

Epistemic foundation of the backward induction paradox*

Geir B. Asheim^a

Thomas Brunnschweiler^b

25 August 2022

Abstract

After having observed a deviation from backward induction, a player might deem the opponent prone to deviate from backward induction again, making it worthwhile to deviate themselves. Such reaction might make the deviation by the opponent worthwhile in the first place—which is the backward induction paradox. This argument against backward induction cannot be made in games where all players move only once. While strategic-form perfect equilibrium yields backward induction in games where players move only once but not necessarily otherwise, no existing non-equilibrium concept captures the backward induction paradox by having these properties. To provide such a concept, we define and epistemically characterize the Independent Dekel-Fudenberg Procedure. Since beliefs are modelled by non-Archimedean probabilities, meaning that some opponent choices might be assigned subjective probability zero without being deemed subjectively impossible, special attention is paid to the formalization of stochastically independent beliefs.

JEL Classification No.: C72.

Keywords: Perfect information games, backward induction paradox, non-Archimedean probabilities, stochastic independence.

*We thank Andrés Perea and Martin Dufwenberg for many helpful suggestions. This paper has been developed from Brunnschweiler's master thesis at the University of Oslo.

^aDepartment of Economics, University of Oslo, P.O. Box 1095 Blindern, 0317 Oslo, Norway
(*Email:* g.b.asheim@econ.uio.no).

^bDepartment of Economics, University of Oslo, P.O. Box 1095 Blindern, 0317 Oslo, Norway
(*Email:* thomas.brunnschweiler@outlook.com).

1 Introduction

After having observed a deviation from backward induction in a finite extensive game, like the *centipede game* or the *finitely repeated prisoners' dilemma*, a player might deem the opponent prone to deviate from backward induction again. If the player believes with sufficient subjective probability in this possibility, it might be worthwhile for the player to deviate from backward induction themselves. In turn, such reaction, if predicted, can provide a reason for the opponent to deviate from backward induction in the first place. This is the *backward induction paradox* as introduced by Basu (1988) and Reny (1985, 1988) and discussed by, among others, Binmore (1987, Section 3), Pettit and Sugden (1989), and Sobel (1993); see also Luce and Raiffa (1957, pp. 80–81) for an early illustration of a related point and Mas-Colell, Whinston, and Green (1995, p. 282) for a textbook treatment.

As pointed out by Dufwenberg and Van Essen (2018, p. 126), this argument against backward induction cannot be made in games where all players move only once. *Strategic-form perfect equilibrium* captures this by yielding backward induction in perfect information games where all players move only once, but not necessarily in games where some player moves more than once. Indeed, Selten (1975) ensures that his concept of *extensive-form perfect equilibrium* leads to backward induction by applying strategic-form perfect equilibrium to the agent-strategic form, where each player moves only once. However, backward induction is not an equilibrium concept, but a procedure that corresponds to increasing levels of reasoning.

Therefore, to offer an epistemic foundation of the backward induction paradox, we provide a non-equilibrium concept, supported by epistemic modeling, that yields the backward induction strategies in perfect information games where all players move only once, but not necessarily the backward induction outcome in games where some player moves more than once. To the best of our knowledge, no previously existing epistemic model solves the backward induction paradox in the sense of yielding such a non-equilibrium concept. In particular, the *Dekel-Fudenberg Procedure* (Dekel and Fudenberg, 1990)—which consists of one round of elimination of weakly dominated strategies, followed by subsequent rounds of elimination of strictly dominated strategies—does not even yield backward induction outcomes in perfect information games where all players only move once, while *sequential/quasi-perfect/proper rationalizability* (Dekel, Fudenberg, and Levine, 1999, 2002; Schuhmacher, 1999; Asheim and Perea, 2005) always yield the backward induction strategies.¹ Moreover, *extensive form rationalizability* (Pearce, 1984) leads to the backward induction outcome in perfect information games with no relevant payoff ties, independently of whether players move more than once (Battigalli, 1997, Thm. 4; Battigalli and Siniscalchi, 2002, Prop. 8).

We define a refinement of the Dekel-Fudenberg Procedure, called the *Independent Dekel-Fudenberg Procedure*, which requires that each player has stochastically independent beliefs about

¹The latter property is also shared by *common belief in future rationality* (Perea, 2014; see also Asheim, 2002), characterized by a backward dominance procedure, and *backward rationalizability* (Perea, 2014; Penta, 2015).




Figure 1: A centipede game and its corresponding strategic form.

the strategy choices of their opponents. As will be illustrated in the game of Figure 2 in Section 2, with such uncorrelated beliefs a player cannot infer anything about the future play of other players by observing the past play of different players.² However, the stochastic independence only concerns “inter-player” inference, not “intra-player” inference, meaning that players can learn about the behavior of opponents, if these opponents are to move more than once. The Independent Dekel-Fudenberg procedure formalizes the backward induction paradox, since it has the feature that, for each player, only the backward induction strategy survives the procedure in games without relevant payoff ties if players move only once, while outcomes incompatible with backward induction might survive the procedure otherwise. Furthermore, we provide an epistemic characterization for the Independent Dekel-Fudenberg Procedure based on common belief of rationality (maximizing expected payoffs given the beliefs about the strategy choices of the opponents), caution (taking into account all strategies of the opponents), and stochastic independence (player i cannot learn anything about the behavior of opponent j by observing the play of different opponent j').

The paper is organized as follows: Section 2 presents the backward induction paradox as well as intuitions for our results in more detail, while the subsequent Section 3 introduces perfect information games. Section 4 specifies the formal meaning of stochastic independence in a context where belief are modelled by non-Archimedean probabilities, meaning that some opponent choices might be assigned subjective probability zero without being deemed subjectively impossible. Section 5 defines the Independent Dekel-Fudenberg Procedure and shows how this concept solves the backward induction paradox, and Section 6 provides its epistemic characterization. Section 7 contains concluding discussion.

2 Backward Induction Paradox

The backward induction paradox can be illustrated in a version of Rosenthal’s (1981) centipede game, as depicted in Figure 1. In this game, the backward induction procedure entails that player 1 chooses D at this player’s second decision node, inducing player 2 to choose d and player 1 to choose Out at their first decision node. However, if player 1 deviates from backward induction by

²In the terminology of Stalnaker (1998), the beliefs of the player about the behavior of two different opponents are epistemically independent.




Figure 2: A centipede game where players move only once.

choosing In, then player 2 weakly prefers c to d if, conditional on being asked to play, this player believes that player 1 will deviate from backward induction also at their second decision node, by choosing C , with at least probability $\frac{1}{3}$. Moreover, player 1 weakly prefers In to Out if this player believes that player 2 will react to being asked to play by choosing c with at least probability $\frac{1}{3}$.

In a game with similar features, namely the finitely repeated prisoners' dilemma, Pettit and Sugden (1989) argue that the backward induction solution, where players choose defect in all rounds, is intuitively implausible. Rather players might choose *cooperate* to signal a willingness to do so also in the future, leading players to adopt a *tit-for-tat* strategy for a while. Indeed, Kreps, Milgrom, Roberts, and Wilson (1982) demonstrate how such behavior can be rational when one player can possibly be committed to a tit-for-tat strategy. This is related to Kreps and Wilson (1982) and Milgrom and Roberts (1982) who show that players might use initial behavior to acquire a reputation for being 'tough' in Selten's (1978) finitely repeated *chain-store game*, leading to a different outcome than that predicted by backward induction in that game.

Reny (1988, 1992b, 1993) and Bicchieri (1989) relate the backward induction paradox to the impossibility of common knowledge of rationality in perfect information games. So, even if the players initially assign subjective probability zero to the event that their opponents do *not* choose best replies to their beliefs given their payoffs—in contrast to the assumptions made in the papers by Kreps, Milgrom, Roberts, and Wilson—the analysis must still allow for such irrationality to be deemed subjectively possible. In the following decades, a series of papers, including Reny (1992a), Aumann (1995), Ben-Porath (1997), Stalnaker (1998), Battigalli and Siniscalchi (2002), Asheim (2002), Asheim and Dufwenberg (2003a,b), Brandenburger (2007), Perea (2007, 2008, 2014), Brandenburger, Friedenberg, and Keisler (2008), Arieli and Aumann (2015), and Battigalli and De Vito (2021), have considered epistemic conditions that lead only to outcomes consistent with backward induction and those that permit also other outcomes. However, their predictions in perfect information games appear not to depend on whether players move more than once.

The Independent Dekel-Fudenberg Procedure, introduced here, yields a prediction which does depend on whether players move more than once. To illustrate how, consider the centipede game of Figure 2, which is a version of the centipede game of Figure 1 where the two agents of player 1 at the first and last decision nodes of the game have been divided into two separate players, 1 and 3, who however have the same payoffs as a function of the outcomes. In this game, the backward induction procedure entails that player 3 chooses D , inducing player 2 to choose d and player 1 to

choose Out. Since D weakly dominates C , only D is a best reply for player 3 to a belief where all opponent strategy profiles are deemed subjectively possible. This implies that C is eliminated in the first round of the Independent Dekel-Fudenberg Procedure, while no strategy is eliminated for players 1 and 2. Turn now to round 2 and player 2. Any belief for player 2 that (i) satisfies that all opponent strategy profiles are deemed subjectively possible, (ii) assigns subjective probability 1 to player 1 and 3 choosing (Out, D) or (In, D), and (iii) is stochastically independent, has the property that the belief of player 2 over the strategies of player 3 conditional on the choice by player 1 assigns subjective probability 1 to D independently of whether player 1 has chosen Out or In. Hence, c is eliminated in the second round of the Independent Dekel-Fudenberg Procedure, while no strategy is eliminated for player 1. Hence, in the third round, player 1 must assign subjective probability 1 to players 2 and 3 choosing (d , D), implying that In is eliminated. In contrast, the elimination stops after the first round if stochastically independent beliefs are not imposed or if the same player chooses at the first and last decision nodes, since then player 2 need not assign subjective probability 1 to the choice of D at the last decision node, conditional on the choice of In at the first decision node. This will be explained in more detail in Section 5.

3 Perfect Information Games

A finite *extensive game form* of almost perfect information with I players and M stages can be described as follows. This description facilitates the proofs while encompassing all game forms associated with both finite perfect information games and finitely repeated games. The sets of *histories* is determined inductively: The set of histories at the beginning of the first stage 1 is $H^1 = \{\emptyset\}$. Let H^m denote the set of histories at the beginning of stages $m \in \{1, 2, \dots, M\}$. At every $h \in H^m$, let, for each player $i \in \mathcal{I} := \{1, 2, \dots, I\}$, i 's nonempty and finite action set be denoted $A_i(h)$, where i is inactive at h if $A_i(h)$ is a singleton. Write $A(h) := A_1(h) \times A_2(h) \times \dots \times A_I(h)$. Define the set of histories at the beginning of stage $m + 1$ as follows: $H^{m+1} := \{(h, a) \mid h \in H^m \text{ and } a \in A(h)\}$. This concludes the induction. Let, for each player $i \in \mathcal{I}$,

$$H_i := \left\{ h \in \bigcup_{m=1}^M H^m \mid A_i(h) \text{ is not a singleton} \right\}$$

denote the set of histories at which player i makes an action choice; H_i is assumed to be nonempty. Then $H := \bigcup_{i=1}^I H_i$ is the set of *subtrees*, and $Z := H^{M+1}$ is the set of *outcomes*.

Let, for each player $i \in \mathcal{I}$, $v_i : Z \rightarrow \mathbb{R}$ denote i 's Bernoulli *utility function*. The combination of the extensive form and the vector (v_1, v_2, \dots, v_I) of utility functions is an *extensive game* Γ with I players. A *pure strategy* for player i is a function s_i that assigns an action in $A_i(h)$ to any $h \in H_i$. Let S_i denote player i 's finite set of pure strategies, and write $S := S_1 \times S_2 \times \dots \times S_I$ and $S_{-i} := S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_I$. Let $z : S \rightarrow Z$ map strategy profiles into outcomes.³

³A pure strategy $s_i \in S_i$ can be viewed as an act on S_{-i} that assigns $z(s_i, s_{-i}) \in Z$ to any $s_{-i} \in S_{-i}$. The set

Let, for each player $i \in \mathcal{I}$, $u_i = v_i \circ z$ denote i 's *payoff function*. Then $G = ((S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}})$ is the *strategic game* derived from Γ .

An extensive game Γ is of *perfect information* if $\{H_1, H_2, \dots, H_I\}$ is a partition of H ; that is, there is no history at which two players choose actions simultaneously. In a perfect information game Γ , let $p : H \rightarrow \{1, 2, \dots, I\}$ be the function that, for each $h \in H$, determines the player who chooses at h . An extensive game Γ has the property that all players choose only once if, for each player $i \in \mathcal{I}$, H_i is a singleton, implying that, for some $h_i \in H$, $H_i = \{h_i\}$ and $S_i = A(h_i)$. An extensive game Γ has *no relevant payoff ties* if, for each player $i \in \mathcal{I}$ and all $s_{-i} \in S_{-i}$, $v_i(z(s'_i, s_{-i})) \neq v_i(z(s''_i, s_{-i}))$ whenever $s'_i, s''_i \in S_i$ lead to different outcomes; that is, $z(s'_i, s_{-i}) \neq z(s''_i, s_{-i})$. In a perfect information game Γ with no relevant payoff ties, the procedure of *backward induction* determines a unique strategy profile s^* through the following inductive procedure: If $h \in H \cap H^M$, then $s_{p(h)}^*(h)$ is the unique action that maximizes $v_{p(h)}(h, a)$ over all $a \in A_{p(h)}$. Assume that s^* has been determined for all $h \in H \cap (H^{m+1} \cup \dots \cup H^M)$, where $m \in \{1, 2, \dots, M-1\}$. If $h \in H \cap H^m$, then $s_{p(h)}^*(h)$ is the unique action that maximizes $v_{p(h)}(h, a, s_{p(h,a)}^*(h, a), \dots)$ over all $a \in A_{p(h)}$. This concludes the induction.

4 Independent Non-Archimedean Probabilities

Analysis of extensive games in the strategic form is facilitated by applying beliefs about opponent behavior where certain actions are deemed subjectively possible although assigned subjective probability zero. This requires so-called non-Archimedean subjective probabilities. Moreover, our foundation of the backward induction paradox requires that such non-Archimedean subjective probabilities be stochastically independent. This section concerns the modeling of stochastically independent non-Archimedean subjective probabilities.

Consider a finite set X . Following Blume, Brandenburger, and Dekel (1991a), a *lexicographic probability system* (LPS) λ on X is a vector $(\mu^1, \mu^2, \dots, \mu^L)$, where μ^ℓ , for $\ell = 1, 2, \dots, L$, are probability (non-negative one-sum) distributions on X . The support of μ^1 , $\text{supp}\mu^1$, is the set of elements in X that are assigned positive subjective probability, while the support of λ , $\text{supp}\lambda = \text{supp}\mu^1 \cup \text{supp}\mu^2 \cup \dots \cup \text{supp}\mu^L$, is the set of elements in X that are deemed subjectively possible.

In the context of a finite strategic game $G = ((S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}})$, player i 's payoff function u_i combined with an LPS $\lambda_i = (\mu_i^1, \mu_i^2, \dots, \mu_i^L)$ on S_{-i} , as a representation of player i 's belief about opponent behavior, determines player i 's preferences over his own strategies $s_i \in S_i$ as follows: s_i is weakly preferred to s'_i given the beliefs λ_i if and only if

$$\left(\Sigma_{u_i}^1(s_i), \Sigma_{u_i}^2(s_i), \dots, \Sigma_{u_i}^L(s_i) \right) \geq_{\mathcal{L}} \left(\Sigma_{u_i}^1(s'_i), \Sigma_{u_i}^2(s'_i), \dots, \Sigma_{u_i}^L(s'_i) \right),$$

of pure strategies S_i is partitioned into equivalent classes of acts since a pure strategy s_i also determines actions in subtrees which s_i prevents from being reached. Each such equivalent class corresponds to a *plan of action* in the sense of Rubinstein (1991). As there is no need here to differentiate between identical acts, the concept of a plan of action suffices. Indeed, in the example of Figure 1, we list only the players' plans of actions.

where $\Sigma_{u_i}^\ell(s_i)$ denotes $\sum_{s_{-i} \in S_{-i}} \mu_i^\ell(s_{-i}) u_i(s_i, s_{-i})$ for $\ell \in \{1, 2, \dots, L\}$, and where $\geq_{\mathcal{L}}$ is defined by, for $a, b \in \mathbb{R}^L$, $a \geq_{\mathcal{L}} b$ if and only if (i) $a_\ell = b_\ell$ for all $\ell \in \{1, 2, \dots, L\}$ or (ii) there exists $\ell \in \{1, 2, \dots, L\}$ such that $a_{\ell'} = b_{\ell'}$ for all $\ell' \in \{1, 2, \dots, \ell - 1\}$ and $a_\ell > b_\ell$. Say that s_i is a *best reply* to λ_i if, for all $s'_i \in S_i$, s_i is weakly preferred to s'_i given the beliefs λ_i . Define i 's *best reply correspondence* β_i from the set of LPSs on S_{-i} to $2^{S_{-i} \setminus \{\emptyset\}}$ as follows: For every LPS λ_i on S_{-i} ,

$$\beta_i(\lambda_i) := \{s_i \in S_i \mid s_i \text{ is a best reply to } \lambda_i\}.$$

To define stochastic independence we impose *strong independence* in the sense of Blume, Brandenburger, and Dekel (1991a, Def. 7.1; 1991b, Sect. 3.3). This version of stochastic independence “requires there to be an equivalent \mathbb{F} -valued probability measure that is a product measure” (Blume, Brandenburger, and Dekel, 1991b, p. 90), where \mathbb{F} is “a non-Archimedean ordered field . . . which is a strict extension of the real number field \mathbb{R} ” (Blume, Brandenburger, and Dekel, 1991a, p. 72), with the notion of ‘a non-Archimedean ordered field’ not being explained in detail and the concept of ‘equivalence’ only being implicitly defined. Therefore, to expound their definition, we introduce the notions of non-standard numbers and non-standard probabilities and refer to literature which analyzes these notions. An *infinitesimal* ε is a positive number with the property that $\varepsilon < a$ for every positive real number $a \in \mathbb{R}$. Following Robertson (1973), Hammond (1994), and Halpern (2010), let $\mathbb{R}(\varepsilon)$ be the smallest field that includes all real numbers and the infinitesimal ε . As shown by Meier and Perea (2020, Sect. 5.1), every finite non-standard number $a \in \mathbb{R}(\varepsilon)$ can uniquely be written as $a = a_1 + a_2\varepsilon + a_3\varepsilon^2 + \dots$, where $a_\ell \in \mathbb{R}$ for every $\ell \in \mathbb{N}$. Let $\text{st}(a) := a_1$ denote the *standard part* of a , which is the real number “closest” to a .

Consider a finite set X . A *non-standard probability distribution* (NPD) on X is a function $\nu : X \rightarrow \mathbb{R}(\varepsilon)$ such that $\nu(x) \geq 0$ for all $x \in X$ and $\sum_{x \in X} \nu(x) = 1$. Following Halpern (2010, Def. 4.1 and Lemma A.7), say that an NPD ν on X is *equivalent* to an LPS $\lambda = (\mu^1, \mu^2, \dots, \mu^L)$ on X if, for all $x \in X$,

$$\nu(x) = \sum_{\ell=1}^L \tilde{\nu}(\ell) \mu^\ell(x),$$

where $\tilde{\nu} : \{1, \dots, L\} \rightarrow \mathbb{R}(\varepsilon)$ is an NPD on $\{1, 2, \dots, L\}$ with the properties that

$$\text{st}\left(\frac{\tilde{\nu}(\ell+1)}{\tilde{\nu}(\ell)}\right) = 0$$

for $\ell \in \{1, 2, \dots, L-1\}$ and $\tilde{\nu}(L) > 0$. To illustrate, let $\lambda = (\mu^1, \mu^2, \mu^3)$ be an LPS on X . Then $\tilde{\nu}$ equal to $(1 - \varepsilon - \varepsilon^2, \varepsilon, \varepsilon^2)$ or $(1 - \varepsilon, \varepsilon - \varepsilon^2, \varepsilon^2)$ or $(1 - 2\varepsilon^2, 2(\varepsilon^2 - 3\varepsilon^3), 6\varepsilon^3)$ are examples of NPDs on $\{1, 2, 3\}$ that can be used to aggregate the LPS λ into an equivalent NPD ν .

In the context of a finite strategic game $G = ((S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}})$, player i 's payoff function u_i , combined with an NPD ν_i on S_{-i} determines player i 's preferences over his own strategies $s_i \in S_i$

as follows: s_i is weakly preferred to s'_i given the beliefs ν_i if and only if

$$\sum_{s_{-i} \in S_{-i}} \nu_i(s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \nu_i(s_{-i}) u_i(s'_i, s_{-i}).$$

Say that s_i is a *best reply* to ν_i if, for all $s'_i \in S_i$, s_i is weakly preferred to s'_i given the beliefs ν_i . If the NPD ν_i on S_{-i} is equivalent to the LPS λ_i on S_{-i} , then the set of best replies coincide:

$$\beta_i(\lambda_i) = \{s_i \in S_i \mid s_i \text{ is a best reply to } \nu_i\}.$$

In fact, as argued by Halpern (2010, footnote 5), this statement holds if we consider the best reply correspondence β_i and the set of best replies as a function of the NPS ν_i for every possible payoff function u_i on $S_i \times S_{-i}$.

An NPD ν_i on S_{-i} is a *product distribution* if there exist NPDs ν_i^j on S_j for $j \in \mathcal{I} \setminus \{i\}$ such that

$$\nu_i(s_{-i}) = \prod_{j \in \mathcal{I} \setminus \{i\}} \nu_i^j(s_j)$$

for all $s_{-i} \in S_{-i}$. An LPS λ_i on S_{-i} is said to be *strongly independent* if there exists an equivalent NPD on S_{-i} that is a product distribution. This concludes our elucidation of the independence concept defined by Blume, Brandenburger, and Dekel (1991a, Def. 7.1; 1991b, Sect. 3.3).

5 Independent Dekel-Fudenberg Procedure

The *Dekel-Fudenberg Procedure* (Dekel and Fudenberg, 1990) eliminates, in the first round, all weakly dominated strategies for all players and, in subsequent rounds, all strictly dominated strategies for all players, until the procedure reaches a round in which no further elimination is possible. In the first subsection we first state an equivalent definition where the eliminated strategies in each round are those that can never be best replies to beliefs where only strategies that are still uneliminated are assigned positive subjective probabilities, but where all opponent strategy profiles are deemed subjectively possible. We then define the *Independent Dekel-Fudenberg Procedure* by imposing the additional requirement that beliefs are strongly independent. In the second subsection we establish three results showing how the Independent Dekel-Fudenberg Procedure can be used to interpret the backward induction paradox, while the ordinary Dekel-Fudenberg Procedure cannot.

5.1 Definitions

Consider first the correspondence $a_i^c : \{S'_{-i} \subseteq S_{-i} \mid S'_{-i} \text{ is a Cartesian product}\} \rightarrow 2^{S_i}$ defined as follows (where superscript c indicates that beliefs are allowed to be correlated):

$$a_i^c(S'_{-i}) := \{s_i \in S_i \mid \text{there exists an LPS } \lambda_i = (\mu_i^1, \dots, \mu_i^L) \text{ on } S_{-i} \text{ with} \\ \text{supp}\mu_i^1 \subseteq S'_{-i} \text{ and } \text{supp}\lambda_i = S_{-i} \text{ such that } s_i \text{ is a best response to } \lambda_i\}$$

for all non-empty Cartesian products S'_{-i} , and $a_i^c(\emptyset) := \emptyset$. By Brandenburger (1992, Prop. 2), the following is an equivalent definition of the Dekel-Fudenberg Procedure.

Definition 1 *The Dekel-Fudenberg Procedure.* Consider the sequence defined by, for all players $i \in \mathcal{I}$, $S_i^0 = S_i$ and, for every $k \geq 1$, $S_i^k = a_i^c(S_1^{k-1} \times \dots \times S_{i-1}^{k-1} \times S_{i+1}^{k-1} \times \dots \times S_I^{k-1})$. A strategy s_i for player i *survives the Dekel-Fudenberg Procedure* if $s_i \in P_i^c := \bigcap_{k=1}^{\infty} S_i^k$.

In particular, for each player $i \in \mathcal{I}$, $S_i^1 = a_i^c(S_{-i})$ is the set of i 's *admissible* strategies, that is, not weakly dominated (Blume, Brandenburger, and Dekel, 1991a, Thm. 4.2), while, for every $k > 1$, $S_i^k = a_i^c(S_{-i}^{k-1})$ is the subset of S_i^1 that are not strictly dominated on S_{-i}^{k-1} (Pearce, 1984, Lemma 3, generalized to I -player games where beliefs are allowed to be correlated). Strategies surviving this procedure are called *permissible* by Brandenburger (1992); hence, the notation P_i^c .

Consider next the correspondence $a_i : \{S'_{-i} \subseteq S_{-i} \mid S'_{-i} \text{ is a Cartesian product}\} \rightarrow 2^{S_i}$ defined as follows:

$$a_i(S'_{-i}) := \{s_i \in S_i \mid \text{there exists a strongly independent LPS } \lambda_i = (\mu_i^1, \dots, \mu_i^L) \text{ on } S_{-i} \text{ with} \\ \text{supp}\mu_i^1 \subseteq S'_{-i} \text{ and } \text{supp}\lambda_i = S_{-i} \text{ such that } s_i \text{ is a best response to } \lambda_i\}$$

for all non-empty Cartesian products S'_{-i} , and $a_i(\emptyset) := \emptyset$. This correspondence can be used to state the following definition.

Definition 2 *The Independent Dekel-Fudenberg Procedure.* Consider the sequence defined by, for all players $i \in \mathcal{I}$, $S_i^0 = S_i$ and, for every $k \geq 1$, $S_i^k = a_i(S_1^{k-1} \times \dots \times S_{i-1}^{k-1} \times S_{i+1}^{k-1} \times \dots \times S_I^{k-1})$. A strategy s_i for player i *survives the Independent Dekel-Fudenberg Procedure* if $s_i \in P_i := \bigcap_{k=1}^{\infty} S_i^k$.

5.2 Results

We start our analysis of the Dekel-Fudenberg and Independent Dekel-Fudenberg Procedures by noting the following helpful result, writing $P_{-i}^c := P_1^c \times \dots \times P_{i-1}^c \times P_{i+1}^c \times \dots \times P_I^c$ and $P_{-i} := P_1 \times \dots \times P_{i-1} \times P_{i+1} \times \dots \times P_I$.

Lemma 1. (a) For each player $i \in \mathcal{I}$, $\emptyset \neq P_i \subseteq P_i^c \subseteq S_i$.

(b) For each player $i \in \mathcal{I}$, $P_i^c = a_i^c(P_{-i}^c)$ and $P_i = a_i(P_{-i})$.

Proof. For all $i \in \mathcal{I}$, a_i^c and a_i are monotone: if S'_{-i} and S''_{-i} are Cartesian products satisfying $\emptyset \neq S'_{-i} \subseteq S''_{-i} \subseteq S_{-i}$, then $\emptyset \neq a_i^c(S'_{-i}) \subseteq a_i^c(S''_{-i}) \subseteq a_i^c(S_{-i})$ and $\emptyset \neq a_i(S'_{-i}) \subseteq a_i(S''_{-i}) \subseteq a_i(S_{-i})$. Hence, since S is finite, both procedures converge in a finite number of rounds to non-empty sets of strategies P_i^c and P_i , respectively, satisfying $P_i^c = a_i^c(P_{-i}^c)$ and $P_i = a_i(P_{-i})$, for all players $i \in \mathcal{I}$. Since, for any Cartesian product $S'_{-i} \subseteq S_{-i}$, $a_i(S'_{-i}) \subseteq a_i^c(S'_{-i})$, a strategy that survives the Independent Dekel-Fudenberg Procedure also survives the Dekel-Fudenberg Procedure, while the converse need not hold. \square

We first show that the Independent Dekel-Fudenberg Procedure determines the profile of backward induction strategies in perfect information games with no relevant payoff ties where all players choose only once.

Proposition 1. *In any perfect information game Γ with no relevant payoff ties and the property that all players choose only once, for each player $i \in \mathcal{I}$, there is a unique strategy that survives the Independent Dekel-Fudenberg Procedure, and this strategy is the player's backward induction strategy.*

Proof. Assume that Γ is a perfect information game with no relevant payoff ties and the property that all players choose only once. Since all players choose only once, for each $i \in \mathcal{I}$, i 's strategy set S_i equals $A(h_i)$, where h_i is the single history after which i makes an action choice. Since the game has no relevant payoff ties, there exists, for each $i \in \mathcal{I}$, a unique strategy $s_i^* \in A(h_i)$ that survives the backward induction procedure. Let \mathcal{I}_m be the set of players that makes an action choice in stage $m \in \{1, 2, \dots, M\}$, implying that $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_M\}$ is a partition of \mathcal{I} . The backward induction procedure has M stages where, for all $k \in \{1, 2, \dots, M-1\}$, $A(h_i)$ is the set strategies surviving k stages for each $i \in \mathcal{I}_1 \cup \dots \cup \mathcal{I}_{M-k}$ and s_i^* is the unique strategy surviving k stages for each $i \in \mathcal{I}_{M-k+1} \cup \dots \cup \mathcal{I}_M$, while s_i^* is the unique strategy surviving M stages for all $i \in \mathcal{I}$. The strategy of proof is to show that, for all $k \in \{1, 2, \dots, M\}$, no strategy but s_i^* survives k stages of the Independent Dekel-Fudenberg Procedure for each $i \in \mathcal{I}_{M-k+1} \cup \dots \cup \mathcal{I}_M$.

We prove this by induction. We initiate the induction by first showing that no strategy but s_i^* survives stage 1 of the Independent Dekel-Fudenberg Procedure for each $i \in \mathcal{I}_M$. This follows since, for each $i \in \mathcal{I}_M$, s_i^* is the unique best response to an LPS λ_i on S_{-i} satisfying $\text{supp}\lambda_i = S_{-i}$.

We next show that no strategy but s_i^* survives stage $k+1$ of the Independent Dekel-Fudenberg Procedure for each $i \in \mathcal{I}_{M-k}$ if no strategy but s_j^* survives k stages of the Independent Dekel-Fudenberg Procedure for each $j \in \mathcal{I}_{M-k+1} \cup \dots \cup \mathcal{I}_M$. Assume that no strategy but s_j^* survives k stages of the Independent Dekel-Fudenberg Procedure for each $j \in \mathcal{I}_{M-k+1} \cup \dots \cup \mathcal{I}_M$, and let $\lambda_i = (\mu_i^1, \dots, \mu_i^L)$ on S_{-i} be a strongly independent LPS on S_{-i} with $\text{supp}\mu_i^l \subseteq S_{-i}^k$ and $\text{supp}\lambda_i = S_{-i}$, where S_{-i}^k is the Cartesian product of opponent strategies that survive k rounds of the Independent Dekel-Fudenberg Procedure. Since λ_i is strongly independent, there exists an equivalent NPD ν_i on S_{-i} which satisfies that $\nu_i(s_{-i}) = \prod_{j \in \mathcal{I} \setminus \{i\}} \nu_i^j(s_j)$, where ν_i^j are NPDs on

S_j for $j \in \mathcal{I} \setminus \{i\}$. Since $S_j^k = \{s_j^*\}$ for all $j \in \mathcal{I}_{M-k+1} \cup \dots \cup \mathcal{I}_M$, it follows from the property that $\text{supp}\mu_i^1 \subseteq S_{-i}^k$ that, for each such j , $\text{st}(v_j(s_j^*)) = 1$. Because of this and the properties of the backward induction procedure, for each $i \in \mathcal{I}_{M-k}$, s_i^* is the unique best response to the LPS λ_i , since $\text{supp}\lambda_i = S_{-i}$. It follows that no strategy but s_i^* survives stage $k + 1$ of the Independent Dekel-Fudenberg Procedure for each $i \in \mathcal{I}_{M-k}$. This concludes the induction.

Since, by Lemma 1(a), for each $i \in \mathcal{I}$, $P_i \neq \emptyset$, it follows that $P_i = \{s_i^*\}$ for each $i \in \mathcal{I}$. \square

Proposition 1 can be illustrated by the centipede game of Figure 2, the version of the centipede game of Figure 1 with three separate players. As explained at the end of Section 2, the Independent Dekel-Fudenberg Procedure leads to the following rounds of elimination in this game:

$$\begin{array}{lll}
S_1^1 = a_1(S_2 \times S_3) = S_1 & S_2^1 = a_2(S_1 \times S_3) = S_2 & S_3^1 = a_3(S_1 \times S_2) = \{D\} \\
S_1^2 = a_1(S_2 \times \{D\}) = S_1 & S_2^2 = a_2(S_1 \times \{D\}) = \{d\} & S_3^2 = a_3(S_1 \times S_2) = \{D\} \\
S_1^3 = a_1(\{d\} \times \{D\}) = \{\text{Out}\} & S_2^3 = a_2(S_1 \times \{D\}) = \{d\} & S_3^3 = a_3(S_1 \times \{d\}) = \{D\} \\
\dots & \dots & \dots \\
S_1^k = a_1(\{d\} \times \{D\}) = \{\text{Out}\} & S_2^k = a_2(\{\text{Out}\} \times \{D\}) = \{d\} & S_3^k = a_3(\{\text{Out}\} \times \{d\}) = \{D\} \\
\dots & \dots & \dots
\end{array}$$

Hence, in the game of Figure 2, the eliminations according to the Independent Dekel-Fudenberg Procedure correspond to the backward induction procedure. Note that the Independent Dekel-Fudenberg Procedure might eliminate faster than the backward induction procedure. This is indeed the case if, in the game of Figure 2, player 1's payoff of Out would have been 6 instead of 2, causing In to be eliminated already in the first round. However, in any case, for a perfect information game with no relevant payoff ties and the property that all players choose only once, only the backward induction strategies survive the procedure. We next show that this is not the case for the ordinary Dekel-Fudenberg Procedure.

Proposition 2. *There exists a perfect information game Γ with no relevant payoff ties and the property that all players choose only once, where an outcome other than the backward induction outcome can be reached even if all players choose strategies that survive the Dekel-Fudenberg Procedure.*

Proof. Consider the game of Figure 2, which is a perfect information game with no relevant payoff ties and the property that all players choose only once. Since D weakly dominates C , only D is a best reply for player 3 to an LPS where all opponent strategy profiles are deemed subjectively possible. Hence, $S_3^1 = a_3^c(S_1 \times S_2) = \{D\}$, implying that C is eliminated in the first round of the Dekel-Fudenberg Procedure, while no strategy is eliminated for players 1 and 2. The Dekel-Fudenberg Procedure allows no further elimination. In particular, c is best reply for player 2 to an LPS λ_2 over $S_1 \times S_3 = \{(\text{Out}, D), (\text{Out}, C), (\text{In}, D), (\text{In}, C)\}$ given by $\lambda_2 = ((1, 0, 0, 0), (0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}))$ since the LPS $\lambda_2|_{\{\text{In}\}}$ conditional on the choice of In by player 1 assigns subjective probability $\frac{1}{2}$ to player 3 choosing C . Note that the LPS λ_2 for player 2 satisfies that

all opponent strategy profiles are deemed subjectively possible and assigns subjective probability 1 to $S_1^1 \times S_3^1 = \{(\text{Out}, D), (\text{In}, D)\}$, but it is not strongly independent. Hence, in addition to the backward induction outcome Out, also the outcomes (In, d) and (In, c, D) can be reached even if all players choose strategies that survive the Dekel-Fudenberg Procedure. \square

Proposition 3. *There exists a perfect information game Γ with no relevant payoff ties, where an outcome other than the backward induction outcome can be reached even if all players choose strategies that survive the Independent Dekel-Fudenberg Procedure. Such a game necessarily involves some player choosing more than once.*

Proof. Consider the game of Figure 1, which is a perfect information game with no relevant payoff ties. Since InD weakly dominates InC, InC cannot be a best reply for player 1 to an LPS where all opponent strategy profiles are deemed subjectively possible. Hence, $S_1^1 = a_1(S_2) = \{\text{Out}, \text{InD}\}$, implying that InC is eliminated in the first round of the Independent Dekel-Fudenberg Procedure, while no strategy is eliminated for player 2. The Independent Dekel-Fudenberg Procedure allows no further elimination. In particular, c is best reply for player 2 to an LPS λ_2 over $S_1 = \{\text{Out}, \text{InD}, \text{InC}\}$ given by $\lambda_2 = ((1, 0, 0), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}))$ since the LPS $\lambda_2|_{\{\text{InD}, \text{InC}\}}$ conditional on the choice of In by player 1 assigns subjective probability $\frac{1}{2}$ to player 1 choosing InC. Note that the LPS λ_2 for player 2 satisfies that all opponent strategy profiles are deemed subjectively possible and assigns subjective probability 1 to $S_1^1 = \{\text{Out}, \text{InD}\}$. It is also trivially strongly independent as the game has only two players. Hence, in addition to the backward induction outcome Out, also the outcomes (In, d) and (In, c, D) can be reached even if both players choose strategies that survive the Independent Dekel-Fudenberg Procedure. \square

6 Epistemic Characterizations

The epistemic analysis builds on the concept of player types, where a type of a player is characterized by an LPS over the others' strategies and types.

6.1 Definitions

For each $i \in \mathcal{I}$, let T_i denote player i 's non-empty and finite type space. The *state space* is defined by $\Omega := S \times T$, where $T := T_1 \times \dots \times T_I$. For each player $i \in \mathcal{I}$, write $\Omega_i := S_i \times T_i$ and $\Omega_{-i} := \Omega_1 \times \dots \times \Omega_{i-1} \times \Omega_{i+1} \times \dots \times \Omega_I$. To each type $t_i \in T_i$ of every player i is associated an LPS $\lambda_i(t_i) = (\mu_i^1(t_i), \mu_i^2(t_i), \dots, \mu_i^{\mathbf{L}(t_i)}(t_i))$ on Ω_{-i} . For each player i , we thus have the player's strategy set S_i , type space T_i and a mapping λ_i that to each of i 's types t_i assigns an LPS $\lambda_i(t_i)$ over the others' strategy choices and types. The structure $((S_i)_{i \in \mathcal{I}}, (T_i)_{i \in \mathcal{I}}, (\lambda_i)_{i \in \mathcal{I}})$ is called an *S-based interactive belief structure*.

For each $i \in \mathcal{I}$, let $\mathbf{s}_i(\omega)$ and $\mathbf{t}_i(\omega)$ denote i 's strategy and type in state $\omega \in \Omega$. In other words, $\mathbf{s}_i : \Omega \rightarrow S_i$ is the projection of the state space to i 's strategy set, assigning to each state $\omega \in \Omega$

the strategy $s_i = \mathbf{s}_i(\omega)$ that i uses in that state. Likewise, $\mathbf{t}_i : \Omega \rightarrow T_i$ is the projection of the state space to i 's type space. For each player $i \in \mathcal{I}$, the *belief operator* B_i maps each event $E \subseteq \Omega$ to the set of states where player i 's type assigns subjective probability 1 to E :

$$B_i(E) := \{\omega \in \Omega \mid \mu_i^1(\mathbf{t}_i(\omega))(E(\omega_i)) = 1\},$$

where $E(\omega_i) := \{\omega_{-i} \in \Omega_{-i} \mid (\omega_i, \omega_{-i}) \in E\}$. This belief operator appears in Hu (2007). The belief operator B_i satisfies $B_i(\emptyset) = \emptyset$, $B_i(\Omega) = \Omega$, $B_i(E') \subseteq B_i(E'')$ if $E' \subseteq E''$ (monotonicity), and $B_i(E) = E$ if $E = (\text{proj}_{\Omega_i} E) \times \Omega_{-i}$. The last property means that each player i always believes their own strategy-type pair. Since also $B_i(E) = (\text{proj}_{\Omega_i} B_i(E)) \times \Omega_{-i}$ for all events $E \subseteq \Omega$, the operator B_i satisfies both positive ($B_i(E) \subseteq B_i(B_i(E))$) and negative ($\neg B_i(E) \subseteq B_i(\neg B_i(E))$) introspection. However, B_i violates the truth axiom, meaning that $B_i(E) \subseteq E$ need not hold for all $E \subseteq \Omega$. Finally, $B_i(E') \cap B_i(E'') \subseteq B_i(E' \cap E'')$ for all $E', E'' \subseteq \Omega$. Say that, at $\omega \in \Omega$, there is *mutual belief* of $E \subseteq \Omega$ if $\omega \in B(E)$, where $B(E) := B_1(E) \cap \dots \cap B_I(E)$. Say that, at $\omega \in \Omega$, there is *common belief* of $E \subseteq \Omega$ if $\omega \in CB(E)$, where $CB(E) := B(E) \cap B(B(E)) \cap B(B(B(E))) \cap \dots$.

We connect types with the payoff functions by, for each player $i \in \mathcal{I}$, defining i 's *choice correspondence* $\mathbf{S}_i : T_i \rightarrow 2^{S_i}$ as follows: For each of i 's types $t_i \in T_i$,

$$\mathbf{S}_i(t_i) := \beta_i(\text{marg}_{S_{-i}} \boldsymbol{\lambda}_i(t_i))$$

consists of i 's best replies when player i is of type t_i . For each player $i \in \mathcal{I}$, write $[\text{rat}_i]$ for the event that player i uses a best reply:

$$[\text{rat}_i] := \{\omega \in \Omega \mid \mathbf{s}_i(\omega) \in \mathbf{S}_i(\mathbf{t}_i(\omega))\}.$$

One may interpret $[\text{rat}_i]$ as the event that i is rational: if $\omega \in [\text{rat}_i]$, then $\mathbf{s}_i(\omega)$ is a best reply to $\text{marg}_{S_{-i}} \boldsymbol{\lambda}_i(\mathbf{t}_i(\omega))$. For each player $i \in \mathcal{I}$, write $[\text{cau}_i]$ for the event that player i has beliefs with full support on the strategy profiles of the others:

$$[\text{cau}_i] := \{\omega \in \Omega \mid \text{supp}(\text{marg}_{S_{-i}} \boldsymbol{\lambda}_i(\mathbf{t}_i(\omega))) = S_{-i}\}.$$

One may interpret $[\text{cau}_i]$ as the event that i is cautious. For each player $i \in \mathcal{I}$, write $[\text{ind}_i]$ for the event that player i has stochastically independent beliefs about the strategy choices of the others:

$$[\text{ind}_i] := \{\omega \in \Omega \mid \text{marg}_{S_{-i}} \boldsymbol{\lambda}_i(\mathbf{t}_i(\omega)) \text{ is strongly independent}\}.$$

Write $[\text{rat}] := [\text{rat}_1] \cap \dots \cap [\text{rat}_I]$, $[\text{cau}] := [\text{cau}_1] \cap \dots \cap [\text{cau}_I]$, and $[\text{ind}] := [\text{ind}_1] \cap \dots \cap [\text{ind}_I]$ for the events that all players, respectively, are rational, are cautious, and have stochastically independent belief about the strategy choices of the others.

6.2 Results

We can now state the following characterization results.

Proposition 4. *For each player $i \in \mathcal{I}$ and any strategy $s_i \in S_i$ for i , s_i survives the Dekel-Fudenberg Procedure if and only if there exists an S -based interactive belief structure $((S_i)_{i \in \mathcal{I}}, (T_i)_{i \in \mathcal{I}}, (\lambda_i)_{i \in \mathcal{I}})$ such that $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in CB([\text{rat}] \cap [\text{cau}])$.*

The proof is deleted, as the result is well-known (Börjers, 1994), and its proof can be obtained by removing the independence requirement from the proof of the following Proposition 5. The epistemic characterization of the Dekel-Fudenberg Procedure in Proposition 4 can be illustrated in the game of Figure 2 by letting $T_1 = \{t_1^{\text{Out}}, t_1^{\text{In}}\}$, $T_2 = \{t_2^d, t_2^c\}$, and $T_3 = \{t_3^D\}$, where $\lambda_1(t_1^{\text{Out}}) = (\mu_1^1(t_1^{\text{Out}}), \mu_1^2(t_1^{\text{Out}}), \mu_1^3(t_1^{\text{Out}}))$ and $\lambda_1(t_1^{\text{In}}) = (\mu_1^1(t_1^{\text{In}}), \mu_1^2(t_1^{\text{In}}))$ are given by:

$$\begin{aligned} \mu_1^1(t_1^{\text{Out}})((d, t_2^d), (D, t_3^D)) &= 1, \\ \mu_1^2(t_1^{\text{Out}})((d, t_2^d), (C, t_3^D)) &= \mu_1^2(t_1^{\text{Out}})((c, t_2^d), (D, t_3^D)) = \frac{1}{2}, \\ \mu_1^3(t_1^{\text{Out}})((c, t_2^d), (C, t_3^D)) &= 1, \\ \mu_1^1(t_1^{\text{In}})((d, t_2^d), (D, t_3^D)) &= \mu_1^1(t_1^{\text{In}})((c, t_2^c), (D, t_3^D)) = \frac{1}{2}, \\ \mu_1^2(t_1^{\text{In}})((d, t_2^d), (C, t_3^D)) &= \mu_1^2(t_1^{\text{In}})((c, t_2^c), (C, t_3^D)) = \frac{1}{2}, \end{aligned}$$

where $\lambda_2(t_2^d) = (\mu_2^1(t_2^d), \mu_2^2(t_2^d), \mu_2^3(t_2^d))$ and $\lambda_2(t_2^c) = (\mu_2^1(t_2^c), \mu_2^2(t_2^c))$ are given by:

$$\begin{aligned} \mu_2^1(t_2^d)((\text{Out}, t_1^{\text{Out}}), (D, t_3^D)) &= 1, \\ \mu_2^2(t_2^d)((\text{Out}, t_1^{\text{Out}}), (C, t_3^D)) &= \mu_2^2(t_2^d)((\text{In}, t_1^{\text{Out}}), (D, t_3^D)) = \frac{1}{2}, \\ \mu_2^3(t_2^d)((\text{In}, t_1^{\text{Out}}), (C, t_3^D)) &= 1, \\ \mu_2^1(t_2^c)((\text{Out}, t_1^{\text{Out}}), (D, t_3^D)) &= 1, \\ \mu_2^2(t_2^c)((\text{Out}, t_1^{\text{Out}}), (C, t_3^D)) &= \mu_2^2(t_2^c)((\text{In}, t_1^{\text{In}}), (D, t_3^D)) = \mu_2^2(t_2^c)((\text{In}, t_1^{\text{In}}), (C, t_3^D)) = \frac{1}{3}, \end{aligned}$$

and where $\lambda_3(t_3^D) = (\mu_3^1(t_3^D), \mu_3^2(t_3^D), \mu_3^3(t_3^D))$ is given by:

$$\begin{aligned} \mu_3^1(t_3^D)((\text{Out}, t_1^{\text{Out}}), (d, t_2^d)) &= 1, \\ \mu_3^2(t_3^D)((\text{Out}, t_1^{\text{Out}}), (c, t_2^d)) &= \mu_3^2(t_3^D)((\text{In}, t_1^{\text{Out}}), (d, t_2^d)) = \frac{1}{2}, \\ \mu_3^3(t_3^D)((\text{In}, t_1^{\text{Out}}), (c, t_2^d)) &= 1. \end{aligned}$$

Then, for each state in $\{(\text{Out}, t_1^{\text{Out}}), (\text{In}, t_1^{\text{In}})\} \times \{(d, t_2^d), (c, t_2^c)\} \times \{(D, t_3^D)\}$, there is common belief of rationality and caution, since $\mathbf{S}_1(t_1^{\text{Out}}) = \{\text{Out}\}$, $\mathbf{S}_1(t_1^{\text{In}}) = \{\text{In}\}$, $\mathbf{S}_2(t_2^d) = \{d\}$, $\mathbf{S}_2(t_2^c) = \{c\}$, and $\mathbf{S}_3(t_3^D) = \{D\}$. This corresponds to the fact that Out and In for player 1, d and c for player 2, and D for player 3 survive the Dekel-Fudenberg Procedure.

Proposition 5. *For each player $i \in \mathcal{I}$ and any strategy $s_i \in S_i$ for i , s_i survives the Independent Dekel-Fudenberg Procedure if and only if there exists an S -based interactive belief structure $((S_i)_{i \in \mathcal{I}}, (T_i)_{i \in \mathcal{I}}, (\lambda_i)_{i \in \mathcal{I}})$ such that $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in CB([\text{rat}] \cap [\text{cau}] \cap [\text{ind}])$.*

Proof. Part 1: For each player $i \in \mathcal{I}$ and any strategy $s_i \in P_i$ (that is, s_i survives the Independent Dekel-Fudenberg Procedure), there exists an S -based interactive belief structure $((S_i)_{i \in \mathcal{I}}, (T_i)_{i \in \mathcal{I}}, (\lambda_i)_{i \in \mathcal{I}})$ such that $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in CB([\text{rat}] \cap [\text{cau}] \cap [\text{ind}])$. For each $i \in \mathcal{I}$ and any $s_i \in P_i$, let $t_i^{s_i}$ denote a type of i for which $s_i \in \mathbf{S}_i(t_i^{s_i})$, $\text{supp}(\text{marg}_{S_{-i}} \mu_i^1(t_i^{s_i})) \subseteq P_{-i}$, $\text{supp}(\text{marg}_{S_{-i}} \lambda_i(t_i^{s_i})) = S_{-i}$, and $\text{marg}_{S_{-i}} \lambda_i(t_i^{s_i})$ is strongly independent. By Lemma 1, such types exist since, for each i , $P_i \neq \emptyset$ and $P_i = a_i(P_{-i})$. Furthermore, assume that, for all $(s_{-i}, t_{-i}) \in \Omega_{-i}$, $\mu_i^1(t_i^{s_i})(s_{-i}, t_{-i}) > 0$ only if, for all $j \neq i$ and $s_j \in P_j$, $t_j = t_j^{s_j}$. Write, for each $i \in \mathcal{I}$, $T_i := \{t_i = t_i^{s_i} \mid s_i \in P_i\}$. The definitions of $[\text{rat}]$, $[\text{cau}]$, and $[\text{ind}]$ imply

$$\{(s_1, \dots, s_I, t_1, \dots, t_I) \mid \text{for all } i \in \mathcal{I}, s_i \in P_i \text{ and } t_i = t_i^{s_i}\} \subseteq CB([\text{rat}] \cap [\text{cau}] \cap [\text{ind}]).$$

Hence, for each player $i \in \mathcal{I}$ and any strategy $s_i \in P_i$, $((S_i)_{i \in \mathcal{I}}, (T_i)_{i \in \mathcal{I}}, (\lambda_i)_{i \in \mathcal{I}})$ has the property that $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in CB([\text{rat}] \cap [\text{cau}] \cap [\text{ind}])$.

Part 2: For each player $i \in \mathcal{I}$, if $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in CB([\text{rat}] \cap [\text{cau}] \cap [\text{ind}])$, where $((S_i)_{i \in \mathcal{I}}, (T_i)_{i \in \mathcal{I}}, (\lambda_i)_{i \in \mathcal{I}})$ is an S -based interactive belief structure, then $s_i \in P_i$. If, for $i \in \mathcal{I}$, $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in [\text{rat}] \cap [\text{cau}] \cap [\text{ind}]$, then $s_i \in a_i(S_{-i})$. Let, for all $i \in \mathcal{I}$,

$$S'_i = \{s_i \in S_i \mid s_i = \mathbf{s}_i(\omega) \text{ for some } \omega \in B^{k-1}([\text{rat}] \cap [\text{cau}] \cap [\text{ind}]), \text{ where } k \in \mathbb{N}\}.$$

Then if, for $i \in \mathcal{I}$, $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in B^k([\text{rat}] \cap [\text{cau}] \cap [\text{ind}])$, then $s_i \in a_i(S'_{-i})$. It now follows from the definition of P_i that $s_i \in P_i$ if $s_i = \mathbf{s}_i(\omega)$ for some $\omega \in CB([\text{rat}] \cap [\text{cau}] \cap [\text{ind}])$. \square

The epistemic characterization of the Independent Dekel-Fudenberg Procedure can also be illustrated in the game of Figure 2, by noting that there is common belief of rationality, caution, and stochastically independent beliefs in the state $((\text{Out}, t_1^{\text{Out}}), (d, t_2^d), (D, t_3^D))$, where the types t_1^{Out} , t_2^d , and t_3^D are defined as above. In particular, $\text{marg}_{S_{-1}} \lambda_1(t_1^{\text{Out}})$, $\text{marg}_{S_{-2}} \lambda_2(t_2^d)$, and $\text{marg}_{S_{-3}} \lambda_3(t_3^D)$ are strongly independent, since by aggregating the three levels of these LPSs by the NPS $\tilde{\nu}$, where $\tilde{\nu}(1) = (1 - \varepsilon)^2$, $\tilde{\nu}(2) = 2(\varepsilon - \varepsilon^2)$, and $\tilde{\nu}(3) = \varepsilon^2$, it follows that the aggregated NPSs are product distributions. This corresponds to the fact that Out for player 1, d for player 2, and D for player 3 survive the Independent Dekel-Fudenberg Procedure. In contrast, $\text{marg}_{S_{-2}} \lambda_2(t_2^c)$, where t_2^c is defined as above, is not strongly independent, reflecting that c for player 2 and In for player 1 do not survive the Independent Dekel-Fudenberg Procedure.

Combined with Propositions 1–3, these results imply that stochastically independent beliefs are an essential ingredient in an epistemic characterization of the backward induction paradox.

7 Discussion

Requiring that beliefs about opponents' choices are stochastically independent in games with more than two players was the traditional view in game theory, as reflected by equilibrium concepts (like Nash equilibrium and strategic-form perfect equilibrium) and non-equilibrium concepts (like rationalizability as originally defined by Bernheim, 1984, and Pearce, 1984). Over the years, however, this view has been challenged with the argument that players can have stochastically dependent beliefs about the choices of opponents even though the opponents choose independently. Moreover, allowing for correlated beliefs leads to the strategies that are never best replies being exactly those that are dominated.

The Dekel-Fudenberg Procedure is an uncontroversial solution concept, as it eliminates only those strategies that cannot be rational if rationality and caution are commonly believed. The question of whether stochastic independence of beliefs about opponents' choices should also be imposed, leading to the Independent Dekel-Fudenberg Procedure, might be made subject to empirical analysis by designing experiments which compare games like those depicted in Figures 1 and 2 as different treatment. We are not aware of any such experiments,⁴ and answering this question is beyond the scope of the present paper. Its purpose has been to point out that this refinement of the Dekel-Fudenberg Procedure can be used to interpret the backward induction paradox (as shown by Proposition 1 and 3), and that its epistemic characterization (Proposition 5) thereby yields an epistemic foundation of this paradox.

Instead of using the Dekel-Fudenberg Procedure as our point of departure, we could have used other concepts that always yield backward induction in 2-player games where the each player moves only once, but which might lead to outcome incompatible with backward induction if players move more than once. The concept of *Fully Permissible Sets* as defined and epistemically characterized by Asheim and Dufwenberg (2003a) for 2-player games and applied to extensive games in Asheim and Dufwenberg (2003b) does have these properties. The concept essentially yields the same prediction as the Dekel-Fudenberg Procedure in the game of Figure 1, while being more restrictive by yielding the backward induction outcome in the game of Reny (1992a, Fig. 1). Asheim and Perea (2019, Def. 9) generalize this concept to games with more than two players without imposing stochastically independent beliefs. If instead stochastic independence is imposed when generalizing Fully Permissible Sets to such games, this concept would yield an alternative interpretation and epistemic foundation of the backward induction paradox.

⁴The experimental results of Dufwenberg and Van Essen (2018) show that backward induction might not obtain even if each player moves only once, in games where the backward induction strategy for each player depends on whether there is an even or odd number of remaining players. This can be interpreted as a test of the common belief assumption rather than the assumption that beliefs are stochastically independent.

References

- Arieli, I. and Aumann, R.J. (2015). ‘The logic of backward induction.’ *Journal of Economic Theory* **159**, 443–464.
- Asheim, G.B. (2002). ‘On the epistemic foundation for backward induction.’ *Mathematical Social Sciences* **44**, 121–144.
- Asheim, G.B. and Dufwenberg, M. (2003a). ‘Admissibility and common belief.’ *Games and Economic Behavior* **42**, 208–234.
- Asheim, G.B. and Dufwenberg, M. (2003b). ‘Deductive reasoning in extensive games.’ *Economic Journal* **113**, 305–325.
- Asheim, G.B. and Perea, A. (2005). ‘Sequential and quasi-perfect rationalizability in extensive games.’ *Games and Economic Behavior* **53**, 15–42.
- Asheim, G.B. and Perea, A. (2019). ‘Algorithms for cautious reasoning in games.’ *International Journal of Game Theory* **48**, 1241–1275.
- Aumann, R.J. (1995). ‘Backward induction and common knowledge of rationality.’ *Games and Economic Behavior* **8**, 6–19.
- Basu, K. (1988). ‘Strategic irrationality in extensive games.’ *Mathematical Social Sciences* **15**, 247–260.
- Battigalli, P. (1997). ‘On rationalizability in extensive games.’ *Journal of Economic Theory* **74**, 40–61.
- Battigalli, P. and De Vito, N. (2021). ‘Beliefs, plans, and perceived intentions in dynamic games.’ *Journal of Economic Theory* **195**, 105283.
- Battigalli, P. and Siniscalchi, M. (2002). ‘Strong belief and forward induction reasoning.’ *Journal of Economic Theory* **106**, 356–391.
- Ben-Porath, E. (1997). ‘Rationality, Nash equilibrium and backwards induction in perfect-information games.’ *Review of Economic Studies* **64**, 23–46.
- Bernheim, B.D. (1984). ‘Rationalizable strategic behavior.’ *Econometrica* **52**, 1007–1028.
- Bicchieri, C. (1989). ‘Self-refuting theories of strategic interaction: A paradox of common knowledge.’ *Erkenntnis* **30**, 69–85.
- Binmore, K. (1987). ‘Modeling rational players: Part i.’ *Economics & Philosophy* **3**, 179–214.
- Blume, L., Brandenburger, A., and Dekel, E. (1991a). ‘Lexicographic probabilities and choice under uncertainty.’ *Econometrica* **59**, 61–79.
- Blume, L., Brandenburger, A., and Dekel, E. (1991b). ‘Lexicographic probabilities and equilibrium refinements.’ *Econometrica* **59**, 81–98.
- Börgers, T. (1994). ‘Weak dominance and approximate common knowledge.’ *Journal of Economic Theory* **64**, 265–276.
- Brandenburger, A. (1992). ‘Lexicographic probabilities and iterated admissibility.’ In *Economic Analysis of Markets and Games* (eds. P. Dasgupta, D. Gale, O. Hart, E. Maskin), pp. 282–290. Cambridge, MA: MIT Press.

- Brandenburger, A. (2007). ‘The power of paradox: some recent developments in interactive epistemology.’ *International Journal of Game Theory* **35**, 465–492.
- Brandenburger, A., Friedenberg, A., and Keisler, H.J. (2008). ‘Admissibility in games.’ *Econometrica* **76**, 307–352.
- Dekel, E. and Fudenberg, D. (1990). ‘Rational behavior with payoff uncertainty.’ *Journal of Economic Theory* **52**, 243–67.
- Dekel, E., Fudenberg, D., and Levine, D. (1999). ‘Payoff information and self-confirming equilibrium.’ *Journal of Economic Theory* **89**, 165–185.
- Dekel, E., Fudenberg, D., and Levine, D. (2002). ‘Subjective uncertainty over behavior strategies: A correction.’ *Journal of Economic Theory* **104**, 473–478.
- Dufwenberg, M. and Van Essen, M. (2018). ‘King of the hill: Giving backward induction its best shot.’ *Games and Economic Behavior* **112**, 125–138.
- Halpern, J.Y. (2010). ‘Lexicographic probability, conditional probability, and nonstandard probability.’ *Games and Economic Behavior* **68**, 155–179.
- Hammond, P.J. (1994). ‘Elementary non-archimedean representations of probability for decision theory and games,’ In *Patrick Suppes: Scientific Philosopher*, Vol. 1 (ed. P. Humphreys), pp. 25–49. Dordrecht: Kluwer.
- Hu, T.-W. (2007). ‘On p -rationalizability and approximate common certainty of rationality.’ *Journal of Economic Theory* **136**, 379–391.
- Kreps, D.M., Milgrom, P., Roberts, J., and Wilson, R. (1982). ‘Rational cooperation in the infinitely repeated prisoners’ dilemma.’ *Journal of Economic Theory* **27**, 245–252.
- Kreps, D.M. and Wilson, R. (1982). ‘Reputation and imperfect information.’ *Journal of Economic Theory* **27**, 253–279.
- Luce, D. and Raiffa, H. (1957). *Games and Decisions*. New York: Wiley.
- Mas-Colell, A., Whinston, M., and Green, J. (1995). *Microeconomic Theory*. Oxford: Oxford University Press.
- Meier, M. and Perea, A. (2020). ‘Reasoning about your own future mistakes.’ Mimeo, EPICENTER Working Paper No. 21.
- Milgrom, P. and Roberts, J. (1982). ‘Predation, reputation, and entry deterrence.’ *Journal of Economic Theory* **27**, 280–312.
- Pearce, D.G. (1984). ‘Rationalizable strategic behavior and the problem of perfection.’ *Econometrica* **52**, 1029–1050.
- Penta, A. (2015). ‘Robust dynamic implementation.’ *Journal of Economic Theory* **160**, 280–316.
- Perea, A. (2007). ‘Epistemic foundations for backward induction: an overview.’ In *Interactive logic. Proceedings of the 7th Augustus de Morgan Workshop Volume 1*, pp. 159–193).
- Perea, A. (2008). ‘Minimal belief revision leads to backward induction.’ *Mathematical Social Sciences* **56**, 1–26.

- Perea, A. (2014). ‘Belief in the opponents future rationality.’ *Games and Economic Behavior* **83**, 231–254.
- Pettit, P. and Sugden, R. (1989). ‘The backward induction paradox.’ *The Journal of Philosophy* **86**, 169–182.
- Reny, P.J. (1985). ‘Rationality, common knowledge, and the theory of games.’ Mimeo, Department of Economics, Princeton University.
- Reny, P.J. (1988). ‘Knowledge and games with perfect information.’ In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Volume Two: Symposia and invited papers*, pp. 363–369.
- Reny, P.J. (1992a). ‘Backward induction, normal form perfection and explicable equilibria.’ *Econometrica* **60**, 627–649.
- Reny, P.J. (1992b). ‘Rationality in extensive-form games.’ *Journal of Economic Perspectives* **6**, 103–118.
- Reny, P.J. (1993). ‘Common belief and the theory of games with perfect information.’ *Journal of Economic Theory* **59**, 257–274.
- Robertson, A. (1973). ‘Function theory on some nonarchimedean fields.’ *The American Mathematical Monthly* **80**, 87–109.
- Rosenthal, R.W. (1981). ‘Games of perfect information, predatory pricing, and the chain-store paradox.’ *Journal of Economic Theory* **25**, 92–100.
- Rubinstein, A. (1991). ‘Comments on the interpretation of game theory.’ *Econometrica* **59**, 909–924.
- Schuhmacher, F. (1999). ‘Proper rationalizability and backward induction.’ *International Journal of Game Theory* **28**, 599–615.
- Selten, R. (1975). ‘Reexamination of the perfectness concept for equilibrium points in extensive games.’ *International Journal of Game Theory* **4**, 25–55.
- Selten, R. (1978). ‘The chain store paradox.’ *Theory and Decision* **9**, 127–159.
- Sobel, J.H. (1993). ‘Backward-induction arguments: A paradox regained.’ *Philosophy of Science* **60**, 114–133.
- Stalnaker, R. (1998). ‘Belief revision in games: forward and backward induction.’ *Mathematical Social Sciences* **36**, 31–56.