Dynamic Consistency in Games without Expected Utility

Andrés Perea*

EPICENTER Working Paper No. 29 (2024)

Abstract

Within dynamic games we are interested in conditions on the players' preferences that imply *dynamic consistency* and the existence of *sequentially optimal strategies*. The latter means that the strategy is optimal at each of the player's information sets, given his beliefs there. These two properties are needed to undertake a meaningful game-theoretic analysis in dynamic games. To explore this we assume that every player holds a *conditional preference relation* – a mapping that assigns to every probabilistic belief about the opponents' strategies a preference relation over his own strategies. We identify sets of very basic conditions on the conditional preference relations that guarantee dynamic consistency and the existence of sequentially optimal strategies, respectively. These conditions are implied by, but are much weaker than, assuming expected utility. That is, to undertake a meaningful game-theoretic analysis in dynamic games we can do with much less than expected utility.

JEL classification: C72, D81

Keywords: Dynamic games, dynamic consistency, sequentially optimal strategy, conditional preference relation, non-expected utility

* *E-mail:* a.perea@maastrichtuniversity.nl *Affiliation:* Maastricht University, Department of Quantitative Economics and EpiCenter, The Netherlands.

1 Introduction

The principle of *dynamic consistency* plays a central role in one-person decision theory. It states that the decision maker's preferences at different points in time must be sufficiently aligned. More precisely, if the decision maker ex-ante ranks two acts that only differ conditional on an event E, then the ranking should not change upon observing that E has been realized. For a detailed account, the reader may consult Machina (1989) and the references therein.

Dynamic consistency is also of key importance to dynamic games, although on a somewhat more implicit basis. In most equilibrium and non-equilibrium concepts for dynamic games, such as sequential equilibrium (Kreps and Wilson (1982)), sequential rationalizability (Dekel, Fudenberg and Levine (1999, 2002) and Asheim and Perea (2005)), backwards rationalizability (Perea (2014) and Penta (2015)) and extensive-form rationalizability (Pearce (1984), Battigalli (1997)), it is assumed that every player possesses strategies that are *sequentially optimal,* that is, optimal at each of his information sets given his conditional beliefs there. The existence of such sequentially optimal strategies relies, in turn, on the dynamic consistency that the players exhibit in the game: If at a certain information set h the player ranks two strategies that only differ conditional on reaching a future information set h', and the player expects h' to be reached with positive probability, then his ranking should not change upon reaching h'.

In one-person decision theory, dynamic consistency has also been explored for preferences that do not conform to expected utility (see, again, Machina (1989) for an overview). This is important because experimental evidence shows that many decision makes deviate from the assumptions of expected utility.

This raises the question: What about dynamic games? If the players are not necessarily assumed to be expected utility maximizers, how much should we pre-suppose so that we can still perform a meaningful game-theoretic analysis? This is the question we wish to explore in this paper.

It is fair to say that such a meaningful analysis is only possible if both dynamic consistency and the existence of sequentially optimal strategies are guaranteed to hold. In this paper we therefore aim for some mild conditions on the players' preferences that imply both of these properties.

Towards this goal we assume that every player in the dynamic game holds a *conditional preference relation* (Gilboa and Schmeidler (2003), Perea (2023)) – a mapping that assigns to every possible probabilistic belief about the opponents' strategies a preference relation over his own strategies. We choose this model because it nicely reflects the game-theoretic principle that the ranking over your own strategies crucially depends on your belief about the behavior of others. And it does so without assuming expected utility. At the same time, it is flexible enough to induce a preference relation for a player at each of his information sets: Simply take his conditional preference relation, take the conditional belief he holds at that information set, and see what preference relation it induces over his own strategies.

One key difference with the more traditional models of Savage (1954) and Anscombe and

Aumann (1963) is that we *assume* that the players are Bayesian, by holding probabilistic beliefs about the opponents' strategies. On the other hand, we do not assume a *unique* belief for the players, as a conditional preference relation specifies a preference relation over strategies for *every possible belief.* The rationale is that in a dynamic game, a player may change his belief throughout his reasoning process, or upon observing new information, and he is typically uncertain about the beliefs held by his opponents.

Within this decision-theoretic framework we identify a set of very basic conditions on conditional preference relations which guarantee dynamic consistency: *preservation of indifference, preservation of strict preference,* and *respect of outcome-equivalent strategies.* The first condition states that for every two beliefs where the player is indifferent between two strategies, he will remain indifferent if he uses any belief on the line segment between these two beliefs. The second condition is similar, but applies to strict preference. The third condition states that if two strategies lead to the same outcome under the opponents' strategy combination s_{-i}, then player i must be indifferent between the two strategies if he assigns probability 1 to s_{-i}. Moreover, to guarantee the existence of sequentially optimal strategies we find that the basic conditions above, together with *transitivity*, are sufficient.

These conditions are implied by, but are much weaker than, expected utility. This is important, since it follows by the axiomatic treatments in Gilboa and Schmeidler (2003) and Perea (2023) that assuming expected utility may be very demanding. In particular, the conditions *three-choice linear preference intensity* and *four-choice linear preference intensity* in Perea (2023), which are needed for expected utility, impose a substantial cognitive burden on behalf of the decision maker. In contrast, the sufficient conditions above are very basic and mild. This paper thus shows that these basic conditions are already enough to undertake a meaningful game-theoretic analysis.

The paper is organized as follows: In Section 2 we lay out the model of a dynamic game, on the basis of which we define strategies and conditional beliefs. In Section 3 we present the decision-theoretic framework based on conditional preference relations. In Section 4 we define dynamic consistency and provide some basic sufficient conditions on the players' conditional preference relations that imply it. In Section 5 we do the same for the existence of sequentially optimal strategies. In Section 6 we provide some concluding remarks. The appendix contains all the proofs.

2 Games, Strategies and Beliefs

2.1 Dynamic Game Forms

In this paper we consider finite dynamic games that allow for simultaneous moves and imperfect information. Formally, a *dynamic game form* is a tuple $D = (I, P, I^a, (A_i, H_i)_{i \in I}, Z)$, where

(a) I is the finite set of *players*;

(b) P is the finite set of *past action profiles*, or *histories;*

(c) the mapping I^a assigns to every history $p \in P$ the (possibly empty) set of *active players* $I^a(p) \subseteq I$ who must choose after history p. If $I^a(p)$ contains more than one player, there are simultaneous moves after p. If $I^a(p)$ is empty, the game terminates after p. By P_i we denote the set of histories $p \in P$ with $i \in I^a(p)$;

(d) for every player i, the mapping A_i assigns to every history $p \in P_i$ the finite set of *actions* $A_i(p)$ from which player i can choose after history p. The objects P, I^a and $(A_i)_{i \in I}$ must be such that the empty history $\emptyset$ is in P, representing the beginning of the game, and the non-empty histories in P are precisely those objects $(p, (a_i)_{i \in I^a(p)})$ where p is a history in P and $a_i \in A_i(p)$ for every $i \in I^a(p)$;

(e) for every player i there is a partition H_i of the set of histories P_i where i is active. Every partition element $h_i \in H_i$ is called an *information set* for player i. In case h_i contains more than one history, the interpretation is that player i does not know at h_i which history in h_i has been reached. The objects A_i and H_i must be such that for every information set $h_i \in H_i$ and every two histories p, p' in h_i, we have that $A_i(p) = A_i(p')$. We can thus write $A_i(h_i)$ for the unique set of available actions at h_i. Moreover, it must be that $A_i(h_i) \cap A_i(h_i') = \emptyset$ for every two distinct information sets $h_i, h_i' \in H_i$;

(f) $Z \subseteq P$ is the collection of histories p where the set of active players $I^a(p)$ is empty. Such histories are called *terminal* histories, or *consequences.*

This definition follows Osborne and Rubinstein (1994), with the difference that we do not specify utilities at the consequences. This is why we call it a dynamic game *form* and not a dynamic game.

Based on this model we can derive the following definitions: We say that a history p *precedes* a history p' (or p' *follows* p) if p' results by adding some action profiles after p. Let $H := \cup_{i \in I} H_i$ be the collection of all information sets for all players. For every two information sets $h, h' \in H$, we say that h *precedes* h' (or h' *follows* h) if there is a history $p \in h$ and a history $p' \in h'$ such that p precedes p'. Two information sets h, h' are *simultaneous* if there is some history p which belongs to both h and h'. We say that h *weakly precedes* h' (or h' *weakly follows* h) if either h precedes h', or h, h' are simultaneous.

The dynamic game form satisfies *perfect recall* (Kuhn (1953)) if every player always remembers which actions he chose in the past, and which information he had about the opponents' past actions. Formally, for every player i, every information set $h_i \in H_i$, and every two histories $p, p' \in h_i$, the sequences of player i actions in p and p' must be the same (and consequently, the collection of player i information sets that p and p' cross must be the same). For the remainder of this paper we will always assume that the dynamic game form satisfies perfect recall.

2.2 Strategies

A strategy for player i assigns an available action to every information set at which player i is active, and that is not excluded by earlier actions in the strategy. Formally, let $\tilde{s}_i$ be a mapping that assigns to *every* information set $h_i \in H_i$ some action $\tilde{s}_i(h) \in A_i(h)$. We call $\tilde{s}_i$ a *complete* *strategy.* Then, a history $p \in P$ is *excluded* by $\tilde{s}_i$ if there is some information set $h_i \in H_i$, with some history $p' \in h_i$ preceding p, such that $\tilde{s}_i(h_i)$ is different from the unique player i action at p' leading to p. An information set $h \in H$ is excluded by $\tilde{s}_i$ if all histories in h are excluded by $\tilde{s}_i$. The *strategy* induced by $\tilde{s}_i$ is the restriction of $\tilde{s}_i$ to those information sets in H_i that are not excluded by $\tilde{s}_i$. A mapping $s_i : \tilde{H}_i \to \cup_{h_i \in \tilde{H}_i} A_i(h_i)$, where $\tilde{H}_i \subseteq H_i$, is a *strategy* for player i if it is the strategy induced by a complete strategy.[1] By S_i we denote the set of strategies for player i, and by $S_{-i} := \times_{j \neq i} S_j$ the set of strategy combinations for i's opponents.

Consider a strategy profile $s = (s_i)_{i \in I}$ in $\times_{i \in I} S_i$. Then, s induces a unique consequence $z(s)$. We say that the strategy profile s *reaches* a history p if p precedes $z(s)$. Similarly, the strategy profile s is said to reach an information set h if s reaches a history in h.

For a given information set $h \in H$ and player i we define the sets

$$\begin{aligned} S(h) &:= \{s \in \times_{i \in I} S_i \mid s \text{ reaches } h\}, \\ S_i(h) &:= \{s_i \in S_i \mid \text{there is some } s_{-i} \in S_{-i} \text{ such that } (s_i, s_{-i}) \in S(h)\}, \text{ and} \\ S_{-i}(h) &:= \{s_{-i} \in S_{-i} \mid \text{there is some } s_i \in S_i \text{ such that } (s_i, s_{-i}) \in S(h)\}. \end{aligned}$$

Intuively, $S_i(h)$ is the set of strategies for player i that allow for information set h to be reached, whereas $S_{-i}(h)$ is the set of opponents' strategy combinations that allow for h to be reached.

2.3 Beliefs

In a dynamic game form, a player holds a belief about the opponents' strategies at every information set where he is active. More precisely, a *conditional belief vector* b_i for player i assigns to every information set $h_i \in H_i$ a conditional probabilistic belief $b_i(h_i) \in \Delta(S_{-i}(h_i))$ about the opponents' strategy combinations that are still possible when h_i is reached. Here we denote, for a finite set X, by $\Delta(X)$ the set of probability distributions on X.

Many concepts for dynamic games require the conditional belief vector to satisfy Bayesian updating. Formally, the conditional belief vector b_i satisfies *Bayesian updating* if for every two information sets $h_i, h_i' \in H_i$ where h_i precedes h_i' and $b_i(h_i)(S_{-i}(h_i')) > 0$, it holds that

$$b_i(h_i')(s_{-i}) = \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i}(h_i'))}$$

for all opponents' strategy combination $s_{-i} \in S_{-i}(h_i')$.

3 Conditional Preference Relations

The ultimate question is: What strategy, or strategies, can a player in a dynamic game plausibly choose? This will depend crucially on the *beliefs* that the player holds about the opponents'

[1] What we call a "strategy" is sometimes called a "plan of action" in the literature (Rubinstein (1991)), and what we call a "complete strategy" is often called a "strategy".

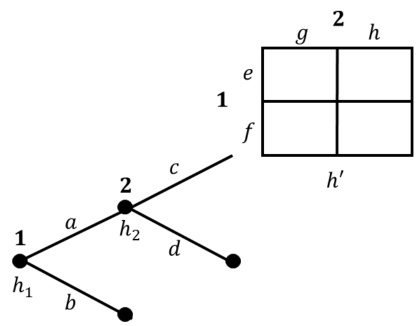

Figure 1: A dynamic game form

strategies: For different beliefs, the player may opt for different strategies. To capture this phenomenon most generally, we assume that the player holds, for *every possible* belief about the opponents' strategies, a preference relation over his own strategies. This is modelled by a *conditional preference relation* (Gilboa and Schmeidler (2003), Perea (2023)), and we take this as the primitive object for our analysis.

Definition 3.1 (Conditional preference relation) *A* ***conditional preference relation*** $\succsim_i$ *for player i specifies for every belief $\beta_i \in \Delta(S_{-i})$ about the opponents' strategy combinations a complete and reflexive preference relation $\succsim_{i,\beta_i}$ over his strategies.*

As an illustration, consider the dynamic game form in Figure 1. Here, h' denotes an information set where players 1 and 2 choose simultaneously. A possible conditional preference relation $\succsim_1$ for player 1 has been depicted in Figure 2. The picture should be read as follows: Every belief for player 1 is a probability distribution over player 2's strategies $(c, g), (c, h)$ and d, and can thus be identified with a point in the triangle. The corner points of the triangle are thus the "opinionated" beliefs that assign probability 1 to one of the three strategies. The picture reveals that for every belief to the left of the curve, player 1 prefers the strategy (a, e) to the strategy (a, f), and the strategy (a, f) to b. For every belief to the right of the curve he prefers (a, f) to (a, e) and (a, e) to b. For every belief on the curve he is indifferent between (a, e) and (a, f), and prefers both strategies to b.

The conditional preference relation $\succsim_1$ above also specifies how player 1 would change the ranking of his strategies when he revises his belief upon reaching a new information set. Suppose, for instance, that player 1 initially holds the belief $(0.5) \cdot (c, g) + (0.5) \cdot d$, where he assigns equal probability to player 2 choosing the strategies (c, g) and d. Figure 2 then tells us that player 1 will initially prefer his strategy (a, e) to (a, f), and his strategy (a, f) to b. Suppose now that, upon reaching his second information set h', he revises his belief by Bayesian updating to (c, g). From Figure 2 we learn that at h' player 1 would prefer (a, f) to (a, e), and (a, e) to b.

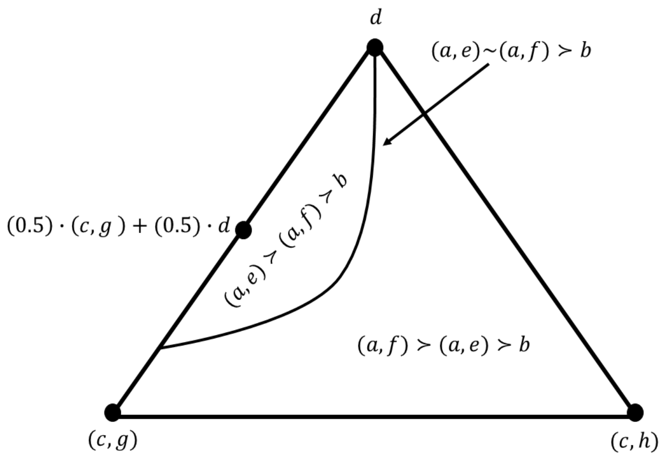

Figure 2: A conditional preference relation for the dynamic game form in Figure 1

This holds in general: If we fix a conditional preference relation $\succsim_i$ for player i, and specify a conditional belief vector b_i, describing what belief player i would have at each of his information sets, then we know for every information set what his preferences over his strategies would be. Indeed, at a given information set $h_i \in H_i$ player i would have the belief $b_i(h_i)$, which in turn induces the preference relation $\succsim_{i,b_i(h_i)}$ over his own strategies.

4 Dynamic Consistency

In this section we first provide a definition of *dynamic concistency* in the context of conditional preference relations, and subsequently lay out some intuitive properties that imply dynamic consistency. At the end we illustrate, by means of an example, that these properties do not require the conditional preference relation to have an expected utility representation.

4.1 Definition

In dynamic decision problems, the term *dynamic consistency* refers to the general idea that the decision maker, as time passes by, should not reverse the ranking between two options "without good reason". More precisely, if the decision maker initially ranks two acts that only differ conditional on an event E, then the decision maker should not change his ranking if he learns that the event E obtains.

Within the context of a conditional preference relation, this idea can be translated as follows: Suppose player i compares two strategies, s_i and t_i, that both can possibly reach an information set $h_i' \in H_i$, and that only differ at information sets that weakly follow h_i'. Now consider an information set $h_i \in H_i$ that precedes h', such that player i believes at h_i that h_i' may be reached with positive probability, and that player i prefers s_i to t_i at h_i. Then, under Bayesian updating player i should still prefer s_i to t_i at h_i'.

The intuition is the following: If the player prefers s_i to t_i at h_i, then apparently player i believes at h_i that the moves of his opponents after, or at, h_i' work in favor of s_i. If the play moves from h_i to h_i', then under Bayesian updating player i will maintain his belief about the opponents' moves after, or at, h_i'. As such, player i should still believe at h_i' that the future moves of his opponents work in favor of s_i.

Definition 4.1 (Dynamic consistency) *A conditional preference relation $\succsim_i$ for player i is* ***dynamically consistent*** *if for every conditional belief vector b_i that satisfies Bayesian updating, every two information sets $h_i, h_i' \in H_i$ where h_i precedes h_i' and $b_i(h_i)(S_{-i}(h_i')) > 0$, and every two strategies $s_i, t_i \in S_i(h_i')$ that only differ at information sets weakly following h_i', and for which*

$$s_i \succsim_{i,b_i(h_i)} t_i,$$

it holds that

$$s_i \succsim_{i,b_i(h_i')} t_i.$$

Note that Bayesian updating is assumed in the definition of dynamic consistency. This is in line with a well-known property in one-person decision theory, stating that dynamic consistency within a Savage-style model requires the decision maker to update his beliefs using Bayesian updating.

It may be verified that the conditional preference relation $\succsim_1$ in Figure 2 violates dynamic consistency. Indeed, consider the conditional belief vector b_1 for player 1 where

$$b_1(h_1) = (0.5) \cdot (c, g) + (0.5) \cdot d \text{ and } b_1(h') = (c, g). \tag{4.1}$$

Then, b_1 satisfies Bayesian updating and $b_1(h_1)(S_2(h')) > 0$. Moreover, the strategies (a, e) and (a, f) only differ at h'. However, according to Figure 2 we have that $(a, e) \succ_{1,b_1(h_1)} (a, f)$ and $(a, f) \succ_{1,b_1(h')} (a, e)$. Hence, dynamic consistency is violated.

4.2 Sufficient Conditions

Why is it that the conditional preference relation in Figure 2 violates dynamic consistency? We will show that it violates two intuitive principles, which we call *preservation of indifference* and *preservation of strict preference.*

In Figure 2 we see that player 1 is indifferent between (a, e) and (a, f) for the belief β_1 that attaches probability 1 to player 2's strategy d, and for a belief β_1' that attaches positive probability to the strategies d and (c, g). But then, it seems reasonable that player 1 will also be indifferent between (a, e) and (a, f) for every belief on the line segment between β_1 and β_1'. This property will be called *preservation of indifference.* However, this property is violated as player 1 prefers (a, e) to (a, f) for all beliefs on the line segment strictly between β_1 and β_1'.

From Figure 2 we also conclude that player 1 prefers (a, f) to (a, e) for the belief β_1'' that assigns probability 1 to the strategy (c, g). As player 1 is indifferent between (a, e) and (a, f) at the belief β_1 above, it seems reasonable that player 1 will prefer (a, f) to (a, e) for all beliefs on the line segment strictly between β_1 and β_1''. This property is called *preservation of strict preference.* Also this property is violated, as player 1 prefers (a, e) to (a, f) for the belief $(0.5) \cdot (c, g) + (0.5) \cdot d$ which is on the line segment strictly between β_1 and β_1''.

To formally define these two properties, we need some further terminology: Take two beliefs $\beta_i, \beta_i' \in \Delta(S_{-i})$ and a number $\lambda \in [0, 1]$. Then, $(1-\lambda)\beta_i + \lambda\beta_i'$ is the belief that assigns to every opponents' strategy combination $s_{-i} \in S_{-i}$ the probability

$$(1-\lambda) \cdot \beta_i(s_{-i}) + \lambda \cdot \beta_i'(s_{-i}).$$

Geometrically, $(1-\lambda)\beta_i + \lambda\beta_i'$ is a belief on the line segment between β_i and β_i'. The following two definitions are adapted from Gilboa and Schmeidler (2003) and Perea (2023).

Definition 4.2 (Preservation of indifference and strict preference) *Consider a conditional preference relation $\succsim_i$. Then,*

(a) $\succsim_i$ satisfies ***preservation of indifference*** *if for every two strategies $s_i, t_i \in S_i$, and every two beliefs $\beta_i, \beta_i' \in \Delta(S_{-i})$ with $s_i \sim_{i,\beta_i} t_i$ and $s_i \sim_{i,\beta_i'} t_i$, it holds that $s_i \sim_{i,(1-\lambda)\beta_i+\lambda\beta_i'} t_i$ for every $\lambda \in (0, 1)$, and*

(b) $\succsim_i$ satisfies ***preservation of strict preference*** *if for every two strategies $s_i, t_i \in S_i$, and every two beliefs $\beta_i, \beta_i' \in \Delta(S_{-i})$ with $s_i \succsim_{i,\beta_i} t_i$ and $s_i \succ_{i,\beta_i'} t_i$, it holds that $s_i \succ_{i,(1-\lambda)\beta_i+\lambda\beta_i'} t_i$ for every $\lambda \in (0, 1)$.*

A last property we need in order to guarantee dynamic consistency is called *respect of outcome-equivalent strategies.* The idea is that if a player believes that two strategies lead to the same outcome, then he should be indifferent between the two strategies. To formally define it, we need an additional definition: For an opponents' strategy combination s_{-i}, we denote by $[s_{-i}]$ the belief that assigns probability 1 to s_{-i}.

Definition 4.3 (Respect of outcome-equivalent strategies) *A conditional preference relation $\succsim_i$* ***respects outcome-equivalent strategies*** *if for every two strategies s_i, t_i and every opponents' strategy combination s_{-i} where (s_i, s_{-i}) leads to the same consequence as (t_i, s_{-i}), it holds that $s_i \sim_{i,[s_{-i}]} t_i$.*

This property represents a weak version of *consequentialism* – a condition in philosophy and decision theory which states that an act should only be evaluated on the basis of its induced consequences and nothing else. See, for instance, the overviews by Sinnott-Armstrong (2023) and Machina (1989, Section 4), and the references therein. For a formulation and discussion of consequentialism in the framework of conditional preference relations for dynamic game forms, as we use it here, the reader may consult Perea (2024).

It may be verified that the conditional preference relation in Figure 2 respects outcome-equivalent strategies. To see this, consider the strategies (a, e) and (a, f), and the opponent's strategy d. Then, $((a, e), d)$ and $((a, f), d)$ lead to the same consequence. At the same time, player 1 is indifferent between (a, e) and (a, f) at the belief $[d]$.

We will now show that the three basic properties above are sufficient to guarantee dynamic consistency.

Theorem 4.1 (Sufficient conditions for dynamic consistency) *Every conditional preference relation that satisfies preservation of indifference, preservation of strict preference and respect of outcome-equivalent strategies is dynamically consistent.*

Note that the three conditions above are relatively basic and mild. As such, this result shows that a collection of weak conditions is enough to guarantee dynamic consistency.

4.3 Expected Utility

In Theorem 4.1 we do not require the conditional preference relation to be induced by a utility function on consequences, as is typically assumed in dynamic games. As an illustration, consider the conditional preference relation $\succsim_1$ for player 1 in Figure 3 for the dynamic game form in Figure 1. It may be verified that this conditional preference relation $\succsim_1$ satisfies preservation of indifference and preservation of strict preference, and that it respects outcome-equivalent strategies. Hence, we conclude in view of Theorem 4.1 that $\succsim_1$ is dynamically consistent.

At the same time, it can be shown that $\succsim_1$ does not have an expected utility representation. Formally, we say that a conditional preference relation $\succsim_i$ has an *expected utility representation* if there is a utility function $u_i : S_i \times S_{-i} \to \mathbf{R}$ such that $s_i \succsim_{i,\beta_i} t_i$ if and only if

$$\sum_{s_{-i} \in S_{-i}} \beta_i(s_{-i}) \cdot u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \beta_i(s_{-i}) \cdot u_i(t_i, s_{-i})$$

for all strategies s_i, t_i and every belief β_i.

To see why $\succsim_1$ in Figure 3 does not have an expected utility representation suppose, on the contrary, that there would be an expected utility representation u_1. Consider the vector v in Figure 3 which is outside the belief simplex. Since the vector v is on the line through the beliefs that yield the same expected utility for the strategies (a, e) and b, we conclude that at the vector v the "expected utility" of (a, e) and b would also be the same. Here, by the "expected utility" of the strategy (a, e) at the vector v we mean

$$\sum_{s_2 \in S_2} v(s_2) \cdot u_1((a, e), s_2),$$

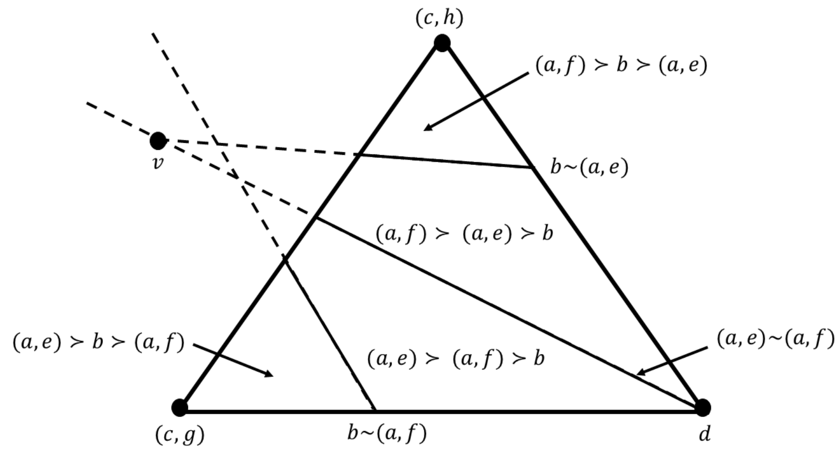

Figure 3: Conditional preference relation that is dynamically consistent, but does not have expected utility representation

where $v(s_2)$ may take negative values. Similarly for the "expected utility" of the strategy b at the vector v.

The vector v is also on the line through the beliefs that yield the same expected utility for the strategies (a, e) and (a, f), which implies that the "expected utility" of (a, e) and (a, f) will also be the same at v. We thus see that at the vector v, the "expected utilities" of $(a, e), (a, f)$ and b are all the same. However, it can be seen from Figure 3 that v is not on the line of beliefs where the expected utility of (a, f) and b are the same, which implies that the "expected utility" of (a, f) and b will not be the same at v. We thus obtain a contradiction. Hence, we conclude that there is no expected utility representation for $\succsim_1$.

It may be verified that the conditional preference relation $\succsim_1$ violates the axiom of *three choice linear preference intensity* in Perea (2023), which is necessary for an expected utility representation. Geometrically, this axiom states the following: Consider three strategies, and for each of the three pairs of strategies consider the corresponding indifference set – the set of beliefs where the player is indifferent between the two strategies involved. If we extend these three indifference sets linearly outside the belief simplex, then *three choice linear preference intensity* requires that these three sets have a common intersection, possibly outside the belief simplex. In Figure 3, these linear extensions are depicted by the dashed lines. Admittedly, *three choice linear preference intensity* is a rather demanding property, but it is needed for an expected utility representation. At the same time, Theorem 4.1 shows that this property is not required for establishing dynamic consistency.

If, on the other hand, we assume that the conditional preference relation $\succsim_i$ *does* have an

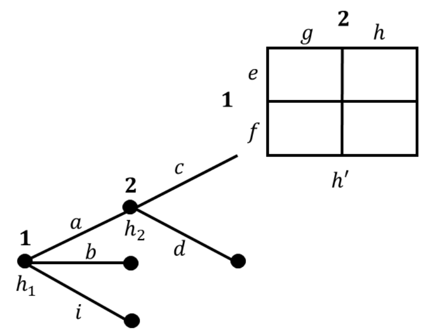

Figure 4: Dynamically consistent conditional preference relation with non-consequentialist expected utility representation

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

Table 1: Dynamically consistent conditional preference relation with non-consequentialist expected utility representation

expected utility representation $u_i : S_i \times S_{-i} \to \mathbf{R}$, then it follows from Gilboa and Schmeidler (2003) and Perea (2023) that $\succsim_i$ satisfies preservation of indifference and preservation of strict preference. If we require, in addition, that $\succsim_i$ respects outcome-equivalent strategies, then we conclude on the basis of Theorem 4.1 that $\succsim_i$ is dynamically consistent. We thus conclude, as a special case, that every conditional preference relation $\succsim_i$ that respects outcome-equivalent strategies and has an expected utility representation u_i, is dynamically consistent.

But even in this case, the conditional preference relation $\succsim_i$ need not be consequentialist. In particular, $\succsim_i$ need not be induced by a utility function on consequences. As an illustration, consider the dynamic game form in Figure 4. The only difference with Figure 1 is that player 1 now has three choices at the beginning. Consider the conditional preference relation $\succsim_1$ for player 1 with the expected utility representation u_1 given by Table 1.

Note that $\succsim_1$ respects outcome-equivalent strategies, because $((a,e),d)$ and $((a,f),d)$ lead to the same consequence whereas, at the same time, $(a,e) \sim_{1,[d]} (a,f)$. As such, Theorem 4.1 guarantees that $\succsim_1$ is dynamically consistent.

However, $\succsim_1$ is non-consequentialist. Indeed, since the consequences induced by strategies

b and i do not depend on player 2's choice, consequentialism implies that player 1's ranking of his strategies b and i should be independent of player 2's choice. But this is not the case, since player 1 prefers b to i if he believes player 2 to choose (c, g), whereas he prefers i to b if he believes player 2 to choose (c, h).

If we assume, on the other hand, that the conditional preference relation $\succsim_i$ has a *consequentialist* expected utility representation u_i which only depends on consequences, as is the case for "traditional" dynamic games, then $\succsim_i$ will always satisfy dynamic consistency. The reason is that in such a case, the conditional preference relation will automatically satisfy *respect of outcome-equivalent strategies.* Hence, our Theorem 4.1 implies that dynamic consistency will always hold for traditional dynamic games where the players' conditional preference relations are given by utility functions at the consequences.

5 Sequentially Optimal Strategies

In this section we show that every conditional preference relation which is dynamically consistent and transitive allows for a strategy that is optimal at every information set, provided the player updates his beliefs by Bayesian updating. Such strategies are called *sequentially optimal.* We start by defining sequentially optimal strategies, after which we state and prove the abovementioned result.

5.1 Definition

We start by defining what it means for a strategy to be optimal at a given information set for a specific conditional belief. Recall that, for a given information set $h_i \in H_i$, we denote by $S_i(h_i)$ the set of strategies that can possibly reach h_i.

Definition 5.1 (Optimal strategy) *Consider a conditional preference relation $\succsim_i$, an information set $h_i \in H_i$, a strategy $s_i \in S_i(h_i)$ that can possibly reach h_i, and a conditional belief $\beta_i \in \Delta(S_{-i}(h_i))$ for player i at h. Then, the strategy s_i is **optimal** for the conditional preference relation $\succsim_i$ at h_i under the conditional belief β_i if*

$$s_i \succsim_{i,\beta_i} s_i' \text{ for every strategy } s_i' \in S_i(h).$$

Next, consider a conditional belief vector b_i that assigns to every information set $h_i \in H_i$ a conditional belief $b_i(h_i) \in \Delta(S_{-i}(h_i))$. Then, a strategy is called *sequentially optimal* for this conditional belief vector if it is optimal at *every* information set that can possibly by reached under the strategy.

Definition 5.2 (Sequentially optimal strategy) *Consider a conditional preference relation $\succsim_i$ and a conditional belief vector b_i. Then, a strategy s_i is **sequentially optimal** for the conditional preference relation $\succsim_i$ under the conditional belief vector b_i if for every information set $h_i \in H_i$ with $s_i \in S_i(h_i)$, the strategy s_i is optimal for $\succsim_i$ under the belief $b_i(h_i)$.*

In general, a sequentially optimal strategy need not exist for a given conditional belief vector that satisfies Bayesian updating. As an illustration, consider the dynamic game form from Figure 1 and the associated conditional preference relation $\succsim_1$ in Figure 2. Consider the conditional belief vector b_1 given by (4.1) which satisfies Bayesian updating. Then, according to Figure 2, only the strategy (a, e) is optimal at h_1, whereas only the strategy (a, f) is optimal at h' under the conditional belief vector b_1. Hence, there is no strategy that is sequentially optimal under the conditional belief vector b_1.

5.2 Sufficient Conditions

For a meaningful analysis it seems necessary that, for every conditional belief vector that satisfies Bayesian updating, there will always be a strategy that is sequentially optimal. As we have seen, this property fails for the conditional preference relation in Figure 2. The question now is: What conditions need to be imposed such that they guarantee the existence of sequentially optimal strategies for all conditional belief vectors satisfying Bayesian updating? The answer is quite simple: The conditions that imply dynamic consistency, together with *transitivity,* are sufficient here.

Definition 5.3 (Transitivity) *A conditional preference relation $\succsim_i$ is* ***transitive*** *if the preference relation $\succsim_{i,\beta_i}$ over strategies is transitive for all beliefs $\beta_i \in \Delta(S_{-i})$.*

If we combine this property with the conditions in Theorem 4.1 that imply dynamic consistency, then this will guarantee the existence of sequentially optimal strategies for all conditional belief vectors satisfying Bayesian updating.

Theorem 5.1 (Existence of sequentially optimal strategies) *Consider a conditional preference relation $\succsim_i$ that satisfies preservation of indifference, preservation of strict preference, respects outcome-equivalent strategies and is transitive. Then, for every conditional belief vector b_i that satisfies Bayesian updating there is a strategy which is sequentially optimal for $\succsim_i$ under b_i.*

As an illustration, consider the conditional preference relation $\succsim_1$ in Figure 3 for the dynamic game form in Figure 1. It may be verified that $\succsim_1$ satisfies preservation of indifference and preservation of strict preference, respects outcome-equivalent strategies and is transitive. Therefore, we conclude on the basis of Theorem 5.1 that every conditional belief vector which satisfies Bayesian updating allows for a sequentially optimal strategy.

At the same time, we have seen earlier that $\succsim_1$ does not have an expected utility representation. This shows that expected utility is not a necessary requirement for the existence of sequentially optimal strategies. It is sufficient, though, when taken together with respect of outcome-equivalent strategies. To see this, take a conditional preference relation that has an expected utility representation and respects outcome-equivalent strategies. Then, it follows from

Gilboa and Schmeidler (2003) and Perea (2023) that it satisfies preservation of indifference and preservation of strict preference, and that it is transitive. Hence, it follows from Theorem 5.1 that every conditional belief vector which satisfies Bayesian updating allows for a sequentially optimal strategy. We thus obtain the following result.

Corollary 5.1 (Expected utility implies existence of sequentially optimal strategies) *Consider a conditional preference relation* $\succsim_i$ *that has an expected utility representation and respects outcome-equivalent strategies. Then, for every conditional belief vector* b_i *that satisfies Bayesian updating there is a strategy which is sequentially optimal for* $\succsim_i$ *under* b_i*.*

This result is known in the decision theoretic and game theoretic literature (see, for instance, Lemma 8.14.1 in Perea (2012)), but the interesting feature is that it follows from a more general result in Theorem 5.1. In particular, expected utility is not needed to guarantee the existence of sequentially optimal strategies – the much more basic conditions in Theorem 5.1 are sufficient to ensure it.

6 Concluding Remarks

Differences with Savage-style framework. The model of conditional preference relations by Gilboa and Schmeidler (2003) employed in this paper *assumes* that the players are Bayesian, as we define a preference relation over the player's strategies for every possible probabilistic belief that he could hold over the opponents' strategies. This is in sharp contrast with the expected utility model of Savage (1954) where the axioms *imply* Bayesianism. Indeed, Savage's axiom system allows us to derive a unique probability measure over states. Other investigations, like Machina and Schmeidler (1992) and Epstein and Le Breton (1993), show that Bayesianism may even be derived from an appropriate set of axioms for scenarios where the expected utility hypothesis is not fulfilled.

On the other hand, we do not consider a *unique* belief for a decision maker in our model, whereas the Savage-style models above typically do. We find this important, as a player in a dynamic game may change his belief upon observing new information, and will in general be inherently uncertain about the beliefs held by his opponents. For a game-theoretic analysis it is therefore important for a player to reason about several possible beliefs for his opponents, and our model is flexible enough to allow for this. Moreover, the notion of a conditional preference relation is capable of describing how the player's preference relation changes when the game moves from one information to another. It can thus be used as a dynamic decision theoretic model.

In a game-theoretic setting, one drawback of the Savage-style models is that many acts (mappings from states to consequences) do not correspond to real choice options of the player, but the model assumes that the player nevertheless ranks all of these acts. The concept of a conditional preference relation circumvents this problem by assuming that for every belief, the player only ranks his strategies in the dynamic game, and nothing else.

In view of the above, we feel that the concept of a conditional preference relation provides a natural decision-theoretic framework for analyzing the behavior of players in a dynamic game. Important is also that it allows for scenarios where the expected utility hypothesis is violated.

Does dynamic consistency imply expected utility? Dynamic consistency is, by its very nature, a notion that applies to dynamic scenarios, where the decision maker updates his preferences upon receiving new information. The general idea is that the decision maker's updated preferences should be sufficiently aligned with his ex-ante preferences. Our version of dynamic consistency has also been defined along those lines: For a given conditional preference relation, the player's updated preference relation upon reaching a new information set must be in line with his preference relation held at the previous information set.

The idea of dynamic consistency has also been explored in static Savage-style scenarios, however. Machina and Schmeidler (1992) and Epstein and Le Breton (1993), for instance, derive for a given event E and sub-act h on the complement of E, an updated preference relation $\succsim_{E,h}$ over sub-acts on E as follows:

$$f \succsim_{E,h} g \text{ if and only if} \begin{bmatrix} f(s), & \text{if } s \in E \\ h(s), & \text{if } s \notin E \end{bmatrix} \succsim \begin{bmatrix} g(s), & \text{if } s \in E \\ h(s), & \text{if } s \notin E \end{bmatrix}.$$

Then, by construction, the updated preference relation $f \succsim_{E,h} g$ is dynamically consistent with the ex-ante preference relation $\succsim$.

But suppose we would only allow for conditioning events that are *observable,* like the events E above, but without the sub-act h. In that case, the updated preference relation $\succsim_{E,h}$ should be independent of the sub-act h, which amounts to imposing the Sure-Thing Principle. Together with the other Savage axioms, this would lead us to expected utility. In that sense, this version of dynamic consistency would imply expected utility.

Note that this is not true for our analysis in this paper: Our sufficient conditions for dynamic consistency do not imply expected utility, as has been shown by the example of Figure 3. Moreover, our notion of dynamic consistency assumes that updated preferences are defined conditional on observable events only. Indeed, the updated preferences are the preferences that a player i holds at each of his information sets h_i, which in turn correspond to the observable events $S_{-i}(h_i)$ where the opponents have made choices that allow h_i to be reached.

Where does this difference come from? First, our decision-theoretic framework is fundamentally different from the Savage-style frameworks, as already discussed above. Hence, assumptions that seem similar at first sight may lead to different conclusions. Moreover, our sufficient conditions that imply dynamic consistency are very basic, and far from yielding expected utility. In turn, the notion of dynamic consistency for Savage-style models as discussed above, where updated preferences are only defined for observable conditioning events, implies the Sure-Thing Principle, which already brings us very close to expected utility.

7 Appendix

7.1 Proof of Section 4

Proof of Theorem 4.1. Consider a conditional preference relation $\succsim_i$ that satisfies preservation of indifference, preservation of strict preference and respect of outcome-equivalent strategies. We will show that $\succsim_i$ is dynamically consistent.

Consider a conditional belief vector b_i that satisfies Bayesian updating, two information sets $h_i, h_i' \in H_i$ where h_i precedes h_i' and $b_i(h_i)(S_{-i}(h_i')) > 0$, and two strategies $s_i, t_i \in S_i(h_i')$ that only differ at information sets weakly following h_i', and where $s_i \succsim_{i,b_i(h_i)} t_i$.

By definition of Bayesian updating, the conditional belief $b_i(h_i') \in \Delta(S_{-i}(h_i'))$ at h_i' is given by

$$b_i(h_i')(s_{-i}) := \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i}(h_i'))} \text{ for every } s_{-i} \in S_{-i}(h_i'). \tag{7.1}$$

We will show that $s_i \succsim_{i,b_i(h_i')} t_i$.

We distinguish two cases: (1) $b_i(h_i)(S_{-i}(h_i')) = 1$, and (2) $b_i(h_i)(S_{-i}(h_i')) < 1$.

Case 1. Suppose that $b_i(h_i)(S_{-i}(h_i')) = 1$. Then, $b_i(h_i) = b_i(h_i')$, and it trivially follows that $s_i \succsim_{i,b_i(h_i')} t_i$ since $s_i \succsim_{i,b_i(h_i)} t_i$.

Case 2. Suppose that $b_i(h_i)(S_{-i}(h_i')) < 1$. Then, $b_i(h_i)(S_{-i} \backslash S_{-i}(h_i')) > 0$. Let $\beta_i \in \Delta(S_{-i} \backslash S_{-i}(h_i'))$ be the belief given by

$$\beta_i(s_{-i}) := \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i} \backslash S_{-i}(h_i'))} \text{ for every } s_{-i} \in S_{-i} \backslash S_{-i}(h_i'). \tag{7.2}$$

Then, in view of (7.1) and (7.2), the belief $b_i(h_i)$ can be written as

$$\begin{aligned} b_i(h_i) &= b_i(h_i)(S_{-i}(h_i')) \cdot b_i(h_i') + b_i(h_i)(S_{-i} \backslash S_{-i}(h_i')) \cdot \beta_i \\ &= b_i(h_i)(S_{-i}(h_i')) \cdot b_i(h_i') + (1 - b_i(h_i)(S_{-i}(h_i'))) \cdot \beta_i. \end{aligned} \tag{7.3}$$

By construction, the belief β_i only assigns positive probability to strategy combinations outside $S_{-i}(h_i')$, and hence we have that

$$\beta_i = \sum_{s_{-i} \in S_{-i} \backslash S_{-i}(h_i')} \beta_i(s_{-i}) \cdot [s_{-i}]. \tag{7.4}$$

Now, take some $s_{-i} \in S_{-i} \backslash S_{-i}(h_i')$. Then, for every history $p \in h_i'$, the strategy combination s_{-i} does not select some of the actions that lead to p. As s_i and t_i only differ at information sets weakly following h_i', we conclude that (s_i, s_{-i}) and (t_i, s_{-i}) lead to the same consequence z which does not follow h_i'. Since $\succsim_i$ *respects outcome-equivalent strategies*, we conclude that

$$s_i \sim_{i,[s_{-i}]} t_i \text{ for every } s_{-i} \in S_{-i} \backslash S_{-i}(h_i'). \tag{7.5}$$

As $\succsim_i$ satisfies *preservation of indifference*, it follows by (7.4) and (7.5) that

$$s_i \sim_{i,\beta_i} t_i. \tag{7.6}$$

Now assume, contrary to what we want to show, that $t_i \succ_{i,b_i(h_i')} s_i$. Since $\succsim_i$ satisfies *preservation of strict preference*, and $b_i(h_i)(S_{-i}(h')) > 0$, we conclude on the basis of (7.3) and (7.6) that $t_i \succ_{i,b_i(h_i)} s_i$. This, however, is a contradiction to our assumption that $s_i \succsim_{i,b_i(h_i)} t_i$. Hence, $t_i \succ_{i,b_i(h_i')} s_i$ cannot be true, which implies that $s_i \succsim_{i,b_i(h_i')} t_i$. Thus, $\succsim_i$ is dynamically consistent. This completes the proof. ■

7.2 Proof of Section 5

Proof of Theorem 5.1. Consider a conditional belief vector b_i that satisfies Bayesian updating. Let

$$H_i^1 = \{h_i \in H_i \mid h_i \text{ is not preceded by any } h_i' \in H_i\}$$

be the collection of first information sets for player i. Consider a first information set $h_i \in H_i^1$, and the induced preference relation $\succsim_{i,b_i(h_i)}$ over strategies there. Since $\succsim_i$ is transitive, we know that the preference relation $\succsim_{i,b_i(h_i)}$ is transitive, and hence there is a strategy $s_i^{1h_i}$ that is optimal at h_i under the belief $b_i(h_i)$.

Now, construct a strategy s_i^1 such that, for every $h_i \in H_i^1$, the strategy s_i^1 coincides with $s_i^{1h_i}$ at all information sets $h_i' \in H_i$ that weakly follow h_i and where $s_i^1 \in S_i(h_i')$. Such a construction is possible since, by perfect recall, we have for every two different information sets $h_i, h_i' \in H_i^1$ that every information set that weakly follows h_i cannot weakly follow h_i'. We will now show, for every $h_i \in H_i^1$, that the strategy s_i^1 is optimal at h_i under the belief $b_i(h_i)$.

Consider an information set $h_i \in H_i^1$, and compare the strategies s_i^1 and $s_i^{1h_i}$ under the belief $b_i(h_i)$. By definition, $b_i(h_i) \in \Delta(S_{-i}(h_i))$, and hence $b_i(h_i)$ can be written as

$$b_i(h_i) = \sum_{s_{-i} \in S_{-i}(h_i)} b_i(h_i)(s_{-i}) \cdot [s_{-i}]. \tag{7.7}$$

Take some $s_{-i} \in S_{-i}(h)$. Then, there is some history $p \in h_i$ such that s_{-i} selects all the actions that lead to p. Since s_i^1 and $s_i^{1h_i}$ coincide at all information sets for player i weakly following h_i, and since there are no choices for player i before h_i, we conclude that (s_i^1, s_{-i}) and $(s_i^{1h_i}, s_{-i})$ lead to the same consequence following p. As $\succsim_i$ *respects outcome-equivalent strategies*, we conclude that

$$s_i^1 \sim_{i,[s_{-i}]} s_i^{1h_i} \text{ for all } s_{-i} \in S_{-i}(h_i). \tag{7.8}$$

Since $\succsim_i$ satisfies *preservation of indifference*, it follows from (7.7) and (7.8) that

$$s_i^1 \sim_{i,b_i(h_i)} s_i^{1h_i}. \tag{7.9}$$

Recall that $s_i^{1h_i}$ is optimal at h_i under the belief $b_i(h_i)$, which means that

$$s_i^{1h_i} \succsim_{i,b_i(h_i)} s_i \text{ for all } s_i \in S_i(h_i).$$

Since the preference relation $\succsim_{i,b_i(h_i)}$ is *transitive*, we conclude on the basis of (7.9) that

$$s_i^1 \succsim_{i,b_i(h_i)} s_i \text{ for all } s_i \in S_i(h_i), \tag{7.10}$$

and hence s_i^1 is optimal at h_i under the belief $b_i(h_i)$.

For every $h_i \in H_i^1$, let

$$H_i^+(h_i) := \{h_i' \in H_i \mid h_i' \text{ follows } h_i \text{ and } b_i(h_i)(S_{-i}(h_i')) > 0\}$$

be the collection of information sets for player i that follow h_i and which, according to the belief at h_i, can possibly be reached with some positive probability. By perfect recall, all of these sets $H_i^+(h_i)$ are disjoint. For every information set $h_i' \in H_i^+(h_i)$, the conditional belief $b_i(h_i') \in \Delta(S_{-i}(h_i'))$ is, by the definition of Bayesian updating, given by

$$b_i(h_i')(s_{-i}) := \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i}(h_i'))} \text{ for every } s_{-i} \in S_{-i}(h_i'). \tag{7.11}$$

Now, take some $h_i' \in H_i^+(h_i)$ such that $s_i^1 \in S_i(h_i')$. We show that s_i^1 is optimal at h_i' for $b_i(h_i')$.

Take some arbitrary strategy $s_i \in S_i(h_i') \backslash \{s_i^1\}$. Then, in particular, $s_i \in S_i(h_i)$ since h_i precedes h_i'. Hence, we know by (7.10) that

$$s_i^1 \succsim_{i,b_i(h_i)} s_i. \tag{7.12}$$

We will show that $s_i^1 \succsim_{i,b_i(h_i')} s_i$.

Let $\tilde{s}_i$ be the strategy that coincides with s_i at all player i information sets that weakly precede or weakly follow h_i', and that coincides with s_i^1 at all other player i information sets h_i'' with $s_i^1 \in S_i(h_i'')$. Note that the belief $b_i(h_i') \in \Delta(S_{-i}(h'))$ can be written as

$$b_i(h_i') = \sum_{s_{-i} \in S_{-i}(h_i')} b_i(h_i')(s_{-i}) \cdot [s_{-i}]. \tag{7.13}$$

Take some $s_{-i} \in S_{-i}(h_i')$. Then, there is a history p in h_i such that s_{-i} selects all the actions that lead to p. Moreover, as s_i and $\tilde{s}_i$ coincide at all player i information sets preceding h_i' it follows by perfect recall that s_i and $\tilde{s}_i$ select all the player i actions leading to p. As s_i and $\tilde{s}_i$ also coincide at all player i information sets weakly following h_i' we conclude that (s_i, s_{-i}) and $(\tilde{s}_i, s_{-i})$ lead to the same consequence. By *respect of outcome equivalent strategies* we then obtain that

$$\tilde{s}_i \sim_{i,[s_{-i}]} s_i. \tag{7.14}$$

As $\succsim_i$ satisfies *preservation of indifference* we conclude from (7.13) and (7.14) that

$$\tilde{s}_i \sim_{i,b_i(h_i')} s_i. \tag{7.15}$$

Remember that $s_i^1, \tilde{s}_i \in S_i(h_i')$. Hence, by perfect recall, s_i^1 and $\tilde{s}_i$ coincide at all player i information sets preceding h_i', which implies that $s_i^1, \tilde{s}_i$ only differ at player i information sets weakly following h_i'. Since $b_i(h_i)(S_{-i}(h_i')) > 0$ and the belief $b_i(h_i')$ is given by (7.11), it follows by Theorem 4.1 that

$$s_i^1 \succsim_{i,b_i(h_i')} \tilde{s}_i. \tag{7.16}$$

As $\succsim_i$ is *transitive* it follows from (7.15) and (7.16) that $s_i^1 \succsim_{i,b_i(h_i')} s_i$. Since this holds for every $s_i \in S_i(h_i')\backslash\{s_i^1\}$, we conclude that the strategy s_i^1 is optimal at h_i' under the belief $b_i(h_i')$.

Let

$$H_i^{1+} := \{h_i' \in H_i \mid h_i' \in H_i^+(h_i) \text{ for some } h_i \in H_i^1\}$$

be the collection of information sets for player i which, according to the beliefs at H_i^1, can possibly be reached with positive probability. On the basis of our insights above, we conclude that the strategy s_i^1 so constructed is optimal at every information set h_i in $H_i^1 \cup H_i^{1+}$ with $s_i^1 \in S_i(h_i)$ under the associated belief $b_i(h_i)$.

Next, define

$$H_i^2 := \{h_i \in H_i\backslash(H_i^1 \cup H_i^{1+}) \mid h_i \text{ not preceded by any } h_i' \in H_i\backslash(H_i^1 \cup H_i^{1+})\}$$

as the collection of first information sets for player i that are not in $H_i^1 \cup H_i^{1+}$. By a similar argument as above, we know that for every $h_i \in H_i^2$ there is a strategy $s_i^{2h_i}$ that is optimal at h_i under the belief $b_i(h_i)$.

Now, construct a strategy s_i^2 that coincides with s_i^1 at all information sets in $H_i^1 \cup H_i^{1+}$, and that, for every $h_i \in H_i^2$, coincides with $s_i^{2h_i}$ at all information sets for player i that weakly follow h_i. In a similar way as above, it can then be shown that for every $h_i \in H_i^2$ with $s_i^2 \in S_i(h_i)$, the strategy s_i^2 is optimal at h_i for the belief $b_i(h_i)$.

We will now show that, for every $h_i' \in H_i^1 \cup H_i^{1+}$ with $s_i^2 \in S_i(h_i')$, the strategy s_i^2 is optimal at h_i' under the belief $b_i(h_i')$. Take some $h_i' \in H_i^1 \cup H_i^{1+}$ with $s_i^2 \in S_i(h_i')$. Then, there is some $h_i \in H_i^1$ such that h_i' weakly follows h_i, and $b_i(h_i')$ is given by (7.11).

By construction, under the belief $b_i(h_i)$ only information sets in $H_i^1 \cup H_i^{1+}$ can possibly be reached with positive probability if player i chooses a strategy in $S_i(h_i)$. But then, it follows by (7.11) that under the belief $b_i(h_i')$, only information sets in $H_i^1 \cup H_i^{1+}$ can possibly be reached with positive probability if player i chooses a strategy in $S_i(h_i')$.

Now, consider an opponents' strategy combination $s_{-i} \in S_{-i}(h_i')$ such that $b_i(h_i')(s_{-i}) > 0$. Since $s_i^2 \in S_i(h_i')$ and s_i^2 coincides with s_i^1 at all information sets in $H_i^1 \cup H_i^{1+}$, we have that $s_i^1 \in S_i(h_i')$ also. But then, we know by the insights above that both (s_i^1, s_{-i}) and (s_i^2, s_{-i}) only reach information sets in $H_i^1 \cup H_i^{1+}$. As s_i^1 and s_i^2 coincide at all information sets in $H_i^1 \cup H_i^{1+}$, we conclude that (s_i^1, s_{-i}) and (s_i^2, s_{-i}) lead to the same consequence. Since $\succsim_i$ *respects outcome-equivalent strategies* we conclude that

$$s_i^2 \sim_{i,[s_{-i}]} s_i^1 \text{ for all } s_{-i} \text{ with } b_i(h_i')(s_{-i}) > 0. \tag{7.17}$$

At the same time, the belief $b_i(h_i')$ can be written as

$$b_i(h_i') = \sum_{s_{-i} \in S_{-i}(h_i'): b_i(h_i')(s_{-i})>0} b_i(h_i')(s_{-i}) \cdot [s_{-i}]. \tag{7.18}$$

Since $\succsim_i$ satisfies *preservation of indifference*, we conclude on the basis of (7.17) and (7.18) that

$$s_i^2 \sim_{i,b_i(h_i')} s_i^1. \tag{7.19}$$

Recall that s_i^1 was optimal at h_i' under the belief $b_i(h_i')$, which means that

$$s_i^1 \succsim_{i,b_i(h_i')} s_i \text{ for all } s_i \in S_i(h_i'). \tag{7.20}$$

As $\succsim_{i,b_i(h_i')}$ is *transitive,* (7.19) and (7.20) imply that

$$s_i^2 \succsim_{i,b_i(h_i')} s_i \text{ for all } s_i \in S_i(h_i')$$

and hence s_i^2 is optimal at h_i' under the belief $b_i(h_i')$.

We thus conclude that, for every $h_i \in H_i^1 \cup H_i^{1+} \cup H_i^2$ with $s_i^2 \in S_i(h_i)$, the strategy s_i^2 is optimal at h_i under the belief $b_i(h_i)$.

For every $h_i \in H_i^2$, define the collection of information sets

$$H_i^+(h_i) := \{h_i' \in H_i \mid h_i' \text{ follows } h_i \text{ and } b_i(h_i)(S_{-i}(h_i')) > 0\}.$$

In the same way as above, it can then be shown that for every $h_i' \in H_i^+(h_i)$ with $s_i^2 \in S_i(h_i')$, the strategy s_i^2 is optimal at h_i' under the belief $b_i(h_i')$.

Let

$$H_i^{2+} := \{h_i' \in H_i \mid h_i' \in H_i^+(h_i) \text{ for some } h_i \in H_i^2\}.$$

Then, we conclude that for every $h_i \in H_i^{2+}$ with $s_i^2 \in S_i(h_i)$, the strategy s_i^2 is optimal at h_i under the belief $b_i(h_i)$.

Altogether, we see that for every $h_i \in H_i^1 \cup H_i^{1+} \cup H_i^2 \cup H_i^{2+}$ with $s_i^2 \in S_i(h_i)$, the strategy s_i^2 is optimal at h_i under the belief $b_i(h_i)$.

We can continue in this fashion until, for some K, every information set for player i is in

$$(H_i^1 \cup H_i^{1+}) \cup (H_i^2 \cup H_i^{2+}) \cup ... \cup (H_i^K \cup H_i^{K+}).$$

Then, the strategy s_i^K so constructed will have the property that, for every $h_i \in H_i$ with $s_i^K \in S_i(h_i)$, the strategy s_i^K is optimal at h_i under the belief $b_i(h_i)$. That is, s_i^K is sequentially optimal for the conditional belief vector b_i. This completes the proof. ■

References

[1] Anscombe, F.J. and R.J. Aumann (1963), A definition of subjective probability, *Annals of Mathematical Statistics* **34,** 199–205.

[2] Asheim, G.B. and A. Perea (2005), Sequential and quasi-perfect rationalizability in extensive games, *Games and Economic Behavior* **53**, 15–42.

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

[4] Dekel, E., Fudenberg, D. and D.K. Levine (1999), Payoff information and self-confirming equilibrium, *Journal of Economic Theory* **89**, 165–185.

[5] Dekel, E., Fudenberg, D. and D.K. Levine (2002), Subjective uncertainty over behavior strategies: A correction, *Journal of Economic Theory* **104,** 473–478.

[6] Epstein, L.G. and M. Le Breton (1993), Dynamically consistent beliefs must be Bayesian, *Journal of Economic Theory* **61,** 1–22.

[7] Gilboa, I. and D. Schmeidler (2003), A derivation of expected utility maximization in the context of a game, *Games and Economic Behavior* **44,** 184–194.

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

[9] Kuhn, H.W. (1953), Extensive games and the problem of information, in H.W. Kuhn and A.W. Tucker (eds.), *Contributions to the Theory of Games,* Volume II (Princeton University Press, Princeton, NJ), pp. 193–216 (*Annals of Mathematics Studies* **28**)

[10] Machina, M.J. (1989), Dynamic consistency and non-expected utility models of choice under uncertainty, *Journal of Economic Literature* **XXVII,** 1622–1668.

[11] Machina, M.J. and D. Schmeidler (1992), A more robust definition of subjective probability, *Econometrica* **60,** 745–780.

[12] Osborne, M.J. and A. Rubinstein (1994), *A Course in Game Theory,* The MIT Press, Cambridge, Massachusetts, London, England.

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

[14] Penta, A. (2015), Robust dynamic implementation, *Journal of Economic Theory* **160,** 280–316.

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

[16] Perea, A. (2014), Belief in the opponents' future rationality, *Games and Economic Behavior* **83,** 231–254.

[17] Perea, A. (2023), Expected utility as an expression of linear preference intensity, Manuscript downloadable at: `www.epicenter.name/Perea/Papers/EU-linear-pref-int.pdf`

[18] Perea, A. (2024), Consequentialism in dynamic games, Epicenter Working Paper No. 30.

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

[20] Savage, L.J. (1954), *The Foundation of Statistics,* Wiley, New York.

[21] Sinnott-Armstrong, W. (2023), Consequentialism, *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), `https://plato.stanford.edu/archives/win2023/entries/consequentialism/`