

Optimal choice for finite and infinite horizons[☆]

Z. Z. Méder^a, J. Flesch^b, R. Peeters^{a,*}

^a*Department of Economics, Maastricht University, The Netherlands*

^b*Department of Quantitative Economics, Maastricht University, The Netherlands*

Abstract

This paper lays down conceptual groundwork for optimal choice of a decision maker facing a finite-state Markov decision problem on an infinite horizon. We distinguish two notions of a strategy being favored on the limit of horizons, and examine the properties of the emerging binary relations. After delimiting two senses of optimality, we define a battery of optimal strategy sets – including the Ramsey-Weizsäcker overtaking criterion – and analyze their relationships and existence properties. We also relate to the work on pointwise limits of strategies by Fudenberg and Levine (1983).

Keywords: optimality, Markov decision problem

JEL: C61

1. Introduction and motivation

Some decisions are clearly limited in time: they only involve decisions and payoffs until a certain temporal point, having no relevant options and effects for the decision maker beyond that point. At one extreme, think of choosing the flavor of a scoop of icecream: the decision is blatantly one-off, and any payoffs from choosing vanilla or chocolate are gained immediately. Other decisions, however, are potentially infinite, in the sense that utility-changing decisions might have to be made after arbitrary long times. Here, one can think of choosing the palette of flavors for an ice-cream company: managers might come and go, today's kids might grow old and die out, but until the company is up and running, the flavor palette will need to be re-considered. Each of these decisions will have both immediate and long-term effects.

The horizon of a decision problem is the temporal distance between the time when the decision maker is confronted with the problem, and the furthest horizon that is payoff-relevant. This paper starts with an infinite-horizon decision problem. However, even when the horizon is infinite, a decision-maker might still act as if it were finite. There are various possible reasons for the decision maker acting in such a way: maybe he (wrongly) believes that the decision problem is actually limited. It could also happen that, due to cognitive limitations, he is unable or unwilling to calculate with any utility that he gains in periods beyond a certain period. In our example, a manager might care only about the short-run profits of the ice-cream company, and postpone R&D expenditures.

Short-sighted attitudes amount to cutting the decision tree so that it becomes finite. The alternative is to take into consideration the entire decision tree, including effects arbitrarily far into the future. If more cognitive resources, right beliefs, or deep commitments – e.g. to an ancient family business – are available, the full decision problem can be tackled. If we let the horizon of the decision-maker grow ad infinitum, i.e. we cut off smaller and smaller parts of the tree, intuition would suggest that decision strategies that are optimal on the full tree can always be approximated by strategies that are optimal for infinitely many such cuts. We show, however, that this is not always the case, and strategies that are optimal on the complete horizon can be suboptimal for all finite horizons. A third and distinct optimality notion emerges by taking the limit of strategies optimal on finite horizons.

In investigating the limit properties of decision strategies, the present work provides a systematic approach towards formulating optimality criteria such as the Ramsey-Weizsäcker overtaking criterion of Brock (1970) or the limit-equilibrium of Fudenberg and Levine (1983). We will see that the lack of a unified framework leads to incompatible interpretations of the overtaking criterion. These criteria, and thus our whole undertaking can be best understood as a robustness analysis for strategies optimal on the complete horizon with regard to changes in the discount factor. More specifically, we ask what happens if the discount factor is perturbed from the left, i.e. the decision maker becomes infinitesimally more impatient, or the stopping probability of the game increases.

Our results provide grounds for optimality refinements. As we will see, it is possible that a strategy that is optimal on the complete horizon is beaten by another strategy on *any* finite horizon. This means that whenever the game ends, the decision maker would have been strictly better off choosing the other strategy. This is, in fact, a good reason not to choose it, and instead go with the strategy that is strictly better for finite horizons.

[☆]The Netherlands Organisation for Scientific Research (NWO) is gratefully acknowledged for its support.

*Corresponding author

Email addresses: z.meder@maastrichtuniversity.nl (Z. Z. Méder),
j.flesch@maastrichtuniversity.nl (J. Flesch),
r.peeters@maastrichtuniversity.nl (R. Peeters)

zons.

The next section outlines our model and notation. In Section 3, we define some relations between strategies, delimiting the strong and weak senses of limit of finite horizons. Section 4 distinguishes two senses of optimality, and analyzes the resulting strategy sets. The limit of strategies approach of Fudenberg and Levine (1983) is considered in Section 5, while the last section concludes.

2. The problem

Our decision maker faces a finite Markov decision problem.

Definition 1. A finite Markov decision problem is given by:

- the set of time periods $\{1, 2, \dots\}$;
- a finite set of states Ω , with $\omega_1 \in \Omega$ as the initial state;
- a finite and nonempty set of pure actions A_ω that the decision maker can choose from in state ω ;
- a payoff function $u_\omega : A_\omega \rightarrow \mathbb{R}$ that assigns a payoff to every action in state ω ;
- transition probabilities $m_\omega : A_\omega \rightarrow \Delta(\Omega)$, with $m_\omega(\omega' | a_\omega)$ denoting the probability to transit from state ω to state ω' when action a_ω is chosen.

In every period, the decision-maker chooses an action from those available to him. We call a path of states and actions that the decision maker can go through a “history”.

Definition 2. A history h has the form $h = (\omega_1, a_{\omega_1}, \dots, \omega_{t-1}, a_{\omega_{t-1}}, \omega_t)$, with:

- $\omega_i \in \Omega$, for $i \in \{1, \dots, t\}$;
- $a_{\omega_i} \in A_{\omega_i}$, for $i \in \{1, \dots, t-1\}$;
- $m_{\omega_{i-1}}(\omega_i | a_{\omega_{i-1}}) > 0$, for $i \in \{1, \dots, t-1\}$.

The length of h or current time at h is denoted by $t = t(h)$. In a similar vein, the function $\omega(h) = \omega_{t(h)}$ indicates the current or end state at history h . We use H to refer to the set of all histories.

Obviously, the well-being of the decision-maker depends not only on his current choice, but also his future actions. Thus, the decision maker needs a decision rule that tells him what to do at each history. Such a decision rule is called a strategy.

Definition 3. A strategy of the decision maker assigns an action to all histories:

$$s : h \mapsto A_{\omega(h)}.$$

The set of all strategies is denoted by S .

When the decision maker chooses strategy s , we denote the probability that history h occurs by $P_s(h)$.

We consider utility functions that are additively separable. The expected utility induced by a strategy s on horizon $T \in \mathbb{N} \cup \{\infty\}$ is:

$$U_T(s) = \sum_{h | t(h) \leq T} r_h(u(s(h))) P_s(h).$$

where $r_h : \mathbb{R} \rightarrow \mathbb{R}$ is continuous and strictly increasing. Moreover, if the maximal absolute payoff in one period is $c = \max_{\omega, a_\omega} |u_\omega(a_\omega)|$, we require that $\sum_{t=1}^{\infty} \max_{h | t(h)=t} r_h(c) < \infty$, to ensure that U_∞ is finite for all possible streams of payoffs.

The horizon of this function is interpreted as the period beyond which the decision maker ignores payoffs. If the decision maker considers all future payoffs, however far they might be, we say that his horizon is infinite.

Two well-known utility functions that satisfy our requirements are the exponential discount function, whereby $r_h(u) = \delta^{t(h)-1} u$, and the quasi-hyperbolic discount function, with $r_h(u) = u$ if $t(h) = 1$ and $r_h(u) = \beta \delta^{t(h)-1} u$ otherwise. For a comprehensive list of discount functions used in various fields see Heilmann (2010), Frederick et al. (2002) or Loewenstein and Read (2003).

3. Relations on a strategy set

We start by defining relations to compare two strategies within S . First, suppose the decision maker has a finite horizon.

Definition 4. A strategy s is (strictly) favored over another strategy s' on horizon T if it induces a (strictly) higher utility on that horizon. More precisely, we write $s \succcurlyeq_T s'$ if $U_T(s) \geq U_T(s')$. Similarly, we write $s \succ_T s'$ whenever $U_T(s) > U_T(s')$.

Next, it is possible that the decision maker considers the complete horizon of infinite length of the decision problem.

Definition 5. A strategy s is (strictly) favored over another strategy s' on the complete horizon if it induces a (strictly) higher utility on that horizon. More precisely, we write $s \succcurlyeq_{CH} s'$ if $U_\infty(s) \geq U_\infty(s')$. Similarly, we write $s \succ_{CH} s'$ whenever $U_\infty(s) > U_\infty(s')$.

Another alternative is to conceive of the infinite horizon as the limit of finite horizons, as the parameter T that determines the length of the horizon goes to infinity. Thus, the third option for an agent when comparing strategies is to look at the expected utility with T being large enough. There are, however, at least two ways for looking at this limiting behavior, in a stronger and in a weaker sense.

Definition 6. A strategy is (strictly) favored over another strategy s' on the limit of finite horizons in the strong sense (LHS) if it induces a (strictly) higher utility for all horizons beyond a certain horizon T' . More precisely, we write $s \succcurlyeq_{LHS} s'$ if there is a T' such that $s \succcurlyeq_T s'$ for all $T \geq T'$. Similarly, we write $s \succ_{LHS} s'$ whenever there is a T' such that $s \succ_T s'$ for all $T \geq T'$.

Definition 7. A strategy is (strictly) favored over another strategy s' on the limit of finite horizons in the weak sense (LHW) if, for any horizon, we can find a longer horizon such that the first strategy induces a (strictly) higher utility on it than the second one. More precisely, we write $s \succ_{\text{LHW}} s'$ if for all T' there is a $T \geq T'$ such that $s \succ_T s'$. Similarly, we write $s \succ_{\text{LHW}} s'$ whenever for all T' there is a $T \geq T'$ such that $s \succ_T s'$.

Note that the latter definition is equivalent to saying that there are infinitely many horizons on which the first strategy induces a (strictly) higher utility.

Theorem 1. The relations of \succ_T , \succ_{CH} , \succ_{LHS} , \succ_{LHW} and their strict versions satisfy the properties given in Tab. 1.

Table 1: Properties of the relations.

	\succ			$>$		
	T/CH	LHS	LHW	T/CH	LHS	LHW
Total	+	-	+	-	-	-
Reflexive	+	+	+	-	-	-
Irreflexive	-	-	-	+	+	+
Symmetric	-	-	-	-	-	-
Asymm.	-	-	-	+	+	-
Antisymm.	-	-	-	+	+	-
Transitive	+	+	-	+	+	-

“+” indicates that the property necessarily holds.

“-” indicates that the property does not hold in general.

R is: *total* if xRy or yRx for all x, y ;

reflexive if xRx for all x ;

irreflexive if xRx for no x ;

symmetric if xRy implies yRx ;

asymmetric if xRy implies not yRx ;

antisymmetric if xRy and yRx imply $x = y$;

transitive if xRy and yRz imply xRz .

Proof.

The relations \succ_T , \succ_{CH} , \succ_T and \succ_{CH} .

The properties of \succ_T and \succ_T simply inherit the properties of the non-strict and strict orders \geq and $>$ on \mathbb{R} , since they are equivalent to single utility comparisons. Similarly, \succ_{CH} and \succ_{CH} inherit the properties of \geq and $>$ on \mathbb{R} .

The relation \succ_{LHS} .

To verify that \succ_{LHS} is not total, see the decision problem in Fig. 1 with \blacksquare as the initial state. The two options of the decision maker are going west (strategy s) and going east (s'). Using exponential discounting with $\delta = 0.5$, we see that

$$U_T(s) = \begin{cases} 0 & \text{if } T = 2k + 1 \\ -(0.5)^T & \text{if } T = 2k \end{cases}$$


and

$$U_T(s') = \begin{cases} 0.5^{T-1} & \text{if } T = 2k + 1 \\ -(0.5)^{T-1} & \text{if } T = 2k. \end{cases}$$

This means that whenever $T = 2k$, then $s \succ_T s'$, so $s' \not\succ_T s$, but whenever $T = 2k + 1$, then $s' \succ_T s$, so $s \not\succ_T s'$. It follows that neither $s \succ_{\text{LHS}} s'$ nor $s' \succ_{\text{LHS}} s$, therefore \succ_{LHS} is not total.

The relation \succ_{LHS} is reflexive (and not irreflexive), as any strategy is weakly better than itself on all horizons.

Figure 1: The relation \succ_{LHS} is not total.



We turn to symmetry: \succ_{LHS} is not symmetric, asymmetric or antisymmetric. Lack of symmetry follows trivially from cases when one strategy generates strictly higher utility on all finite horizons. Due to its reflexivity, it is not asymmetric. To see that it is not antisymmetric either, take two different strategies $s \neq s'$ that induce the same expected utility on all horizons. For these, we have $s \succ_{\text{LHS}} s'$ and $s' \succ_{\text{LHS}} s$ with $s \neq s'$. Thus \succ_{LHS} is not antisymmetric.

To show that \succ_{LHS} is transitive, suppose $s \succ_{\text{LHS}} s'$ and $s' \succ_{\text{LHS}} s''$. This means that there is a T' such that $s \succ_T s'$ for all $T \geq T'$, and that there is a T'' such that $s' \succ_T s''$ for all $T \geq T''$. Let $T''' = \max\{T', T''\}$. Consequently, $s \succ_T s''$ whenever $T \geq T'''$. It follows that $s \succ_{\text{LHS}} s''$.

The relation \succ_{LHW} .

In contrast with \succ_{LHS} , the relation \succ_{LHW} is total. Choosing any s and s' , for any T , either $s \succ_T s'$ or $s' \succ_T s$. But the set of possible choices for T is infinite. Therefore, $s \succ_{\text{LHW}} s'$ for infinitely many T or $s' \succ_{\text{LHW}} s$ for infinitely many T (or possibly, both).

Next, \succ_{LHW} is reflexive for the same reason as \succ_{LHS} , and thus not irreflexive.


The symmetry properties of \succ_{LHW} are also the same as those of \succ_{LHS} , again for the same reasons.

Fig. 2 shows a decision problem where \succ_{LHW} is not transitive. Suppose again that \blacksquare is the starting state and discounting is exponential with $\delta = 0.5$. The decision maker has three options: going southwest (strategy s), south (s') or southeast (s''). The payoffs for these strategies on different horizons are summarized in Tab. 2. We see that whenever $T = 3k$, then $s \succ_T s'$, so $s \succ_{\text{LHW}} s'$. Moreover, whenever $T = 3k + 2$, then $s' \succ_T s''$, so $s' \succ_{\text{LHW}} s''$. But $U_T(s) < 0 = U_T(s'')$ for all T , so there is no T such that $s \succ_T s''$. It follows that $s \not\succ_{\text{LHW}} s''$, therefore \succ_{LHW} is not transitive.

Table 2: Utility comparison for Fig. 2.

	$U_T(s)$	$U_T(s')$	$U_T(s'')$
$T = 3k + 1$	-0.5^{T-1}	0	0
$T = 3k + 2$	-0.5^{T-1}	0.5^{T-2}	0
$T = 3k$	-0.5^{T-1}	-0.5^{T-2}	0

Figure 2: The relations \succsim_{LHW} and $>_{\text{LHW}}$ are not transitive.



The relation $>_{\text{LHS}}$.

By comparing any strategy s to itself, or two distinct strategies that generate the same payoffs on all finite horizons, we can see that $>_{\text{LHS}}$ is not total.

Contrary to its non-strict version, $>_{\text{LHS}}$ is irreflexive, since no strategy can strictly beat itself on any horizon. Therefore, it is also not reflexive.

The case when one strategy yields higher payoff than the other for all horizons shows that $>_{\text{LHS}}$ is not symmetric. It is, however, both asymmetric and antisymmetric: if $s >_{\text{LHS}} s'$, then there exists T' such that $s >_T s'$ for all $T \geq T'$. But then it is not possible that also a T'' exists for which $s' >_T s$ for all $T \geq T''$, since this would imply $s >_T s'$ and $s' >_T s$ for $T \geq \max\{T', T''\}$, which is impossible.

To show that $>_{\text{LHS}}$ is transitive, we can use the argument for the transitivity of its non-strict counterpart, replacing non-strict inequalities with strict ones everywhere.

The relation $>_{\text{LHW}}$.

For similar reasons as $>_{\text{LHS}}$, the relation $>_{\text{LHW}}$ is not total.

This relation is irreflexive (and not reflexive), and we can again use the same argument as for $>_{\text{LHW}}$.

The analogy also carries over for non-symmetry. However, $>_{\text{LHW}}$ is not asymmetric or antisymmetric: the example in Fig. 1 shows that it is possible that $s >_T s'$ for infinitely many T and also $s' >_T s$ for infinitely many T .

It is easy to see that the example on Fig. 2 also shows the non-transitivity of $>_{\text{LHW}}$, since all of the strategies in Tab. 2 are strictly best on the respective horizons. \square

4. Optimal strategies

The decision maker is looking for an optimal strategy. However, he might – for any reason – decide to restrict his choice to a nonempty subset G of the full strategy space S . For example the decision maker might focus on stationary strategies.

Given a relation R on S , a strategy in G can be said to be optimal in two senses. First, a strategy can be regarded as optimal if no further improvements can be made on it, i.e. if there is no strategy in G that is favored to it. One can think of a former definition for being a world champion in chess, which used to be not being beaten in a championship final – in case of a draw, the incumbent would keep his title. Formally, a strategy $s^* \in G$ is optimal according to relation R in the “not-beaten” (NB) sense, if there is no other strategy $s \in G$ such that sRs^* .

Alternatively, a strategy in G can be said to be optimal if it is favored to all other strategies in G . Here, one can think of horse

races, where in case of a tie bookmakers consider all the horses in the dead heat as “winners”. Formally, a strategy $s^* \in G$ is optimal according to relation R in the “beat-all” (BA) sense if s^*Rs for all other strategies $s \in G$.

For a total and transitive relation, a simple duality emerges: a strategy is not beaten by the strict order if and only if it is favored to every other strategy according to the weak order, and a strategy is not beaten by the weak order if and only if it is favored to every other strategy according to the strict order. Therefore, we have the following definitions for optimal strategies for finite horizons and the complete horizon.

Definition 8. The set of strategies in G that are optimal on horizon T is given by:

$$G_T = \{s \in G_s \mid \nexists s' \in G_s, s' >_T s\} = \{s \in G \mid \forall s' \in G, s \succsim_T s'\}.$$

The set of strategies in G that are uniquely optimal on horizon T is given by:

$$\overline{G}_T = \{s \in G \mid \nexists s' \in G_s, s' \succsim_T s\} = \{s \in G \mid \forall s' \in G_s, s >_T s'\},$$

with $G_s = G \setminus \{s\}$ here and henceforth.

Definition 9. The set of strategies in G that are optimal on the complete horizon is given by:

$$G_{\text{CH}} = \{s \in G \mid \nexists s' \in G_s, s' >_{\text{CH}} s\} = \{s \in G \mid \forall s' \in G_s, s \succsim_{\text{CH}} s'\}.$$

The set of strategies in G that are uniquely optimal on the complete horizon is given by:

$$\overline{G}_{\text{CH}} = \{s \in G \mid \nexists s' \in G_s, s' \succsim_{\text{CH}} s\} = \{s \in G \mid \forall s' \in G_s, s >_{\text{CH}} s'\}.$$

As its name suggests, a uniquely optimal strategy set contains at most one element.

This simple duality of the two senses of optimality fails however in the case of the limit of finite horizon relations, since these are either not total (\succsim_{LHS}) or not transitive (\succsim_{LHW}). Thus, we have a total of eight possible definitions for the concept of “optimal strategy on the limit of finite horizon”:

Definition 10. The sets of strategies s in G that are “not-beaten-optimal” or “beat-all-optimal” on the limit of finite horizons in the strong or weak sense, according to the respective strict or non-strict relations, are defined as in Tab. 3:

Table 3: Definitions of optimal strategy sets.

	LHS	LHW
NB, \succsim	$\{\nexists s' \in G_s, s' \succsim_{\text{LHS}} s\}$	$\{\nexists s' \in G_s, s' \succsim_{\text{LHW}} s\}$
NB, $>$	$\{\nexists s' \in G_s, s' >_{\text{LHS}} s\}$	$\{\nexists s' \in G_s, s' >_{\text{LHW}} s\}$
BA, \succsim	$\{\forall s' \in G_s, s \succsim_{\text{LHS}} s'\}$	$\{\forall s' \in G_s, s \succsim_{\text{LHW}} s'\}$
BA, $>$	$\{\forall s' \in G_s, s >_{\text{LHS}} s'\}$	$\{\forall s' \in G_s, s >_{\text{LHW}} s'\}$

For example, we denote $\{s \in G \mid \nexists s', s' \succsim_{\text{LHS}} s\}$ by $G_{\text{LHS}}^{\text{NB}\succsim}$. The other sets are denoted in a consistent manner.

Theorem 2. For any Markov decision problem and $G \subseteq S$, the different notions of optimality induce the following relations on the resulting sets of optimal strategies:

$$\begin{array}{ccc} & G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ} & \\ & \supseteq & \\ G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ} & & G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ} \subseteq G_{CH} \\ & \supseteq & \\ & G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ} & \end{array}$$

Proof.

$$G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ}, G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ}, G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ}, G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ}.$$

Strategy s is in $G_{LHW}^{NB \succ}$ if and only if there is no $s' \in G_s$ such that $s' \succ_T s$ for infinitely many T . This is true if and only if $U_T(s') \geq U_T(s)$ for only finitely many T , which is again equivalent to there being a $T'(s')$ such that $U_T(s) > U_T(s')$ for all $T \geq T'(s)$. The latter condition is the definition of $G_{LHS}^{BA \succ}$, so indeed $G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ}$. Similarly, the definitions of all the other pairs of sets are logically equivalent.

$$G_{LHS}^{BA \succ} \subseteq G_{LHS}^{BA \succ}, G_{LHW}^{BA \succ} \subseteq G_{LHW}^{BA \succ}.$$

These cases are trivial, since $s \succ_T s' \Rightarrow s \succ_T s'$.

$$G_{LHS}^{BA \succ} \subseteq G_{LHW}^{BA \succ}.$$

Choose any $s \in G_{LHS}^{BA \succ}$. Thus, for all $s' \in G$, there is a T' such that $U_T(s) > U_T(s')$ for all $T \geq T'$. Obviously then, for all s' there are infinitely many T' for which $U_T(s) \geq U_T(s')$, so we have $s \in G_{LHW}^{BA \succ}$.

$$G_{LHS}^{BA \succ} \subseteq G_{LHW}^{BA \succ}.$$

This case is analogous to the previous one, by replacing \succ with \succ .

$$G_{LHW}^{BA \succ} \subseteq G_{CH}.$$

Fix $s \in G_{LHW}^{BA \succ}$ and $s' \in G$. Since $s \succ_{LHW} s'$ for infinitely many T , $U_\infty(s) = \lim_{T \rightarrow \infty} U_T(s) \geq \lim_{T \rightarrow \infty} U_T(s') = U_\infty(s')$. Therefore, $s \in G_{CH}$. \square

We know that for any closed and compact $G \subseteq S$ strategies optimal on the complete horizon exist, i.e. G_{CH} is nonempty for closed, compact G .

The set $G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ}$ contains those strategies that are weakly favored to any other strategy on infinitely many horizons. For this reason, we call them “repeatedly” optimal strategies, and denote the resulting set of strategies by G_R . We conjecture that for any Markov decision problem and closed, compact $G \subseteq S$, the set of repeatedly optimal strategies is nonempty. Fig. 4 shows an example where transfinite induction, a natural attempt of proving the nonemptiness of G_R , fails.

Members of the set $G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ}$ are those strategies for which, for every other strategy s' , there is a certain horizon – depending on s' – after which s is weakly favored to s' for all further horizons. For this reason, we call them “uniformly” optimal strategies, and denote the resulting set of optimal strategies by G_U . Uniformly optimal strategies are not guaranteed to exist. For the decision problems in Figs. 1 and 2, there are only two and three possible strategies, those defined in Thm. 1. Since $s, s' \notin S_U$ for the first problem, and $s, s', s'' \notin S_U$ for

the second problem, there are no uniformly optimal strategies in these decision problems for $G = S$.

It can easily be seen that $G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ}$ might be empty: one can think of a decision problem where two strategies induce the same payoffs over all horizons.

The definition of $G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ}$ corresponds to the Ramsey-Weizsäcker overtaking criterion (Brock (1970)). We denote the set of strategies optimal according to the overtaking criterion by G_{OT} . From our previous results, it is obvious that a decision problem might have no optimal strategy according to the overtaking criterion. Moreover, if G_{OT} is nonempty, then it contains exactly one element.

Originally, the role of the overtaking criterion is to compare infinite utility streams where the stakeholders – the ones “gaining” the utility – are distinct in each period, as in a non-overlapping generations model. However, the lack of conceptual clarity led to misinterpretations of the criterion. For example, Rubinstein (1979)’s definition makes the overtaking criterion equivalent to optimality on the complete horizon in our framework, due to the presence of discounting.

5. Pointwise limits

Since we can only conjecture the existence of repeatedly optimal strategies, the question arises whether we can generate repeatedly optimal strategies as the limit of optimal strategies on finite horizons? A natural way to attempt this is through the pointwise limit of a sequence of strategies.

Definition 11. Suppose $\mathfrak{s} = (s_T)_{T \in \mathbb{N}}$ is an infinite sequence of strategies. We say that strategy $s \in S$ is a pointwise limit of \mathfrak{s} if there is an index set $I \subseteq \mathbb{N}$ of infinite cardinality so that for every history h there is a horizon $T(h)$ such that for every $T' \geq T(h)$ in I , $s(h) = s_{T'}(h)$.

It is easy to provide an algorithm for “constructing” pointwise limits of sequences of strategies. First, order all histories in a way that shorter histories always precede longer ones. Then, go through the histories step by step: at each history, choose any action that is taken by infinitely many strategies, and eliminate the remaining strategies. After each step of choice and elimination, the number of remaining strategies will be infinite. Thus the algorithm will never halt, and at least one pointwise limit will exist for any infinite sequence of strategies.

The pointwise limit approach is also taken by Fudenberg and Levine (1983), who show the following theorem.

Theorem 3. Suppose \mathfrak{s} is a sequence of strategies optimal on all finite horizons: $\mathfrak{s} = (s_T)$ with $s_T \in G_T$ for each $T \in \{1, 2, \dots\}$, and let $s \in S$ be a pointwise limit of \mathfrak{s} . Then s is optimal on the complete horizon: $s \in G_{CH}$. We denote the set of all such strategies with G_{FL} .

Since they use a game-theoretic approach, and their formalism is somewhat different, we attach a short proof.


Proof. Suppose the contrary. Then, there must be an $s \in G_{FL}$ and $s' \in G$ such that $U_\infty(s') = U_\infty(s) + \delta$ for some $\delta > 0$. From

the finiteness condition $\sum_{t=1}^{\infty} \bar{r}(t) < \infty$ we get $\lim_{t \rightarrow \infty} \bar{r}(t) = 0$. Therefore, there exists a T' such that $U_T(s') > U_T(s)$ for all $T > T'$.

We know that for some infinite index set I , strategy s is the pointwise limit of (g_i) , with $g_i \in G_i$ for each $i \in I$. This implies that there is some horizon $T'' \in I$ and a strategy $g_{T''}$ optimal on $T'' > T'$ such that $s(h) = g_{T''}(h)$ for all h with $t(h) < T''$. Consequently, s is also optimal on T'' , but this contradicts $U_T(s') > U_T(s)$ for all $T > T'$. \square

A strategy generated through pointwise limits is not necessarily repeatedly optimal. To see this, consider the decision problem in Fig. 3. Suppose that \blacksquare is the starting state and discounting is exponential with $\delta = 0.5$. Let g_t be the following strategy for $t > 1$: go southeast, stay there until t , go east in that period, then stay there.


Figure 3: Strategy generated through pointwise limits is dominated for every horizon



First we show that g_t is the unique optimal strategy on horizon t . On this horizon, it induces the payoff sequence $(0, \dots, 0, 8)$ and hence a utility of $8 \cdot 0.5^{t-1}$. There are only three types of alternative strategies to it: taking less time or more time to get the payoff of 8, or going southwest. Since the payoff of 8 is taken at the last possible moment, any strategy that takes more time will induce a utility of 0 on horizon t . On the other hand, any strategy that takes less time will end up with a utility smaller or equal to 0, because of the repeated payoff of -16 on every round after taking the payoff of 8. Last, going southwest induces a utility of 0.5^{t-1} . Thus g_t is uniquely optimal on horizon t .

Since g_t is uniquely optimal on horizon t , it can easily be verified that G_{FL} is a singleton set, only consisting of the strategy s of going southeast, then staying at that point forever. This strategy earns 0 on all horizons. However, we have seen that going southwest earns 0.5^{t-1} on all finite horizons, therefore $s' >_{LHS} s$, and s is not repeatedly optimal.

Figure 4: Blablabla

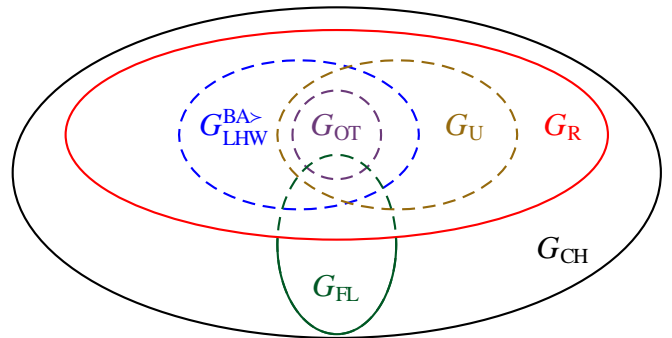


6. Discussion

The previous sections have shown that the different optimality notions are related as illustrated in Fig. 5. It should be noted that our examples concern non-generic decision problems. However, we do not regard this as a serious drawback, as many interesting decision problems are non-generic, as indeed many existing (equilibrium) refinements. In the generic case, all of these sets collapse into a single one.

Figure 5: The different optimality notions.

A dashed boundary indicates that the set might be empty.



The existence of repeatedly optimal strategies remains an intriguing open question, although we conjecture that for closed subsets of the strategy space, repeatedly optimal strategies will necessarily exist for any discount function. If the existence of repeatedly optimal strategies can be shown, we have good reasons to use it as a refinement of optimal strategies on the complete horizon. Namely, the choice of a repeatedly optimal strategy will guarantee that for infinitely many periods, the decision maker has reason not to feel any regret over his strategy choice. If a uniformly optimal strategy can be found, there is even more reason to rejoice: compared with every other strategy, any regret for choosing the uniformly optimal strategy will fully dissipate after just finitely many periods.

While our model is presented in a decision-theoretic framework, the game-theoretic extensions are straightforward. Results concerning the relationships of various strategy sets will carry over, but the existence properties of various strategy sets will have to be readdressed for games.

References

References

Brock, W.A. (1970): An axiomatic basis for the Ramsey-Weizsäcker overtaking criterion. *Econometrica* 38/6, 927–929.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002): Time discounting and time preference: A critical review. *Journal of Economic Literature* 40/2, 351–401.

Fudenberg, D and Levine, D (1983): Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory* 31, 251–268.

Garcia, A. (2004): Forecast horizons for a class of dynamic games. *Journal of Optimization Theory and Applications* 122/3, 471–486.

Heilmann, C. (2010): Rationality and time. Thesis submitted to the London School of Economics.

- Loewenstein, G. and Read, D. (2003): *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. Russell Sage.
- Rubinstein, A. (1979): Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory* 21, 1–9.