ as the ratio of players unaware of D who have never been unaware of both C and E . Simple algebra shows that $\lambda_{<CE>}^t = t\pi_C^n \pi_E^n + o(\pi_C^n \pi_E^n) = o(\pi_D^n)t$, whereas $\lambda_{<D>}^t = t\pi_D^n + o(\pi_D^n)$. So that, for $\pi_D^n \rightarrow 0$ and any t , $\lambda_{<CE>}^t \rightarrow 0$, and $f^*(A, C) = 1$. On the other hand, if¹⁸ $\pi_D^n / (\pi_C^n \pi_E^n) \rightarrow K \in \mathbb{R}^{++}$, then the population play will slowly drift from (A, C) to a mixture of (A, C) and (B, C) . Also any Nash equilibrium of the form $(\alpha A + (1 - \alpha)B, C)$, with $\alpha \in (0, 1)$, is F -destabilized. $\diamond$

To simplify the statement of the following results (proved in the Appendix), we call σ° *always F -stable* if it is F -stable for any sequence of forgetfulness matrices $\pi^n > 0$, $\pi^n \rightarrow 0$.

Theorem 5 *If σ° is a strict equilibrium then it is always F -stable.¹⁹ If σ° is always F -stable, then it is trembling hand perfect.*

We also show that trembling hand perfection coincides with F -stability in 2×2 games we call *double-sided*,²⁰ i.e. games where there is not any player i such that $\forall \sigma_{-i}, u_i(a_i, \sigma_{-i})$ is independent of a_i .

Corollary 6 *For any double-sided 2×2 game, σ° is a trembling hand perfect equilibrium if and only if it is always F -stable.*

4 Games with Perfect Information

4.1 Extending the Model

Consider a perfect-information game $\Gamma = (X, Z, I, \iota, A, u_1, \dots, u_n)$, the formal definition is in Appendix. To represent a player's perception of the game while she is playing, it is

¹⁸This is essentially equivalent to require that each player in population 1 forgets action E whenever she forgets action C , instead of forgetting them independently. As her optimal choice against action C and E is the same, she may collapse them in a single opponent's "action".

¹⁹Actually the statement can be strengthened to include all quasi-strict isolated equilibria.

²⁰We suspect that these games have already been classified with a different terminology, but we were not able to find a definition established in the literature.

necessary to allow different mental models at different decision nodes. In fact, suppose player i is unaware of an action a at her first decision node x_1 . In case her opponent takes action a , player i may find herself at a decision node x_2 , of which she was unaware when taking her choice at node x_1 . She will then need a description of the game different from the one she had at x_1 .

Definition 3 *Each player in population i holds a mental frame $R^i \subseteq A$ s.t. $\forall x \in X$, $R^i \cap A(x) \neq \emptyset$, and $A_i \subset R^i$. At node y , she holds the mental model²¹ R^y defined as follows:*

1. *Let A^y be the union of R^i and the path between x_0 and y . That is,*

$$A^y = R^i \cup \{a \in A | x' = a(x), x_0 \preceq x, x' \preceq y\}.$$

2. *Let the pair (Y^y, R^y) be the largest tree with root x_0 , and contained in (Y, A^y) . I.e,*

$$Y^y = \{x_0\} \cup \{y' \in Y | \exists n, \{a_1, \dots, a_n\} \subseteq A^y \text{ s.t. } x_1 = a_1(x_0), x_2 = a_2(x_1), \dots, y' = a_n(x_{n-1})\}, \text{ and } R^y = A^y|_{Y^y}.$$

We extend the definition of ABE to account for perfection, and obtain the concept of Awareness Perfect Bayesian Equilibrium. For brevity, we directly present *simple* APBE.

As different types of opponents may take different actions, there could be a “separating” equilibrium in which a player identifies the types of players she is matched with, by observing their play at predecessor nodes. To account for that, we allow each player to Bayes update her beliefs (with respect to her opponents’ types), at every decision node. For any $x, y \in Y$, and $\sigma \in \times_X \Delta(A(x))$, define $\mu(y|\sigma, x) = \prod_{a: x' \rightarrow^a x'', x \preceq x', x'' \preceq y} \sigma_{x'}(a)$ the probability of reaching y from x under the behavioral strategy profile σ . For any node x , and mental model R^x define by $\beta_{T^j}[R^x]$ the conjecture of any player at x holding R^x with respect to the strategy of type T^j .

²¹For imperfect-information games, mental models R^h are defined on each information set h , and they satisfy the following requirement. Consider the predecessor set h' : if h is not in $\Gamma|_{R^{h'}}$, then R^h must include at least a path from $R^{h'}$ to h . Because of that an awareness type must not only include forgotten actions, but also actions that would be recalled in case an opponent took an action the type is unaware of. The problem is solved in Squintani (1999).

Definition 4 Given game Γ and state λ , the behavioral strategy profile $\sigma \in \times_I(\times_{X_i}\Delta(A(x_i))^{T^i})$ is a simple awareness perfect Bayesian equilibrium whenever for any role i , any type T^i , any decisional node $x : \iota(x) = i$,

1. (Bayesian Update) for any opponents' type profile $T^{-i} \in \mathcal{T}^{-i}$, define $\psi_{R^x}(T^{-i})$ such that

$$\psi_{R^x}(T^{-i}) = \frac{\mu(x|\beta_T[R^x], x_0)\lambda^{-i}(T^{-i})}{E^{\lambda^{-i}}[\mu(x|\beta_T[R^x], x_0)]} \quad \text{"whenever" possible,}$$

2. (Sequential Rationality) For any $a \in A(x)$, $\sigma_{T^i}(a) > 0$,

$$a \in \arg \max_{A(x)} E^{\psi_{R^x}} \left[\sum_Z \mu(z|\beta_T[R^x], a(x)) u_i(z) \right]$$

3. (Conjectures) For any j , and T^j , $\text{supp}(\beta_{T^j}[R^x]) \subseteq R_j^x$ and

$$\beta_{T^j}[R^x](a_j) = \sigma_{T^j}(a_j) + \left[1 - \sum_{\hat{a}_j \in R_j^x} \sigma_{T^j}(\hat{a}_j) \right] \hat{\beta}_{R^x}(a_j)$$

Where $\hat{\beta}_{R^x}$ is a backward induction solution²² strategy for the game $\Gamma|_{Y^x}$.

Given any type profile T , the terminal-node path distribution $\zeta_T \in \Delta(z)$ is such that $\forall z$, $\zeta_T(z) = \mu(z|\sigma_T, x_0)$. The aggregate path $f \in \Delta(z)$ is such that $\forall z$, $f(z) := E^\lambda[\zeta_T(z)]$.

For any mental model R^i the *forgettable actions* set is $\hat{F}(R^i) := \{a \notin A_i \mid \forall x, [R^i \cap A(x)] \setminus \{a\} \neq \emptyset\}$. The forgotten actions set $F(R^i)$ is defined as for normal form games. Each player of type T^i , who observes path z and forgets $F(R^i)$ will change to type $\hat{T}^i = ((R^i \setminus F(R^i)) \cup R^z, \sigma_{T^i})$ when observed actions cannot be forgotten, and to type $\hat{T}^i = ((R^i \cup R^z) \setminus F(R^i), \sigma_{T^i})$ when they can be forgotten. Each player i 's individual transition depends on her type R^i , on her random set $F(R^i)$, and on the observed path z . The latter depends on the type of opponents she is matched with, through ζ_T . The transition ξ is constructed by aggregating individual random transitions, across different types T^i , and across different types opponents' type matches T^{-i} . Inertia, and the definition of F -stability are treated similarly to the normal-form case.

²²Battigalli (1997) shows that backward induction coincide with his extensive form rationalizability concept in terms of path of play.

4.2 Forgetfulness of Opponents' Unobserved Actions

Example 1 in the second section shows that the backward induction path may be destabilized by infinitesimal forgetfulness. It must be pointed out, however, that the BI path of Example 1 may be F -destabilized only if $\pi_H^n = 0$, for n large enough. So one can only conclude that F -destabilization may occur for some sequence of *weakly positive* matrices π^n , where $\pi^n \rightarrow 0$. That is a somewhat weak result, because it seems to suggest that forgetfulness destabilizes the BI path only if it is restricted to some particular actions. To clear up the issue, we modify the Example 1 and show that F -destabilization may occur for some sequence of *strictly* positive matrices π^n that converge to 0.

Example 6 Consider the game illustrated in the figure 2.

The backward induction solution is (NA, ZC) . Let $\pi^n \rightarrow 0, \pi^n > 0, \forall n$ such that $\lim \pi_A^n / \pi_F^n > b > 1$. The formal analysis is lengthy and thus relegated to the Appendix. Intuitively, some players from population 2 will forget A as it is off path, and so play H . When the ratio of players unaware of A becomes larger than $1/2$, all the players aware of H will play T instead of N . Now Z is off-path, so some players from population 1 can forget it. At the same time, some players from population 1 will forget H if faced with players from population 2 that did not forget A . One obtains that at the steady state, $\rho_Z^* > 0, \rho_H^* > 0, \rho_Z^* + \rho_H^* = 1$, and $0 < \rho_F^* < 1/2, \rho_F^* + \rho_A^* = 1$. The limit average path will be a mixture: $f^*(NZ) = 1 - \rho_Z^*, f^*(TH) = \rho_Z^* \rho_A^*, f^*(TCA) = \rho_Z^* \rho_F^*$. This is not Nash and not even unitary self-confirming. Note, though, that the result does not hold for all vanishing sequences of strictly positive matrices. If one requires that $\forall n, \pi_A^n < \pi_F^n$, the backward induction path is stable. $\diamond$

Example 6 differs from Example 1 because of the addition of player 2's actions Q and Z after player 1's action N . The addition of the action Z is trivial, as the payoff for the path N in Example 1 is equal to the payoff for the path NZ in Example 6, moreover the addition of the action Q is irrelevant for the backward induction solution, as it is conditionally strictly




Figure 2: Pseudo-Tripede

dominated by Z at player 2's decision node.²³ Yet the difference in terms of forgetfulness is substantial. In fact, the BI path of Example 1 is F -destabilized only if $\pi_H^n = 0$ for n large enough, whereas that of Example 6, is F -destabilized for any sequence of matrices $\pi^n \rightarrow 0$ such that $\forall n, \pi_A^n > \pi_F^n$, regardless of the magnitude of π_H^n . As the two examples can be used to represent the same economic situations, an apparent contradiction arises. In fact, the comparison shows that the sequence along which forgetfulness vanishes is what really matters. If the modeler deems that $\pi_A > \pi_F$, she would be mistaken concluding that forgetfulness does not matter when $\pi_H > 0$, on the basis of Example 1. She should instead consider the “complete” model represented in Figure 2.

As already pointed out in the second section, the action F can be interpreted as a non-credible threat nested in the credible threat C . We will define all games displaying such structure *games with nested threats*, and show that the backward induction outcome of a game Γ can be F -destabilized if and only if Γ is a game of nested threats. In particular, the BI path is F -destabilized if the non-credible threat can be forgotten with “high” infinitesimal probability, and the action alternative to the credible threat cannot be forgotten at all. The characterization of games where the BI path is F -destabilized when *all* the actions can be forgotten can be simply derived from the games of nested threats, as shown in the previous discussion comparing Examples 1 and 6.

²³An action $a \in A(x)$ is conditionally strongly dominated, if there exist $\hat{a} \in A(x)$ such that for any $\mu \in \Delta(\{z : a(x) \prec z\})$ and any $\hat{\mu} \in \Delta(\{z : \hat{a}(x) \prec z\})$, $\sum_Z u_{\iota(x)}(z)\mu(z) < \sum_Z u_{\iota(x)}(z)\hat{\mu}(z)$. See Fudenberg and Tirole (1991).

We will first introduce our formal characterization and then interpret it in terms of nested threats. Consider a generic perfect-information game. Let a_x° be the backward induction choice at node x , and let $u^\circ(x)$ be the backward induction value of x . Denote by $a^\circ(\Gamma)$ the backward-induction path of game Γ . For any player i and any node x , we will also define as *backward-inductive maximin* value $\bar{u}_i(x)$, the value obtained by i if at any choice following (and including) x player i maximizes her utility and her opponents minimize it (the formal definitions is in the Appendix). Finally, for notational ease, let the pair (x, a) be called a *deviation* whenever $a \in A(x)$, and $a \neq a_x^\circ$. As customary, subscripts of actions and nodes denote the player assigned the move, and to avoid trivialities, we will restrict attentions to games such that $\forall x, \#A(x) > 1$ and $\forall a \in A(x), \iota(a(x)) \neq \iota(x)$.

Definition 5 *A perfect-information generic extensive-form game is a game of nested threats if there exists deviations (x_i, a_i) [x_i on path], and (x_j, a_j) [$a_i(x_i) \preceq x_j$] such that:*

1. $u_i^\circ(a_j(x_j)) > u_i^\circ(x_i)$
2. $\forall x : x_i \prec x \preceq x_j, [u_{\iota(x)}^\circ(a_j(x_j)) \geq \bar{u}_{\iota(x)}(a'(x)), \forall a' \in A(x)]$.

The mathematical interpretation in terms of nested threats is as follows. For any i, x , and $a_i \in A(x)$, call the outcome z a *threat* to (x, a_i) if $a_i(x) \prec z$ and $u_i(z) < u_i^\circ(x)$. The threat z to (x, a_i) is *trivial* if it can be reached only when player i does not maximize her utility at some node,²⁴ the threat z to (x, a_i) is *credible* if $u^\circ(a_i(x)) = u(z)$, if a threat is neither credible, nor trivial, it is *non-credible*. Definition 5 requires²⁵ that there exists z , credible threat to some (x_i, a_i) , x_i on path, and that there exists z' non-credible threat to some $(x_j, a_{x_j}^\circ)$, where $a_i(x_i) \preceq x_j$. In fact, by assumption, x_i is on path and there exist x_j such that $a_i(x_i) \preceq x_j$. Denote by z the terminal node reached from $a_i(x_i)$ along BI solution. Since a_i is not on path, $u_i^\circ(x_i) > u_i^\circ(a_i(x_i)) = u_i(z)$: we have established that z is

²⁴See Appendix for the formal definition.

²⁵The remaining requirements in Definition 5 are only imposed so that the nested threats are relevant for all players' choice. For instance, if it were not the case that $u_i^\circ(a_j(x_j)) > u_i^\circ(x_i)$, then a deviation of player j from $a_{x_j}^\circ$ to a_j would be irrelevant for player i .

a credible threat to (x_i, a_i) . Also, applying Requirement (2) to x_j and $a_{x_j}^\circ$, one obtains that $\bar{u}_i(a_{x_j}^\circ(x_j)) < u_j^\circ(a_j(x_j)) < u_j^\circ(a_{x_j}^\circ(x_j))$. Let z' be the terminal node reached from $a_{x_j}^\circ(x_j)$ along BI-maximin on j : we have established that z' is a non-credible threat to (x_j, a_j) .

Theorem 7 *The BI solution $a^\circ(\Gamma)$ is F -destabilized for some weakly positive sequence $\pi^n \rightarrow 0$ if and only if Γ is a game of nested threats.*

Sufficiency is established setting π^n as follows. For any x_k on the path from x_i to x_j (x_j included), consider all actions a_k that do not lead into x_j (or that differ from x_j). Allow forgetfulness of all and only the actions alternative to the BI-maximin path starting at any $a_k(x_k)$. Definition 5 requires i to deviate from the BI path when all such actions have been forgotten.

We prove necessity through three different claims. Since all the actions on path cannot be forgotten, the first claim establishes that for the BI path to be F -destabilized, a deviation a_j must occur in the continuation of an action a_i , that is itself alternative to an action on path. Suppose now that the BI path is F -destabilized, and the new path goes through a_i : the second claim requires that at least a deviation (x_j, a_j) (in the continuation of a_i) occurs because the BI value of a_j is larger than the BI-maximin value of $a_{x_j}^\circ$. In fact, along the new path, players may be reminded of actions of which they were previously unaware. So if players deviate in the wrong belief to obtain more than the BI value, sooner or later, an action that deludes their expectations will be taken, and thus at next period they will avoid taking the deviation. For the backward induction path to be destabilized through action a_i , and the new path to reach the last deviation a_j , the third claim simply requires it to be the case that forgetfulness of actions leads all the choices from x_i onwards into a_j .

The second part of our characterization will point out that whenever a game is rich enough, the BI path may be F -destabilized. Given a generic game $\Gamma = (X, Z, \prec, I, \iota, A, u)$, we construct a generic game Γ' that contains the game Γ , and such that the BI path of Γ' coincides with that of Γ on their common nodes. Formally the game Γ' is a *backward-induction-irrelevant expansion* of Γ if $X \subset X'$, $(I, \iota, \prec, A) = (I', \iota'|_X, \prec'|_X, A'|_X)$ and

$\forall x \in X, u^\circ(x) = u^\circ(x)'$. In terms of full memory analysis, the games Γ' and Γ represent the same strategic interaction, and Γ' is a richer, more complete, model than Γ . Intuitively, a BI-irrelevant expansion can be obtained by attaching trivial or conditionally dominated actions to the original game. Trivial actions are payoff-irrelevant, and so they do not change the BI solution, they are usually interpreted as “inaction”. By definition, conditionally strictly dominated actions are never chosen (under full memory), and so they are irrelevant for backward induction. Proposition 8 points out that we can construct a game with the payoff structure of Example 6, around the last choice on the BI path of Γ . That choice is thus upset by forgetfulness, and also upstream BI choices may be upset in a “domino effect”. Let $a^\circ(\Gamma)$ be the BI path of the game Γ , and let $f^*(\Gamma')|_A$ be the long-run prediction for the game Γ' restricted to the action set A .

Proposition 8 *For any n -player game Γ with $n > 1$, there exist a BI-irrelevant expansion Γ' and a sequence of matrices $\pi^n > 0, \pi^n \rightarrow 0$, such that $a^\circ(\Gamma')$ is F -destabilized, $f^*(\Gamma')|_A \neq a^\circ(\Gamma)$, and $f^*(\Gamma')$ Pareto dominates $a^\circ(\Gamma)$.*

Theorem 7 seems to restrict the F -destabilization result to a class of games with somewhat convoluted strategic properties. Nevertheless, Proposition 8 proves that forgetfulness may matter in all games that are rich enough. The analysis underlines that the fundamental ingredient in dealing with the issue of forgetfulness is not the game it is applied to, but the relative probability of forgetting different actions in the game. Any game can be expanded without modifying the full-memory BI solution, so that there exist a sequence of forgetfulness matrices for which the BI solution is F -destabilized. However, the exercise is irrelevant, if the modeler believes players forget actions of the expanded game according to relative forgetfulness probabilities, for which the BI solution is not F -destabilized.

4.3 Forgetfulness of Observed Actions

To introduce the characterization for the case in which also forgetfulness of opponents' observed actions is allowed, consider the following example.




Figure 3: Chain-Store Game

Example 7 Consider the game in Figure 3: the only BI solution is IA . The players in population 1 forget with probability π_F the action F and so play I or forget with probability π_A the action A and then play O , in which case they will never observe A as they make it off-path. Thus the system evolves as

$$\begin{cases} \rho_A^{t+1} = \pi_A(1 - \rho_A^t - \rho_F^t) + \rho_A^t \\ \rho_F^{t+1} = \pi_F(1 - \rho_A^t - \rho_F^t) + \rho_F^t \end{cases}$$

the only stationary state of the system is $\rho_A^* > 0, \rho_F^* > 0, \rho_A^* + \rho_F^* = 1$. The limit stationary distribution is $f^*(O) = \rho_A^*, f^*(IA) = \rho_F^*$. $\diamond$

The main feature of Example 7 is that the action F is a non-credible threat to the action I on path. We will extend the characterization of Theorem 7, to show that the BI solution of a game Γ is F -destabilized for some sequence of forgetfulness matrices also if Γ is a game with a non-credible threat to an action on path. Moreover, given a game Γ whose BI path cannot be F -destabilized, the addition of just one conditionally strictly dominated action yields a BI-equivalent game Γ' whose BI path is F -destabilized.

Proposition 9 *The backward induction path $a^\circ(\Gamma)$ is F -destabilized for some weakly positive sequence $\pi^n \rightarrow 0$ if and only if Γ has a deviation (x_j, a_j) such that $\forall x \preceq x_j, [u_{\iota(x)}^\circ(a_j(x_j)) \geq \bar{u}_{\iota(x)}(a'(x)), \forall a' \in A(x)]$.*

The proof is analogous to that of Theorem 7 and is thus omitted. It can be easily checked that the characterization includes games of nested threats from Definition 5, and that the

condition in Proposition 9 is implied by non-credible threats to nodes on path. Say, in fact, that z is a non-credible threat to $(x_i, a_{x_i}^\circ)$ x_i on path, then $\bar{u}_i(a_{x_i}^\circ(x_i)) \square u_i(z) < u_i^\circ(a_{x_i}^\circ(x_i))$. The remaining requirements of Proposition 9 are only imposed to make sure that such non-credible threat is relevant for i 's behavior. They essentially mean that if i believes to reach z when playing $a_{x_i}^\circ$, she will prefer to permanently deviate to a different path.

At the same time, note that, any game satisfying the characterization, such that x_j is the last node $x \preceq x_j$, x is on path, has a credible threat to a node on path. In fact, $u_j^\circ(a_{x_j}^\circ(x_j)) > u_j^\circ(a_j(x_j)) > \bar{u}_j(a_{x_j}^\circ(x_j))$, let z be the terminal node reached by BI-maximin path starting at $a_{x_j}^\circ(x_j)$. As $a_{x_j}^\circ(x_j) \prec z$ and $u_j(z) = \bar{u}_j(a_{x_j}^\circ(x_j)) < u_j^\circ(a_{x_j}^\circ(x_j))$, it follows that z is a non-credible threat to $(x_j, a_{x_j}^\circ)$.

5 Conclusion and Possible Extensions

In this paper we have shown that infinitesimal forgetfulness of unobserved opponents' actions may destabilize standard solutions. The unique backward induction path drifts into a non-Nash self-confirming equilibrium in any perfect-information game of nested threats. The class of 'games of nested threats' includes all games that are rich enough. In one-shot simultaneous-move games, forgetfulness allows for a simple characterization of stable equilibria in 2×2 games. In general, while all pure-strategy equilibria are stable, some "bad" proper equilibria are destabilized, and all long-run predictions are Nash equilibria. Allowing forgetfulness of all opponents' actions, we obtain a normal-form refinement equivalent to trembling hand perfection in 2×2 games, and generally stronger.

A meaningful extension consists of a general dynamic analysis of the model. The assumption of initial full awareness is here made only to highlight our destabilization results, and to focus the exposition. 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 cycles, and of their attraction sets, would then complete the analysis.

Another direction in which to extend this work is to allow players' awareness of their own possible actions to evolve over time. If one extends the analysis presented in this paper in that direction, she faces the problem that, once a player has forgotten the possibility to play a certain action, she may never recall that option again. In Squintani (1999), in particular, we show that, even if forgetfulness is restricted to own unobserved actions, some non-self-confirming-equilibrium outcomes may be stable even in *normal-form* games. One possibility to let players regain awareness of forgotten actions is to consider repeated matching with changing roles instead of fixed ones. However, Squintani (1999) shows that if one allows for role mobility, some long run predictions are not self-confirming equilibrium even when ruling out forgetfulness of own actions.

A Appendix: Normal Form Games

Proof of Proposition 1. Let a° be a pure strategy equilibrium of $G : \forall i, \forall a_i \in A_i, u_i(a^\circ) \geq u_i(a_i, a_{-i}^\circ)$. We claim that $\forall t, \forall i$, if $\forall T^i \in \text{supp}(\lambda_i^{t-1}), \sigma_{T^i}^{t-1}(a_i^\circ) = 1$ then $\forall T^i \in \text{supp}(\lambda_i^t), \sigma_{T^i}^t(a_i^\circ) = 1$, too. In fact, for any i , since observed actions may not be forgotten, $\forall R^i \in \text{supp}(\rho_i^t), a_{-i}^\circ \in R^i$. Suppose each player j , and type T^j plays a_j° , then $\beta[T^i](a_{-i}^\circ) = 1$, since $a_{-i}^\circ \in R^i$. So $\forall R^i \in \text{supp}(\rho_i^t), \sum_{A_{-i}} u_i(a_i, a_{-i}) \beta[T^i](a_{-i}) = u_i(a_i, a_{-i}^\circ), \forall a_i \in A_i$, and thus profile a° is concluded to be a ABE. By inertia, $\forall T^i \in \text{supp}(\lambda_i^t), \sigma_{T^i}^t(a_i^\circ) = 1$. Countable induction proves the first claim.

Take σ° full support isolated mixed strategy NE. For any i , define $BT^i := \{T^i = (R^i, \hat{\sigma}_i) | R^i \neq A\}$ [the set of unaware types], and let $GT^i = \mathcal{T}^i \setminus BT^i$. Take any distribution λ such that $\forall i, \lambda^i(BT^i) < \min\{\sigma_i^\circ(a_i) | i \in I, a_i \in A_i\}$. Fix any strategy profile $\sigma_\sim := ((\sigma_{T^i})_{BT^i})_I$. For any i , consider the system $E^{\lambda_i}[\sigma_{T^i}|BT^i] + E^{\lambda_i}[\sigma_{T^i}|GT^i] = \sigma_i^\circ$. Since it is linear and non-degenerate, it has a set of solutions $(\sigma_{T^i})_{GT^i} =: \sigma_\star^i$. Since $\forall i, \lambda_i(BT^i) < \min\{\sigma_i^\circ(a_i) | i \in I, a_i \in A_i\}$, it follows that $\forall a_i, E^{\lambda_i}[\sigma_{T^i}(a_i)|BT^i] < \sigma_i^\circ(a_i)$; so that there are some $\sigma_\star^i \in \Delta(A_i)^{\#(GT)}$. Since for any $i, \varphi_{-i}(\lambda, (\sigma_\star, \sigma_\sim)) = \sigma_{-i}^\circ$, and σ° is an full-support equilibrium, $u_i(a_i, \varphi_{-i}(\lambda, (\sigma_\star, \sigma_\sim)))$ is constant across all $a_i \in A_i$. Thus optimal strategy of any type in GT^i is any $\sigma^i \in \Delta(A_i)$. So any $\sigma_\star^i$ is an ABE strategy profile against $(\sigma_\star^{-i}, \sigma_\sim^{-i})$, for any i . Since $\sigma_\sim$ was arbitrary, there is a set of profiles $(\sigma_\star, \sigma_\sim)$

which are ABE. As σ° is isolated, if $\lambda_i(BT^i)$ is small enough, inertia requires all players of any type $T^i \in GT^i$, in any population i to play a profile σ_* that belongs to one of those ABE: that profile satisfies $\varphi_i(\lambda, (\sigma_*, \sigma_\sim)) = \sigma_i^\circ$. To prove that $\forall t, \forall i, \lambda_i(BT^i)$ is small enough, as ξ is continuous, one can use a first order expansion, together with an induction argument. So suppose that, for any $i, \forall \tau < t, \varphi(\lambda^\tau, \sigma^\tau) = \sigma_{-i}^\circ$, then $\forall (j, a'_j \in A_j), \lambda_i^t\{(R^i, \sigma_i) : a_j \notin R_j^i\} \approx \pi_{a_j}^i \sum_{\tau=1}^t (1 - \sigma_j^\circ(a_j))^\tau \rightarrow \pi_{a_j}^i / \sigma_j^\circ(a_j)$. Since $1/\sigma^\circ(a_j)$ is bounded, for $\pi_{a_j}^i \rightarrow 0, \lambda_i^t\{(R^i, \sigma_i) : a_j \notin R_j^i\} \rightarrow 0$, independently of t . That concludes the induction argument. ■

Proof of Proposition 2. To show that the F -destabilization is robust with respect to the addition of strictly dominated strategies, take $a_i \in A_i$ s.t. $\forall \sigma_{-i} \in \Delta(A_{-i}), \exists \hat{a}_i \in A_i : u_i(\hat{a}_i, \sigma_{-i}) > u_i(a_i, \sigma_{-i})$. Because own actions cannot be forgotten, $\forall t, i, T^i \in \text{supp}(\lambda_i^t)$, where $T^i = (R^i, \hat{\sigma}_i), \forall \sigma_{-i} \in \Delta(R^i), \exists \hat{a}_i \in A_i : u_i(\hat{a}_i, \sigma_{-i}) > u_i(a_i, \sigma_{-i})$. So $\sigma_{T^i}(a_i) = 0$ for any ABE σ . Therefore, $\forall t$, the presence of the strategy a_i does not influence the inertial ABE σ^t . ■

Proof of Theorem 3. Consider a 2×2 game and label actions $A_1 = \{A, B\}$ and $A_2 = \{C, D\}$. We know that if σ° is either a pure strategies or an isolated full-support strategies NE, then it is F -stable. Let us pick σ° s.t. $\sigma_B^\circ = 1, \sigma_C^\circ \in (0, 1)$ to represent any non-full-support mixed-strategy equilibrium. For σ° to be a NE, it must be the case that $u_2(B, D) = u_2(B, C)$.

Consider the first the case that $u_2(A, D) \neq u_2(A, C)$, and say without loss of generality that $u_2(A, D) > u_2(A, C)$. Now, for σ° s.t. $\sigma_B^\circ = 1$ to be a NE, B cannot be weakly dominated. Say B dominates A , then at any t , any type R^1 plays B . Inertia then implies that $f_C^t = \sigma_C^\circ$ and $f_B^t = 1$ for any t , so that σ° is F -stable. So we are left with the case that neither A nor B dominate each other. We deal with the case for entry games in the main body of the paper. Now we need to consider the case when $u_1(A, C) \square u_1(B, C)$ and $u_1(B, D) = u_1(A, D)$. In such case, σ° is F -stable because the $(1 - \sigma_C^\circ)\pi_C$ fraction of players in population 1 who forgets C , will still play B , by inertia.

Then, we repeat the analysis for the case $u_2(A, D) = u_2(A, C)$. Again, if B dominates A , then σ° is F -stable, and the same occurs if $u_1(A, C) \square u_1(B, C)$ and $u_1(B, D) = u_1(A, D)$. Moreover, for $u_1(A, C) \square u_1(B, C)$ and $u_1(A, D) > u_1(B, D)$, σ° is F -stable, because all aware players in population 2 are always indifferent between C and D , and, because of inertia, they will stick to strategy σ_C° , thereby keeping the ratio of unaware players in

population 1 infinitesimal for $\pi \rightarrow 0$.

Finally, we deal with full support, non-isolated equilibria. In 2×2 games, there exists a full support, non-isolated equilibrium only if $\exists i : u_i(\sigma) = u_i(\sigma') \forall \sigma, \sigma' \in \Delta(A)$, and neither action of player j dominates the other. Ruling out the case in which $u_1(\sigma) = u_1(\sigma')$ and $u_2(\sigma) = u_2(\sigma') \forall \sigma, \sigma' \in \Delta(A)$, we are left with games that belong to the entry-game set.

Say without loss of generality, that $u_1(A, C) \square u_1(B, C)$ and $u_1(A, D) > u_1(B, D)$, and $u_2(\sigma) = u_2(\sigma') \forall \sigma, \sigma' \in \Delta(A)$. Given any full-support equilibrium σ° , some players in population 1 may forget C and play A , some others may forget D and play B . Since $u_2(\sigma) = u_2(\sigma') \forall \sigma, \sigma' \in \Delta(A)$, by inertia, all players in population 2 maintain strategy σ_C° . Against such strategy all aware players in population 1 are indifferent (since σ° is a full-support equilibrium), and so, by inertia, they maintain σ_A° . At next period unaware types will be reminded, but again by inertia they will play last period strategy. As long as $\pi_C^n \neq \pi_D^n$, $f_A^* \neq \sigma_A^\circ$ so that $f^* \neq \sigma^\circ$. ■

Proof for Proposition 4. First consider the transition ξ for $\pi = 0$ at the stationary point λ^* . In the remainder of the proof, for notational ease, take λ to denote the measure associated with the distribution λ . For any i, R^i , since $F(R^i) = \emptyset$, it follows that $\lambda_i^*[T^i | R^i = \text{supp}(\phi(G, \lambda^*))] = 1, \forall i$, or else $\lambda^* = \gamma(\Gamma, \pi, \lambda^*, \phi(\Gamma, \lambda^*))$ is violated. That implies that, for any pair T^i, T^j in $\text{supp}(\lambda^*)$, it follows that $\beta[T^i] = \phi(\Gamma, \lambda^*)_{-i}$. Since λ^* is a product measure, $E^*(G, \lambda^*)$ is a set of Nash equilibria of the game G . Pick a sequence $\{\lambda_n\}_{n \geq 1}, \lambda_n \in S_n, \forall n$. By upper-hemi-continuity of the ABE correspondence, $E^*(G, \lambda^n) \subseteq E^*(G, \lambda^*)$ for n large enough. ■

Proof of Theorem 5. To show that strict equilibrium are not F -destabilized, take a° (pure-strategy) strict NE. For any j , denote by $GT^j := \{(R^j, \hat{\sigma}_j) | a_{-j}^\circ \in R^j\}$ (good types) and $BT^j := \{(R^j, \hat{\sigma}_j) | a_{-j}^\circ \notin R^j\}$ (bad types). Define $GT = \times_{j \in I} GT^j$ and $BT = \times_{j \in I} BT^j$.

Similarly to the proof for Theorem 1 on full-support mixed strategies, we first claim that for any profile σ_{BT} , if for any j , $\lambda_j(BT^j)$ is small enough, then there exists a ABE s.t. $\sigma_{T^j}(a^\circ) = 1, \forall T^j \in GT^j, \forall j$. In fact, as A is a finite set, there exists a $\delta > 0$ s.t. $\min_{\{i \in I, a_i \in A_i\}} [u_i(a^\circ) - u_i(a_i, a_{-i}^\circ)] > \delta, \forall a_i \neq a_i^\circ$. Therefore for any i , $\lambda_{-i}^t(GT^{-i})[u_i(a^\circ) - u_i(a_i, a_{-i}^*)] > E^{\rho^t}_{-i} \left[\sum_{a_{-i} \in A_{-i}} \sigma_{T^{-i}}(a_{-i}) [u_i(a_i^*, a_{-i}) - u_i(a_i, a_{-i})] \right]$, as long as for any $j \neq i$, $\lambda_j(BT^j)$ is small enough. Now set $\pi = \max_{i,j} \pi_{a_j^i}$. Suppose that $\forall (\tau \square t, i), \lambda_i^\tau(BT^i) \square \pi + o(\pi)$, and $\sigma_{T^j}^\tau(a^\circ) = 1, \forall T^j \in GT^j, \forall j$. For π small enough, by the above argument, for any σ_{BT} , there is a ABE s.t. $\sigma_{T^j}^t(a_j^*) = 1, \forall T^j \in GT^j, \forall j$. One of these ABE is inertial, and

so we may select it for our argument. At time $t+1$, all types in BT^j will be reminded of a_j° with probability of at least $\lambda_{-j}^\tau(GT^{-j})$. Therefore, $\lambda_i^\tau(BT^i) \square (\pi + o(\pi))(1 - (1 - \pi)^{I-1}) + \pi(1 - (\pi + o(\pi))) = \pi + o(\pi)$. The result then follows by induction.

Let us then prove that if σ° is not THP, then it will be F -destabilized for some $\pi^n > 0, \pi^n \rightarrow 0$. Consider the transition ξ at $t = 1$. Because all opponents actions can be forgotten, for any i and any model R^i , it is the case that $\rho_i^1(R^i) = \prod_{(a_j \notin R_j^i, j \in I)} \pi_{a_j}^i \prod_{(a_j \in R_j^i, j \in I)} (1 - \pi_{a_j}^i) > 0$. Say that σ^1 is an inertial ABE, and for any $j \in I$, let $\hat{f}_j^1 = E^{\lambda_j} [\sigma_{(R^j, \hat{\sigma}_j)} | R^j \neq A^{-j}]$. If $\forall j \in I, \text{supp}(\hat{f}_j^1) = A_j$, then $\exists i : \sigma_{i*} \neq \sigma_i^\circ$. If not, as $f_{-i}^1 = \rho^{-i}(A^{-i})E[\sigma_{-i*}] + \hat{f}_{-i}^1$, for any $j, f_{-i}^1 \rightarrow \sigma^\circ$ for $\pi \rightarrow 0$, and since $\sigma_i^\circ \in BR_i(f_{-i}^1)$ and i is arbitrary, we would contradict the hypothesis that σ° is not THP (cf. Fudenberg and Tirole 1991, pg. 352). Moreover, as f^1 , the distribution at time 1 is full support, then for any t , also f^t is full support. Suppose that there were a $f^t \rightarrow \sigma^\circ$ for $\pi \rightarrow 0$. Since σ° is not THP, $\exists i, a_i \notin \text{supp}(\sigma_i^\circ)$ s.t. $u_i(a_i, f_{-i}^t) > u_i(\sigma_i^\circ, f_{-i}^t)$. Since own actions may not be forgotten, it cannot be the case that $f^t \rightarrow \sigma^\circ$.

Even in the case that f^1 is not full support, it may be still be the case that, for $\pi \rightarrow 0$, $f^1 \not\rightarrow \sigma^\circ$, as long as $\exists i, a_i \notin \text{supp}(\sigma_i^\circ)$ s.t. $u_i(a_i, f_{-i}^1) > u_i(\sigma_i^\circ, f_{-i}^1)$. However, for $t > 1$, the relative ratios of unaware types for $k \neq j$ may be very different than for $t = 1$, so we cannot guarantee that f^t stays away from σ° for small forgetfulness. On the other hand, we can show that (selecting some particular sequences $\pi^n > 0, \pi^n \rightarrow 0$), σ° will be destabilized even if $\forall i, \sigma_i^\circ \in BR_i(f_{-i}^{(n),1})$. As such arguments holds also if $\forall i, \sigma_i^\circ \in BR_i(f_{-i}^{(n),t})$ and $f^{(n),t} \rightarrow \sigma^\circ$, it takes care also of the case just mentioned. Since σ° is not THP, for any full-support $\sigma^n \rightarrow \sigma^\circ$, there is an i and an $a_i \notin \text{supp}(\sigma_i^\circ)$ s.t. $u_i(a_i, \sigma_{-i}^n) > u_i(\sigma_i^\circ, \sigma_{-i}^n)$, as by definition of NE, $u_i(a_i, \sigma_{-i}^\circ) \square u_i(\sigma_i^\circ, \sigma_{-i}^\circ)$, it is also the case that $u_i(a_i, \sigma_{-i}^\circ) = u_i(\sigma_i^\circ, \sigma_{-i}^\circ)$. Moreover, there is a j s.t. $u_i(a_i, \sigma_j^n, \sigma_{-ij}^\circ) > u_i(\sigma_i^\circ, \sigma_j^n, \sigma_{-ij}^\circ)$; and so there is an a_j s.t. $u_i(a_i, a_j, \sigma_{-ij}^\circ) > u_i(\sigma_i^\circ, a_j, \sigma_{-ij}^\circ)$. Therefore, even if $f_{-ij} = \sigma_{-ij}^\circ$, any best reply $\sigma_{\bar{T}^i}$ by types $\bar{T}^i$ with model $\bar{R}^i, := (A_{ij} \times \{s_j\})$ is s.t. $\text{supp}(\sigma_{\bar{T}^i}) \cap \text{supp}(\sigma_i^\circ) = \emptyset$. At times $t \geq 2$, each player with model $\bar{R}^i$ may be reminded of actions $\hat{a}_j \neq a_j$, denote with GT^i , the set of types so generated. Even in the case that $f_{-i} = \sigma_{-i}^\circ$, however, a_i is an inertial ABE strategy for any $T^i \in GT^i$. In fact, $\text{supp}(\sigma_{\bar{T}^i}) \cap \text{supp}(\sigma_i^\circ) = \emptyset$ implies that $\|\sigma_{\bar{T}^i} - \sigma_i^\circ\| \geq \|\sigma_{\bar{T}^i} - a_i\|$ (with a minor notational violation, we let a_i denote the mixed strategy that gives all mass to a_i) and we know that $u_i(a_i, \sigma_{-i}^\circ) = u_i(\sigma^\circ)$. For any t , $\Delta \rho_i^t(\bar{R}^i) + \Delta \lambda_i^t(GT^i) > 0$. Set $\pi_{a_j}^i = o(\prod_{\hat{a}_j \neq a_j} \pi_{\hat{a}_j}^i) = \pi_{a_k}^l$ for any $(l, k) \neq (i, j)$. Then, at any

t , for any type $(R^i, \hat{\sigma}_i) \notin GT^i$, $[R^i \neq \bar{R}^i$, and $R^i \neq A^{-i}]$, $\lambda_i^t(R^i, \hat{\sigma}_i) = o(\rho_i^t(\bar{R}^i) + \lambda_i^t(GT^i))$, so that at steady state $\rho_i^*(\bar{R}^i) + \lambda_i^*(GT^i) \approx 1$. Therefore, even if for any t , $f_{-i}^t = \sigma_{-i}^\circ$, still for $t \rightarrow \infty$, $f_i^t \rightarrow \sigma_i^\circ$. F -destabilization occurs in the form of a slow drift. ■

Proof of Corollary 6. We want to show that in 2×2 games, all THP equilibria are F -stable. Say that σ° is THP, then it is undominated (in 2-player games). For both i , there is an a_{-i} s.t. $u_i(a_i, a_{-i})$ is not constant in a_i . Therefore, for σ° to be undominated, it must be either strict or (isolated) full-support. In both cases, as we know, it is F -stable. ■

B Appendix: Perfect Information Games

Definitions A perfect-information game $\Gamma = (X, Z, I, \iota, A, u_1, \dots, u_n)$ is defined as follows. The set X represents decisional nodes, and the set Z terminal nodes, $I = \{1, \dots, I\}$ is the set of players, (there is a population continuum of size 1 for each $i \in I$), $\iota : X \rightarrow I$ labels the decision nodes to players. The functions $u_i : Z \rightarrow \mathbb{R}$ are the utilities obtained by reaching a terminal node. They are extended on $\Delta(Z)$ according to the multilinear expansion formula. The action space A is introduced as follows. Let $Y = X \cup Z$, and introduce $<$, an irreflexive, acyclic ordering on Y satisfying the following requirements: $[\forall z \in Z, \nexists y \in Y \text{ s.t. } z < y]$; $[\exists ! x_0 \in Y \text{ s.t. } \forall y \in Y, y \not< x_0]$; and $[\forall y \in Y \neq x_0, \exists ! x \in X : x < y]$. We will often denote by $\prec$ the transitive closure of $<$, and $x \preceq y$ will mean that either $x \prec y$ or $x = y$. Let A the subset of Y^2 generated by $<$: it is customary to write $x \rightarrow^a y$ to mean $a = (x, y)$. The pair (Y, A) is thus a tree²⁶ with root $x_0 \in X$ and leaves Z . A path from node x to node y on a tree $G = (Y, A)$ from node x to node y is the unique set of actions $\{a_0, \dots, a_n\}$ s.t. $x \rightarrow^{a_0} y_1, y_1 \rightarrow^{a_1} y_2, \dots, y_n \rightarrow^{a_n} y$. We partition A into $\{A(x) : x \in X\}$ where $\forall x, A(x)$ is the set of actions exiting x .

Given a perfect information game Γ , let $Y_0 = Z$, and for any $j \geq 1$, set $Y_j = \{x \in Y \setminus (\cup_{k=0}^{j-1} Y_k) | \nexists x' \in Y \setminus (\cup_{k=0}^{j-1} Y_k), x \prec x'\}$. The *backward-induction solution* a° is defined as follows. Set $u^\circ(z) = u(z), \forall z \in Z$, and for any $j \geq 1$, and $\forall x \in Y_j$, let $a_x^\circ = \arg \max_{a \in A(x)} u_{\iota(x)}^\circ(a(x))$ and $u^\circ(x) = u^\circ(a_x^\circ(x))$. For any player i , the *backward-inductive maximin* values $\bar{u}_i$ are defined as follows. Set $\bar{u}_i(z) = u_i(z), \forall z \in Z$, and for any

²⁶In fact, (Y, A) is an orientation of connected, acyclic graph, and $\forall y \neq x_0$, the number of edges entering y is 1, and $\forall z \in Z$ the number exiting edges is 0. See Diestel (1997), first chapter.

$j \geq 1$, $\forall x \in Y_j$, let $\bar{a}_x = \arg \max_{a \in A(x)} \bar{u}_i(a(x))$ if $\iota(x) = i$, $\bar{a}_x = \arg \min_{a \in A(x)} \bar{u}_i(a(x))$ if $\iota(x) \neq i$, and $\bar{u}_i(x) = u_i(\bar{a}_x(x))$. For any player i , to define her trivial threats proceed as follows. Let A_{-i} be the set of all a_{-i} , selections of an action $a_x \in A(x)$ for each $x : \iota(x) \neq i$. Fix a_{-i} , let $\hat{u}_i(z) = u_i(z)$, $\forall z \in Z$, and for any $j \geq 1$, $\forall x \in Y_j$, let $\hat{a}_x = \arg \max_{a \in A(x)} \bar{u}_i(a(x))$ if $\iota(x) = i$, $\hat{a}_x = A_{-i}(x)$ if $\iota(x) \neq i$, and $\hat{u}_i(x) = u_i(\hat{a}_x(x))$. Let $\hat{z}(a_{-i}, x)$ be the terminal node reached by the path from x along the $\hat{a}$ solution. Let $\hat{Z}_i(x) := \{\hat{z}(a_{-i}, x) \text{ for some } a_{-i}\}$ be the set of i -non-trivial outcomes from node x . A *threat* z to (x_i, a_i) is *trivial* if $z \notin \hat{Z}_i(a_i(x_i))$.

Formal Analysis for Example 6. The list of the types and APBE strategies in the population is as follows: for player 1, $\sigma_{[A]} = HZ$, $\sigma_{[F]} = CH = \sigma_{[\star]}^1$; for player 2, $\sigma_{[QH]} = \sigma_{[H]} = NA$, $\sigma_{[QC]} = \sigma_{[ZH]} = \sigma_{[ZC]} = \sigma_{[C]} = \sigma_{[Z]} = TA$, and $\sigma_{[Q]} = NA = \sigma_{[\star]}^2$ if $f^1(HZ) > f^1(CZ)$ and TA vice-versa.

Instead of keeping track of all these types, we will first simplify the system. Consider the transition of types $[A]$ and $[C]$ when $f^1(HZ) > f^1(CZ)$. The fully aware types in population 1 will forget C because they play N and cannot see C on path. The types $[C]$ play T and thus observe A on the path TCA unless they face type $[A]$. The fully aware types in population 2 will forget A unless they play against types $[C]$ or not $[Z]$ (or any mixture with one of the two). The types $[A]$ play H and thus will never observe A on path and be reminded of it anymore. Therefore, the system for (ρ_A, ρ_C) follows: $\rho_A^{t+1} = \rho_A^t + \pi_A(1 - \rho_A^t - \rho_F^t)(\rho_H^t + \rho_Q^t + \rho_{QH}^t + \rho_2^t)$ (where the latter is the amount of fully aware types in population 2) and $\rho_C^{t+1} = \rho_C^t \rho_A^t + \pi_C(1 - \rho_C^t - \rho_Z^t - \rho_Q^t)$. As $\Delta^t(\rho_A) > 0$ and $\rho_A^1 > 0$, this system has only one attractor state which displays $\rho_A^* > 0$, and $\rho_C^* = \pi_C \frac{1 - \rho_Z^* - \rho_Q}{1 - \rho_A^* + \pi_C} \approx 0$ with $\pi_C \approx 0$. Moreover, $\forall t, \Delta^t(\rho_C) = \rho_C^t(\rho_A^t - 1 - \pi_C) + \pi_C(1 - \rho_Z^t - \rho_Q^t)$, that is, for π_C small enough, $\Delta(\rho_C^1) > 0$ and $\Delta(\rho_C^{t+1}) < 0, \forall t > 1$. Therefore, we can conclude that $\rho_C^t \approx 0, \forall t$, as $\rho_C^1 = \pi_C \approx 0$, for π small enough.

After dropping types $[C]$, from the analysis, will drop types $[QH]$, $[QC]$, $[ZC]$, and $[ZH]$. Note that $\sigma_{[ZH]} = \sigma_{[ZC]} = \sigma_{[Z]} = TA$ and that type $[H]$ cannot forget Z and we ruled out type $[C]$, whereas type $[Z]$ can forget either H or C , and finally Z is off-path for type $[ZC]$, and $[ZH]$, thus the latter cannot be recalled of Z , instead they can be recalled of H or C . So we can group $[ZH]$ and $[ZC]$ with $[Z]$ and the modification of the system is irrelevant for the results. Analogously, group $[QH]$ and $[CH]$ with $[H]$.

As long as $\rho_A^t \square \frac{1}{2}$, the system evolves now as follows.

$$\begin{cases} \rho_A^{t+1} \approx \rho_A^t + \pi_A(1 - \rho_A^t - \rho_F^t)(1 - \rho_Z^t) \\ \rho_F^{t+1} = \rho_F^t + \pi_F(1 - \rho_A^t - \rho_F^t) \\ \rho_H^{t+1} \approx \rho_H^t + \pi_H(1 - \rho_H^t)(1 - \rho_A^t) \\ \rho_Q^{t+1} = \rho_Q^t(1 - \pi_H) + \pi_Q(1 - \rho_Q^t - \rho_H^t - \rho_Z^t) \\ \rho_Z^{t+1} \approx 0 \end{cases}$$

All the steady states of this system display $\rho_A^* + \rho_F^* = 1$, and as $\pi_A/\pi_F > b > 1$, it is the case that $\forall t$, $\Delta(\rho_A^t) > \Delta(\rho_F^t) > 0$ and $\rho_A^t > \rho_F^t$. Thus the only attractor is s.t. $\rho_A^* > 1/2 > \rho_F^*$.

Nevertheless, consider τ s.t. $\rho_A^{\tau-1} \square 1/2 < \rho_A^\tau$: at time τ , all the types A^1 and Q will switch to T . Thus the system will now evolve as follows.

$$\begin{cases} \rho_A^{t+1} \approx \rho_A^t + \pi_A(1 - \rho_A^t - \rho_F^t)\rho_H^t \\ \rho_F^{t+1} = \rho_F^t + \pi_F(1 - \rho_A^t - \rho_F^t) \\ \rho_H^{t+1} \approx \rho_H^t + \pi_H(1 - \rho_H^t - \rho_Z^t)(1 - \rho_A^t) \\ \rho_Q^{t+1} = \rho_Q^t(1 - \pi_H) + \pi_Q(1 - \rho_H^t - \rho_Q^t - \rho_Z^t) \\ \rho_Z^{t+1} \approx \rho_Z^t + \pi_Z(1 - \rho_H^t - \rho_Q^t - \rho_Z^t) \end{cases}$$

The only attractor steady state for the system is as follows: for player 1, $\rho_Q^* = 0, \rho_H^* > \rho_H^\tau + \rho_Q^\tau, \bar{\rho}_Z > 0, \rho_Z^* + \rho_H^* = 1$, for player 2, $0 < \rho_F^* < 1/2, \rho_F^* + \rho_A^* = 1$. Thus for any π^n the average path will be a mixture: $f^*(NZ) = 1 - \rho_Z^*, f^*(TH) = \rho_Z^*\rho_A^*, f^*(TCA) = \rho_Z^*\rho_F^*$.

It is only left to show that f^* is not a Nash. As the game is generic and 2-player, every unitary self-confirming equilibrium is Nash, (cf. Fudenberg-Levine (1993a), pg. 541.) The path f^* can be originated by a Nash equilibrium σ^* only if $\sigma_N^* > 0, \sigma_{TA}^* > 0, \sigma_{TF}^* = 0$ and $\sigma_{HZ}^* > 0, \sigma_{CZ}^* > 0, \sigma_{HQ}^* = 0, \sigma_{CQ}^* = 0$. Thus to show that f^* is not a Nash, it is enough to show that it is not Nash in the reduced normal form given by $\{N, TA\} \times \{HZ, CZ\}$. It can be easily checked that in such reduced normal form CZ weakly dominates HZ , so that it cannot be a best reply against any mixed belief. $\diamond$

Proof of Theorem 7. The proof will be divided in two separate parts.

Part 1 Necessity.

Proof. For any node x , type T^i with frame R^i , model R^x , beliefs (ψ, μ) , and action $a \in A(x)$, let $u_{R^i}(a|\psi, \mu) := \sum_{T^{-i}} \psi_{R^x}(T^{-i}) \sum_Z \mu(z|\sigma_T[R^x], a(x)) u_i(z)$, $A_{T^i}(x|\psi, \mu) =$

$\arg \max_{A(x)} u_{R^i}(a|\psi, \mu)$, and $a_{T^i}(x|\psi, \mu)$ as the element of the previous set when it is a singleton. Also, with a minor notational violation, we define f_x^* as the distribution on the actions $A(x)$ induced by f^* .

We shall prove necessity through 3 different claims. The first one establishes that if the backward induction is F -destabilized, a deviation must occur in the continuation of an action a_i alternative to an action on path.

Claim 1 *The BI path $a^\circ(\Gamma)$ is F -destabilized only if $[\exists a_i \in A(x_i) : u^\circ(x_i) = u^\circ(x_0), a_i \neq a_{x_i}^\circ]$, $[\exists a \in A(x) : a_i(x_i) \preceq x, a \neq a_x^\circ]$, and $[\exists z' : a(x) \preceq z', u^\circ(z') > u^\circ(x_i)]$.*

Proof. Enumerate the backward induction nodes x_j in natural number index, maintaining the order of the sets Y_j they belong to.

Take x_1 (the last decision node on the BI path), and denote $i := \iota(x_1)$. Suppose that $f_x^*(a_x^\circ) = 1$ for all nodes $x \prec x_1$ [i.e. the deviation from the BI path induced by F -destabilization occurs at node x_1]. For any $a_i \in A(x_1)$, either $a_i(x_1) \in Z$ or $\iota(a_i(x_1)) \neq i$, since the game has no trivial actions. By construction, $a_{x_1}^\circ(x_1) \in Z$. By assumption, $\forall(t, T^i \in \text{supp}(\lambda_i^t)), A(x_1) \subset R^{x_1}$. So if $a_i(x_1) \in Z$, then $u_{T^i}(a_i|\psi, \mu) = u_i^\circ(a_i(x_1)), \forall \psi, \mu$. In case $a_i(x_1) \notin Z$, if $\forall(t, T^i \in \text{supp}(\lambda_i^t)), u_{T^i}(a_i|\psi, \mu) \square u_i^\circ(a_{x_1}^\circ)$, then the backward induction path is not destabilized by a_i and we can take another action. If instead, at some time $t, \exists T^i \in \text{supp}(\lambda_i^t)$ s.t. $u_{T^i}(a_i|\psi, \mu) > u_i^\circ(a_{x_1}^\circ)$, then $\exists z' : a_i(x_1) \prec z' : \mu(z') > 0, u_i(z') > u_i^\circ(a_{x_1}^\circ) = u_i^\circ(x_1)$. Moreover, as $u_i^\circ(a_i(x_i)) < u_i(z')$, there exist $z \neq z'$ s.t. $u_i^\circ(z) = u_i^\circ(a_i(x_i))$. Let x as the last node s.t. $x \prec z$ and $x \prec z'$ and let $a \in A(x)$ s.t. $a(x) \prec z$, it follows that $a \neq a_x^\circ$. We have shown that if $f_x^*(a_x^\circ) = 1$ for all nodes $x \prec x_1$, then $[\exists a_i \in A(x_1) : a_i \neq a_{x_1}^\circ]$, $[\exists a \in A(x) : a_i(x_1) \preceq x, a \neq a_x^\circ]$, and $[\exists z' : a(x) \preceq z', u^\circ(z') > u^\circ(x_i)]$.

Consider now x_2 let $\iota(x_2) =: j$, and suppose that $f_x^*(a_x^\circ) = 1$ for all nodes $x \prec x_2$. If $\forall(t, T^j \in \text{supp}(\lambda_j^t)), a_{x_1}^\circ \in R^{x_2}$, then the argument goes as above. To let there be a $t, T^j \in \text{supp}(\lambda_j^t) : a_{x_1}^\circ \notin R^{x_2}$, there must exist $\tau < t$, and a node $x_l \prec x_2, [\iota(x_l) =: k]$ and a type $T^k \in \text{supp}(T_k^\tau)$ s.t. $A_{T^k}(x_l|\mu, \psi) \neq \{a^\circ(x_k)\}$. If not, as $\rho_\star^0 = 1, f^0(a^\circ) = 1$, it follows that the BI path is played from x_0 to x_2 . As for any $T^j \in \text{supp}(\lambda_j^\tau), a_{T^j}(x_2|\mu, \psi) = a_{x_2}^\circ$, and for any $R^i \in \text{supp}(\rho_i^\tau), a_{T^j}(x_1|\mu, \psi) = a_{x_1}^\circ$, it follows that f^t reaches $a_{x_1}^\circ$. As $a_{x_1}^\circ \in Z$, we conclude that $\zeta_{T^j}^{t-1}(a_{x_1}^\circ) = 1$ [type T^j observed $a_{x_1}^\circ$ at time $t-1$] and thus $a_{x_1}^\circ \in R^j, \forall R^j \in \text{supp}(\rho_k^t)$. That is to say: to have a deviation from BI path at node x_2 at time t , there must

have been a deviation on node $x_k \prec x_j$ at time $\tau < t$. Thus node x_2 is irrelevant and we can proceed directly with node x_3 .

Let $k := \iota(x_3)$. If $\forall (R^j, \hat{\sigma}_j) \in \text{supp}(\psi), a_{x_1}^\circ \in R^j$ (player k believes to play against a type T^j s.t. $a_{x_1}^\circ \in R^j$), then we can apply the same argument used for node x_2 to the two cases $a_{x_1}^\circ \notin R^{x_3}$, and $a_{x_2}^\circ \notin R^{x_3}$, to conclude that node x_3 is irrelevant and move further up. If instead $\exists (R^j, \hat{\sigma}_j) \in \text{supp}(\psi)$ s.t. $a_{x_1}^\circ \notin R^j$, such types do not play $a_{x_2}^\circ$, and may have enough weight to make k deviate from $a_{x_3}^\circ$. However, let τ' be the first time in which $\exists (R^j, \hat{\sigma}_j)^j \in \text{supp}(\psi)$ s.t. $a_{x_1}^\circ \notin R^j$. As ψ is generated by $\lambda^{\tau'}$, $\exists (R^j, \hat{\sigma}_j) \in \text{supp}(\lambda^{\tau'})$ s.t. $a_{x_1}^\circ \notin R^j$. Analogously as before, there must exist a time $\tau < \tau'$ and a node $x_l \prec x_2$, ($\iota(x_l) =: l$) and a type $T^l \in \text{supp}(\lambda_l^\tau)$ s.t. $A_{T^l}(x_l | \mu, \psi) \neq \{a^\circ(x_l)\}$. In case $x_l = x_3$, since τ' is the *first* time at which $\exists (R^j, \hat{\sigma}_j) \in \text{supp}(\psi)$ s.t. $a_{x_1}^\circ \notin R^j$, k cannot have deviated from $a_{x_3}^\circ$ because she believed to play against a type R^j unaware of $a_{x_1}^\circ$. Again the argument showed for node x_1 holds. If $x_l \prec x_3$, then again x_3 is irrelevant, and we can move further up.

Repeating the argument by (backward) induction on $x_j : u^\circ(x_j) = u^\circ(x_0)$, the first claim is proved. ■

The second claim shows that for the backward-induction path to be F -destabilized, and the new path to go through a_i , there must exist a deviation (in the continuation of a_i) that leads into a path where no further deviations from backward induction occur.

Claim 2 *Suppose that the BI path $a^\circ(\Gamma)$ is F -destabilized, and $f_x^*(a_x^\circ) = 1, \forall x \prec x_i$, but $f_{x_i}^*(a_i) > 0$, x_i on path, $a_i \neq a_{x_i}^\circ$. Then $\exists a_j \in A(x_j)$ [$a_i(x_i) \preceq x_j$ and $a_j \neq a_{x_j}^\circ$] s.t. $u_j^\circ(a_j(x_j)) > \bar{u}_j(a_{x_j}^\circ(x_j))$.*

Proof. To show that if there is not any $a_j \in A(x_j)$ [$a_i(x_i) \preceq x_j$ and $a_j \neq a_{x_j}^\circ$] s.t. $u_j^\circ(a_j(x_j)) > \bar{u}_j(a_{x_j}^\circ(x_j))$, then $\#(t \sqcap T : f_{x_i}^*(a_i) = 0) / T \rightarrow 1$, for $\pi^n \rightarrow 0$.

We will consider all nodes $x_j : a_i(x_i) \preceq x_j$, and look for times t and types $T^j \in \text{supp}(\lambda_j^t)$ and ψ, μ s.t. $A_{T^j}(x_j | \psi, \mu) \neq \{a_{x_j}^\circ\}$. If we could not find any such x_j , by claim 1, it would follow that $f_{x_i}^t(a_i) = 0$. First note that since own actions cannot be forgotten, $\forall [x_j, T^j, \psi, \mu, a \in A(x_j)], u_{T^j}(a(x_j) | \psi, \mu) \geq \bar{u}_j(a(x_j))$.

Start with any $x_1 \in Y_1, a_i(x_i) \preceq x_1$ [the last decision nodes in the continuation of a_i], denote $l := \iota(x_1)$. Again, as own action cannot be forgotten, $\forall [t, T^l \in \text{supp}(\lambda_l^t)], A(x_1) \subset R^{x_1}$. Since $a_l(x_1) \in Z$, $u_{T^l}(a_1 | \psi, \mu) = u_i^\circ(a_1(x_1)), \forall \psi, \mu$. Thus $a_{T^l}(x_1 | \psi, \mu) = a_{x_1}^\circ, \forall \psi, \mu$. So

that in any case, x_1 is irrelevant and we can proceed with other nodes $x_1 \in Y_1$ and with other sets $Y_k, k = 2, \dots, K$. Consider now any $x_2 \in Y_2, a_i(x_i) \prec x_2$ let $\iota(x_2) = j$. If $\forall a_j \in A(x_2) [a_j \neq a_{x_2}^\circ], u_{T^j}(a_j|\psi, \mu) < \bar{u}_j(a_{x_2}^\circ(x_2)) \square u_{T^j}(a_{x_2}^\circ(x_2)|\psi, \mu)$, then $a_{T^j}(x_2|\psi, \mu) = a_{x_2}^\circ$. So if $\forall [t, T^j \in \text{supp}(\lambda_j^t), \mu, \psi, a_j \in A(x_2)], u_{T^j}(a_j|\psi, \mu) < \bar{u}_j(a_{x_2}^\circ(x_2))$ the node x_2 is irrelevant and we can proceed with other nodes in Y_2 and with other sets $Y_k, k = 3, \dots, K$. So say that $\exists [t, T^j \in \text{supp}(\lambda_j^t), \mu, \psi]$, s.t. $a_j \in A(x_2), u_{T^j}(a_j|\psi, \mu) > u_{T^j}(a_{x_2}^\circ(x_2)|\psi, \mu) \geq \bar{u}_j(a_{x_2}^\circ(x_2))$. If $u_{T^j}(a_j|\psi, \mu) \neq u_j^\circ(a_j(x_2))$, then $u_j^\circ(a_j(x_2)) > \bar{u}_j(a_{x_2}^\circ(x_2))$, and the claim is proved taking $x_j = x_2$. So suppose that $u_{T^j}(a_j|\psi, \mu) > u_j^\circ(a_j(x_2))$, and that $u_j^\circ(a_j(x_2)) < \bar{u}_j(a_{x_2}^\circ(x_2))$. For clarity, we remind that we want to see whether a_j can upset the backward induction path through a_i .

Say that $a_j \in A_{T^j}(x_2|\psi, \mu)$, and that $\exists T^i \in \text{supp}(\lambda_i^t)$, s.t. $a_{T^i}(x_i|\psi, \mu) = a_i$, where μ and ψ are induced by the APBE strategy $\sigma_{T^j}(x_2|\psi, \mu)$, and by rational choices at nodes $x_k : x_i \prec x_k \prec x_j$. There are two cases: either x_j is reached by f^t , or there exists $x_k, T^k \in \text{supp}(\lambda_k^t)$, s.t. $a_{T^k}(x_k|\psi, \mu) \neq a_{x_k}^\circ$, where μ and ψ are induced by $a_{T^j}(x_2|\psi, \mu) = a_j$, and by rational choices at nodes $x : x_k \prec x \prec x_j$. The latter case is equivalent to the situation in which $R^{x_k} \cap A(x_2) = \{a_j\}$ and $A(x) \subset R^{x_k}, \forall x : x_k \prec x \prec x_j$. Therefore, it can be subsumed in the analysis for node x_k : it is *as if* x_2 did not cause f^t to deviate along a_i so that we can proceed with another node. Whenever x_2 is reached by the path of play f^t , any type $T^j \in \text{supp}(\lambda_j^t)$ that is called to play at x_2 and plays a_j will observe the choice at node $a_j(x_2)$ [let $\iota(a_j(x_2)) = l$]. Since $a_j(x_2) \in Y_1$, as proved before, for any $[T^l \in \text{supp}(\lambda_l^t), \mu, \psi]$, $a_{T^l}(a_j(x_2)|\psi, \mu) = a_{a_j(x_2)}^\circ$. So that $\zeta_{T^j}^t(a_{a_j(x_2)}^\circ) > 0$, and with positive probability players of type T^j , assume at time $t+1$ the model $\hat{R}^j : a_{a_j(x_2)}^\circ \in \hat{R}^{x_2}$, so as to be grouped in type $\hat{T}^j$. Since $u_{\hat{T}^j}(a_j|\psi, \mu) = u_j^\circ(a_j(x_2)) < \bar{u}_j(a_{x_2}^\circ(x_2))$, $a_j \notin A_{\hat{T}^j}(x_k|\psi, \mu)$.

Now, let $T[a_j] := \{R^j : a_{T^j}(x_k|\psi, \mu) = a_j, \text{ for some } \psi, \mu\}$. For π^n small enough, the transition function ξ_τ implies that $\rho_j^\tau(R_{a_j})$ is decreasing in $\bar{\tau} > t$, and, if x_2 were to be reached on path forever, $\lambda_j^\tau(T[a_j]) \rightarrow 0$, for $\bar{\tau} \rightarrow \infty$. Since the game is generic, there exists an $\varepsilon > 0$ s.t. if $\lambda_j^\tau(T[a_j])_{(n)} < \varepsilon$, the effect of the R_{a_j} types is irrelevant on the payoff of choices $a_k [a_k(x_k) \preceq x_2]$. Thus there exists a finite $\bar{\tau}$ s.t. a_j cannot upset the BI path at time $\bar{\tau}$. There could be other $\hat{a}_j$ s.t. $\hat{a}_j \in A_{\hat{T}^j}(x_2|\psi, \mu)$. Yet, repeating the analysis for action a_j , one sees that at time $t+2$, each player of type $\hat{T}^j$ would assume model $\tilde{R}^j : a_{\hat{a}_j(x_2)}^\circ \in \tilde{R}^{x_2}$. Therefore, for π^n small enough, there exists a finite time $\bar{\tau}$ s.t. $\lambda_j^\tau(GT) \geq \prod_{A(x_2)} (1 - \pi_a) \approx 1$ where $GT := \{T^j : A_{T^j}(x_j|\psi, \mu) \neq \{a_{x_j}^\circ\}, \forall \psi, \mu\}$. Fix π^N

small, and call $\tau(N)$ the first (finite) time s.t. x_2 is reached on path. If $\tau(N)$ does not exist for any N , we can move to a different node. If $\tau(N)$ exists, for any fixed $t < \tau(N)$, the transition function γ implies that $\lambda_j^t(GT)_{(n)} \rightarrow 1$ for $N < n \rightarrow \infty$ ($\pi^n \rightarrow 0$ with n). Since the game is generic, there exists an $\varepsilon > 0$ s.t. if $\lambda_j^t(GT)_{(n)} < \varepsilon$, the effect of the RG types is irrelevant on the payoff of choices a_k ($a_k(x_k) \preceq x_2$). Thus, $\tau(N) \rightarrow \infty$ for $\pi^N \rightarrow 0$.

Denote by $\tau_2(N)$ the second time s.t. x_2 is reached on path. Above we proved that for small π^N , at time $\tau(N) + T$, the ratio of “bad” types $\lambda_j^{\bar{\tau}}(GT) \approx 1$. That means that at time $\tau(N) + \bar{\tau}$, a_j does not perturb the backward induction path. Thus we can substitute time $\tau(N) + \bar{\tau}$ for time 0, and time $\tau_2(N)$ for time $\tau(N)$, to see that $\tau_2(N) - \tau(N) \rightarrow \infty$ for $\pi^N \rightarrow 0$. Going on, for any k , $\tau_k(N) - \tau_{k-1}(N) \rightarrow \infty$. Thus, if a_j were to perturb a° , that would be a small-time deviation: $\#(\{t \sqcap \bar{\tau} : f^t \neq a^\circ\})/\bar{\tau} \rightarrow 0$, for $\pi^n \rightarrow 0$.

Having completed with the nodes in Y_2 , take a node $x_3 \in Y_3, a_i(x_i) \prec x_3$ let $\iota(x_3) = k$, and consider $a_k \in A(x_3)$. The argument seen on node x_2 may be repeated for x_3 until considering the case when $u_{\hat{T}^k}(a_k|\psi, \mu) > u^\circ(a_k(x_3))$ and $f^t(a_k(x_3)) > 0$. With a minor notational violation, let $x_2 := a_k(x_3)$, $x_1 := a_{x_2}^\circ(x_2)$, $z = a_{x_1}^\circ(x_1)$, and $j := \iota(x_2)$. If $\zeta_{T^k}(z) > 0$, the last part of the argument for node x_2 may be repeated without further ado. If instead $\zeta_{T^k}(z) = 0$, then note that $\forall [T^j \in \text{supp}(\lambda_j^t), \mu, \psi]$, $A_{T^j}(x_2|\mu, \psi) \neq \{a_{x_2}^\circ\}$. As $a_k \in A_{T^k}(x_3|\mu, \psi)$, and x_3 is reached, then also x_2 is reached. Since $x_2 \in Y_2$, $\zeta_{T^j}(z) > 0$, and we have proved above that in finite time the BI choice at x_2 is reestablished. Yet, whenever $\exists T^j \in \text{supp}(\lambda_j^t), \mu, \psi$, s.t. $A_{T^j}(x_2|\mu, \psi) = \{a_{x_2}^\circ\}$, $\zeta_{T^k}(z) > 0$, and so we can reestablish in finite time also the BI choice at x_3 . Repeating the argument by (backward) induction on $x_l \in Y_l$, the second claim is proved. ■

The last claim simply means that if the backward induction path is destabilized through action a_i , to reach the last deviation a_j , it must be the case that all the choices from x_i onwards lead to a_j .

Claim 3 *Suppose that the BI path $a^\circ(\Gamma)$ is F -destabilized as follows. $f_x^\circ(a_x^\circ) = 1, \forall x \prec x_i$, $f_{x_i}^*(a_i) > 0$, x_i on path, $a_i \neq a_{x_i}^\circ$, $f_{x_j}^*(a_j) > 0$, $a_i(x_i) \preceq x_j$, $a_j \neq a_{x_j}^\circ$, and $f_x^*(a_x^\circ) = 1, \forall x : a_j(x_j) \prec x$. Then $u_i^\circ(a_j(x_j)) > u_i^\circ(x_i)$ and $\forall x : x_i \prec x \preceq x_j$, $[u_{\iota(x)}^\circ(a_j(x_j)) \geq \bar{u}_{\iota(x)}(a'(x)), \forall a' \in A(x)]$.*

Proof. If it is not the case that $u_i^\circ(a_j(x_j)) > u_i^\circ(x_i)$ then $\forall [T^i \in \text{supp}(\lambda_i^t), \mu, \psi]$, $u_{T^i}(a_{x_i}^\circ|\psi, \mu) > u_{T^i}(a_i|\psi, \mu)$ and thus the BI path does not deviate at node x_i . Pick

any $x : x_i \prec x \preceq x_j$. Pick $a \in A(x)$ s.t. $a(x) \preceq a_j(x_j)$. If it is not the case that $u_{\iota(x)}^\circ(a_j(x_j)) \geq \bar{u}_{\iota(x)}(a'(x)), \forall a' \in A(x)$, then $\exists a' \in A(x)$ s.t. $\forall [T^i \in \text{supp}(\lambda_i^t), \mu, \psi]$, $u_{T^i}(a'|\psi, \mu) > u_{T^i}(a|\psi, \mu)$. In which case, f^* cannot reach $a_j(x_j)$. ■

That concludes the proof of Necessity ■

Part 2 Sufficiency

Proof. Pick one arbitrary pair of deviations $(x_i, a_i), (x_j, a_j)$ as characterized in Definition 5. We shall construct an appropriate sequence of matrices π^n and an offsetting path of simple inertial ABE where each player with the same frame take the same strategy.

For any node $x : a_i(x_i) \preceq x \preceq x_j$, let $A'(x) := \{a \in A(x) | a(x) \not\preceq a_j(x_j)\}$, [the set of actions that do not lead into $a_j(x_j)$], and for any $a' \in A'(x)$, let $\bar{a}(a'(x))$ be the $\iota(x)$ BI-maximin path starting at node $a'(x)$. Let $A''(x) := \{a \in A(x) | \exists a' \in A'(x), a'(x) \prec a(x'), \iota(x') \neq \iota(x), \text{ and } a \notin \bar{a}(a'(x))\}$ [the actions alternative to the BI-maximin path of actions in $A'(x)$]. Enumerate the nodes $x : a_i(x_i) \preceq x \preceq x_j$, in natural number index reverting the $\prec$ order, let node $x_j := x_1$ and let K denote the index of node $a_i(x_i)$. For any k , let $m_k := \sum_{0 \leq l < k} \#(A''(x_k))$ (with $A''(x_0) := \emptyset$). For any n construct π^n as follows. For any $k : 1 \leq k \leq K$, let $\pi_a^{\iota(x_k)} := 1/n^{1+m_k}$ for any $a \in A''(x_k)$. For any action $a \in A \setminus (\cup_{k=1}^K A''(x_k))$ set $\pi_a^i = 0, \forall i \in I$. By Definition 5, all actions a above characterized are off the BI path a° . As $f^0 = a^\circ$, at time 0, they are unobserved and may be forgotten.

First we analyze the behavior of the population when $f_{x_i}^t(a_i) = 0$ (the nodes in the continuation of a_i are not reached). For any n , the awareness types distribution at node x_j evolves as follows: $\forall A''' \subseteq A''(x_j), \rho^1(A \setminus A''') = \prod_{a \in A'''} \pi_a, \rho^1(A) = 1 - \sum_{A''' \subseteq A''(x_j)} \rho^1(A \setminus A''')$, and $\forall t \geq 1, \rho^{t+1}(A \setminus A''') = \rho^t(A \setminus A''') + \sum_{A''' \subset A'''' \subseteq A''(x_j)} \rho^t(A \setminus A''') \prod_{a \in A'''' \setminus A'''} \pi_a$. So $\rho_j^1(A \setminus A''') = o(1/n^{\#(A''')-1})$, but $\exists \tau_1(n)$ s.t. $\forall t > \tau_1(n), \rho_j^t(A \setminus A''(x_1)) \approx 1$. By Requirement (2) in Definition 5, $\forall (\mu, \psi), a_{(A \setminus A''(x_1), \mu, \psi)}^\circ(x_j) = a_j$.

Consider now x_2 . $\forall a \in A''(x_2), \pi_a^{\iota(x_2)} := 1/n^{1+m_2}$, It follows that $\forall A''' \subseteq A''(x_2) \rho_{\iota(x_2)}^{\tau_1(n)}(A \setminus A''') = o(1/n^{l-1})$ where $l \geq \#(A''') + 1$. For n large enough, that is to say, $\rho_{\iota(x_2)}^{\tau_1(n)}(A) \approx 1$. Yet, the evolution of the distribution at node x_2 is the same as at node x_j , so there exists a time $\tau_2(n)$ s.t. $\forall t > \tau_2(n), \rho_j^t(A \setminus A''(x_1)) \approx 1$. Above we proved that at such t , nearly all types $R^j \in \rho_j$ play a_j at node x_j , so, by Requirement (2) in Definition 5, and by the

genericity of Γ , it follows that $a_{(A \setminus A''(x_2), \mu, \psi)}^\circ(x_2) = a_2 : a_2(x_2) = x_j$, for ψ induced by $\rho_j^t(A \setminus A''(x_1)) \approx 1$ and for μ induced by $a_{(A \setminus A''(x_1), \mu, \psi)}^\circ(x_j) = a_j$.

Repeating the analysis for all nodes $x_3, \dots, x_K$ we establish that for n large enough, there exist a collection of finite times $(\tau_1(n), \tau_2(n), \dots, \tau_k(n))$ s.t. that $\forall k, \forall t > \tau_k(n)$, $\rho_{\iota(x_k)}^t(A \setminus A''(x_k)) \approx 1$ and $a_{(A \setminus A''(x_k), \mu, \psi)}^\circ(x_k) = a_k : a_k(x_k) \preceq a(x_j)$. For any time $t > \tau_K(n)$, that determines μ and ψ s.t. by Requirement (1) in Definition 5, $a_{(A, \mu, \psi)}^\circ(x_i) = a_i$. So $f^t \neq a^\circ$, in particular, let z the terminal node reached by the backward induction path starting at node $a_j(x_j)$ it is the case that $\forall x_k, \zeta_{A \setminus A''(x_k)}^t(z) \approx 1$. As by construction none of the actions $a \in A''(x_k)$ is s.t. $a(x) \prec z$, none of the forgotten actions will ever be observed. So for any k , $\rho_{\iota(x_k)}^t(A \setminus A''(x_k))$ is non-decreasing in t . Thus the deviation at f^t is stable: $f^*(z) > 0$, which concludes that $f^* \neq a^\circ$.

Now we consider the case when some node x_k is reached by f^t before $\tau_K(n)$. Start with node x_j , and say that f^t reaches x_j before than $\tau_K(n)$. Let $\bar{\tau}$ the first time that f^t reaches x_j . If $\bar{\tau} > \tau_1(n)$, then the event has no effect on the analysis for node x_j , as $\rho_j^t(A \setminus A''(x_j)) \approx 1$ and $\zeta_{A \setminus A''(x_j)}^t(z) \approx 1$. So say that $\bar{\tau} < \tau_1(n)$. For any $k > 1$, $\rho_{\iota(x_k)}^{\bar{\tau}}(A) \approx 1$, therefore $f^{\bar{\tau}}$ may reach x_j , only because $\exists A''' \subset A''(x_j)$ s.t. $a_{(A \setminus A''', \mu, \psi)}^\circ(x_j) =: a'_j \neq a_j^\circ$, and such types induce ψ, μ so as to make all players at nodes $x : x_i \preceq x \prec x_j$ to switch to action $a : a(x) \preceq x_j$. Now, if $\exists \tau > \bar{\tau}$ s.t. $f^\tau = a^\circ$, then our conclusion is reached.

Let z' be the terminal node reached by $f^{\bar{\tau}}$. To let there be $\tau > t : f^\tau = a^\circ$, it must be the case that $\exists a \in A''(x_j)$ s.t. $a_{(A \setminus (A''' \cup \{a\}), \mu, \psi)}^\circ(x_j) = a_j^\circ$. As $\zeta_{A \setminus A'''}^{\bar{\tau}}(z) \approx 1$, moreover, it is the case that such $\tau = \bar{\tau} + 1$ and that $\rho_j^{\bar{\tau}}(A \setminus A''') \approx 0$. The node x_j may be reached several more times $(\bar{\tau}_2, \bar{\tau}_3, \dots)$, and again $\rho_j^{\bar{\tau}_k+1}(A \setminus A''') \approx 0$. However, for any $t \neq \bar{\tau}_k + 1$, $\Delta(\rho_j^{\bar{\tau}}(A \setminus A''(x_j))) > (1 - \rho_j^{\bar{\tau}}(A \setminus A''(x_j))) \prod_{a \in A''(x_j)} \pi_a$, and for any $t = \bar{\tau}_k + 1$, $\Delta(\rho_j^{\bar{\tau}}(A \setminus A''(x_j))) \geq 0$. That directly implies that, for n large enough $\exists \tau'_1(n)$ s.t. $\forall t > \tau'_1(n)$, $\rho_j^t(A \setminus A''(x_j)) \approx 1$. Therefore, if f^t is to reach x_j for $t > \tau'_1(n)$, then f^t will also reach $a_j(x_j)$. Analogously as the argument for $\tau_1(n)$, also $\rho_{\iota(x_2)}^{\tau'_1(n)}(A) \approx 1$.

Proceeding analogously at node x_j , we find a $\tau'_2(n)$, and iterating for $k > 2$, we find $\tau'_k(n)$, that satisfy the same conditions of the $\tau_k(n)$ nodes found at previous passage. ■ ■

Proof of Proposition 8. We say that the game Γ' is obtained from Γ by the addition of the *trivial* action $a \in A'$ if and only if $\exists z \in Z, x' \in X', (X, Z \setminus \{z\}, \prec |_{X \cup Z \setminus \{z\}}, I, \iota, A, u|_{Z \setminus \{z\}}) = (X' \setminus \{x'\}, Z' \setminus \{a(x')\}, \prec' |_{X' \setminus \{x'\} \cup Z' \setminus \{a(x')\}}, I', \iota'|_{X' \setminus \{x'\}}, A' \setminus \{a\})$,

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

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 and only if $(X, Z, \prec, I, \iota, A, u) = (X', Z' \setminus \{a'(x)\}, \prec' |_{X' \cup Z' \setminus \{a'(x)\}}, I', \iota', A \setminus \{a'\}, u'|_{Z \setminus \{a'(x)\}})$ 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 $\{\Gamma_k\}_{k=1}^K$ s.t. $\forall k, \Gamma_k$ can be obtained from Γ_{k-1} by the addition of a trivial or conditionally dominated action, and $\Gamma_1 = \Gamma, \Gamma_K = \Gamma'$, then the game Γ' is a *backward-induction-irrelevant expansion* of Γ .

Take any game Γ , consider the BI set ordering: take x_i s.t. $u^\circ(x_i) = u^\circ(x_0)$ and $\forall x_j : u^\circ(x_j) = u^\circ(x_0), x_j \in Y_{k_j}, k_i < k_j$ (x_i is the last node on BI path). Relabel $a^\circ(x_i)$ as B and $\iota(x_i)$ as player 1. Consider the sequence of actions $\{Z, Q, T, H, C, A, F\}$ and add them sequentially to the game Γ , in the following manner.

For the addition Z , say that $x_1 \xrightarrow{B} x_2 : \iota(x_2) = 2, x_2 \xrightarrow{Z} z_Z, u(z_Z) = u(a^\circ(x_i))$, thus Z is a trivial addition. For $Q : x_2 \xrightarrow{Q} z_Q, u(z_Q) < u(a^\circ(x_i))$, so Q is a strictly dominated (at node x_2) addition. Let $T : x_1 \xrightarrow{B} z_T, u_1(z_T) < u_1(a^\circ(x_i))$ and $u_2(z_T) > u_2(a^\circ(x_i))$, thus T is strictly dominated addition. Let $C : x_1 \xrightarrow{T} x'_2 : \iota(x'_2) = 2, x'_2 \xrightarrow{C} z_C, u(z_C) = u(z_T)$, so C is a trivial addition. Let $H : x'_2 \xrightarrow{H} z_H, u_1(z_H) > u_1(a^\circ(x_i))$ and $u_2(a^\circ(x_i)) < u_2(z_H) < u_2(z_C)$, so H is strictly dominated addition. Let $A : x'_2 \xrightarrow{Y} x'_1 : \iota(x'_1) = 1, x'_1 \xrightarrow{A} z_A, u(z_A) = u(z_C)$, thus A is a trivial addition. Let $F : x'_1 \xrightarrow{F} z_F, u(z_F) < u(z_A)$ so F is strictly dominated addition.

Then proceed exactly as in Example 6. ■

References

- [1] Alos-Ferrer [1999]: “Dynamical Systems with a Continuum of Randomly Matched Agents”, *Journal of Economic Theory* **86**: 245-267
- [2] Aumann R. and A. Brandeburger [1995]: “Epistemic Conditions for Nash Equilibrium” *Econometrica* **63**: 1161-80
- [3] Barsalou L. W. [1992]: *Cognitive Psychology - An Overview for Cognitive Scientists* Hillsdale NJ, Lawrence Erlbaum As. Pu.

- [4] Battigalli P. [1997]: “On Rationalizability in Extensive Games” *Journal of Economic Theory* **74**: 40-61
- [5] Bernheim D. [1984]: “Rationalizable Strategic Behaviour” *Econometrica* **52**: 1007-1028
- [6] Boylan R. [1992]: “Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals” *Journal of Economic Theory* **57**: 473-504
- [7] Dekel E. Fudenberg D. and D. Levine [1999]: “Payoff Information and Self-Confirming Equilibrium” *Tel Aviv University* WP 9-99
- [8] Dekel E. Lippman B. and A. Rustichini [1998]: “Standard State-Space Models Preclude Unawareness” *Econometrica* **66**: 159-173
- [9] Diestel R. [1997]: *Graph Theory* Springer, New York
- [10] Dudley R. M. [1989]: *Real Analysis and Probability* Chapman and Hall, New York
- [11] Durrett R. [1996]: *Probability Theory and Examples* Belmont CA, Wadsworth Pu.
- [12] Fudenberg D. Kreps D. M. and D. K. Levine [1988]: “On the Robustness of Equilibrium Refinements” *Journal of Economic Theory* **44**: 354-380
- [13] Fudenberg D. and D.K. Levine [1993a]: “Self-Confirming Equilibrium” *Econometrica* **61**: 523-545
- [14] Fudenberg D. and D.K. Levine [1993b]: “Steady State Learning and Nash Equilibrium” *Econometrica* **61**: 547-573
- [15] Fudenberg D. and D.K. Levine [1998]: *The Theory of Learning in Games*, MIT Press
- [16] Fudenberg D. and Tirole J. [1991]: *Game Theory* Cambridge MA, MIT press
- [17] Hale J. [1969]: *Ordinary Differential Equations*, Wiley, NY
- [18] Harsanyi J. [1967]: “Games with Incomplete Information Played by ‘Bayesian’ Players, part I, II, III” *Management Science* **14**: 159-182, 320-334, 486-502

- [19] Jackson M. and E. Kalai [1995]: “Social Learning in Recurring Games” *Kellogg G. S. M. at Northwestern University* mimeo
- [20] Kohlberg E. and J.F. Mertens [1986]: “On the Strategic Stability of Equilibria” *Econometrica* **54**: 1003-1038
- [21] Kreps D. [1990] “Corporate Culture and Economic Theory”, in *Perspectives on Positive Political Economy*, J. Alt and K. Shepsle eds. Cambridge University Press
- [22] Kreps D. and Wilson R. [1982] “Reputation and Imperfect Information” *Journal of Economic Theory* **27**: 253-279
- [23] Milgrom P. and J. Roberts [1982] “Limit Pricing and Entry under Incomplete Information”, *Econometrica* **53**: 889-904
- [24] Modica S, A. Rustichini and J.M. Tallon [1998]: “Unawareness and Bankruptcy: A General Equilibrium Model” *Economic Theory* **12**: 259-92
- [25] Myerson R. [1978]: “Refinements of the Nash Equilibrium Concept” *International Journal of Game Theory* **7**: 73-80
- [26] Nash J. [1950]: “Equilibrium Points in n-Persons Games” *Proceedings of the National Academy of Science* **36**: 48-49
- [27] Pearce D. [1984]: “Rationalizable Strategic Behavior and the Problem of Perfection” *Econometrica* **52**: 1029-1050
- [28] Schlag K. [1998]: “Why Imitate, and If So, How?” *Journal of Economic Theory* **78**: 130-156
- [29] Samuelson L. [1997]: *Evolutionary Games and Equilibrium Selection* MIT Press
- [30] Selten R. [1975]: “A Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games” *International Journal of Game Theory* **4**: 25-55
- [31] Squintani F. [1999]: “The Evolution of Awareness of Actions with Infinitesimal Forgetfulness.” *Northwestern University* mimeo
- [32] Weibull J. W. [1995]: *Evolutionary Game Theory* Cambridge MA, MIT press