J. Econ. Theory 229 (2025) 106067

Contents lists available at ScienceDirect

Journal of Economic Theory

journal homepage: www.elsevier.com/locate/jet

Dynamic consistency in games without expected utility ☆

Andrés Perea

Maastricht University, Department of Quantitative Economics and EpiCenter, the Netherlands

ARTICLE INFO

JEL classification:
C72
D81

Keywords:
Dynamic games
Dynamic consistency
Sequentially optimal strategy
Conditional preference relation
Non-expected utility

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. To explore these properties we assume, following Gilboa and Schmeidler (2003) and Perea (2025a), 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. Moreover, it is shown that non-expected utility is compatible with dynamic consistency and consequentialism in our framework.

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 et al. (1999, 2002) and Asheim and Perea (2005)), backwards rationalizability (Perea (2014), Penta (2015) and Catonini and Penta (2025)) 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'.

At the same time, the game-theoretic literature also offers plausible concepts that do *not* insist on dynamic consistency. An example is Battigalli et al. (2019) which investigates players with ambiguity averse smooth ambiguity preferences (Klibanoff et al. (2005)) who repeatedly play a dynamic game. Their preferences are dynamically inconsistent, yet the proposed concept of self-confirming equilibrium assumes that the players are "sophisticated" – they correctly anticipate their future preferences and actions, and choose

☆ I would like to thank the editor, associate editor, and three anonymous referees for their valuable comments and suggestions, which led to a substantial improvement of the paper. I am also grateful to some anonymous referees at the TARK 2023 conference for their useful feedback on an earlier version of this paper.

E-mail address: a.perea@maastrichtuniversity.nl.

https://doi.org/10.1016/j.jet.2025.106067
Received 30 August 2024; Received in revised form 12 August 2025; Accepted 13 August 2025

Available online 21 August 2025

their current actions by "folding back" planning. Battigalli et al. (2023) refer to this as *one-step optimality*. Battigalli et al. (2019) also discuss other game-theoretic concepts where players are dynamically inconsistent, and the interested reader is referred to their overview.

In one-person decision theory, dynamic consistency has also been explored for preferences that do not conform to expected utility (see, again, Machina (1989)). This is important because experimental evidence shows that many decision makers deviate from the assumptions of expected utility. A natural question that arises is: What about dynamic games? If the players are not necessarily assumed to be expected utility maximizers, how much should we presuppose so that we obtain dynamically consistent preferences? This is the question we wish to explore in this paper.

Towards this goal we assume that every player in the dynamic game holds a *conditional preference relation* (Gilboa and Schmeidler (2003), Perea (2025a)) – 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 probabilistically motivated, 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}. It thus assumes a weak form of *consequentialism*, as player i only cares about the outcome that is reached under these two strategies and his belief, and not about the choices he would have made at information sets he does not expect to be reached. Such counterfactual choices may play a role if the player exhibits preferences in the spirit of psychological games, as formalized in Battigalli and Dufwenberg (2009), for instance.[1]

Finally, 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 relevant, since it follows by the axiomatic treatments in Gilboa and Schmeidler (2003) and Perea (2025a) 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 (2025a), 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.

In particular, it is shown that within our framework, non-expected utility preferences are compatible with dynamic consistency and (some form of) consequentialism. This seemingly contradicts a "folk theorem" in decision theory, which states that the combination of dynamic consistency and consequentialism leads to expected utility. However, it is argued towards the end of the paper that our version of dynamic consistency is substantially weaker than the typical notion in the decision-theoretic literature, since we do not impose dynamic consistency for every possible conditioning event, but only for those that correspond to the information sets in the dynamic game at hand.

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. We also provide an economically meaningful example of a conditional preference relation that does not have an expected utility representation, yet is dynamically consistent, satisfies a form of consequentialism stronger than respect of outcome-equivalent strategies, and allows for sequentially optimal strategies. In Section 6 we discuss the implications of dynamic consistency and consequentialism in the decision-theoretic literature, and compare these to ours. We start by providing an overview of papers in decision theory which show that (some versions of) dynamic consistency and consequentialism lead to expected utility under the rules of conditional probability. We then explain why these results do not apply to our setting, followed by a discussion of papers in decision theory that make non-expected utility compatible with either dynamic consistency or consequentialism, by relaxing the other property of the two. We finally discuss how the rules of conditional probability also follow from diachronic Dutch book arguments. The appendix contains the proof of our second result.

[1] Strictly speaking, the model in Battigalli and Dufwenberg (2009) allows the player's preferences to depend on counterfactual choices by his opponents at unreached information sets, but not on counterfactual choices by himself. The model could naturally be extended to situations where the player's preferences would depend on his own counterfactual choices as well.

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, the set $I^a(p)$ is non-empty, 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), denoted by $p \lhd p'$, 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), denoted by $h \lhd h'$, 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), denoted by $h \unlhd h'$, 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 *complete strategy* $\tilde{s}_i$ for player i assigns to every information set $h_i \in H_i$ an available action $\tilde{s}_i(h) \in A_i(h)$. Let $\tilde{S}_i$ be the set of complete strategies for player i, and $\tilde{S}_{-i} := \times_{j \neq i} \tilde{S}_j$ the set of opponents' complete strategy combinations. Every combination of complete strategies $(\tilde{s}_i)_{i \in I}$ induces a consequence $z((\tilde{s}_i)_{i \in I}) \in Z$. By

$$H_i(\tilde{s}_i) := \{h_i \in H_i \mid \exists \tilde{s}_{-i} \in \tilde{S}_{-i}, p \in h_i \text{ s.t. } p \lhd z(\tilde{s}_i, \tilde{s}_{-i})\}$$

we denote the collection of player i information sets that can be reached by $\tilde{s}_i$. By $r_i(\tilde{s}_i)$ we denote the restriction of $\tilde{s}_i$ to information sets in $H_i(\tilde{s}_i)$, and it is called the *reduced strategy* induced by $\tilde{s}_i$. By $S_i := r_i(\tilde{S}_i)$ we denote the set of *reduced strategies* for player i. In the sequel, when we say *strategy*[2] we always mean a reduced strategy. Every combination of strategies $(s_i)_{i \in I}$ reaches a consequence $z((s_i)_{i \in I}) \in Z$.

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

$$S(h) := \{s \in \times_{i \in I} S_i \mid \exists p \in h \text{ s.t. } p \lhd z(s)\},$$

$$S_i(h) := \{s_i \in S_i \mid \exists s_{-i} \in S_{-i} \text{ s.t. } (s_i, s_{-i}) \in S(h)\}, \text{ and}$$

$$S_{-i}(h) := \{s_{-i} \in S_{-i} \mid \exists s_i \in S_i \text{ s.t. } (s_i, s_{-i}) \in S(h)\}.$$

Intuitively, $S_i(h)$ 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] 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".

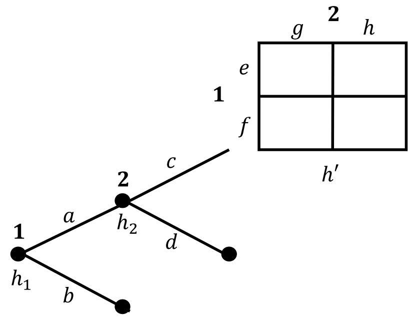

Fig. 1. A dynamic game form.

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 the rules of conditional probability when the game moves from one player i information set to another. Battigalli et al. (2023) refer to this condition as *forward consistency.*[3] Formally, the conditional belief vector b_i is *forward consistent* 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' strategies: For different beliefs, the player may opt for different strategies. To capture this phenomenon 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 (2025a)), 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 Fig. 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 Fig. 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 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. Fig. 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 in a forward consistent way to (c, g). From Fig. 2 we learn that at h' player 1 would prefer (a, f) to (a, e), and (a, e) to b.

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.

[3] In many papers this property is called *Bayesian updating.*

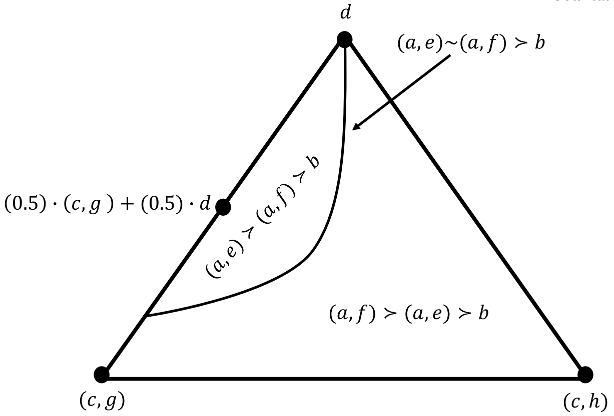

Fig. 2. A conditional preference relation for the dynamic game form in Fig. 1.

4. Dynamic consistency

In this section we first provide a definition of *dynamic consistency* 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_i', 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 forward consistency, 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 forward consistency 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 forward consistent conditional belief vector b_i, 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

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

It may be verified that the conditional preference relation $\succsim_1$ in Fig. 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 is forward consistent and $b_1(h_1)(S_2(h')) > 0$. Moreover, the strategies (a,e) and (a,f) only differ at h'. However, according to Fig. 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 Fig. 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 Fig. 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 Fig. 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 (2025a).

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 $\lambda\in(0,1)$.

These conditions are tightly related to the axiom of *betweenness* in Dekel (1986), at least from a mathematical point of view. Dekel (1986) considers a preference relation over lotteries, and states that it satisfies *betweenness* if for every two lotteries L,L' the following two properties hold: (i) if $L\sim L'$ then $L\sim(1-\lambda)L+\lambda L'\sim L'$ for all $\lambda\in(0,1)$, and (ii) if $L\succ L'$, then $L\succ(1-\lambda)L+\lambda L'\succ L'$ for all $\lambda\in(0,1)$. Property (i) thus states that the indifference sets are convex, whereas (ii) guarantees that these indifference sets cannot be "thick", unless the decision maker (DM) is indifferent between all lotteries. Similar properties are implied by conditions (a) and (b) above: Condition (a) states that the set of beliefs for which player i is indifferent between two strategies is convex, whereas condition (b) states that this set of beliefs cannot be "thick" unless player i is indifferent between these two strategies for all beliefs.

Not only is the axiom of betweenness similar to our axioms of preservation of indifference and preservation of strict preference, also the underlying decision-theoretic frameworks are related. To see this, consider for player i a strategy s_i and a probabilistic belief $\beta_i\in\Delta(S_{-i})$ in our setting. Together, they induce a probability distribution $\mathbb{P}_{(s_i,\beta_i)}$ over consequences, which may be identified with a lottery $L_{(s_i,\beta_i)}$. The conditional preference relation $\succsim_i$ thus induces a restricted binary relation over lotteries $L_{(s_i,\beta_i)}$ and $L_{(t_i,\beta_i)}$ that only differ by the choice of strategy, but not by the choice of belief. More precisely, $L_{(s_i,\beta_i)}\succsim_i L_{(t_i,\beta_i)}$ precisely when $s_i\succsim_{i,\beta_i}t_i$. This induced preference ranking of special pairs of lotteries brings us close to the decision-theoretic framework of von Neumann and Morgenstern (1947) and Dekel (1986), where the decision maker holds a binary relation over *all* possible lotteries on a given set of consequences.

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}. That is, it represents the Dirac measure on 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 is closely related to the *relevance* axiom in Myerson (1991), which is formulated for a framework where the decision maker holds for every event E (i.e. a set of states) a preference relation $\succsim_E$ over Anscombe-Aumann acts. *Relevance* states that whenever two acts f,g behave identically at every state in an event E, then $f\sim_E g$. Our condition above can be viewed as an instance of relevance, where f,g correspond to the strategies s_i,t_i, and the event E is the set $\{s_{-i}\}$.

Respect of outcome-equivalent strategies 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.

Perea (2025b) introduces the notion of *preference-based consequentialism*[4] for our framework of conditional preference relations, as follows. For every strategy s_i and belief β_i, recall from above that $\mathbb{P}_{(s_i,\beta_i)}$ is the induced probability distribution on consequences. *Preference-based consequentialism* states that for every four strategies s_i,s_i',t_i,t_i' and every two beliefs β_i,β_i' where $\mathbb{P}_{(s_i,\beta_i)}=\mathbb{P}_{(s_i',\beta_i')}$ and $\mathbb{P}_{(t_i,\beta_i)}=\mathbb{P}_{(t_i',\beta_i')}$, it should hold that $s_i\succsim_{i,\beta_i}t_i$ precisely when $s_i'\succsim_{i,\beta_i'}t_i'$. It is easily seen that preference-based consequentialism implies

[4] The name *preference-based* consequentialism is chosen as to distinguish it from *utility-based* consequentialism – a stronger notion that is also defined in Perea (2025b).

respect of outcome-equivalent strategies, by setting $s_i' = t_i' = s_i$ and $\beta_i = \beta_i' = [s_{-i}]$. The other implication, however, is not true. This confirms our interpretation of respect of outcome-equivalent strategies as a weak form of consequentialism.

It may be verified that the conditional preference relation in Fig. 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 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.*

Proof. 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 forward consistent conditional belief vector b_i, 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 forward consistency, the conditional belief $b_i(h_i') \in \Delta(S_{-i}(h_i'))$ at h_i' is given by

$$b_i(h_i')(s_{-i}) := \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i}(h_i'))} \ \forall s_{-i} \in S_{-i}(h_i'). \tag{4.2}$$

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

$$\beta_i(s_{-i}) := \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i} \backslash S_{-i}(h_i'))} \ \forall s_{-i} \in S_{-i} \backslash S_{-i}(h_i'). \tag{4.3}$$

Then, in view of (4.2) and (4.3), the belief $b_i(h_i)$ can be written as

$$\begin{aligned} b_i(h_i) &= b_i(h_i)(S_{-i}(h_i')) \cdot b_i(h_i') + b_i(h_i)(S_{-i} \backslash S_{-i}(h_i')) \cdot \beta_i \\ &= b_i(h_i)(S_{-i}(h_i')) \cdot b_i(h_i') + (1 - b_i(h_i)(S_{-i}(h_i'))) \cdot \beta_i. \end{aligned} \tag{4.4}$$

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

$$\beta_i = \sum_{s_{-i} \in S_{-i} \backslash S_{-i}(h_i')} \beta_i(s_{-i}) \cdot [s_{-i}]. \tag{4.5}$$

Here, the sum represents the convex combination of Dirac measures. 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

$$s_i \sim_{i,[s_{-i}]} t_i \ \forall s_{-i} \in S_{-i} \backslash S_{-i}(h_i'). \tag{4.6}$$

As $\succsim_i$ satisfies *preservation of indifference*, it follows by (4.5) and (4.6) that

$$s_i \sim_{i,\beta_i} t_i. \tag{4.7}$$

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 (4.4) and (4.7) 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. □

4.3. Expected utility

In Theorem 4.1 we do not require the conditional preference relation to have an expected utility representation, as is typically assumed in dynamic games. As an illustration, consider the conditional preference relation $\succsim_1$ for player 1 in Fig. 3 for the dynamic game form in Fig. 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. In fact, $\succsim_1$ does not only respect outcome-equivalent strategies, it even satisfies the stronger notion of *preference-based consequentialism* (Perea (2025b)), as the two notions are equivalent in the particular dynamic game form of Fig. 1.

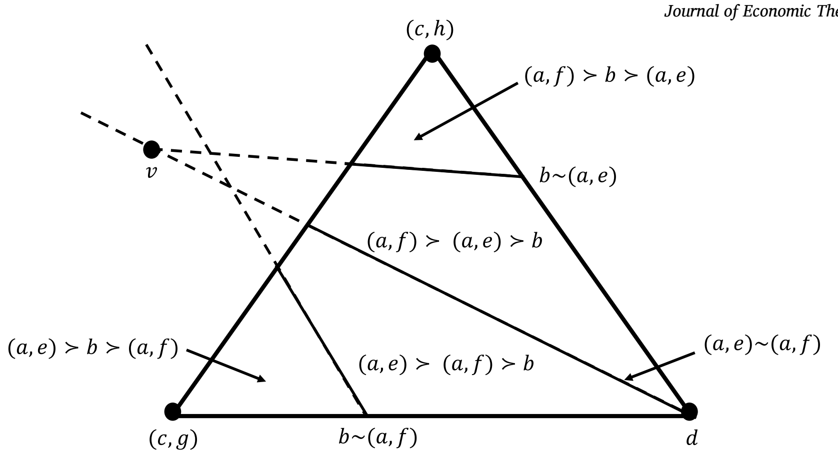

Fig. 3. Conditional preference relation that is dynamically consistent, but does not have expected utility representation.

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 Fig. 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 Fig. 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

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

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 Fig. 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 $\succsim_1$.

Fig. 3 bears some resemblance with the second triangle in Figure 1 of Gul and Lantto (1990), which depicts the indifference curves of a preference relation over lotteries that satisfies the *betweenness* axiom, but that does not have an expected utility representation. Recall that the betweenness axiom is mathematically similar to our axioms of preservation of indifference and preservation of strict preference. The indifference curves in the Gul-Lantto triangle are non-parallel lines, whereas the indifference curves of an expected utility preference relation over lotteries should all be parallel lines. An important difference is that in our framework, the lines are sets of *beliefs* for which the player is indifferent between two of his strategies, and not sets of *lotteries* between which the decision maker is indifferent. Also, a conditional preference relation having an expected utility representation in our framework would not lead to parallel lines of indifference beliefs, but rather to lines that intersect at the same point (possibly outside the belief triangle).

Also, our Fig. 2 is similar to the last triangle in Figure 1 of Gul and Lantto (1990), which represents the non-linear indifference curves of a preference relation over lotteries that violates the betweenness axiom. In the same fashion, our Fig. 2 represents a conditional preference relation that violates our axiom of preservation of indifference, leading to a non-linear set of beliefs for which the decision maker is indifferent between two of his strategies.

It may be verified that the conditional preference relation $\succsim_1$ in Fig. 3 violates the axiom of *three-choice linear preference intensity* in Perea (2025a), 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 Fig. 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 expected utility representation $u_i : S_i \times S_{-i} \to \mathbf{R}$, then it follows from Gilboa and Schmeidler (2003) and Perea (2025a) that $\succsim_i$ satisfies preservation of indifference and

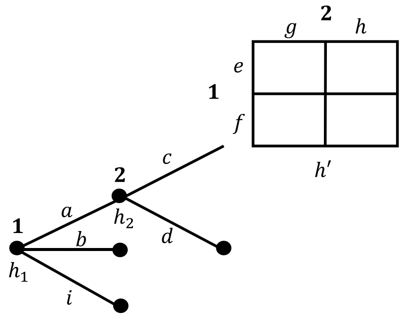

Fig. 4. Dynamically consistent conditional preference relation with non-consequentialist expected utility representation.

Table 1
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

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 *preference-based consequentialist* (Perea (2025b)), as defined in Section 4.2. As an illustration, consider the dynamic game form in Fig. 4. The only difference with Fig. 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 not preference-based consequentialist. Indeed, since the consequences induced by strategies b and i do not depend on player 2's choice, preference-based 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).

Such an expected utility representation that violates preference-based consequentialism is allowed, however, by the psychological games model of Battigalli and Dufwenberg (2009). In that model, player 1's preferences between two strategies are allowed to depend directly on player 2's strategy, no matter whether this strategy affects the consequence or not. That is, player 1 may care about hypothetical choices by player 2 at information sets he expects not be reached.

If we assume, on the other hand, that the conditional preference relation $\succsim_i$ is preference-based consequentialist and has an expected utility representation u_i, 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 preference-based consequentialist and have an expected utility representation.

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 in a forward consistent way. 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

$$s_i \succsim_{i,\beta_i} s_i' \quad \forall s_i' \in S_i(h).$$

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.[5]

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 forward consistency. As an illustration, consider the dynamic game form from Fig. 1 and the associated conditional preference relation $\succsim_1$ in Fig. 2. Consider the conditional belief vector b_1 given by (4.1) which satisfies forward consistency. Then, according to Fig. 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

As we have seen in Fig. 2, there are conditional preference relations that do not allow for sequentially optimal strategies under a forward consistent conditional belief vector. The question now is: What conditions need to be imposed on conditional preference relations such that they *do* guarantee the existence of sequentially optimal strategies for all forward consistent conditional belief vectors? 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[6] 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 forward consistent conditional belief vectors.

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 forward consistent conditional belief vector b_i there is a strategy which is sequentially optimal for $\succsim_i$ under b_i.*

The proof can be found in the appendix. As an illustration of the theorem, consider the conditional preference relation $\succsim_1$ in Fig. 3 for the dynamic game form in Fig. 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 forward consistent conditional belief vector 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 (2025a) 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 forward consistent conditional belief vector 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 forward consistent conditional belief vector b_i there is a strategy which is sequentially optimal for $\succsim_i$ under b_i.*

For consequentialist expected utility representations, this result is known in the decision-theoretic and game-theoretic literature (see, for instance, Lemma 8.14.1 in Perea (2012)). Interestingly, our corollary implies that it also holds for non-consequentialist expected utility representations, like the one in Table 1 for the game form in Fig. 4. Our Theorem 5.1, in turn, states that an expected

[5] In the literature, this type of optimality is often called *weak* sequential rationality, as to stress that it only requires optimality at those information sets that are not precluded by the strategy itself. It is the natural form of optimality if one uses strategies as *plans of action* (Rubinstein (1991)), as we do.

[6] Here, we mean the usual transitivity of a binary relation on a set.

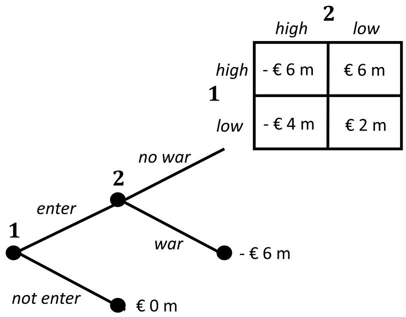

Fig. 5. Competition between an incumbent and a loss averse entrant.

Table 2
Utilities used by the entrant to compare different pairs of strategies.

u	*war*	(*no war, high*)	(*no war, low*)
(*enter, high*)	**−12**	**−12**	6
not enter	0	0	0
v	*war*	(*no war, high*)	(*no war, low*)
(*enter, low*)	**−12**	**−8**	2
not enter	0	0	0
w	*war*	(*no war, high*)	(*no war, low*)
(*enter, high*)	−6	−6	6
(*enter, low*)	−6	−4	2

utility representation is not needed at all to guarantee the existence of sequentially optimal strategies – the much more basic conditions in Theorem 5.1 are sufficient to ensure it.

5.3. Economic example

In Sections 4.3 and 5.2 we have argued that in our framework, dynamic consistency, and even the existence of sequentially optimal strategies, are compatible with non-expected utility preferences that exhibit a weak version of consequentialism, as embodied by respect of outcome-equivalent strategies. We will now illustrate this by means of an economic example.

Example: Competition between an incumbent and a loss averse entrant.

Consider a potential entrant (firm 1) that must decide whether or not to enter the market that is currently dominated by an incumbent monopolist (firm 2). Upon entering, the incumbent must decide whether or not to start a price war in order to move the entrant out. If there is no price war, both firms must independently decide whether to invest a high or low amount in improving the production technology. This situation can be modeled by the dynamic game in Fig. 5, where the numbers at the consequences are the profits (positive or negative) for the entrant, in millions of euros. Hence, high investment corresponds to a *high-risk high-gain* action, whereas low investment embodies a *low-risk low-gain* action.

Suppose that the entrant exhibits loss aversion (Kahneman and Tversky (1979)) if, and only if, it compares an uncertain act with a certain outcome. In this case, the entrant compares every consequence that may be obtained under the uncertain act with the certain outcome, and experiences every relative *loss* "twice as intensely" as the relative gains. More concretely, if the entrant compares the strategy (*enter, high*) to the strategy *not enter*, which yields the certain outcome 0, then it evaluates both strategies on the basis of the utilities u displayed in the first matrix of Table 2. In the columns we have put the three strategies of the incumbent. Note that the relative loss of 6 million, which results from the strategy combinations ((*enter, high*), *war*) and ((*enter, high*), (*no war, high*)), is translated into a utility of −12 (in bold), representing the loss aversion of the entrant.

Similarly, if the entrant compares the strategy (*enter, low*) to the strategy *not enter*, then it uses the utilities v from the second matrix in Table 2. Again, the loss aversion components of the utilities are in bold. Finally, if the entrant compares the strategies (*enter, high*) and (*enter, low*), then there is no certain outcome involved, and the utilities w that are being used are assumed to be equal to the profits, as displayed in the third matrix of Table 2.

Formally, the conditional preference relation $\succsim_1$ of the entrant is such that for every belief $\beta_1 \in \Delta(S_2)$ we have that (i) (*enter, high*) $\succsim_{1,\beta_1}$ *not enter* if and only if $Eu((\textit{enter, high}), \beta_1) \geq Eu(\textit{not enter}, \beta_1)$, (ii) (*enter, low*) $\succsim_{1,\beta_1}$ *not enter* if and only if $Ev((\textit{enter, low}), \beta_1) \geq Ev(\textit{not enter}, \beta_1)$, and (iii) (*enter, high*) $\succsim_{1,\beta_1}$ (*enter, low*) if and only if $Ew((\textit{enter, high}), \beta_1) \geq Ew((\textit{enter, low}), \beta_1)$. Here, $Eu(\cdot, \beta_1)$ represents the expected utility with respect to the utility function u and the belief β_1, and similarly for $Ev(\cdot, \beta_1)$ and $Ew(\cdot, \beta_1)$. The conditional preference relation $\succsim_1$ can be visualized by means of Fig. 6.

Here, we have abbreviated the strategies for the entrant to n, (e, h) and (e, l), respectively. Moreover, the vector (0, 2/3, 1/3) describes the belief that assigns probability 0 to *war*, probability 2/3 to (*no war, high*), and probability 1/3 to (*no war, low*), and similarly for the other beliefs in the picture.

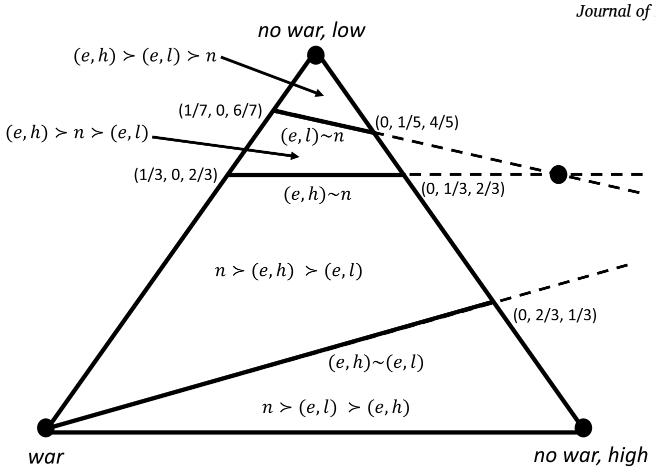

Fig. 6. Conditional preference relation for entrant.

It may be verified that the conditional preference relation $\gtrsim_1$ satisfies preservation of indifference, preservation of strict preference, respect of outcome-equivalent strategies and transitivity. To see why it satisfies respect of outcome-equivalent strategies, note that the strategies (*enter, high*) and (*enter*, *low*) lead to the same consequence under the incumbent's strategy *war*, and that (*enter, high*) $\sim_{1,[war]}$ (*enter*, *low*). In fact, $\gtrsim_1$ even satisfies the stronger requirement of preference-based consequentialism, as it is equivalent to respect of outcome-equivalent strategies in this dynamic game form. In the light of Theorems 4.1 and 5.1 we may thus conclude that $\gtrsim_1$ is dynamically consistent, and that every forward consistent conditional belief vector allows for a sequentially optimal strategy.

At the same time it may be verified that the entrant's conditional preference relation $\gtrsim_1$ does not have an expected utility representation. This is most easily seen from Fig. 6. Consider the three sets of beliefs where the entrant is indifferent between (*enter, high*) and *not enter*, between (*enter*, *low*) and *not enter*, and between (*enter, high*) and (*enter*, *low*), respectively. If we linearly extend these three sets of beliefs outside the belief simplex, then the three dashed lines so obtained do not meet at the same point. By the same argument as for Fig. 3 above, it can then be concluded that there is no expected utility representation for $\gtrsim_1$. The reason is that the conditional preference relation $\gtrsim_1$ violates the rather demanding axiom of *three-choice linear preference intensity* in Perea (2025a), which is needed for expected utility.

Note, however, that the utility functions u, v and w in Table 2 provide an expected utility representation for the comparison between (*enter, high*) and *not enter*, between (*enter*, *low*) and *not enter*, and between (*enter, high*) and (*enter*, *low*), respectively. In spite of this, there is no expected utility representation that works for the comparison of all three pairs of strategies simultaneously. The reason is that the entrant only exhibits loss aversion when it compares (*enter, high*) to *not enter*, and when it compares (*enter*, *low*) to *not enter*, but not when it compares (*enter, high*) to (*enter*, *low*). Nevertheless, the conditional preference relation for the entrant is dynamically consistent, satisfies preference-based consequentialism, and allows for sequentially optimal strategies.

6. Combining dynamic consistency and consequentialism

In the decision-theoretic literature, there is a "folk theorem" which states that combining dynamic consistency and consequentialism leads to expected utility under the rules of conditional probability. We start with an overview of some papers that prove this folk theorem, based on their own specific versions of dynamic consistency and consequentialism. At the same time, the examples in Fig. 3 and Fig. 6 show that our framework allows for dynamically consistent and preference-based consequentialist conditional preference relations which do not have an expected utility representation, thus "contradicting" the folk theorem above. In the second part we explain why the folk theorem does not apply to our setting. We subsequently discuss some papers in decision theory which reconcile a specific form of non-expected utility maximization with either dynamic consistency or consequentialism, by relaxing the other property of the two. In the last part we investigate the rules of conditional probability as implied by dynamic consistency and consequentialism in the "folk theorem", and discuss how these rules can alternatively be justified by *diachronic Dutch book* arguments.

6.1. When the combination leads to expected utility

Using the words of Siniscalchi (2011) and Ghirardato (2002), there is a "folk theorem" in decision theory which states that dynamic consistency and consequentialism, together with some more basic axioms, lead to expected utility in combination with the rules of conditional probability. The main step in proving this result is showing that dynamic consistency and consequentialism together lead to Savage's (1954) *sure thing principle* – the main axiom that brings us close to expected utility already. Here, the sure thing principle states that for every four acts f, g, h and h', and every event E, the DM prefers fEh to gEh precisely when he prefers fEh' to gEh', where fEh is the act that agrees with f on states in E, and agrees with h on states outside E. Supplemented with some more basic axioms, the sure thing principle leads to expected utility and the rules of conditional probability.

Although different papers in the decision theory literature vary in their formalization of dynamic consistency, the main message is that if the decision maker (DM) ex-ante prefers an act f to an act g, where f and g agree on states outside an event E, then upon observing the precise event E the DM should still prefer f to g. The papers also provide different definitions of consequentialism, but a shared feature of these is that the DM, conditional on observing an event E, should base his preference between two acts solely on their behavior at states within E, and not on their counterfactual behavior outside E.

A verbal, and necessarily imprecise, argument for the "folk theorem" is the following: Suppose the DM satisfies dynamic consistency and consequentialism, and prefers fEh to gEh. Then, by dynamic consistency, he will prefer f to g conditional on E. As, by consequentialism, his preferences conditional on E do not depend on the behavior of f and g outside E, it follows again by dynamic consistency that he prefers fEh' to gEh' for every other act h'. Hence, the sure thing principle follows.

The literature has offered various proofs of this "folk theorem", each with their own version of dynamic consistency and consequentialism. Myerson (1991) considers a model where the DM, for every possible event E, holds a conditional preference relation $\succsim_E$ over Anscombe-Aumann acts (Anscombe and Aumann (1963)). His *subjective substitution* axiom states that for two disjoint events E and F we must have that $f \succsim_{E\cup F} g$ whenever $f \succsim_E g$ and $f \succsim_F g$. It corresponds to the "reverse" direction of dynamic consistency as it implies, for every act h, that $fEh \succsim gEh$ whenever $f \succsim_E g$, where $\succsim$ is the "unconditional" preference relation obtained by conditioning on the full state space. To be precise, its relation with dynamic consistency can only be established if we also assume some form of consequentialism, which guarantees that $\succsim_E$ only depends on the behavior of acts at states in E.

The axiom of *state neutrality*, in turn, states that if f behaves identically on the states x and y, and g behaves identically on x and y also, then the DM prefers f to g conditional on $\{x\}$ precisely when he prefers f to g conditional on $\{y\}$. In other words, the preferences conditional on a single-state event only depend on the behavior of the acts at that particular state, and nothing else. It may thus be viewed as an instance of consequentialism. Myerson (1991) shows that these two axioms, together with some more basic axioms, allow for an expected utility representation under the rules of conditional probability.

Ghirardato (2002), like Myerson (1991), also assumes that the DM holds preferences $\succsim_E$ conditional on every event E, but this time the preferences are over Savage-acts (Savage (1954)), not Anscombe-Aumann acts. His definition of *dynamic consistency* states that for every two acts f, g and every non-null event E, we have that $f \succsim_E g$ precisely when $fEg \succsim g$. As above, $\succsim$ represents the unconditional preference relation. *Consequentialism* in this paper means that for every non-null event E, and every two acts f and g that agree on E, it must hold that $f \sim_E g$. It is then shown that these two axioms imply Savage's sure thing principle.

In Siniscalchi (2011), it is assumed that the DM's information partition concerning the true state becomes more and more fine-grained over time, and that the DM holds, at every period t and state ω, a preference relation $\succsim_{t,\omega}$ over the decision trees following (t,ω). This, in turn, implies a preference relation over plans and actions following (t,ω). By this modelling choice, *consequentialism* is built in as a structural property, as the decision trees following (t,ω) only concern actions that can indeed be taken after (t,ω), and not counterfactual actions that could have been taken at foregone subtrees. *Dynamic consistency* states that if at time $t+1$, act a is preferred to act b at every state ω' that is deemed possible at (t,ω), then a must be preferred to b at (t,ω). In that sense, it is related to Myerson's (1991) subjective substitution axiom. It is shown that these two properties imply Savage's sure thing principle (adapted to the present framework) for the ex-ante preferences.

The model in Hammond (1988) is different. There, it is assumed that the DM is characterized by a behavioral norm β which specifies, for every decision tree T and every decision node n within T, a set of optimal actions $\beta(T,n)$. The consequences in the decision tree T may involve lotteries, and the set of possible consequences is assumed to depend on the state.

The *consistency* axiom states that for the subtree T' of T that starts at n, we must have that $\beta(T',n) = \beta(T,n)$. In other words, if the DM plans an action that he believes to be optimal at a later decision node n, then he would still find this action optimal upon reaching n. It thus reveals a form of dynamic consistency. The *consequentialism* axiom states that the optimal actions in decision tree T at decision node n should only depend on the possible consequences that can be reached in T after n, and nothing more. It is shown that these two axioms, together with continuity, imply that there is a preference ordering on consequences with an *expected utility representation*, such that the DM will always select those actions that yield the highest conditional expected utility. That is, the two axioms above lead to expected utility and the rules of conditional probability.

6.2. Why it does not apply to our model

Above we have discussed several papers which show that (some versions of) dynamic consistency and consequentialism lead to expected utility. At the same time, we have seen that our conditional preference relations from Fig. 3 and Fig. 6 satisfy dynamic consistency and preference-based consequentialism, yet do not have an expected utility representation. In other words, our notions of dynamic consistency and consequentialism do *not* lead to expected utility. Where does this difference come from?

One reason for the discrepancy above is that our decision theoretic model, which builds on Gilboa and Schmeidler (2003) and Perea (2025a), is fundamentally different from the models that are most often used in the decision theory literature. Recall our notion of a conditional preference relation $\succsim_i$ for player i, which specifies for every belief β_i about the states (opponents' strategy combinations) a preference relation $\succsim_{i,\beta_i}$ over i's strategies. Compare this to the models of updated preferences in, for instance, Myerson (1991) and Ghirardato (2002), where the DM holds for every possible event E a new preference relation $\succsim_E$ over his acts. Hence, in our model the conditioning events are the possible beliefs that the DM can hold, whereas in the latter models they correspond to the possible sets of states the DM can observe. In particular, we *assume* that the DM holds a probabilistic belief over states, without pinning this belief down, whereas the other models *show*, based on a system of axioms, that the DM's preferences are governed by a *unique* probabilistic belief.

Another difference between our model and the ones above lies in the sets of acts being considered: In our model we only consider those acts that correspond to the *strategies* of player i in the particular dynamic game at hand, whereas the other models consider all possible mappings from states to (lotteries over) consequences, similarly to Savage (1954) and Anscombe and Aumann (1963). That is, in our model we consider preferences over much smaller sets compared to the other models mentioned above.

A different reason why in our framework dynamic consistency and consequentialism do not lead to expected utility is that, in spite of the difference in modelling outlined above, our version of dynamic consistency is arguably *weaker* than the versions that are typically employed in the decision-theoretic literature. Indeed, within the latter strand of literature, dynamic consistency typically requires that for *every possible event* E, and for every two acts f, g that agree outside E and where f is preferred to g ex-ante, the DM should still prefer f to g upon observing E. Our notion of dynamic consistency only requires such property for events E that correspond to *information sets* in the particular dynamic game at hand. If we identify the states in the dynamic game with the opponents' strategy combinations, then there are typically much more possible events E than information sets. As such, our notion of dynamic consistency requires significantly less than the typical notion in decision theory.

6.3. Weakening dynamic consistency or consequentialism

Through the examples in Fig. 3 and 6 we have shown that, in our framework, non-expected utility preferences are compatible with consequentialism and some version of dynamic consistency that is weaker than the versions typically employed in the decision-theoretic literature. This positions our paper within a strand of literature in decision theory where the authors attempt to reconcile models of non-expected utility with either dynamic consistency or consequentialism, by relaxing the other property of the two. We will now give an overview of some of the papers in that literature. For a more detailed discussion and overview of this matter, the reader may consult Machina (1989) and Siniscalchi (2009).

Machina (1989) considers a dynamic framework for decision making under risk where the uncertainty about the state is gradually released over time. More precisely, the DM faces a decision tree consisting of decision moves and chance moves, the latter with known probabilities. Assume that ex-ante, the DM holds a preference relation $\succsim$ over lotteries on consequences that does not conform to expected utility. Suppose that a decision node n is reached, and that L is the counterfactual lottery that would have resulted in the part of the decision tree that does not follow n, given some behavior of the DM there. Then, one can consider the conditional preference relation $\succsim_{n,L}$ where for every two lotteries L_1 and L_2 following n, we have that $L_1 \succsim_{n,L} L_2$ precisely when $(L_1, L) \succsim (L_2, L)$. Here, (L_1, L) denotes the composite lottery that agrees with L_1 after n, and agrees with L otherwise. By construction, the resulting system of conditional preferences is dynamically consistent. However, $\succsim_{n,L}$ is clearly non-consequentialist, as it requires the DM to evaluate his behavior, and that of the chance moves, at unreached parts of the decision tree.

Machina and Schmeidler (1992) and Epstein and Le Breton (1993) proceed similarly, but now within a Savage-style framework for decision making under uncertainty. They start from non-expected utility preferences $\succsim$ over Savage-acts, and extend these to a dynamically consistent system of conditional preferences, as follows. For every event E and act h, define the conditional preferences $\succsim_{E,h}$ by setting $f \succsim_{E,h} g$ precisely when $fEh \succsim gEh$. By definition, the resulting system of conditional preferences is dynamically consistent. At the same time, consequentialism is violated as the conditional preferences $\succsim_{E,h}$ require the DM to reason about the behavior of act h at states outside E. Both Machina and Schmeidler (1992) and Epstein and Le Breton focus on non-expected utility preferences that are *probabilistically sophisticated*, that is, the DM holds a probabilistic belief about the states, and evaluates an act only on the basis of the induced probability distribution on consequences with respect to this belief.

Gul and Lantto (1990) consider a model that is similar to Hammond (1988) and Siniscalchi (2011). The DM is characterized by a strategy correspondence that assigns to every decision tree with two periods a set of optimal strategies, or plans, for that tree. As in Hammond (1988), the consequences may involve lotteries. The key axioms are *consistency* and *weak consequentialism* which, when taken together, are weaker than the combination of dynamic consistency and consequentialism.

Consistency states that there is preference relation on consequences such that the DM always selects the strategy that yields the most preferred consequence. It has a consequentialist component, as it reveals that only induced consequences matter for evaluating different strategies. At the same time, it also has a dynamic consistency flavor, since it requires that decisions at different points in time are governed by the same preference relation over consequences.

Weak consequentialism states that, if the DM optimally chooses lottery p in period 1, and is faced in period 2 with a decision problem where p is again available, but now within an even smaller set of lotteries, then the DM should again optimally choose p in period 2. It may be viewed as a weak version of dynamic consistency, as it links the optimal choices in period 2 to the optimal choices in period 1.

It is shown that consistency and weak consequentialism, in combination with monotonicity and continuity, guarantee that the preference relation over consequences that governs the DM's optimal strategies satisfies the *betweenness property* (Dekel (1986)). Since the betweenness property is weaker than imposing expected utility, it follows that the axioms above are compatible with non-expected utility preferences. This is possible because consistency (in this paper) and weak consequentialism, when taken together, are weaker than the combination of consistency and consequentialism as defined in Hammond (1988). Recall from our discussion above that the latter two conditions led to expected utility.

Recall the model and the axioms of Siniscalchi (2011) as discussed in Section 6.1, where we have seen that his version of dynamic consistency, together with consequentialism, lead to expected utility ex-ante preferences. In the same paper, Siniscalchi (2011) proposes a weaker version of dynamic consistency, called *constant-act dynamic consistency*, which only poses restrictions on the DM's preferences between acts and constant acts. It is shown that this weak version of dynamic consistency, together with consequentialism, allows for conditional preferences that are in line with *max-min expected utility* (Gilboa and Schmeidler (1989).

Epstein and Schneider (2003) employ the same model as Siniscalchi (2011) and the same (strong) version of dynamic consistency discussed in Section 6.1. Although Siniscalchi (2011) has shown that this version of dynamic consistency, together with consequentialism, leads to expected utility *ex-ante* preferences, Epstein and Schneider (2003) prove that the DM's *conditional* preferences may be in line with max-min expected utility.

Klibanoff and Hanany (2007) show how a system of conditional max-min expected utility preferences can be dynamically consistent, if we allow it to violate consequentialism. Within an Anscombe-Aumann setting, consider a DM with ex-ante max-min expected utility preferences $\succsim$, relative to a set of priors C. Take an event E, and an act f that is ex-ante optimal within a set B of acts. From the set of priors C, consider those priors p for which the expected utility of f is maximal among the acts in B that coincide with f outside E. Let this set of priors be $Q^{E,f,B}$. From $Q^{E,f,B}$, choose a prior p^* such that the Bayesian update p^*_E yields the lowest expected utility for f among all priors in $Q^{E,f,B}$. Finally, from C only update those priors p that are in $Q^{E,f,B}$ and such that the expected utility of f under the Bayesian update p_E is at least as high as the expected utility of f under p^*_E. This leads to a conditional set of priors $C_{E,f,B}$ which, in turn, leads to conditional max-min expected utility preferences $\succsim_{E,f,B}$. We thus obtain a system of conditional maxmin expected utility preferences.

Their definition of *dynamic consistency* states that if f is ex-ante optimal within a set B of acts, and act g coincides with f outside an event E, then we must have that $f \succsim_{E,f,B} g$. It is shown that the system of conditional max-min expected utility preferences constructed above is dynamically consistent. At the same time, it violates consequentialism as the conditional preferences $\succsim_{E,f,B}$ require the DM to evaluate the behavior of the act f at states outside E. If, alternatively, one would update *all* priors in C one-by-one, then the induced system of conditional max-min expected utility preferences will typically be dynamically inconsistent.

In Hanany and Klibanoff (2009), the authors extend the approach above to a broader class of ambiguity averse preferences that includes max-min expected utility as a special case. When zooming in on *ambiguity averse smooth ambiguity preferences*, or KMM-preferences for short (Klibanoff, Marinacci and Mukerji (2005)), it is shown how the DM's ex-ante KMM-preferences can be extended to a dynamically consistent system of conditional KMM-preferences. The construction is similar to the approach above: Consider the ex-ante belief μ over beliefs over states upon which the DM's ex-ante KMM-preferences $\succsim$ are based. Take an event E, and an act f that is ex-ante optimal within a set B of acts. Then, transform μ into a conditional belief $\mu_{E,f,B}$ in a "dynamically consistent way", which in general will be different from taking the Bayesian update of μ. As above, also this system of conditional preferences will be non-consequentialist, since constructing the conditional belief $\mu_{E,f,B}$ requires evaluating the behavior of act f at states outside E.

Epstein and Zin (1989) investigate a DM who holds preferences over infinite horizon consumption lotteries. Every such lottery can be identified with a pair (c_0, m), where c_0 is the initial consumption, and m is a probability measure over infinite horizon consumption lotteries. They focus on preferences that have a *recursive* utility function V, which means that for every lottery (c_0, m) we have that $V(c_0, m) = W(c_0, \mu(V[m]))$, where W is an aggregator function, μ is some certainty equivalent, and $V[m]$ is the probability measure on utilities induced by m under V. Recursion may be viewed as a weak version of dynamic consistency, since it states that at every period the lotteries should be evaluated by means of the same utility function V. It is shown that there are interesting preferences which do not have an expected utility representation, yet *are* represented by a recursive utility function. As these preferences are, by construction, consequentialist, it follows that consequentialism and a weak version of dynamic consistency, as embodied by recursion, allow for non-expected utility preferences.

6.4. Diachronic Dutch book arguments

In the first part above we have discussed papers which show that (some versions of) dynamic consistency and consequentialism lead to expected utility under the rules of conditional probability. In particular, Bayes' rule of conditionalization follows from imposing dynamic consistency and consequentialism. In the literature there is an alternative, and very different, justification for the rules of conditional probability, by means of *diachronic Dutch book* arguments.

A *diachronic Dutch book* is a system of conditional bets that the DM is willing to accept, conditional on every event he can possibly observe, yet yields a sure loss from the ex-ante point of view. In Teller (1973) there is an argument by David Lewis showing that a DM who violates the rules of conditional probability will always be vulnerable to a diachronic Dutch book. Skyrms (1987) proves the converse, by showing that a DM who adheres to the rules of conditional probability will never be vulnerable to a diachronic Dutch book. In other words, the rules of conditional probability are both necessary and sufficient for being immune against diachronic Dutch books.

Note that in our definition of dynamic consistency we *impose* the rules of conditional probability by means of the built-in forward consistency requirement. This is fundamentally different from the typical dynamic consistency versions in the decision-theoretic literature, where the rules of conditional probability are not presupposed.

7. Concluding remarks

Our contribution in this paper is two-fold: Within a framework of conditional preference relations (Gilboa and Schmeidler (2003)) for dynamic games, we first present collections of weak conditions that lead to dynamic consistency, and the existence of sequentially optimal strategies, respectively. Second, we show that these conditions are weak enough such that they can be reconciled with non-expected utility. As our axiom of *respect of outcome-equivalent strategies* represents a weak instance of consequentialism, it follows that within our framework, and given our specific versions of dynamic consistency and consequentialism, the conditions of dynamic consistency and consequentialism are compatible with non-expected utility preferences. This is in sharp contrast with some of the papers in decision theory, which show that their specific versions of dynamic consistency and consequentialism necessarily lead to expected utility.

There are several reasons for this discrepancy. First, our versions of dynamic consistency and consequentialism are arguably weaker than the typical versions considered in the decision-theoretic literature. Indeed, our version of dynamic consistency only applies to events that correspond to the information sets in the specific dynamic game at hand, whereas the typical versions in the literature apply to all possible events (i.e. all possible sets of states). Also, our version of consequentialism (as embodied by respect of outcome-equivalent strategies) only requires the player to be indifferent between two strategies under a given probability 1 belief if these strategies lead to the same consequence under that belief. In turn, the typical versions of consequentialism require that any two acts should only be compared on the basis of the probability distributions over consequences they induce, which is much stronger.

Moreover, the decision-theoretic framework of conditional preference relations (Gilboa and Schmeidler (2003)) we use is very different from the typical settings that are considered in the decision-theoretic literature, which are mostly based on the models by von Neumann-Morgenstern, Savage and Anscombe-Aumann. Although we have seen in Section 4.2 that our setting bears some similarities with von Neumann-Morgenstern, some fundamental differences remain which complicate a direct and detailed comparison with most of the decision-theoretic papers on dynamic consistency and consequentialism.

Declaration of competing interest

None.

Appendix A

Proof of Theorem 5.1. Consider a forward consistent conditional belief vector b_i. Let

$$H_i^1 = \{h_i \in H_i \mid \nexists h_i' \in H_i \text{ s.t. } h_i' \lhd h_i\}$$

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 $\succsim_{i,b_i(h_i)}$ on S_i is transitive. As the set S_i of strategies is finite, and it is well-known that every transitive preference relation on a finite set allows for a maximal element in that set, it follows that there is a strategy $s_i^{1h_i}$ that is optimal in S_i for the preference relation $\succsim_{i,b_i(h_i)}$. In other words, $s_i^{1h_i}$ 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

$$b_i(h_i) = \sum_{s_{-i} \in S_{-i}(h_i)} b_i(h_i)(s_{-i}) \cdot [s_{-i}]. \tag{A.1}$$

Here, the sum represents a convex combination of Dirac measures. 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

$$s_i^1 \sim_{i,[s_{-i}]} s_i^{1h_i} \quad \forall s_{-i} \in S_{-i}(h_i). \tag{A.2}$$

Since $\succsim_i$ satisfies *preservation of indifference*, it follows from (A.1) and (A.2) that

$$s_i^1 \sim_{i,b_i(h_i)} s_i^{1h_i}. \tag{A.3}$$

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

$$s_i^{1h_i} \succsim_{i,b_i(h_i)} s_i \quad \forall s_i \in S_i(h_i).$$

Since the preference relation $\succsim_{i,b_i(h_i)}$ is *transitive*, we conclude on the basis of (A.3) that

$$s_i^1 \succsim_{i,b_i(h_i)} s_i \quad \forall s_i \in S_i(h_i), \tag{A.4}$$

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

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

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 forward consistency, given by

$$b_i(h_i')(s_{-i}) := \frac{b_i(h_i)(s_{-i})}{b_i(h_i)(S_{-i}(h_i'))} \ \forall s_{-i} \in S_{-i}(h_i'). \tag{A.5}$$

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 (A.4) that

$$s_i^1 \succsim_{i,b_i(h_i)} s_i. \tag{A.6}$$

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

$$b_i(h_i') = \sum_{s_{-i} \in S_{-i}(h_i')} b_i(h_i')(s_{-i}) \cdot [s_{-i}]. \tag{A.7}$$

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

$$\tilde{s}_i \sim_{i,[s_{-i}]} s_i. \tag{A.8}$$

As $\succsim_i$ satisfies *preservation of indifference* we conclude from (A.7) and (A.8) that

$$\tilde{s}_i \sim_{i,b_i(h_i')} s_i. \tag{A.9}$$

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 (A.5), it follows by Theorem 4.1 that

$$s_i^1 \succsim_{i,b_i(h_i')} \tilde{s}_i. \tag{A.10}$$

As $\succsim_i$ is *transitive* it follows from (A.9) and (A.10) 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

$$H_i^{1+} := \{h_i' \in H_i \mid \exists h_i \in H_i^1 \text{ s.t. } h_i' \in H_i^+(h_i)\}$$

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

$$H_i^2 := \{h_i \in H_i\backslash(H_i^1 \cup H_i^{1+}) \mid \nexists h_i' \in H_i\backslash(H_i^1 \cup H_i^{1+}) \text{ s.t. } h_i' \lhd h_i\}$$

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 (A.5).

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 (A.5) 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

$$s_i^2 \sim_{i,[s_{-i}]} s_i^1 \quad \forall s_{-i} \text{ s.t. } b_i(h_i')(s_{-i}) > 0. \tag{A.11}$$

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

$$b_i(h_i') = \sum_{s_{-i} \in S_{-i}(h_i') : b_i(h_i')(s_{-i})>0} b_i(h_i')(s_{-i}) \cdot [s_{-i}]. \tag{A.12}$$

Since $\succsim_i$ satisfies *preservation of indifference*, we conclude on the basis of (A.11) and (A.12) that

$$s_i^2 \sim_{i,b_i(h_i')} s_i^1. \tag{A.13}$$

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

$$s_i^1 \succsim_{i,b_i(h_i')} s_i \quad \forall s_i \in S_i(h_i'). \tag{A.14}$$

As $\succsim_{i,b_i(h_i')}$ is *transitive*, (A.13) and (A.14) imply that

$$s_i^2 \succsim_{i,b_i(h_i')} s_i \quad \forall s_i \in S_i(h_i')$$

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

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

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

$$H_i^{2+} := \{h_i' \in H_i \mid \exists h_i \in H_i^2 \text{ s.t. } h_i' \in H_i^+(h_i)\}.$$

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

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

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. □

Data availability

No data was used for the research described in the article.

References

Anscombe, F.J., Aumann, R.J., 1963. A definition of subjective probability. Ann. Math. Stat. 34, 199–205.
Asheim, G.B., Perea, A., 2005. Sequential and quasi-perfect rationalizability in extensive games. Games Econ. Behav. 53, 15–42.
Battigalli, P., 1997. On rationalizability in extensive games. J. Econ. Theory 74, 40–61.
Battigalli, P., Catonini, E., Lanzani, G., Marinacci, M., 2019. Ambiguity attitudes and self-confirming equilibrium in sequential games. Games Econ. Behav. 115, 1–29.
Battigalli, P., Catonini, E., Manili, J., 2023. Belief change, rationality, and strategic reasoning in sequential games. Games Econ. Behav. 142, 527–551.
Battigalli, P., Dufwenberg, M., 2009. Dynamic psychological games. J. Econ. Theory 144, 1–35.
Catonini, E., Penta, A., 2025. Backward induction reasoning beyond backward induction. Am. Econ. J. Microecon. Forthcoming.
Dekel, E., 1986. An axiomatic characterization of preferences under uncertainty: weakening the independence axiom. J. Econ. Theory 40, 304–318.
Dekel, E., Fudenberg, D., Levine, D.K., 1999. Payoff information and self-confirming equilibrium. J. Econ. Theory 89, 165–185.
Dekel, E., Fudenberg, D., Levine, D.K., 2002. Subjective uncertainty over behavior strategies: a correction. J. Econ. Theory 104, 473–478.
Epstein, L.G., Le Breton, M., 1993. Dynamically consistent beliefs must be Bayesian. J. Econ. Theory 61, 1–22.
Epstein, L.G., Schneider, M., 2003. Recursive multiple priors. J. Econ. Theory 113, 1–31.
Epstein, L.H., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. Econometrica 57, 937–969.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique priors. J. Math. Econ. 18, 141–153.
Gilboa, I., Schmeidler, D., 2003. A derivation of expected utility maximization in the context of a game. Games Econ. Behav. 44, 184–194.
Ghirardato, P., 2002. Revisiting Savage in a conditional world. Econ. Theory 20, 83–92.
Gul, F., Lantto, O., 1990. Betweenness satisfying preferences and dynamic choice. J. Econ. Theory 52, 162–177.

Hammond, P.J., 1988. Consequentialist foundations for expected utility. Theory Decis. 25, 25–78.
Hanany, E., Klibanoff, P., 2009. Updating ambiguity averse preferences. B.E. J. Theor. Econ. 9 (1), 37.
Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, 263–291.
Klibanoff, P., Hanany, E., 2007. Updating preferences with multiple priors. Theor. Econ. 2, 261–298.
Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. Econometrica 73, 1849–1892.
Kreps, D.M., Wilson, R., 1982. Sequential equilibria. Econometrica 50, 863–894.
Kuhn, H.W., 1953. Extensive games and the problem of information. In: Kuhn, H.W., Tucker, A.W. (Eds.), Contributions to the Theory of Games, Volume II. In: Annals of Mathematics Studies, vol. 28. Princeton University Press, Princeton, NJ, pp. 193–216.
Machina, M.J., 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. J. Econ. Lit. XXVII, 1622–1668.
Machina, M.J., Schmeidler, D., 1992. A more robust definition of subjective probability. Econometrica 60, 745–780.
Myerson, R.B., 1991. Game Theory: Analysis of Conflict. Harvard University Press, Cambridge, Massachusetts, London, England.
Osborne, M.J., Rubinstein, A., 1994. A Course in Game Theory. The MIT Press, Cambridge, Massachusetts, London, England.
Pearce, D.G., 1984. Rationalizable strategic behavior and the problem of perfection. Econometrica 52, 1029–1050.
Penta, A., 2015. Robust dynamic implementation. J. Econ. Theory 160, 280–316.
Perea, A., 2012. Epistemic Game Theory: Reasoning and Choice. Cambridge University Press.
Perea, A., 2014. Belief in the opponents' future rationality. Games Econ. Behav. 83, 231–254.
Perea, A., 2025a. Expected utility as an expression of linear preference intensity. Theory Decis., 1–38.
Perea, A., 2025b. Consequentialism in dynamic games. Econ. Philos. https://doi.org/10.1017/S0266267125100588. Forthcoming.
Rubinstein, A., 1991. Comments on the interpretation of game theory. Econometrica 59, 909–924.
Savage, L.J., 1954. The Foundation of Statistics. Wiley, New York.
Skyrms, M., 1987. Dynamic coherence and probability kinematics. Philos. Sci. 54, 1–20.
Siniscalchi, M., 2009. Two out of three ain't bad: a comment on "The ambiguity aversion literature: a critical assessment". Econ. Philos. 25, 335–356.
Siniscalchi, M., 2011. Dynamic choice under ambiguity. Theor. Econ. 6, 379–421.
Sinnott-Armstrong, W., 2023. Consequentialism. In: Zalta, Edward N., Nodelman, Uri (Eds.), The Stanford Encyclopedia of Philosophy, Winter 2023 edition. https://plato.stanford.edu/archives/win2023/entries/consequentialism/.
Teller, P., 1973. Conditionalization and observation. Synthese 26, 218–258.
von Neumann, J., Morgenstern, O., 1947. Theory of Games and Economic Behavior, 2nd edition. Princeton University Press, Princeton, NJ.