Finite Reasoning Procedures for Dynamic Games

Prepared for edited volume on “Strategic reasoning in dynamic games”

Andrés Perea
Maastricht University

This Version: January 2012

Abstract

In this chapter we focus on the epistemic concept of *common belief in future rationality* (Perea (2011a)), which describes a backward induction type of reasoning for general dynamic games. It states that a player always believes that his opponents will choose rationally now and in the future, always believes that his opponents always believe that their opponents choose rationally now and in the future, and so on, *ad infinitum.* It thus involves infinitely many conditions, which might suggest that this concept is demanding for real players in a game. In this chapter we show, however, that this is not true. For finite dynamic games we present some *finite* reasoning procedures that a player can use to reason his way towards common belief in future rationality.

1. Introduction

If you make a choice in a game, then you must realize that the final outcome does not only depend on your own choice, but also on the choices of your opponents. It is therefore natural that you first *reason* about your opponents in order to form a *plausible belief* about their choices, before you make your own choice. Now, how can we formally model such reasoning procedures about your opponents? And how do these reasoning procedures affect the choice you will eventually make in the game? These questions naturally lead to *epistemic game theory* – a modern approach to game theory which takes seriously the fact that the players in a game are *human beings* who reason before they reach their final decision.

In our view, the most important idea in epistemic game theory is *common belief in rationality* (Tan and Werlang (1988), see also Brandenburger and Dekel (1987)). It states that a player, when making his choice, chooses *optimally* given the belief he holds about the opponents' choices. Moreover, the player also believes that his opponents will choose optimally as well, and that their opponents believe that the other players will also choose optimally, and so on, *ad infinitum.* This idea really constitutes the basis for epistemic game theory, as most – if not all – concepts within epistemic game theory can be viewed as some variant of *common belief in rationality.*

For dynamic games there is a backward induction analogue to common belief in rationality, namely *common belief in future rationality* (Perea (2011a)). This concept states that a player, at each of his information sets, believes that his opponents will choose rationally *now and in the future.* However, it does not require a player to believe that his opponents have chosen rationally in the past! On top of this, the concept states that a player also always believes that his opponents, at each of their information sets, believe that their opponents will choose rationally now and in the future, and so on, *ad infinitum.*

For dynamic games with perfect information, various authors have used some variant of the idea of *common belief in future rationality* as a possible foundation for backward induction. See Asheim (2002), Baltag, Smets and Zvesper (2009), Feinberg (2005) and Samet (1996). Among these contributions, the concept of *stable belief in dynamic rationality* in Baltag, Smets and Zvesper (2009) matches completely the idea of *common belief in future rationality,* although they restrict attention to non-probabilistic beliefs. Perea (2007a) provides an overview of the various epistemic foundations for backward induction that have been offered in the literature.

Some people have critized *common belief in rationality* because it involves *infinitely* many conditions, and hence – they argue – it will be very difficult for a player to meet each of these infinitely many conditions. The same could be said about *common belief in future rationality.* The main purpose of this chapter will be to show that this critique is actually not justified, provided we stick to *finite* games. We will show, namely, that in dynamic games with finitely many information sets, and finitely many choices at every information set, *common belief in future rationality* can be achieved by reasoning procedures that use *finitely* many steps only!

Let us be more precise about this statement. Suppose a player in a dynamic game holds not only conditional beliefs about his opponents' strategies, but also conditional beliefs about his opponents' conditional beliefs about the other players' strategies, and so on, *ad infinitum.* That is, this player holds a full *belief hierarchy* about his opponents – an object that is needed in order to formally define *common belief in future rationality.* Such belief hierarchies can be efficiently encoded within an *epistemic model with types,* in which every type holds, at each of his information sets, a conditional belief about the opponents' strategies *and types.* Within such an epistemic model, we can then *derive* for every type a *full hierarchy of conditional beliefs* about the opponents' strategies and beliefs. This construction is based on Harsanyi's (1967–1968) seminal way of encoding belief hierarchies for games with incomplete information. If a belief hierarchy can be derived from an epistemic model with *finitely many types* only, we say that this belief hierarchy is *finitely generated.* Such finitely generated belief hierarchies will play a central role in this chapter, as we will show that they are "sufficient" when studying *common belief in future rationality* in finite dynamic games.

Let us now come back to the question whether *common belief in future rationality* can be achieved by finite reasoning procedures. As a first step, we show in Section 3 that for a *finitely generated* belief hierarchy, it only takes *finitely* many steps to verify whether this given belief hierarchy expresses *common belief in future rationality* or not. So, although *common belief in future rationality* involves infinitely many conditions, checking these conditions can be reduced to a *finite* procedure whenever we consider belief hierarchies that are finitely generated.

Next, we study the question whether there is an easy way to *generate* a belief hierarchy that expresses *common belief in future rationality.* We answer that question in Section 4, where we present an easy, finite reasoning procedure that a player can use to produce *some* belief hierarchy that expresses *common belief in future rationality.* The procedure starts by specifying, at every information set, a choice for the active players at that information set, and some history of past choices that leads to this information set. We call this a *choice-history cobination.* In the next round this initial choice-history combination generates a *new* choice-history combination, essentially by specifying, at every information set, choices that are optimal given the *future* choices in the *initial* choice-history combination. By repeating these steps, the procedure generates a chain of choice-history combinations. As long as the game is finite, this chain of choice-history combinations will eventually enter into a cycle after finitely many steps. We then show that this cycle of choice-history combinations can easily be transformed into a finite epistemic model in which every type expresses *common belief in future rationality.* In particular, this procedure yields for every player *some* belief hierarchy that expresses *common belief in future rationality.* As this procedure involves only finitely many steps, we obtain an easy, finite reasoning procedure that a player can use to reason his way towards *some* belief hierarchy that expresses *common belief in future rationality.*

In Section 5 our goal is more ambitious, as we wish to develop a finite reasoning procedure that leads a player towards *all* strategies he can rationally make under *common belief in future rationality.* This reasoning procedure is based on the *backward dominance procedure* (Perea (2011a)), which is a recursive elimination procedure that delivers all strategies that can rationally be made under *common belief in future rationality.* For finite dynamic games this algorithm terminates after finitely many steps. We then show how to construct – on the basis of this algorithm – a finite epistemic model in which every type expresses *common belief in future rationality.* Moreover, for every remaining strategy in the algorithm there will be a type within that epistemic model for which that strategy is optimal. Hence, for every strategy that can rationally be chosen under *common belief in future rationality*, this reasoning procedure yields a finitely generated belief hierarchy, expressing *common belief in future rationality*, for which this strategy is optimal. As this procedure ends within finitely many steps, we thus offer a *finite* reasoning procedure that will lead the player to *all* strategies he can rationally make under *common belief in future rationality.* Moreover, this procedure shows that for *common belief in future rationality* it is enough to consider *finitely generated* belief hierarchies. Namely, every strategy that is optimal for *some* belief hierarchy – not necessarily finitely generated – that expresses *common belief in future rationality,* is also optimal for some *finitely generated* belief hierarchy that expresses *common belief in future rationality.*

Overall, we thus see that in finite dynamic games the concept of *common belief in future rationality* can be characterized by *finite* reasoning procedures. Moreover, for such games it is sufficient to use only *finitely generated* belief hierarchies when focusing on *common belief in future rationality.* So, although *common belief in future rationality* involves infinitely many conditions, these conditions can be "reduced" to finitely many steps in these games.

Static games are just a special case of dynamic games, where every player only makes a choice once, and where all players choose simultaneously. It is clear that in static games, the concept of *common belief in future rationality* reduces to the basic concept of *common belief in rationality.* As such, the results in this chapter immediately carry over to *common belief in rationality* as well. Hence, also the concept of *common belief in rationality* in finite static games can be characterized by *finite* reasoning procedures, just by applying the reasoning procedures in this chapter to the special case of static games.

This chapter can therefore be seen as an answer to the critique that epistemic concepts like *common belief in rationality* and *common belief in future rationality* would be too demanding because of the infinitely many conditions. We believe this critique is not justified. Similar conclusions can be drawn for various *other* epistemic concepts in the literature – apart from *common belief in (future) rationality* – as we will see in Section 6.

The outline of the chapter is as follows. In Section 2 we formally define the idea of *common belief in future rationality* within an epistemic model. Section 3 presents a finite reasoning procedure to verify whether a finitely generated belief hierarchy expresses *common belief in future rationality* or not. In Section 4 we develop an easy, finite procedure that a player can use to reason his way towards *some* belief hierarchy that expresses *common belief in future rationality.* Section 5 presents a finite reasoning procedure – based on the *backward dominance procedure* (Perea (2011a)) – which yields, for *every* strategy that can rationally be chosen under *common belief in future rationality*, some finitely generated belief hierarchy expressing *common belief in future rationality* which supports that strategy. We conclude the chapter with a discussion in Section 6. For simplicity, we stick to two-player games throughout this chapter. However, all ideas and results can easily be extended to games with more than two players.

2. Common Belief in Future Rationality

In this section we present the idea of *common belief in future rationality* (Perea (2011a)), and show how it can be formalized within an epistemic model with types.

2.1. Main Idea

Common belief in future rationality (Perea (2011a)) reflects the idea that you believe, at each of your information sets, that your opponent will choose rationally *now* and in the *future,* but not necessarily that he chose rationally in the *past.* In fact, in some dynamic games it is simply *impossible* to believe, at certain information sets, that your opponent has chosen rationally in the past, as this information set can only be reached through a suboptimal choice by the opponent. But it is *always* possible to believe that your opponent will choose rationally now and in the future. On top of this, *common belief in future rationality* also states that you always believe that your opponent reasons in precisely this way as well. That is, you always believe that your opponent, at each of his information sets, believes that you will choose rationally *now* and in the *future*. By iterating this thought process *ad infinitum* we eventually arrive at *common belief in future rationality.*

2.2. Dynamic Games

We now wish to formalize the idea of *common belief in future rationality.* As a first step, we formally introduce *dynamic games.* As already announced in the introduction, we will restrict attention to two-player games for simplicity, although everything in this chapter can easily be generalized to games with more than two players. At the same time, the model of a dynamic game presented here is a bit more general than usual, as we explicitly allow for *simultaneous* choices by players at certain stages of the game. In our setup, a *dynamic game* is a tuple $\Gamma = (I, X, Z, (X_i, C_i, H_i, u_i)_{i \in I})$ where

(a) $I = \{1, 2\}$ is the set of *players;*

(b) X is the set of *non-terminal histories.* Every non-terminal history $x \in X$ represents a situation where one or more players must make a choice;

(c) Z is the set of *terminal histories.* Every terminal history $z \in Z$ represents a situation where the game ends;

(d) $X_i \subseteq X$ is the set of histories at which player i must make a choice. At every history $x \in X$ at least one player must make a choice, that is, for every $x \in X$ there is at least some i with $x \in X_i$. However, for a given history x there may be various players i with $x \in X_i$. This models a situation where various players simultaneously choose at x. For a given history $x \in X$, we denote by $I(x) := \{i \in I : x \in X_i\}$ the set of *active players* at x;

(e) C_i assigns to every history $x \in X_i$ the set of *choices* $C_i(x)$ from which player i can choose at x;

(f) H_i is the collection of *information sets* for player i. Formally, $H_i = \{h_i^1, ..., h_i^K\}$ where $h_i^k \subseteq X_i$ for every k, the sets h_i^k are mutually disjoint, and $X_i = \cup_k h_i^k$. The interpretation of an information set $h \in H_i$ is that at h player i knows that some history in h has been realized, without knowing precisely which one;

(g) u_i is player i's *utility function,* assigning to every terminal history $z \in Z$ some utility $u_i(z)$ in $\mathfrak{R}$.

Throughout this chapter we assume that all sets above are finite. The histories in X and Z consist of finite sequences of choice-combinations

$$((c_i^1)_{i \in I^1}, (c_i^2)_{i \in I^2}, ..., (c_i^K)_{i \in I^K}),$$

where $I^1, ..., I^K$ are nonempty subsets of players, such that

(a) $\emptyset$ (the empty sequence) is in X,

(b) if $x \in X$ and $(c_i)_{i \in I(x)} \in \prod_{i \in I(x)} C_i(x)$, then $(x, (c_i)_{i \in I(x)}) \in X \cup Z$,

(c) if $z \in Z$, then there is no choice combination $(c_i)_{i \in \hat{I}}$ such that $(z, (c_i)_{i \in \hat{I}}) \in X \cup Z$,

(d) for every $x \in X \cup Z$, $x \neq \emptyset$, there is a unique $y \in X$ and $(c_i)_{i \in I(y)} \in \prod_{i \in I(y)} C_i(y)$ such that $x = (y, (c_i)_{i \in I(y)})$.

Hence, a history $x \in X \cup Z$ represents the sequence of choice-combinations that have been made by the players until this moment.

Moreover, we assume that the collections H_i of information sets are such that

(a) two histories in the same information set for player i have the same set of available choices for player i. That is, for every $h \in H_i$, and every $x, y \in h$, it holds that $C_i(x) = C_i(y)$. This condition must hold since player i is assumed to know his set of available choices at h. We can thus speak of $C_i(h)$ for a given information set $h \in H_i$;

(b) two histories in the same information set for player i must pass through exactly the same collection of information sets for player i, and must hold exactly the same past choices for player i. This condition guarantees that player i has *perfect recall,* that is, at every information set $h \in H_i$ player i remembers the information he possessed before, and the choices he made before.

Say that an information set h *follows* some other information set h' if there are histories $x \in h$ and $y \in h'$ such that $x = (y, (c_i^1)_{i \in I^1}, (c_i^2)_{i \in I^2}, ..., (c_i^K)_{i \in I^K})$ for some choice-combinations $(c_i^1)_{i \in I^1}, (c_i^2)_{i \in I^2}, ..., (c_i^K)_{i \in I^K}$. The information sets h and h' are called *simultaneous* if there is some history x with $x \in h$ and $x \in h'$. Finally, we say that information set h *weakly follows* h' if either h follows h', or h and h' are simultaneous.

2.3. Strategies

A strategy for player i is a complete choice plan, prescribing a choice at each of his information sets that can possibly be reached by this choice plan. Formally, for every $h, h' \in H_i$ such that h precedes h', let $c_i(h, h')$ be the choice at h for player i that leads to h'. Note that $c_i(h, h')$ is unique by perfect recall. Consider a subset $\hat{H}_i \subseteq H_i$, not necessarily containing all information sets for player i, and a function s_i that assigns to every $h \in \hat{H}_i$ some choice $s_i(h) \in C_i(h)$. We say that s_i *possibly reaches* an information set h if at every $h' \in \hat{H}_i$ preceding h we have that $s_i(h') = c_i(h', h)$. By $H_i(s_i)$ we denote the collection of player i information sets that s_i possibly reaches. A *strategy* for player i is a function s_i, assigning to every $h \in \hat{H}_i \subseteq H_i$ some choice $s_i(h) \in C_i(h)$, such that $\hat{H}_i = H_i(s_i)$.

Notice that this definition slightly differs from the standard definition of a strategy in the literature. Usually, a strategy for player i is defined as a mapping that assigns to *every* information set $h \in H_i$ some available choice – also to those information sets h that cannot be reached by s_i. The definition of a strategy we use corresponds to what Rubinstein (1991) calls a *plan of action.* One can also interpret it as the equivalence class of strategies (in the classical sense) that are outcome-equivalent. Hence, taking for every player the set of strategies as we use it corresponds to considering the pure strategy reduced normal form. However, for the concepts and results in this chapter it does not make any difference which notion of strategy we use.

For a given information set h, denote by $S_i(h)$ the set of strategies for player i that possibly reach h. By $S(h)$ we denote the set of strategy profiles $(s_i)_{i \in I}$ that reach some history in h.

2.4. Epistemic Model

We say that a strategy is *rational* for you at a certain information set if it is optimal at that information set, *given your conditional belief there* about the opponent's strategy choice. In order to believe that your *opponent* chooses rationally at a certain information set, you must therefore not only hold conditional beliefs about the opponent's strategy choice, but also conditional beliefs about the opponent's conditional beliefs about your strategy choice. This is what we call a *second-order* belief. Moreover, if we go one step further and want to model the event that you believe that your opponent believes that you choose rationally, we need not only your belief about the opponent's beliefs about your strategy choice, but also your belief about the opponent's beliefs about your beliefs about the opponent's strategy choice – that is, your *third-order* belief. Consequently, formally defining the idea of *common belief in future rationality* requires us to consider *infinite belief hierarchies,* specifying your conditional beliefs about the opponent's strategy choice, your conditional beliefs about the opponent's conditional beliefs about your strategy choice, and so on *ad infinitum.*

A problem with infinite belief hierarchies is that writing them down explicitly is an impossible task, since we would need to write down infinitely many beliefs. So is there a way to efficiently *encode* such infinite belief hierarchies without writing too much? The answer is "yes", as we will see right now. Note that your belief hierarchy specifies first-order beliefs about the opponent's strategy choice, second-order beliefs about the opponent's first-order beliefs about your strategy choice, third-order beliefs about the opponent's second-order beliefs, and so on. Hence we conclude that your belief hierarchy specifies conditional beliefs about the opponent's strategy choice *and the opponent's belief hierarchy.* Now let us call every belief hierarchy a *type.* Then, every type can be identified with its conditional beliefs about the opponent's strategy choice *and the opponent's type.* This elegant and powerful idea goes back to Harsanyi (1967–1968), who used it to model infinite belief hierarchies in games with incomplete information.

Let us now implement this idea of encoding infinite belief hierarchies formally. Fix some finite dynamic game Γ with two players.

Definition 2.1 (Epistemic model). *A finite* ***epistemic model*** *for the game* Γ *is a tuple* $M = (T_1, T_2, b_1, b_2)$ *where*

(a) T_i *is the finite set of types for player* i, *and*

(b) b_i *assigns to every type* $t_i \in T_i$ *and every information set* $h \in H_i$ *some probabilistic belief* $b_i(t_i, h) \in \Delta(S_j(h) \times T_j)$ *about opponent* j*'s strategy-type pairs.*

Remember that $S_j(h)$ denotes the set of strategies for opponent j that possibly reach h. By $\Delta(S_j(h) \times T_j)$ we denote the set of probability distributions on $S_j(h) \times T_j$. So, within an epistemic model every type holds at each of his information sets a conditional belief about the opponent's strategy choice and the opponent's type, as we discussed above. For every type $t_i \in T_i$ we can now *derive* its complete belief hierarchy from the belief functions b_i and b_j. Namely, type t_i holds at information set $h \in H_i$ a conditional belief $b_i(t_i, h)$ on $S_j(h) \times T_j$. By taking the marginal of $b_i(t_i, h)$ on $S_j(h)$ we obtain t_i's *first-order* belief at h on j's strategy choice. Moreover, t_i holds at information set $h \in H_i$ a conditional belief about j's possible types. As each of j's types t_j holds first-order conditional beliefs on i's strategy choices, we can thus derive from b_i and b_j the *second-order* conditional belief that t_i holds at $h \in H_i$ about j's first-order beliefs about i's strategy choice. By continuing this procedure we can thus deduce, for every type t_i in the model, each of its belief levels by making use of the belief functions b_i and b_j. In this way, the epistemic model above can be viewed as a short and convenient way to *encode* the infinite belief hierarchy of a player.

By means of this epistemic model we can in particular model the *belief revision* of players during the game. Consider two different information sets h and h' for player i, where h' comes after h. Note that type t_i's conditional belief at h' about j's strategy choice may be different from his conditional belief at h, and hence a type t_i may *revise* his belief about j's strategy choice as the game moves from h to h'. Moreover, t_i's conditional belief at h' about j's *type* may be different from his conditional belief at h, and hence a type t_i may *revise* his belief about j's *type* – and hence j's conditional beliefs – as the game moves from h to h'. So, all different kinds of belief revisions – about the opponent's strategy, but also about the opponent's beliefs – can be captured within this epistemic model.

Note that in our definition above we require the sets of types to be *finite.* This imposes a restriction on the possible belief hierarchies we can encode, since not every belief hierarchy can be derived from a type within an epistemic model with *finite* sets of types. For some belief hierarchies we would need *infinitely* many types to encode them. Belief hierarchies that *can* be derived from a *finite* epistemic model will be called *finitely generated.*

Definition 2.2 (Finitely generated belief hierarchy). *A belief hierarchy β_i for player i is* ***finitely generated*** *if there is some finite epistemic model $M = (T_1, T_2, b_1, b_2)$, and some type $t_i \in T_i$ in that model, such that β_i is the belief hierarchy induced by t_i within M.*

Throughout this chapter we will restrict attention to *finite* epistemic models, and hence to *finitely generated* belief hierarchies. We will see in Section 5 that this is not a serious restriction within the context of *common belief in future rationality*, as every strategy that is optimal for *some* belief hierarchy – not necessarily finitely generated – that expresses *common belief in future rationality*, is also optimal for some *finitely generated* belief hierarchy that expresses *common belief in future rationality.* Moreover, finitely generated belief hierarchies are much easier to work with than those that are not finitely generated.

2.5. Common Belief in Future Rationality: Formal Definition

Remember that *common belief in future rationality* states that you always believe that the opponent chooses rationally now and in the future, you always believe that the opponent always believes that you choose rationally now and in the future, and so on, *ad infinitum.* Within an epistemic model we can state these conditions formally.

We first define what it means for a strategy s_i to be optimal for a type t_i at a given information set h. Consider a type t_i, a strategy s_i and an information set $h \in H_i(s_i)$ that is possibly reached by s_i. By $u_i(s_i, t_i \mid h)$ we denote the expected utility from choosing s_i under the conditional belief that t_i holds at h about the opponent's strategy choice.

Definition 2.3 (Optimality at a given information set). *Consider a type t_i, a strategy s_i and a history $h \in H_i(s_i)$. Strategy s_i is optimal for type t_i at h, if $u_i(s_i, t_i \mid h) \geq u_i(s_i', t_i \mid h)$ for all $s_i' \in S_i(h)$.*

Remember that $S_i(h)$ is the set of player i strategies that possibly reach h. We can now define belief in the opponent's future rationality.

Definition 2.4 (Belief in the opponent's future rationality). *Consider a type t_i and an information set $h \in H_i$. Type t_i believes at h in j's future rationality if $b_i(t_i, h)$ only assigns positive probability to j's strategy-type pairs (s_j, t_j) where s_j is optimal for t_j at every $h' \in H_j(s_j)$ that weakly follows h. Type t_i believes in the opponent's future rationality if he does so at every information set $h \in H_i$.*

So, to be precise, a type that believes in the opponent's future rationality believes that the opponent chooses rationally now (if the opponent makes a choice at a simultaneous information set), and at every information set that follows. As such, the correct terminology would be "belief in the opponent's *present* and future rationality", but we stick to "belief in the opponent's future rationality" as to keep the name short.

We can then formally define the other conditions in common belief in rationality in an inductive manner.

Definition 2.5 (Common belief in future rationality). *Consider a finite epistemic model $M = (T_1, T_2, b_1, b_2)$.*

(Induction start) A type $t_i \in T_i$ is said to express 1-fold belief in future rationality if t_i believes in j's future rationality.

(Induction step) For every $k \geq 2$, a type $t_i \in T_i$ is said to express k-fold belief in future rationality if at every information set $h \in H_i$, the belief $b_i(t_i, h)$ only assigns positive probability to j's types t_j that express $(k-1)$-fold belief in future rationality.

Type $t_i \in T_i$ is said to express ***common belief in future rationality*** *if it expresses k-fold belief in future rationality for all k.*

Finally, we define those strategies that can rationally be chosen under *common belief in future rationality.* We say that a strategy s_i is *rational* for a type t_i if s_i is optimal for t_i at every $h \in H_i(s_i)$. In the literature, this is often called *sequential rationality.* We say that strategy s_i can *rationally be chosen under common belief in future rationality* if there is some epistemic model $M = (T_1, T_2, b_1, b_2)$, and some type $t_i \in T_i$, such that t_i expresses *common belief in future rationality*, and s_i is rational for t_i.

3. Checking Common Belief in Future Rationality

Some people have criticized the concept of *common belief in rationality*, because one has to verify *infinitely many* conditions in order to conclude that a given belief hierarchy expresses *common belief in rationality.* The same could be said about *common belief in future rationality.* We will show in this section that this is *not* true for *finitely generated* belief hierarchies. Namely, verifying whether a *finitely generated* belief hierarchy expresses *common belief in future rationality* or not only requires checking *finitely* many conditions, and can usually be done very quickly. To that purpose we present a reasoning procedure with finitely many steps which, for a given finitely generated belief hierarchy, tells us whether that belief hierarchy expresses *common belief in future rationality* or not.

Consider an epistemic model $M = (T_1, T_2, b_1, b_2)$ with finitely many types for both players. For every type $t_i \in T_i$, let $T_j(t_i)$ be the set of types for player j that type t_i deems possible at some of his information sets. That is, $T_j(t_i)$ contains all types $t_j \in T_j$ such that $b_i(t_i, h)(c_j, t_j) > 0$ for some $h \in H_i$ and some $c_j \in C_j$. We recursively define the sets of types $T_j^k(t_i)$ and $T_i^k(t_i)$ as follows.

Algorithm 3.1 (Relevant types for t_i). *Consider a finite dynamic game* Γ *with two players, and a finite epistemic model* $M = (T_1, T_2, b_1, b_2)$ *for* Γ. *Fix a type* $t_i \in T_i$.

(Induction start) Let $T_i^1(t_i) := \{t_i\}$.

(Induction step) For every even round $k \geq 2$, *let* $T_j^k(t_i) := \cup_{t_i \in T_i^{k-1}(t_i)} T_j(t_i)$. *For every odd round* $k \geq 3$, *let* $T_i^k(t_i) := \cup_{t_j \in T_j^{k-1}(t_i)} T_i(t_j)$.

So, $T_j^2(t_i)$ contains all opponent's types that t_i deems possible, $T_i^3(t_i)$ contains all types for player i which are deemed possible by some type t_j that t_i deems possible, and so on. This procedure eventually yields the sets of types

$$T_i^*(t_i) = \cup_k T_i^k(t_i), \text{ and } T_j^*(t_i) = \cup_k T_j^k(t_i).$$

These sets $T_i^*(t_i)$ and $T_j^*(t_i)$ contain precisely those types that enter t_i's belief hierarchy in some of its levels, and we will call these the *relevant* types for t_i. Since there are only finitely many types in M, there must be some round K such that $T_j^*(t_i) = T_j^K(t_i)$, and $T_i^*(t_i) = T_i^{K+1}(t_i)$. That is, this procedure must stop after finitely many rounds.

Now, suppose that type t_i expresses *common belief in future rationality.* Then, in particular, t_i must believe in j's future rationality. Moreover, t_i must only consider possible opponent's types t_j that believe in i's future rationality, that is, every type in $T_j^2(t_i)$ must believe in the opponent's future rationality. Also, t_i must only consider possible types for j that only consider possible types for i that believe in j's future rationality. In other words, all types in $T_i^3(t_i)$ must believe in the opponent's future rationality. By continuing in this fashion, we conclude that all types in $T_i^*(t_i)$ and $T_j^*(t_i)$ believe in the opponent's future rationality. So, we see that every type t_i that expresses *common belief in future rationality*, must have the property that all types in $T_i^*(t_i)$ and $T_j^*(t_i)$ believe in the opponent's future rationality.

However, we can show that the opposite is also true! Consider, namely, a type t_i within a finite epistemic model $M = (T_1, T_2, b_1, b_2)$ for which all types in $T_i^*(t_i)$ and $T_j^*(t_i)$ believe in the opponent's future rationality. Then, in particular, every type in $T_i^1(t_i)$ believes in j's future rationality. As $T_i^1(t_i) = \{t_i\}$, it follows that t_i believes in j's future rationality. Also, every type in $T_j^2(t_i)$ believes in the opponent's future rationality. As $T_j^2(t_i)$ contains exactly those types for j that t_i deems possible, it follows that t_i only deems possible types for j that believe in i's future rationality. By contuining in this way, we conclude that t_i expresses *common belief in future rationality.*

The two insights above lead to the following theorem.

Theorem 3.2 (Checking common belief in future rationality). *Consider a finite dynamic game Γ with two players, and a finite epistemic model $M = (T_1, T_2, b_1, b_2)$ for Γ. Then, a type t_i expresses common belief in future rationality, if and only if, all types in $T_i^*(t_i)$ and $T_j^*(t_i)$ believe in the opponent's future rationality.*

Note that checking whether all types in $T_i^*(t_i)$ and $T_j^*(t_i)$ believe in the opponent's future rationality can be done within finitely many steps. We have seen above, namely, that the sets of relevant types for t_i – that is, the sets $T_i^*(t_i)$ and $T_j^*(t_i)$ – can be derived within finitely many steps, and only contain finitely many types. So, within a finite epistemic model, checking for *common belief in future rationality* only requires finitely many reasoning steps. Consequently, if we take a finitely generated belief hierarchy, then it only takes finitely many steps to verify whether it expresses *common belief in future rationality* or not.

4. Reasoning Towards One Belief Hierarchy

Suppose you are a player in a finite dynamic game. Is there a reasoning procedure that eventually leads you to *some* belief hierarchy that expresses *common belief in future rationality*? We will see that there is, and such a reasoning procedure only involves finitely many steps.

4.1. Procedure

The procedure works as follows. We start by prescribing a *choice-history combination* $(c^1, x^1) = (c_i^1(h), x_i^1(h))_{i \in I, h \in H_i}$ where, for both players i and every information set $h \in H_i$, we select a choice $c_i^1(h)$ for player i at h, and some history $x_i^1(h)$ in h. Remember that a history in h is some sequence of past choice-combinations that leads to h. The interpretation of the choice-history combination (c^1, x^1) is that player i, at every information set $h \in H_i$, believes that opponent j has chosen according to $x_i^1(h)$ in the past, and will choose according to c^1 now and in the future.

Moreover, we select the histories in such a way that, whenever h, h' are two information sets for player i, information set h' follows h, and the history $x_i^1(h)$ together with the choices in c^1 after $x_i^1(h)$ lead to h', then $x_i^1(h')$ contains exactly the past choice-combinations in $x_i^1(h)$, and the choices in c^1 between $x_i^1(h)$ and h'. We say that the choice-history combination (c^1, x^1) satisfies *Bayesian updating,* as it corresponds exactly to the well-known Bayesian updating condition for belief revision when we restrict to probability 1 beliefs. Clearly, for any given choice-combination c^1, the histories x^1 can always be chosen in such a way that Bayesian updating is satisfied.

Subsequently we create a new choice-history combination (c^2, x^2) as follows. For player i we start at his ultimate information sets – that is, his information sets $h \in H_i$ that are not followed by any other information set $h' \in H_i$. At such an ultimate information set $h \in H_i$ we select a choice $c_i^2(h)$ that is optimal for player i at h, given the opponent's choices $c_j^1(h')$ at information sets $h' \in H_j$ weakly following h, and given the past choices leading to history $x_i^1(h)$ in h.

We then look at the pen-ultimate information sets for player i – that is, the information sets $h \in H_i$ that are followed by at least one other information set $h' \in H_i$, but where all such h' are ultimate information sets for player i. For such a pen-ultimate information set h for player i, we select a choice $c_i^2(h)$ that is optimal for player i at h, given the opponent's choices $c_j^1(h')$ at information sets $h' \in H_j$ weakly following h, given the past choices leading to history $x_i^1(h)$ in h, and *given his own future choices* $c_i^2(h')$ at ultimate information sets h' for player i that follow h.

By continuing this procedure, we finally obtain a choice-combination $(c_i^2(h))_{i \in I, h \in H_i}$ where, for both players i and all information sets $h \in H_i$, the choice $c_i^2(h)$ is optimal for player i at h, given the opponent's choices $c_j^1(h')$ at information sets $h' \in H_j$ weakly following h, given i's own choices $c_i^2(h')$ at information sets $h' \in H_i$ following h, and given the past choices leading to history $x_i^1(h)$ in h. The procedure to achieve this choice combination will be called the *backward induction method,* as it proceeds backwards from the ultimate information sets of player i to the first information sets of player i in the game.

To conclude step 2, we specify for both players i, and every information set $h \in H_i$, some history $x_i^2(h)$ in h, such that the new choice-history combination (c^2, x^2) satisfies Bayesian updating.

In step 3 we then take a new choice-history combination (c^3, x^3) which satisfies Bayesian updating, and where, for both players i and every information set $h \in H_i$, the choice $c_i^3(h)$ is optimal for player i at h, given the opponent's choices $c_j^2(h')$ at information sets $h' \in H_j$ weakly following h, given i's own choices $c_i^3(h')$ at information sets $h' \in H_i$ following h, and given the past choices leading to history $x_i^2(h)$ in h. Here, the choice combination c^3 is found by applying the backward induction method above. And so on. This reasoning procedure thus generates a chain of choice-history combinations, and can be formally defined as follows.

Algorithm 4.1 (Choice-history algorithm). *Consider a finite dynamic game Γ with two players.*

(Induction start) Let $(c^1, x^1) = (c_i^1(h), x_i^1(h))_{i \in I, h \in H_i}$ be an arbitrary choice-history combination that satisfies Bayesian updating.

(Induction step) For every round $k \geq 2$, apply the backward induction method to find a choice-combination $c^k = (c_i^k(h))_{i \in I, h \in H_i}$ where $c_i^k(h)$ is optimal for player i at h, given the opponent's choices $c_j^{k-1}(h')$ at information sets $h' \in H_j$ weakly following h, given i's own choices $c_i^k(h')$ at information sets $h' \in H_i$ following h, and given the past choices leading to history $x_i^{k-1}(h)$ in h. Then, choose for every information set $h \in H_i$ the history $x_i^k(h)$ in h in such a way that the resulting choice-history combination (c^k, x^k) satisfies Bayesian updating.

This procedure generates a chain of choice-history combinations $((c^1, x^1), (c^2, x^2), ...)$. Since there are only finitely many information sets, choices and histories in the game, there must be a round m such that $(c^m, x^m) = (c^{m+l}, x^{m+l})$ for some $l \geq 1$. But then, the chain of the choice-history combinations between rounds m and $m + l$ will simply repeat itself over and over again after round $m + l$. That is, the chain of choice-history combinations will eventually enter into a cycle

$$(c^m, x^m) \to (c^{m+1}, x^{m+1}) \to (c^{m+2}, x^{m+2}) \to ... \to (c^{m+l}, x^{m+l}) = (c^m, x^m).$$

From this cycle we will now construct a belief hierarchy for every player i that expresses *common belief in future rationality.* In fact, for every choice-history combination (c^k, x^k) in this cycle we will construct a type t_i^k for both players i.

The conditional beliefs for the types about the opponent's strategy-type pairs are as follows. Fix some choice-history combination (c^k, x^k) with $k \in \{m + 1, ..., m + l\}$. Remember that $(c^{m+l}, x^{m+l}) = (c^m, x^m)$.

- At every information set $h \in H_i$, type t_i^k believes that, with probability 1, opponent j is of type t_j^{k-1}.
- At every information set $h \in H_i$, type t_i^k believes that, with probability 1, opponent j has made the choices in $x_i^{k-1}(h)$ in the past, and will make the choice $c_j^{k-1}(h')$ at every player j information set h' that weakly follows h.

By construction of the algorithm, player i's choice $c_i^k(h)$ is optimal at h if he believes that j chooses according to c^{k-1} now and in the future, if he believes that j has chosen according to $x_i^{k-1}(h)$ in the past, and given his own future choices in c^k. As type t_i^k has precisely this belief about player j, we may conclude that at every information set $h \in H_i$, the choice $c_i^k(h)$ is optimal for type t_i^k, given i's own future choices in c^k. We say that the choice combination $(c_i^k(h))_{h \in H_i}$ is *locally optimal* for type t_i^k.

Moreover, it is clear that type t_i^k also satisfies *Bayesian updating* when revising his beliefs, as his conditional beliefs are derived from the choice-history combination (c^{k-1}, x^{k-1}) which satisfies Bayesian updating by construction. But we know from Hendon, Jacobsen and Sloth (1996) and Perea (2002) that, if a choice combination $(c_i^k(h))_{h \in H_i}$ is *locally* optimal for some conditional belief vector, and the conditional belief vector satisfies Bayesian updating, then the choice combination is also *globally* optimal. By the latter we mean that at every information set $h \in H_i$ for player i, there is no other choice combination that as a whole would give him a higher expected utility than $(c_i^k(h))_{h \in H_i}$. We may thus conclude that, at every information set $h \in H_i$, the choice combination $(c_i^k(h))_{h \in H_i}$ generated by the procedure above is *globally* optimal for type t_i^k.

Note that type t_i^k believes at every information set $h \in H_i$ that opponent j is of type t_j^{k-1}, and that opponent j chooses according to the choice combination $(c_j^{k-1}(h''))_{h'' \in H_j}$ at information sets $h' \in H_j$ that weakly follow h. We know from above that at every such information set $h' \in H_j$ weakly following h, the choice combination $(c_j^{k-1}(h''))_{h'' \in H_j}$ is globally optimal for type t_j^{k-1}. This means, however, that type t_i^k believes at every information set $h \in H_i$ that opponent j chooses optimally at every $h' \in H_j$ weakly following h. In other words, type t_i^k believes in the opponent's future rationality.

Since this applies to every type t_i^k in the epistemic model above, we thus conclude that every type in the epistemic model above believes in the opponent's future rationality. But then, it follows immediately from Theorem 3.2 that each of the types also expresses *common belief in future rationality.* We thus have found a reasoning procedure that leads you to *some* belief hierarchy that expresses *common belief in future rationality.*

Theorem 4.2 (Reasoning towards one belief hierarchy). *Consider a finite dynamic game* Γ *with two players.*

Suppose we apply the choice-history algorithm until we reach a cycle

$$(c^m, x^m) \to (c^{m+1}, x^{m+1}) \to (c^{m+2}, x^{m+2}) \to ... \to (c^{m+l}, x^{m+l}) = (c^m, x^m)$$

of choice-history combinations.

Then, define for every choice-history combination (c^k, x^k) *in this cycle, and both players* i*, some type* t_i^k *such that*

(a) at every information set $h \in H_i$*, type* t_i^k *believes that, with probability 1, opponent* j *is of*

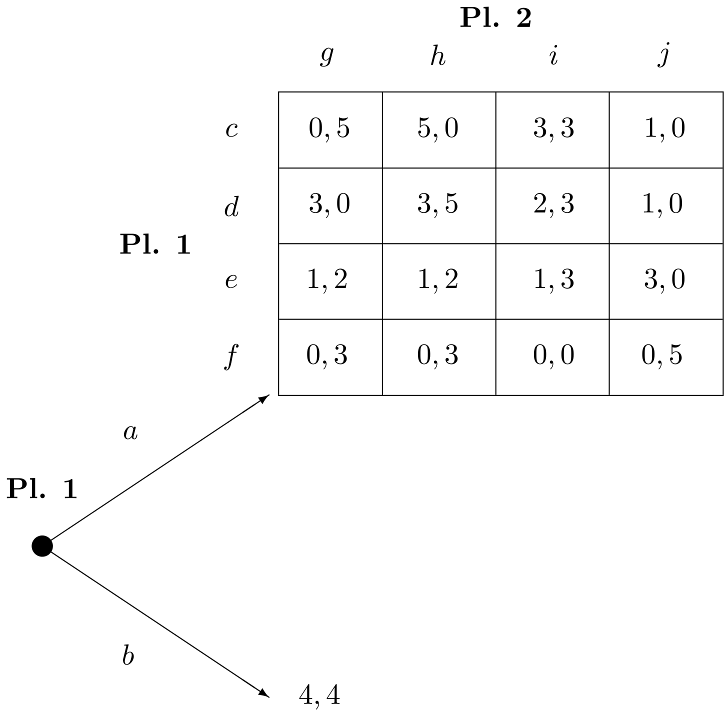

Figure 1: Example of a dynamic game

---

type t_j^{k-1}, *and*

(b) at every information set $h \in H_i$, *type* t_i^k *believes that, with probability 1, opponent* j *has made the choices in* $x_i^{k-1}(h)$ *in the past, and will make the choice* $c_j^{k-1}(h')$ *at every player* j *information set* h' *that weakly follows* h.

Then, every type so generated expresses common belief in future rationality.

Note that this is a *finite* reasoning procedure, since the choice-history algorithm will enter into a cycle after finitely many steps, after which we can start the construction of types above. So, a player in a finite dynamic game can always reason himself in *finitely* many steps towards *some* belief hierarchy that expresses *common belief in future rationality.*

4.2. Example

We will now illustrate the reasoning procedure above by means of an example. Consider the dynamic game in Figure 1. At the beginning, player 1 can choose between a and b. If he chooses b, the game ends, and the utilities for players 1 and 2 will be $(4,4)$. If he chooses a the game continues, and players 1 and 2 must simultaneously choose from $\{c,d,e,f\}$ and $\{g,h,i,j\}$, respectively. The utilities for both players in that case can be found in the table following choice a. Let us denote the beginning of the game by $\emptyset$, and the information set following choice a by h_1. Hence, $\emptyset$ and h_1 are the two information sets in the game. At $\emptyset$ only player 1 makes a choice, whereas both players 1 and 2 are active at h_1.

We will now use the reasoning procedure above to construct some belief hierarchies for players 1 and 2 that express *common belief in future rationality.* We start by applying the choice-history algorithm to generate a chain of choice-history combinations until we reach a cycle.

Let us begin with an arbitrary choice-history combination, say $(c^1,x^1)=((a,f,j),a)$. Here, the choice combination $c^1=(a,f,j)$ specifies the choice a at $\emptyset$, and the choices f and j at h_1. The history $x^1=a$ specifies at h_1 the history a, stating that player 2 believes at h_1 that player 2 has chosen a in the past. Note that this is the only history possible, as at h_1 it is evident that player 1 has chosen a in the past.

Within the choice-history algorithm, this initial choice-history combination (c^1,x^1) induces a new choice-history combination (c^2,x^2) as follows. We start at the ultimate information set h_1. Given player 2's choice j in (c^1,x^1) at h_1, it is optimal for player 1 to choose e at h_1. Conversely, given player 1's choice f in (c^1,x^1) at h_1, and given the history a in (c^1,x^1) at h_1, it is optimal for player 2 to choose j at h_1. So, the new choice combination c^2 specifies the choices e and j at h_1. We then move to the pen-ultimate information set $\emptyset$, which is the beginning of the game. Given player 2's future choice j in the *previous* choice-history combination (c^1,x^1) at h_1, and given player 1's own future choice e in the *current* choice-history combination (c^2,x^2) at h_1, it is optimal for player 1 to choose b at $\emptyset$. So, the new choice combination c^2 specifies the choice b at $\emptyset$. As a is the only history possible at h_1, it follows that the new choice-history combination is $(c^2,x^2)=((b,e,j),a)$.

By continuing in this fashion, the choice-history algorithm yields the following chain of choice-history combinations,

$$\begin{aligned}((a,f,j),a) \;&\rightarrow\; ((b,e,j),a)\rightarrow((b,e,i),a)\rightarrow((b,c,i),a)\rightarrow((b,c,g),a)\rightarrow\\ &\rightarrow\; \underbrace{((b,d,g),a)\rightarrow((b,d,h),a)\rightarrow((a,c,h),a)\rightarrow((a,c,g),a)\rightarrow((b,d,g),a)}_{\text{cycle}}\end{aligned}$$

which eventually enters into the cycle

$$\underbrace{((b,d,g),a)}_{(c^6,x^6)}\rightarrow\underbrace{((b,d,h),a)}_{(c^7,x^7)}\rightarrow\underbrace{((a,c,h),a)}_{(c^8,x^8)}\rightarrow\underbrace{((a,c,g),a)}_{(c^9,x^9)}\rightarrow\underbrace{((b,d,g),a)}_{(c^6,x^6)}.$$

The next step in the reasoning procedure is to transform this cycle of choice-history combinations into an epistemic model in which every type expresses *common belief in future rationality.* Every choice-history combination $(c^6,x^6),(c^7,x^7),(c^8,x^8),(c^9,x^9)$ in the cycle will generate a type for both players, so the sets of types are given by

$$T_1 = \{t_1^6, t_1^7, t_1^8, t_1^9\} \text{ and } T_2 = \{t_2^6, t_2^7, t_2^8, t_2^9\}.$$

According to the construction in Theorem 4.2, the conditional beliefs for the types are

$$\begin{aligned}
b_1(t_1^6, \emptyset) &= b_1(t_1^6, h_1) = (g, t_2^9), \\
b_1(t_1^7, \emptyset) &= b_1(t_1^7, h_1) = (g, t_2^6), \\
b_1(t_1^8, \emptyset) &= b_1(t_1^8, h_1) = (h, t_2^7), \\
b_1(t_1^9, \emptyset) &= b_1(t_1^9, h_1) = (h, t_2^8),
\end{aligned}$$

$$\begin{aligned}
b_2(t_2^6, h_1) &= ((a, c), t_1^9), \\
b_2(t_2^7, h_1) &= ((a, d), t_1^6), \\
b_2(t_2^8, h_1) &= ((a, d), t_1^7), \\
b_2(t_2^9, h_1) &= ((a, c), t_1^8).
\end{aligned}$$

Here, $b_1(t_1^6, \emptyset) = b_1(t_1^6, h_1) = (g, t_2^9)$ means that player 1's type t_1^6 believes, at $\emptyset$ and h_1, that player 2 will choose g and that player 2 is of type t_2^9. Similarly for the other types of player 1. Moreover, $b_2(t_2^6, h_1) = ((a, c), t_1^9)$ means that player 2's type t_2^6 believes at h_1 that player 1 chooses the strategy (a, c), and that player 1 is of type t_1^9. Note that player 2 *must* believe at h_1 that player 1 chooses a strategy that starts with a, as at h_1 it is evident that player 1 has chosen a at $\emptyset$. Similarly for the other types of player 2.

In this way, we have constructed an epistemic model M with finitely many types, in which every type expresses *common belief in future rationality.* Moreover, the epistemic model has the special property that every type, at each of his information sets, assigns probability 1 to *one* type and *one* strategy for the opponent. In other words, every belief hierarchy so constructed only deems possible *one* opponent's strategy and *one* opponent's belief hierarchy at every information set.

Note that strategy (a, c) is optimal for player 1's types t_1^8 and t_1^9, whereas strategy b is optimal for the types t_1^6 and t_1^7. Moreover, strategy g is optimal for player 2's types t_2^6 and t_2^9, whereas strategy h is optimal for the types t_2^7 and t_2^8. Hence, the reasoning method above generates the strategies $(a, c), b, g$ and h as strategies that can rationally be chosen under *common belief in future rationality.*

5. Reasoning Towards a "Sufficient" Set of Belief Hierarchies

In the previous section we have presented a finite reasoning procedure that leads a player to *some* belief hierarchies that expresses *common belief in future rationality.* However, in general not *every* strategy that can rationally be chosen under *common belief in future rationality* will be supported by some belief hierarchy generated by this procedure!

In the dynamic game of Figure 1, for instance, the strategy i for player 2 *can* rationally be chosen under *common belief in future rationality,* but is not supported by any belief hierarchy generated by the procedure above. To see this, note that within the epistemic model constructed above, player 2's type t_2^6 believes at h_1 that player 1 chooses (a, c) while being of type t_1^9, whereas type t_2^7 believes at h_1 that player 1 chooses (a, d) while being of type t_1^6. Moreover, both types believe in the opponent's future rationality. But then, we can add a new type t_2 for player 2 that at h_1 assigns probability 0.5 to player 1 choosing (a, c) while being of type t_1^9, and that at h_1 assigns probability 0.5 to player 1 choosing (a, d) while being of type t_1^6. Then, also t_2 will believe in the opponent's future rationality. Therefore, within the new enlarged epistemic model, every type will believe in the opponent's future rationality, and hence, by Theorem 3.2, all types within the new epistemic model will express *common belief in future rationality*. In particular, the new type t_2 will express *common belief in future rationality.* As strategy i is optimal for type t_2, it follows that player 2 can rationally choose strategy i under *common belief in future rationality.*

But strategy i is not supported by any belief hierarchy generated by the procedure in the previous section. Indeed, for the types t_2^6 and t_2^9 only strategy g is optimal, whereas for the types t_2^7 and t_2^8 only strategy h is optimal. One could argue that this maybe depends on the initial choice-history combination $((a, f, j), a)$ we have chosen for this procedure. Perhaps by chosing some other initial choice-history combination (c^1, x^1), we could eventually generate some different epistemic model in which there *would* be some belief hierarchy expressing *common belief in future rationality* that supports the strategy i for player 2. But this could never happen! Namely, under *common belief in future rationality,* player 2 can at h_1 only assign positive probability to strategies (a, c) and (a, d). To see this, note that at h_1, strategy (a, f) is always worse for player 1 than (a, e), and hence player 2 must believe at h_1 that player 1 does not chose (a, f). But then, strategy j will always be worse for player 2 at h_1 than i, and hence player 1 must believe that player 2 will not choose j at h_1. But then, strategy (a, e) will always be worse for player 1 at h_1 than (a, d), and hence player 2 must believe at h_1 that player 1 does not choose (a, e) or (a, f). In other words, player 2 must at h_1 only assign positive probability to (a, c) and (a, d).

But the reasoning method in the previous section generates very special belief hierarchies, which at every information set assign probability 1 to one particular opponent's strategy. So, no matter which initial choice-history combination we choose, the procedure above will always generate belief hierarchies for player 2 that at h_1 either assign probability 1 to (a, c), or assign probability 1 to (a, d). But then, strategy i can never be optimal for such a belief hierarchy! Hence, strategy i can never be supported by the reasoning method of the previous section, although it can rationally be chosen under *common belief in future rationality.*

In this section our goal is more ambitious, in that we wish to find a reasoning procedure that covers *all* strategies you can rationally choose under *common belief in future rationality* in a given dynamic game. More precisely, the procedure must generate a finite set of belief hierarchies such that, for *every* strategy that can rationally be chosen under *common belief in future rationality,* there will be a belief hierarchy in this set which supports that strategy. In that sense, the reasoning procedure must yield a *sufficient* set of belief hierarchies.

5.1. Procedure

To see how this reasoning procedure works, let us start with exploring the consequences of "believing in the opponent's future rationality". Suppose player i believes at some information set $h \in H_i$ that opponent j chooses rationally now and in the future. Then, player i will at h only assign positive probability to strategies s_j for player j that are optimal, at every $h' \in H_j$ weakly following h, for *some* belief that j can hold at h' about i's strategy choice.

Consider such a future information set $h' \in H_j$, and let $\Gamma^0(h') = (S_j(h'), S_i(h'))$ be the *full decision problem* for player j at h', at which he can only choose strategies in $S_j(h')$ that possibly reach h', and believes that player i can only choose strategies in $S_i(h')$ that possibly reach h'. From Lemma 3 in Pearce (1984), we know that a strategy s_j is optimal for player j at h' for some belief about i's strategy choice, if and only if, s_j is *not strictly dominated* within the full decision problem $\Gamma^0(h') = (S_j(h'), S_i(h'))$ by a *randomized* strategy r_j. Here, a *randomized strategy* r_j for player j is a probability distribution on j's strategies in $S_j(h')$ – that is, j selects each of his strategies s_j' with probability $r_j(s_j')$. And we say that the strategy s_j is *strictly dominated* by the randomized strategy r_j in $\Gamma^0(h')$ if r_j always yields a higher expected utility than s_j against any strategy s_i of player i in $\Gamma^0(h')$.

Putting these things together, we see that if i believes at h in j's future rationality, then i assigns at h only positive probability to j's strategies s_j that are not strictly dominated within any full decision problem $\Gamma^0(h')$ for player j that weakly follows h. Or, put differently, player i assigns at h probability zero to any opponent's strategy s_j that *is* strictly dominated at some full decision problem $\Gamma^0(h')$ for player j that weakly follows h. That is, we eliminate any such opponent's strategy s_j from player i's full decision problem $\Gamma^0(h) = (S_i(h), S_j(h))$ at h.

We thus see that, if player i believes in j's future rationality, then player i eliminates, at each of his full decision problems $\Gamma^0(h)$, those opponent's strategies s_j that are strictly dominated within some full decision problem $\Gamma^0(h')$ for player j that weakly follows h. Let us denote by $\Gamma^1(h)$ the reduced decision problem for player i at h that remains after eliminating such opponent's strategies s_j from $\Gamma^0(h)$.

Next, suppose that player i does not only believe in j's future rationality, but also believes that j believes in i's future rationality. Take an information set h for player i, and an arbitrary information set h' for player j that weakly follows h. As i believes that j believes in i's future rationality, player i believes that player j, at information set h', believes that player i will only choose strategies from $\Gamma^1(h')$. Moreover, as i believes in j's future rationality, player i believes that j will choose rationally at h'. Together, these two insights imply that player i believes at h that j will only choose strategies s_j that are not strictly dominated within $\Gamma^1(h')$. Or, equivalently, player i eliminates from his decision problem $\Gamma^1(h)$ all strategies s_j for player j that are strictly dominated within $\Gamma^1(h')$. As this holds for every player j information set h' that weakly follows h, we see that player i will eliminate, from each of his decision problems $\Gamma^1(h)$, all opponent's strategies s_j that are strictly dominated within some decision problem $\Gamma^1(h')$ for player j that weakly follows h.

Hence, if player i expresses up to 2-fold belief in future rationality, then he will eliminate, from each of his decision problems $\Gamma^1(h)$, all opponent's strategies s_j that are strictly dominated within some decision problem $\Gamma^1(h')$ for player j that weakly follows h. Let us denote by $\Gamma^2(h)$ the reduced decision problem for player i that remains after eliminating such opponent's strategies s_j from $\Gamma^1(h)$.

By continuing in this fashion, we conclude that if player i expresses up to k-fold belief in future rationality – that is, expresses 1-fold, 2-fold, ... until k-fold belief in future rationality – then he believes at every information set $h \in H_i$ that opponent j will only choose strategies from the reduced decision problem $\Gamma^k(h)$. This leads to the following reasoning procedure, known as the *backward dominance procedure* (Perea (2011a)). The procedure is closely related to Penta's (2009) *backwards rationalizability* procedure, and is equivalent to Chen and Micali's (2011) *backward robust solution.*

Algorithm 5.1 (Backward dominance procedure). *Consider a finite dynamic game Γ with two players.*

(Induction start) For every information set h, let $\Gamma^0(h) = (S_1(h), S_2(h))$ be the full decision problem at h.

(Induction step) For every $k \geq 1$, and every information set h, let $\Gamma^k(h) = (S_1^k(h), S_2^k(h))$ be the reduced decision problem which is obtained from $\Gamma^{k-1}(h)$ by eliminating, for both players i, those strategies s_i that are strictly dominated at some decision problem $\Gamma^{k-1}(h')$ weakly following h at which i is active.

Suppose that h is an information set at which player i is active. Then, the interpretation of the reduced decision problem $\Gamma^k(h) = (S_1^k(h), S_2^k(h))$ is that at round k of the procedure, player i believes at h that opponent j chooses some strategy in $S_j^k(h)$. As the sets $S_j^k(h)$ become smaller as k becomes bigger, the procedure thus puts more and more restrictions on player i's conditional beliefs about j's strategy choice. However, since in a finite dynamic game there are only finitely many information sets and strategies, this procedure must stop after finitely many rounds! Namely, there must be some round K such that $S_1^{K+1}(h) = S_1^K(h)$ and $S_2^{K+1}(h) = S_2^K(h)$ for all information sets h. But then, $S_j^k(h) = S_j^K(h)$ for all information sets h and every $k \geq K+1$, and hence the procedure will not put more restrictions on i's conditional beliefs about j's strategy choice after round K. This reasoning procedure is therefore a *finite* procedure, guaranteed to end within finitely many steps.

Above we have argued that if player i reasons in accordance with *common belief in future rationality*, then his belief at information set h about j's strategy choice will only assign positive probability to strategies in $S_j^K(h)$. As a consequence, he can only rationally choose a strategy s_i that is optimal, at every information set $h \in H_i$, for such a conditional belief that only considers

j's strategy choices in $S_j^K(h)$. But then, by Lemma 3 in Pearce (1984), strategy s_i must not be strictly dominated at any information set $h \in H_i$ if we restrict to j's strategy choices in $S_j^K(h)$. That is, s_i must be in $S_i^K(\emptyset)$, where $\emptyset$ denotes the beginning of the game. We can thus conclude that every strategy s_i that can rationally be chosen under *common belief in future rationality*, must be in $S_i^K(\emptyset)$ – that is, must survive the backward dominance procedure at the beginning of the game.

We can show, however, that the converse is also true! That is, every strategy in $S_i^K(\emptyset)$ can be supported by a belief hierarchy that expresses *common belief in future rationality*. Suppose, namely, that player i has performed the backward dominance procedure in his mind, which has left him with the strategies $S_i^K(h)$ and $S_j^K(h)$ at every information set h of the game. Then, by construction, every strategy $s_i \in S_i^K(h)$ is not strictly dominated on $S_j^K(h')$, for every information set h' weakly following h at which i is active. Thus, by Lemma 3 in Pearce (1984), every strategy $s_i \in S_i^K(h)$ is optimal, at every $h' \in H_i$ weakly following h, for some probabilistic belief $b_i^{s_i,h}(h') \in \Delta(S_j^K(h'))$. Similarly, every strategy $s_j \in S_j^K(h)$ will be optimal, at every $h' \in H_j$ weakly following h, for some probabilistic belief $b_j^{s_j,h}(h') \in \Delta(S_i^K(h'))$.

Now we can construct the sets of types

$$\begin{aligned} T_i &= \{t_i^{s_i,h} : h \in H \text{ and } s_i \in S_i^K(h)\} \text{ and} \\ T_j &= \{t_j^{s_j,h} : h \in H \text{ and } s_j \in S_j^K(h)\}, \end{aligned}$$

where H denotes the collection of all information sets in the game. We define the conditional beliefs of the types about the opponent's strategy-type pairs to be

$$b_i(t_i^{s_i,h}, h')(s_j, t_j) = \begin{cases} b_i^{s_i,h}(h')(s_j), & \text{if } t_j = t_j^{s_j,h'} \\ 0, & \text{otherwise} \end{cases}$$

for every $h' \in H_i$, and

$$b_j(t_j^{s_j,h}, h')(s_i, t_i) = \begin{cases} b_j^{s_j,h}(h')(s_i), & \text{if } t_i = t_i^{s_i,h'} \\ 0, & \text{otherwise} \end{cases}$$

for all $h' \in H_j$.

This yields an epistemic model M. Hence, every type $t_i^{s_i,h}$ for player i, at every information set $h' \in H_i$, only considers possible strategy-type pairs $(s_j, t_j^{s_j,h'})$ where $s_j \in S_j^K(h')$, and his conditional belief at h' about j's strategy choice is given by $b_i^{s_i,h}(h')$. By construction, strategy $s_i \in S_i^K(h)$ is optimal for $b_i^{s_i,h}(h')$ at every $h' \in H_i$ weakly following h. As a consequence, strategy $s_i \in S_i^K(h)$ is optimal for type $t_i^{s_i,h}$ at every $h' \in H_i$ weakly following h. The same holds for player j. Since type $t_i^{s_i,h}$, at every information set $h' \in H_i$, only considers possible strategy-type pairs $(s_j, t_j^{s_j,h'})$ where $s_j \in S_j^K(h')$, it follows that type $t_i^{s_i,h}$, at every information set $h' \in H_i$, only considers possible strategy-type pairs $(s_j, t_j^{s_j,h'})$ where strategy s_j is optimal for $t_j^{s_j,h'}$ at every $h'' \in H_j$ weakly following h'. That is, type $t_i^{s_i,h}$ believes in the opponent's future rationality.

Since this holds for every type $t_i^{s_i,h}$ in this epistemic model M, it follows directly from Theorem 3.2 that every type in the epistemic model M above expresses *common belief in future rationality.*

Now, take some strategy $s_i \in S_i^K(\emptyset)$, which survives the backward dominance procedure at the beginning of the game. Then, we know from our insights above that s_i is optimal for the type $t_i^{s_i,\emptyset}$ at every $h \in H_i$ weakly following $\emptyset$ – that is, at *every* $h \in H_i$ in the game. As the type $t_i^{s_i,\emptyset}$ expresses *common belief in future rationality*, we thus see that every strategy $s_i \in S_i^K(\emptyset)$ can rationally be chosen by some type $t_i^{s_i,\emptyset}$ that expresses *common belief in future rationality.* In other words, for every strategy $s_i \in S_i^K(\emptyset)$ that survives the backward dominance procedure at $\emptyset$, there is a belief hierarchy expressing c*ommon belief in future rationality* – namely the belief hierarchy induced by $t_i^{s_i,\emptyset}$ in the epistemic model M – for which s_i is optimal. This insight thus leads to the following theorem.

Theorem 5.2 (Reasoning towards a sufficient set of belief hierarchies). *Consider a finite dynamic game Γ with two players.*

Suppose we apply the backward dominance procedure until it terminates at round K. That is, $S_1^{K+1}(h) = S_1^K(h)$ and $S_2^{K+1}(h) = S_2^K(h)$ for all information sets h.

For every information set h, both players i, every strategy $s_i \in S_i^K(h)$, and every information set $h' \in H_i$ weakly following h, let $b_i^{s_i,h}(h') \in \Delta(S_j^K(h'))$ be a probabilistic belief on $S_j^K(h')$ for which s_i is optimal.

For both players i, define the set of types $T_i = \{t_i^{s_i,h} : h \in H \text{ and } s_i \in S_i^K(h)\}$, and for every type $t_i^{s_i,h}$ and every $h' \in H_i$ define the conditional belief

$$b_i(t_i^{s_i,h}, h')(s_j, t_j) = \begin{cases} b_i^{s_i,h}(h')(s_j), & \text{if } t_j = t_j^{s_j,h'} \\ 0, & \text{otherwise} \end{cases}$$

about j's strategy-type pairs. Then, all types in this epistemic model M express common belief in future rationality. Moreover,

(1) for every strategy $s_i \in S_i^K(\emptyset)$ that survives the backward dominance procedure at $\emptyset$ there is a belief hierarchy in M expressing common belief in future rationality for which s_i is optimal at all $h \in H_i$ possibly reached by s_i – namely the belief hierarchy induced by $t_i^{s_i,\emptyset}$,

(2) for every strategy $s_i \notin S_i^K(\emptyset)$ that does not survive the backward dominance procedure at $\emptyset$, there is no belief hierarchy whatsoever expressing common belief in future rationality for which s_i is optimal at all $h \in H_i$ possibly reached by s_i.

So, whenever a strategy s_i *is* optimal for some belief hierarchy that expresses *common belief in future rationality*, this reasoning procedure generates one. In that sense, we can say that this reasoning procedure yields a "sufficient" set of belief hierarchies. Note also that this is a reasoning procedure with *finitely* many steps, as the backward dominance procedure terminates after finitely many rounds, after which we only have to construct finitely many types – one for each information set h and each surviving strategy $s_i \in S_i^K(h)$ at h.

The theorem above also shows that *finitely generated* belief hierarchies are sufficient when it comes to exploring *common belief in future rationality* within a finite dynamic game. Suppose, namely, that some strategy s_i is optimal, at all $h \in H_i$ possibly reached by s_i, for *some* belief hierarchy – not necessarily finitely generated – that expresses *common belief in future rationality.* Then, according to part (2) in the theorem, strategy s_i must be in $S_i^K(\emptyset)$. But in that case, the procedure above generates a finitely generated belief hierarchy for which the strategy s_i is optimal – namely the belief hierarchy induced by the type $t_i^{s_i,\emptyset}$ within the finite epistemic model M. So we see that, whenever a strategy s_i is optimal for *some* belief hierarchy – not necessarily finitely generated – that expresses *common belief in future rationality*, then it is also optimal for a *finitely generated* belief hierarchy that expresses *common belief in future rationality.*

Corollary 5.3 (Finitely generated belief hierarchies are sufficient). *Consider a finite dynamic game Γ with two players. If a strategy s_i is optimal for* ***some*** *belief hierarchy – not necessarily finitely generated – that expresses common belief in future rationality, then it is also optimal for a* ***finitely generated*** *belief hierarchy that expresses common belief in future rationality.*

Here, whenever we say that s_i is optimal for some belief hierarchy, we mean that it is optimal for this belief hierarchy at every information set $h \in H_i$ possibly reached by s_i. This corollary thus states that working with finitely generated belief hierarchies is sufficient for the concept of *common belief in future rationality* in finite dynamic games.

5.2. Example

We shall now illustrate the reasoning procedure above by means of an example. Consider again the dynamic game from Figure 1. We will first run the backward dominance procedure for this example, and then build an epistemic model on the basis of that procedure, following the construction in Theorem 5.2.

Recall that there are two information sets in this game, namely the beginning of the game $\emptyset$, at which only player 1 is active, followed by the information set h_1 at which both player 1 and player 2 are active. The full decision problems at both information sets are given in Table 1. We will now start the backward dominance procedure. In round 1, we see that within the full decision problem $\Gamma^0(\emptyset)$ at the beginning of the game, the strategies $(a, d), (a, e)$ and (a, f) are strictly dominated for player 1 by b. So, we eliminate $(a, d), (a, e)$ and (a, f) from $\Gamma^0(\emptyset)$, but not – yet – from $\Gamma^0(h_1)$, as h_1 follows $\emptyset$. Moreover, within the full decision problem $\Gamma^0(h_1)$ at h_1,

$\Gamma^0(\emptyset)$

	g	h	i	j
(a,c)	$0,5$	$5,0$	$3,3$	$1,0$
(a,d)	$3,0$	$3,5$	$2,3$	$1,0$
(a,e)	$1,2$	$1,2$	$1,3$	$3,0$
(a,f)	$0,3$	$0,3$	$0,0$	$0,5$
b	$4,4$	$4,4$	$4,4$	$4,4$

$\Gamma^0(h_1)$

	g	h	i	j
(a,c)	$0,5$	$5,0$	$3,3$	$1,0$
(a,d)	$3,0$	$3,5$	$2,3$	$1,0$
(a,e)	$1,2$	$1,2$	$1,3$	$3,0$
(a,f)	$0,3$	$0,3$	$0,0$	$0,5$

Table 1: Full decision problems in the game of Figure 1

player 1's strategy (a,f) is strictly dominated by (a,d) and (a,e), and hence we eliminate (a,f) from $\Gamma^0(h_1)$ and $\Gamma^0(\emptyset)$. Note, however, that we already eliminated (a,f) at $\Gamma^0(\emptyset)$, so we only need to eliminate (a,f) from $\Gamma^0(h_1)$ at that step. For player 2, no strategy is strictly dominated within $\Gamma^0(\emptyset)$ or $\Gamma^0(h_1)$, so we cannot yet eliminate any strategy for player 2. This leads to the reduced decision problems $\Gamma^1(\emptyset)$ and $\Gamma^1(h_1)$ in Table 2.

We now turn to round 2. Within $\Gamma^1(h_1)$, player 2's strategy j is strictly dominated by i. Hence, we can eliminate strategy j from $\Gamma^1(h_1)$, but also from $\Gamma^1(\emptyset)$, as h_1 follows $\emptyset$. No other strategies can be eliminated at this round. This leads to the reduced decision problems $\Gamma^2(\emptyset)$ and $\Gamma^2(h_1)$ in Table 3.

In round 3, player 1's strategy (a,e) is strictly dominated by (a,d) within $\Gamma^2(h_1)$, and hence we can eliminate (a,e) from $\Gamma^2(h_1)$. This leads to the final decision problems in Table 4, from which no further strategies can be eliminated. Note, for instance, that strategy i is not strictly dominated for player 2 within $\Gamma^3(h_1)$, as it is optimal for the belief that assigns probability 0.5 to (a,c) and (a,d).

We will now build an epistemic model on the basis of the final decision problems $\Gamma^3(\emptyset)$ and $\Gamma^3(h_1)$, using the construction in Theorem 5.2. At $\emptyset$, the surviving strategies are (a,c) and b for player 1, and g, h and i for player 2. That is, $S_1^3(\emptyset) = \{(a,c), b\}$ and $S_2^3(\emptyset) = \{g,h,i\}$. Moreover, at h_1 the surviving strategies are given by $S_1^3(h_1) = \{(a,c),(a,d)\}$ and $S_2^3(h_1) = \{g,h,i\}$. Moreover, these strategies are optimal, at $\emptyset$ and/or h_1, for the following beliefs:

$\Gamma^1(\emptyset)$

	g	h	i	j
(a,c)	$0,5$	$5,0$	$3,3$	$1,0$
b	$4,4$	$4,4$	$4,4$	$4,4$

$\Gamma^1(h_1)$

	g	h	i	j
(a,c)	$0,5$	$5,0$	$3,3$	$1,0$
(a,d)	$3,0$	$3,5$	$2,3$	$1,0$
(a,e)	$1,2$	$1,2$	$1,3$	$3,0$

Table 2: Reduced decision problems after round 1 of backward dominance procedure

---

$\Gamma^2(\emptyset)$

	g	h	i
(a,c)	$0,5$	$5,0$	$3,3$
b	$4,4$	$4,4$	$4,4$

$\Gamma^2(h_1)$

	g	h	i
(a,c)	$0,5$	$5,0$	$3,3$
(a,d)	$3,0$	$3,5$	$2,3$
(a,e)	$1,2$	$1,2$	$1,3$

Table 3: Reduced decision problems after round 2 of backward dominance procedure

---

$\Gamma^3(\emptyset)$

	g	h	i
(a,c)	$0,5$	$5,0$	$3,3$
b	$4,4$	$4,4$	$4,4$

$\Gamma^3(h_1)$

	g	h	i
(a,c)	$0,5$	$5,0$	$3,3$
(a,d)	$3,0$	$3,5$	$2,3$

Table 4: Final decision problems in the backward dominance procedure

- strategy $(a,c) \in S_1^3(\emptyset)$ is optimal, at $\emptyset$, for the belief $b_1^{(a,c),\emptyset}(\emptyset) \in \Delta(S_1^3(\emptyset))$ that assigns probability 1 to h,
- strategy $(a,c) \in S_1^3(\emptyset)$ is optimal, at h_1 following $\emptyset$, for the belief $b_1^{(a,c),\emptyset}(h_1) \in \Delta(S_1^3(h_1))$ that assigns probability 1 to h,
- strategy $b \in S_1^3(\emptyset)$ is optimal, at $\emptyset$, for the belief $b_1^{b,\emptyset}(\emptyset) \in \Delta(S_1^3(\emptyset))$ that assigns probability 1 to g,
- strategy $(a,d) \in S_1^3(h_1)$ is optimal, at h_1, for the belief $b_1^{(a,d),h_1}(h_1) \in \Delta(S_1^3(h_1))$ that assigns probability 1 to g,
- strategy $g \in S_2^3(\emptyset)$ is optimal, at h_1 following $\emptyset$, for the belief $b_2^{g,\emptyset}(h_1) \in \Delta(S_1^3(h_1))$ that assigns probability 1 to (a,c),
- strategy $h \in S_2^3(\emptyset)$ is optimal, at h_1 following $\emptyset$, for the belief $b_2^{h,\emptyset}(h_1) \in \Delta(S_1^3(h_1))$ that assigns probability 1 to (a,d),
- strategy $i \in S_2^3(\emptyset)$ is optimal, at h_1 following $\emptyset$, for the belief $b_2^{i,\emptyset}(h_1) \in \Delta(S_1^3(h_1))$ that assigns probability 0.5 to (a,c) and probability 0.5 to (a,d).

On the basis of these beliefs we can now construct an epistemic model as in Theorem 5.2. So, for both players i, both information sets h, and every strategy $s_i \in S_i^3(h)$, we construct a type $t_i^{s_i,h}$, resulting in the type sets

$$T_1 = \{t_1^{(a,c),\emptyset}, t_1^{b,\emptyset}, t_1^{(a,c),h_1}, t_1^{(a,d),h_1}\} \text{ and } T_2 = \{t_2^{g,\emptyset}, t_2^{h,\emptyset}, t_2^{i,\emptyset}, t_2^{g,h_1}, t_2^{h,h_1}, t_2^{i,h_1}\}.$$

The conditional beliefs for the types about the opponent's strategy-type pairs can then be based on the beliefs above. By using the construction in Theorem 5.2, this yields the following beliefs for the types:

$$
\begin{aligned}
b_1(t_1^{(a,c),\emptyset}, \emptyset) &= (h, t_2^{h,\emptyset}), \qquad b_1(t_1^{(a,c),\emptyset}, h_1) = (h, t_2^{h,h_1}), \\
b_1(t_1^{b,\emptyset}, \emptyset) &= (g, t_2^{g,\emptyset}), \qquad b_1(t_1^{b,\emptyset}, h_1) = (g, t_2^{g,h_1}), \\
b_1(t_1^{(a,c),h_1}, \emptyset) &= (h, t_2^{h,\emptyset}), \qquad b_1(t_1^{(a,c),h_1}, h_1) = (h, t_2^{h,h_1}), \\
b_1(t_1^{(a,d),h_1}, \emptyset) &= (g, t_2^{g,\emptyset}), \qquad b_1(t_1^{(a,d),h_1}, h_1) = (g, t_2^{g,h_1}),
\end{aligned}
$$

$$
\begin{aligned}
b_2(t_2^{g,\emptyset}, h_1) &= ((a,c), t_1^{(a,c),h_1}), \\
b_2(t_2^{g,h_1}, h_1) &= ((a,c), t_1^{(a,c),h_1}), \\
b_2(t_2^{h,\emptyset}, h_1) &= ((a,d), t_1^{(a,d),h_1}), \\
b_2(t_2^{h,h_1}, h_1) &= ((a,d), t_1^{(a,d),h_1}), \\
b_2(t_2^{i,\emptyset}, h_1) &= (0.5) \cdot ((a,c), t_1^{(a,c),h_1}) + (0.5) \cdot ((a,d), t_1^{(a,d),h_1}), \\
b_2(t_2^{i,h_1}, h_1) &= (0.5) \cdot ((a,c), t_1^{(a,c),h_1}) + (0.5) \cdot ((a,d), t_1^{(a,d),h_1}).
\end{aligned}
$$

Here, $b_2(t_2^{i,\emptyset}, h_1) = (0.5) \cdot ((a,c), t_1^{(a,c),h_1}) + (0.5) \cdot ((a,d), t_1^{(a,d),h_1})$ means that type $t_2^{i,\emptyset}$ assigns at h_1 probability 0.5 to the event that player 1 chooses (a,c) while being of type $t_1^{(a,c),h_1}$, and assigns probability 0.5 to the event that player 1 chooses (a,d) while being of type $t_1^{(a,d),h_1}$.

By Theorem 5.2 we know that all types so constructed express *common belief in future rationality*, and that for every strategy that can rationally be chosen under *common belief in future rationality* there is a type in this model for which that strategy is optimal. Indeed, the backward dominance procedure delivers the strategies $(a,c), b, g, h$ and i at $\emptyset$, at hence we know from Perea (2011a) that these are exactly the strategies that can rationally be chosen under *common belief in future rationality.* Note that

- strategy (a,c) is optimal, at $\emptyset$ and h_1, for the type $t_1^{(a,c),\emptyset}$,
- strategy b is optimal, at $\emptyset$, for the type $t_1^{b,\emptyset}$,
- strategy g is optimal, at h_1, for the type $t_2^{g,\emptyset}$,
- strategy h is optimal, at h_1, for the type $t_2^{h,\emptyset}$, and
- strategy i is optimal, at h_1, for the type $t_2^{i,\emptyset}$.

So, for every strategy that can rationally be chosen under *common belief in future rationality,* we have constructed – by means of the epistemic model above – a finitely generated belief hierarchy that expresses *common belief in future rationality,* and that supports this strategy.

Note, however, that there is some redundancy in the epistemic model above. Namely, it is easily seen that the types $t_1^{(a,c),\emptyset}$ and $t_1^{(a,c),h_1}$ have identical belief hierarchies, and so do $t_1^{b,\emptyset}$ and $t_1^{(a,d),h_1}$. The same holds for $t_2^{g,\emptyset}$ and t_2^{g,h_1}, for $t_2^{h,\emptyset}$ and t_2^{h,h_1}, and also for $t_2^{i,\emptyset}$ and t_2^{i,h_1}. Hence, we can substitute $t_1^{(a,c),\emptyset}$ and $t_1^{(a,c),h_1}$ by a single type $t_1^{(a,c)}$, and we can substitute $t_1^{b,\emptyset}$ and $t_1^{(a,d),h_1}$ by a single type t_1^b. Similarly, we can substitute $t_2^{g,\emptyset}$ and t_2^{g,h_1} by a single type t_2^g, we can substitute $t_2^{h,\emptyset}$ and t_2^{h,h_1} by a single type t_2^h, and $t_2^{i,\emptyset}$ and t_2^{i,h_1} by t_2^i. This eventually leads to the smaller – yet equivalent – epistemic model with type sets

$$T_1 = \{t_1^{(a,c)}, t_1^b\} \text{ and } T_2 = \{t_2^g, t_2^h, t_2^i\}$$

and beliefs

$$\begin{aligned} b_1(t_1^{(a,c)}, \emptyset) &= b_1(t_1^{(a,c)}, h_1) = (h, t_2^h) \\ b_1(t_1^b, \emptyset) &= b_1(t_1^b, h_1) = (g, t_2^h) \\ \\ b_2(t_2^g, h_1) &= ((a,c), t_1^{(a,c)}), \\ b_2(t_2^h, h_1) &= ((a,d), t_1^b), \\ b_2(t_2^i, h_1) &= (0.5) \cdot ((a,c), t_1^{(a,c)}) + (0.5) \cdot ((a,d), t_1^b). \end{aligned}$$

This redundancy is typical for the construction of the epistemic model in Theorem 5.2. In most games, the epistemic model constructed in this way will contain types that are "duplicates" of each other, as they generate the same belief hierarchy.

6. Discussion

6.1. Algorithms as Reasoning Procedures

In this chapter we have presented two algorithms that lead to belief hierarchies expressing *common belief in future rationality.* The first algorithm can be viewed as an instance of a *best-reply dynamics,* in which you start with a choice-history combination (c^1, x^1), then take a choice-history combination (c^2, x^2) that is a "best reply" against (c^1, x^1), then take a choice-history combination (c^3, x^3) that is a "best reply" against (c^2, x^2), and so on. The real novelty here is that we interpret this algorithm as a *reasoning procedure* that some player – say player i – performs in his mind to generate an "easy" belief hierarchy that expresses *common belief in future rationality.* So, the complete procedure takes place in the mind of a *single person*, and no real interaction is needed between the two players in order for player i to reason according to this procedure. So, player i can perform the best-reply dynamics in his mind without there being a real dynamics between him and the opponent.

The second algorithm is based on the backward dominance procedure proposed in Perea (2011a). The difference is that in this chapter we interpret this algorithm not as a computational tool for the analyst, but rather as a finite reasoning procedure that some player *inside* the game can use (a) to verify which strategies he can rationally choose under *common belief in future rationality*, and (b) to support each of these strategies by a belief hierarchy expressing *common belief in future rationality.*

Hence, one of the main messages in this chapter is that the two algorithms above for *common belief in future rationality* do not only serve as a computational tool for the analyst, but can also be used by a player inside the game as an intuitive reasoning procedure. Compare this to the concepts of Nash equilibrium (Nash (1950, 1951)) for static games, and sequential equilibrium (Kreps and Wilson (1982)) for dynamic games. There is no easy, finite iterative procedure to find one Nash equilibrium – let alone *all* Nash equilibria – in a game. In particular, there is no clear reasoning procedure that a player inside the game can use to reason his way towards a Nash equilibrium. Besides, we believe that Nash equilibrium imposes some implausible conditions on a player's belief hierarchy, as it requires a player to believe that his opponent is *correct* about the actual beliefs he holds (see Perea (2007b), Aumann and Brandenburger (1995), Polak (1999), Brandenburger and Dekel (1989), and Asheim (2006, p.5). In view of all this, we think that Nash equilibrium is not a very appealing concept if we wish to describe the reasoning of players about their opponents. The same actually holds for the concept of sequential equilibrium.

6.2. Simple Existence Proof for Common Belief in Future Rationality

The reasoning procedure in Theorem 4.2 shows, in particular, that for a finite dynamic game it is always possible to hold a belief hierarchy that expresses *common belief in future rationality.* That is, *common belief in future rationality* is always possible in every finite dynamic game. For static games there is an analogous result which states that *common belief in rationality* is always possible in every finite static game. The traditional proof for this result, which is typically used in textbooks, makes use of the existence of Nash equilibrium. The argument runs as follows: For a given static game with finitely many choices we know, by the work of Nash (1950, 1951), that a Nash equilibrium in randomized choices (or mixed strategies) exists. Then, this Nash equilibrium can be used to define an epistemic model with *one* type for both players, say t_i and t_j, in which the belief of type t_i about j's choice is given by j's randomized choice in the Nash equilibrium, and similarly for type t_j. Then, it can be shown that both types t_i and t_j express *common belief in rationality.* However, this existence proof heavily relies of the existence proof for Nash equilibrium, and is therefore not an elementary proof.

In a similar way we could prove the "existence" of *common belief in future rationality* for finite dynamic games, making use of the existence of sequential equilibrium (Kreps and Wilson (1982)). A *sequential equilibrium* consists of a behavioral strategy for both players, specifying at each of their information sets a probability distribution over the available choices, and a system of beliefs, specifying for both players at each of their information sets a probability distribution over the histories at that information set. For a given finite dynamic game we know, by the work of Kreps and Wilson (1982), that a sequential equilibrium exists. Then, this sequential equilibrium can be used to define an epistemic model with *one* type for both players, say t_i and t_j, in which the conditional beliefs of type t_i about j's strategy choice are given by i's system of beliefs and j's behavioral strategy in the sequential equilibrium, and similarly for type t_j. In fact, i's system of beliefs in the sequential equilibrium generates i's conditional beliefs about j's *past* choices, whereas j's behavioral strategy generates i's conditional beliefs about j's *future* choices. Then, it can be shown that both types t_i and t_j express *common belief in future rationality.* See the literature section of Chapter 8 in Perea (2012) for more details on this line of argument. However, this existence proof heavily relies on the existence proof for sequential equilibrium, and is therefore not an elementary proof.

Our reasoning procedure in Theorem 4.2, on the other hand, *does* provide a completely *elementary* existence proof for *common belief in future rationality.* There is no reference to any equilibrium concept in this procedure, and the key insight that makes the procedure work is that every chain of choice-history combinations must eventually enter into a cycle – something that is immediately clear since there are only finitely many choices and histories in the game. In that sense, the reasoning procedure in Theorem 4.2 can be used as a simple and elementary proof for the "existence" of *common belief in future rationality* in finite games. Moreover, we believe that this existence proof is also much more instructive than the one based on the existence of sequential equilibrium, as it also illustrates *how* to construct a belief hierarchy that expresses *common belief in future rationality.* The other existence proof *assumes* a sequential equilibrium, but does not explain *how* a player reasons his way towards this sequential equilibrium. Such an existence proof would therefore be much more opaque.

6.3. Finitely Generated Belief Hierarchies

In this chapter we have restricted our attention to *finitely generated* belief hierarchies – that is, belief hierarchies that can be derived from an epistemic model with *finitely* many types. By doing so we actually exclude some belief hierarchies, as not every belief hierarchy can be generated within a finite epistemic model. If we wish to include *all* possible belief hierarchies in our model, then we must necessarily look at *complete* type spaces for dynamic games as constructed in Battigalli and Siniscalchi (1999).

But for our purposes here it is actually sufficient to concentrate on finitely generated belief hierarchies. Theorem 5.2 implies, namely, that whenever a strategy s_i is optimal for *some* belief hierarchy – not necessarily finitely generated – that expresses *common belief in future rationality*, then s_i is also optimal for some *finitely generated* belief hierarchy that expresses *common belief in future rationality.* Moreover, finitely generated belief hierarchies have the advantage that they are particularly easy to work with, and that checking for *common belief in future rationality* can be done within finitely many steps, as is shown in Theorem 3.2.

6.4. Other Concepts

We have concentrated on the concept of *common belief in future rationality* in this chapter. For this concept we have provided a particularly easy, finite reasoning procedure that a player can use to generate *some* belief hierarchy that expresses *common belief in future rationality* (see Theorem 4.2), and another finite reasoning procedure that yields, for *every* strategy s_i that can rationally be chosen under *common belief in future rationality*, a belief hierarchy expressing *common belief in future rationality* that supports this strategy s_i (see Theorem 5.2). Moreover, the first reasoning procedure yields a belief hierarchy with a very simple structure, as it assigns at every information set probability 1 to one particular strategy and one particular belief hierarchy for the opponent.

We could also look at *other* concepts in the epistemic game theory literature, and see whether we can find associated finite reasoning procedures like the ones we found for *common belief in future rationality*. These epistemic concepts include the basic notion of *common belief in rationality* (Tan and Werlang (1988), Brandenburger and Dekel (1987)) which is based on the concept of *rationalizability* (Bernheim (1984) and Pearce (1984)), and refinements of common belief in rationality like *permissibility* (Brandenburger (1992). Börgers (1994)), *proper rationalizability* (Schuhmacher (1999), Asheim (2001)) and *common assumption of rationality* (Brandenburger, Friedenberg and Keisler (2008)) for static games with lexicographic beliefs, and *common strong belief in rationality* (Battigalli and Siniscalchi (2002)) for dynamic games.

For each of the epistemic concepts above there exists a *finite* recursive procedure that yields all choices (or strategies, if we have a dynamic game) that can rationally be chosen under the concept. We list these procedures, with their references, in Table 5. Among these procedures, iterated elimination of weakly dominated choices is an old algorithm with a long tradition in game theory, and it is not clear where this procedure has been described for the first time in the literature. The procedure already appears in early books by Luce and Raiffa (1957) and Farquharson (1969). An overview of these epistemic concepts and their associated recursive procedures can be found in Perea (2012).

In fact, each of these recursive procedures can be interpreted not only as algorithmic methods for the analyst, but can also be understood as *reasoning procedures* that a player inside the game can use to reason his way towards the choices or strategies he can rationally make under the epistemic concept at hand. Namely, some of the references in Table 5 contain proofs which show that the recursive procedures select *precisely* those choices or strategies that are optimal for *some* belief hierarchy that meets the requirements in the associated epistemic concept. In particular, these proofs construct, for every choice or strategy surviving the procedure, some belief hierarchy that supports this choice or strategy, and which satisfies all conditions imposed by the epistemic concept – similarly to Theorem 5.2 in this chapter. A player inside the game can then use the recursive procedure, together with this construction, to reason his way towards a "sufficient set of belief hierarchies" for this particular epistemic concept.

Moreover, for most of the epistemic concepts in Table 5 there is an easy recursive procedure

Epistemic concept	Recursive Procedure
Common belief in rationality (Tan and Werlang (1988)	Iterated elimination of strictly dominated choices (based on Pearce (1984), Tan and Werlang (1988))
Permissibility (Brandenburger (1992), Börgers (1994))	Dekel-Fudenberg procedure (Dekel and Fudenberg (1990))
Proper rationalizability (Schuhmacher (1999), Asheim (2001))	Iterated addition of preference restrictions (Perea (2011b))
Common assumption of rationality (Brandenburger, Friedenberg and Keisler (2008))	Iterated elimination of weakly dominated choices
Common belief in future rationality (Perea (2011a))	Backward dominance procedure (Perea (2011a))
Common strong belief in rationality (Battigalli and Siniscalchi (2002))	Iterated conditional dominance procedure (Shimoji and Watson (1998), based on Pearce (1984), Battigalli (1997))

Table 5: Overview of epistemic concepts and their recursive procedures

---

that generates *some* belief hierarchy meeting the epistemic conditions, by constructing a *chain of choices* or *beliefs* – similarly to the chain of choices and histories constructed in Theorem 4.2 – and showing that this chain enters into a cycle. This chain can then be transformed into a finite epistemic model where all types satisfy the epistemic conditions imposed by the concept, similarly to our construction for *common belief in future rationality* in Theorem 4.2.

In fact, such a construction works for all epistemic concepts in the table, except for *common assumption of rationality* and *common strong belief in rationality* – see Perea (2012) for an overview. This method does not work for *common assumption of rationality* and *common strong belief in rationality* because for these concepts, in order to verify that a given belief hierarchy meets the epistemic conditions in the concept, it is not sufficient to only look at *this particular* belief hierarchy, but we must consider sufficiently many *alternative* belief hierarchies as well. That is, the conditions in *common assumption of rationality* and *common strong belief in rationality* cannot be checked locally, but need a global perspective instead – something that cannot be achieved by constructing a chain of beliefs similarly to Theorem 4.2. For each of the other concepts in Table 5 such a local perspective *is* sufficient. Namely, for the other concepts, in order to verify that a belief hierarchy meets the epistemic conditions imposed by the concept, it is sufficient to only look at this single belief hierarchy.

Similarly to what we have seen for *common belief in future rationality*, the constructions above – generating chains of choices or beliefs – can serve as easy and elementary existence proofs for the concepts of *common belief in rationality, permissibility* and *proper rationalizability.*

There is, however, no such easy existence proof available for the concepts of *common assumption of rationality* and *common strong belief in rationality.*

References

[1] G.B. Asheim, 'Proper rationalizability in lexicographic beliefs', *International Journal of Game Theory,* 30 (2001), 453–478

[2] G.B. Asheim, 'On the epistemic foundation for backward induction', *Mathematical Social Sciences,* 44 (2002), 121–144

[3] G.B. Asheim, *The Consistent Preferences Approach to Deductive Reasoning in Games* (Theory and Decision Library, Springer, Dordrecht, The Netherlands, 2006)

[4] R.J. Aumann and A. Brandenburger, 'Epistemic conditions for Nash equilibrium', *Econometrica,* 63 (1995), 1161–1180

[5] A. Baltag, S. Smets and J.A. Zvesper, 'Keep 'hoping' for rationality: a solution to the backward induction paradox', *Synthese,* 169 (2009), 301–333 (*Knowledge, Rationality and Action,* 705–737)

[6] P. Battigalli, 'On rationalizability in extensive games', *Journal of Economic Theory,* 74 (1997), 40–61

[7] P. Battigalli and M. Siniscalchi, 'Hierarchies of conditional beliefs and interactive epistemology in dynamic games', *Journal of Economic Theory,* 88 (1999), 188–230

[8] P. Battigalli and M. Siniscalchi, 'Strong belief and forward induction reasoning', *Journal of Economic Theory,* 106 (2002), 356–391

[9] B.D. Bernheim, 'Rationalizable strategic behavior', *Econometrica,* 52 (1984), 1007–1028

[10] T. Börgers, 'Weak dominance and approximate common knowledge', *Journal of Economic Theory,* 64 (1994), 265–276

[11] A. Brandenburger, 'Lexicographic probabilities and iterated admissibility' in P. Dasgupta *et al.* (eds.), *Economic Analysis of Markets and Games* (MIT Press, Cambridge, MA, 1992), pp. 282–290

[12] A. Brandenburger and E. Dekel, 'Rationalizability and correlated equilibria', *Econometrica,* 55 (1987), 1391–1402

[13] A. Brandenburger and E. Dekel, 'The role of common knowledge assumptions in game theory', in F. Hahn (ed.), *The Economics of Missing Markets, Information and Games* (Oxford University Press, Oxford, 1989), pp. 46–61

[14] A. Brandenburger, A. Friedenberg and J. Keisler, 'Admissibility in games', *Econometrica,* 76 (2008), 307–352

[15] J. Chen and S. Micali, 'The robustness of extensive-form rationalizability', Working paper (2011)

[16] E. Dekel and D. Fudenberg, 'Rational behavior with payoff uncertainty', *Journal of Economic Theory,* 52 (1990), 243–267

[17] R. Farquharson, *Theory of Voting* (Yale University Press, New Haven, CT, 1969)

[18] Y. Feinberg, 'Subjective reasoning - dynamic games', *Games and Economic Behavior,* 52 (2005), 54–93

[19] J.C. Harsanyi, 'Games with incomplete information played by "bayesian" players, I–III', *Management Science,* 14 (1967–1968), 159–182, 320–334, 486–502

[20] E. Hendon, H.J. Jacobsen and B. Sloth, 'The one-shot-deviation principle for sequential rationality', *Games and Economic Behavior,* 12 (1996), 274–282

[21] R.D. Luce and H. Raiffa, *Games and Decisions* (John Wiley and Sons, New York, 1957)

[22] D.M. Kreps and R. Wilson, 'Sequential equilibria', *Econometrica,* 50 (1982), 863–894

[23] J.F. Nash, 'Equilibrium points in N-person games', *Proceedings of the National Academy of Sciences of the United States of America,* 36 (1950), 48–49

[24] J.F. Nash, 'Non-cooperative games', *Annals of Mathematics,* 54 (1951), 286–295

[25] D. Pearce, 'Rationalizable strategic behavior and the problem of perfection', *Econometrica,* 52 (1984), 1029–1050

[26] A. Penta, 'Robust dynamic mechanism design', Working paper (2009)

[27] A. Perea, 'A note on the one-deviation property in extensive games', *Games and Economic Behavior,* 40 (2002), 322–338

[28] A. Perea, 'Epistemic foundations for backward induction: An overview' in J. van Benthem, D. Gabbay and B. Löwe (eds.), *Interactive Logic Proceedings of the 7th Augustus de Morgan Workshop, London. Texts in Logic and Games 1* (Amsterdam University Press, 2007a), pp. 159–193

[29] A. Perea, 'A one-person doxastic characterization of Nash strategies', *Synthese,* 158 (2007b), 251–271 (*Knowledge, Rationality and Action* 341–361)

[30] A. Perea, 'Belief in the opponents' future rationality', Working paper (2011a)

[31] A. Perea, 'An algorithm for proper rationalizability', *Games and Economic Behavior,* 72 (2011b), 510–525

[32] A. Perea, *Epistemic Game Theory: Reasoning and Choice* (Cambridge University Press, 2012), Forthcoming

[33] B. Polak, 'Epistemic conditions for Nash equilibrium, and common knowledge of rationality', *Econometrica,* 67 (1999), 673–676

[34] A. Rubinstein, 'Comments on the interpretation of game theory', *Econometrica,* 59 (1991), 909–924

[35] D. Samet, 'Hypothetical knowledge and games with perfect information', *Games and Economic Behavior,* 17 (1996), 230–251

[36] F. Schuhmacher, 'Proper rationalizability and backward induction', *International Journal of Game Theory,* 28 (1999), 599–615

[37] M. Shimoji and J. Watson, 'Conditional dominance, rationalizability, and game forms', *Journal of Economic Theory,* 83 (1998), 161–195

[38] T. Tan and S.R.C. Werlang, 'The bayesian foundations of solution concepts of games', *Journal of Economic Theory,* 45 (1988), 370–391