

Optimal choice for finite and infinite horizons

Citation for published version (APA):

Méder, Z. Z., Flesch, J., & Peeters, R. J. A. P. (2012). *Optimal choice for finite and infinite horizons*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 024 <https://doi.org/10.26481/umamet.2012024>

Document status and date:

Published: 01/01/2012

DOI:

[10.26481/umamet.2012024](https://doi.org/10.26481/umamet.2012024)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Zsombor Méder, János Flesch,
Ronald Peeters

Optimal choice for finite and infinite horizons

RM/12/024

METEOR

Maastricht University School of Business and Economics
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands

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-Weizäcker 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, overtaking criterion

JEL: C61

1. Introduction and motivation

Some decision problems 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 come and go, today's kids grow old and die out, but until the company is up and running, the flavor palette will need to be reconsidered. 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,

[☆]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)

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. One 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. However, this is not always the case, and strategies that are optimal on the complete horizon can be suboptimal for *all* finite truncations of the horizon. A third and distinct optimality notion emerges by taking pointwise limits of strategies optimal on finite horizons.

In investigating the limit properties of decision strategies, the present work provides a systematic approach towards formulating and comparing optimality criteria such as the Ramsey-Weizsäcker overtaking criterion (von Weizsäcker, 1965) or the limit-equilibrium of Fudenberg and Levine (1983).

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 up to any period, the decision maker would be strictly better off in expectation with the other strategy. In such cases, there might be good intuitive reasons to go with one strategy rather than the other, albeit they induce the same expected payoff on the complete horizon.

The next section outlines our model and notation. In Section 3, we define some relations between strategies, delimiting the strong and weak senses of taking limits 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. Then, we consider an approach to prove the non-emptiness of one of our optimality refinements, while the last section provides some concluding remarks.

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 on his future actions. Thus, the decision maker needs a decision rule that tells him what to do at any history. Such a decision rule is called a strategy.

Definition 3. A strategy of the decision maker is a function s that assigns to every history h an action $s(h) \in A_{\omega(h)}$. The set of all strategies is denoted by S .

We suppose that discounting is exponential, so the expected utility induced by a strategy s on horizon $T \in \mathbb{N}$ is:

$$U_T(s) = \sum_{h | t(h) \leq T} \delta^{t(h)-1} u_{\omega(h)}(s(h)) P_s(h),$$

where $u_{\omega(h)}(s(h))$ is the payoff received from taking the action $s(h)$ at history h , while $P_s(h)$ is the probability that history h occurs when the decision maker chooses strategy s . The horizon T 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. The utility induced by s on the infinite horizon is $U_\infty(s) = \lim_{T \rightarrow \infty} U_T(s)$.

3. Comparing strategies

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 a finite horizon T if it induces a (strictly) higher utility on that horizon. More precisely, we write $s \succ_T s'$ 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 $s \succ_{CH} s'$ 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. This option involves comparing strategies by their induced expected utility for large enough T . 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 s 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 T beyond a certain horizon T' . More precisely, we write $s \succ_{\text{LHS}} s'$ if there is a T' such that $s \succ_T s'$ for all $T \geq T'$. Similarly, we write $s >_{\text{LHS}} s'$ whenever there is a T' such that $s >_T s'$ for all $T \geq T'$.

Definition 7. A strategy s is (strictly) favored over another strategy s' on the limit of finite horizons in the weak sense (LHW) if, for any horizon T' , we can find a longer horizon T such that s induces a (strictly) higher utility on T than s' . 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 $s >_{\text{LHW}} s'$ whenever for all T' there is a $T \geq T'$ such that $s >_T s'$.

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.

Following Ramsey (1928), von Weizsäcker (1965) investigates both $>_{\text{LHS}}$ and $>_{\text{LHW}}$, although in a different setting than ours. Later, the optimal growth literature adopts the stronger version as the Ramsey-Weizsäcker overtaking criterion (Gale, 1967; Brock, 1970; Brock and Haurie, 1976). Within a game-theoretic framework, Rubinstein (1979) considers an even stronger version of this relation.

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	-	-	-	-	-	-
Asymmetric	-	-	-	+	+	-
Antisymmetric	-	-	-	+	+	-
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} , $>_T$ and $>_{\text{CH}}$.

The properties of \succ_T and $>_T$ 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 $>_{\text{CH}}$ inherit the properties of \geq and $>$ on \mathbb{R} .

The relation \succ_{LHS} .

To verify that \succ_{LHS} is not total, consider the decision problem in Fig. 1 with \odot as the initial state. The two options of the decision maker are going west (strategy s) and going east (s'). Suppose the discount factor is $\delta = 0.5$. We see that

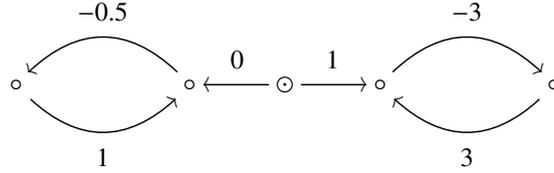
$$U_T(s) = \begin{cases} 0 & \text{if } T = 2k + 1 \\ -(0.5)^T & \text{if } T = 2k \end{cases}$$

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

Figure 1: The relation \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.

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, $s \succ_T s'$ for infinitely many T or $s' \succ_T s$ 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 \odot is the starting state and the discount factor is $\delta = 0.5$. 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 is not favored over s'' according to \succ_{LHW} , therefore \succ_{LHW} is not transitive.

Figure 2: The relations \succ_{LHW} and \succ_{LHS} are not transitive.

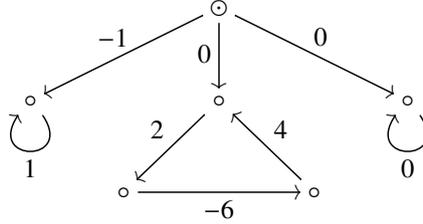


Table 2: Utility comparison for Fig. 2.

	$U_T(s)$	$U_T(s')$	$U_T(s'')$
$T = 3k + 1$	-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

The relation \succ_{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 \succ_{LHS} is not total.

Contrary to its non-strict version, \succ_{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 \succ_{LHS} is not symmetric. It is, however, both asymmetric and antisymmetric: if $s \succ_{LHS} s'$, then there exists T' such that $s \succ_T s'$ for all $T \geq T'$. But then it is not possible that also a T'' exists for which $s' \succ_T s$ for all $T \geq T''$, since this would imply $s \succ_T s'$ and $s' \succ_T s$ for $T \geq \max\{T', T''\}$, which is impossible.

To show that \succ_{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 \succ_{LHW} .

For similar reasons as \succ_{LHS} , the relation \succ_{LHW} is not total.

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

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

It is easy to see that the example on Fig. 2 also shows the non-transitivity of \succ_{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 restrict his attention to 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 other 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 tie, the incumbent would keep his title. Formally, a strategy $s^* \in G$ is optimal in G 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 $s \in G$.

The orders defined on finite and the complete horizons involve single utility comparison. Therefore, 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. Similarly, 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.

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

$$G_T = \{s \in G_s \mid \nexists s' \in G_s, s' \succ_T s\} = \{s \in G \mid \forall s' \in G_s, s \succcurlyeq_T s'\},$$

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

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

$$\overline{G}_T = \{s \in G \mid \nexists s' \in G_s, s' \succcurlyeq_T s\} = \{s \in G \mid \forall s' \in G_s, s \succ_T s'\}.$$

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

$$G_{CH} = \{s \in G \mid \nexists s' \in G_s, s' \succ_{CH} s\} = \{s \in G \mid \forall s' \in G_s, s \succcurlyeq_{CH} s'\}.$$

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

$$\overline{G}_{CH} = \{s \in G \mid \nexists s' \in G_s, s' \succcurlyeq_{CH} s\} = \{s \in G \mid \forall s' \in G_s, s \succ_{CH} s'\}.$$

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

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

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 given by the conditions in Tab. 3:

	LHS	LHW
NB, \succcurlyeq	$\nexists s' \in G_s, s' \succcurlyeq_{LHS} s$	$\nexists s' \in G_s, s' \succcurlyeq_{LHW} s$
NB, \succ	$\nexists s' \in G_s, s' \succ_{LHS} s$	$\nexists s' \in G_s, s' \succ_{LHW} s$
BA, \succcurlyeq	$\forall s' \in G_s, s \succcurlyeq_{LHS} s'$	$\forall s' \in G_s, s \succcurlyeq_{LHW} s'$
BA, \succ	$\forall s' \in G_s, s \succ_{LHS} s'$	$\forall s' \in G_s, s \succ_{LHW} s'$

We denote $\{s \in G \mid \nexists s' \in G_s, s' \succcurlyeq_{LHS} s\}$ by $G_{LHS}^{NB \succcurlyeq}$. 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:

$$\begin{array}{c} G_{\text{LHW}}^{\text{NB}\succ} = G_{\text{LHS}}^{\text{BA}\succ} \\ \subseteq \\ \bar{G}_{\text{CH}} \subseteq G_{\text{LHW}}^{\text{NB}\succ} = G_{\text{LHS}}^{\text{BA}\succ} \\ \subseteq \\ G_{\text{LHS}}^{\text{NB}\succ} = G_{\text{LHW}}^{\text{BA}\succ} \\ \subseteq \\ G_{\text{LHS}}^{\text{NB}\succ} = G_{\text{LHS}}^{\text{BA}\succ} \\ \subseteq \\ G_{\text{LHS}}^{\text{NB}\succ} = G_{\text{LHW}}^{\text{BA}\succ} \\ \subseteq \\ G_{\text{LHS}}^{\text{NB}\succ} = G_{\text{LHW}}^{\text{BA}\succ} \subseteq G_{\text{CH}}. \end{array}$$

Proof.

$$\bar{G}_{\text{CH}} \subseteq G_{\text{LHS}}^{\text{BA}\succ}.$$

Take $s \in \bar{G}_{\text{CH}}$. Then, $U_{\infty}(s) > U_{\infty}(s')$ for all $s' \in G_s$, hence there exists a $T'(s')$ so that for all subsequent horizons $T \geq T'(s')$, we get $U_T(s) > U_T(s')$. Thus, $s \in G_{\text{LHS}}^{\text{BA}\succ}$.

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

Strategy s is in $G_{\text{LHW}}^{\text{NB}\succ}$ 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 $G_{\text{LHS}}^{\text{BA}\succ}$, so indeed $G_{\text{LHW}}^{\text{NB}\succ} = G_{\text{LHS}}^{\text{BA}\succ}$. Similarly, the definitions of all the other pairs of sets are logically equivalent.

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

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

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

Choose any $s \in G_{\text{LHS}}^{\text{BA}\succ}$. 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 $s \in G_{\text{LHW}}^{\text{BA}\succ}$.

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

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

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

Fix $s \in G_{\text{LHW}}^{\text{BA}\succ}$ and $s' \in G$. Since $s \succ_{\text{LHW}} s'$ 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, $s \in G_{\text{CH}}$. \square

We know that whenever $G \subseteq S$ is nonempty and closed for pointwise limits, strategies optimal on the complete horizon exist, i.e. G_{CH} is nonempty (cf. Thm. 3).

The set $G_{\text{LHS}}^{\text{NB}\succ} = G_{\text{LHW}}^{\text{BA}\succ}$ contains those strategies that are non-strictly 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 whenever G is nonempty and closed for pointwise limits, the set of repeatedly optimal strategies is also non-empty. In Section 6 we deal with an attempt to prove the non-emptiness of G_{R} .

Members of the set $G_{\text{LHW}}^{\text{NB}\succ} = G_{\text{LHS}}^{\text{BA}\succ}$ are those strategies s for which, for every other strategy s' , there is a certain horizon – depending on s' – after which s is non-strictly favored to s' for all further horizons. This is exactly the set of strategies that von Weizsäcker (1965) calls optimal. Therefore, we call these “overtakingly” optimal strategies, and denote the resulting set by G_{OT} . Such strategies are not guaranteed to exist. Indeed, for the decision problems in

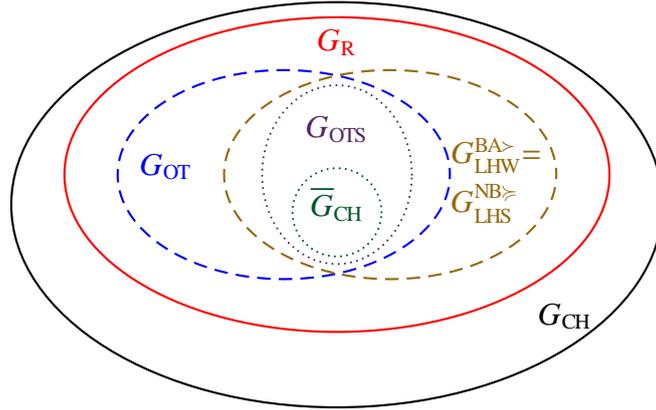
Figs. 1 and 2, there are only two, respectively three possible strategies, those defined in Thm. 1. Since $s, s' \notin S_U$ for the first problem, and $s, s', s'' \notin S_U$ for the second problem, there are no overtakingly optimal strategies in these decision problems for $G = S$. The seminal paper of von Weizsäcker (1965) remarks – in a different setting – that it might be possible for a strategy to be repeatedly, but not overtakingly optimal; however, he does not provide an example.

It is easily seen that $G_{LHS}^{NB \succ} = G_{LHW}^{BA \succ}$ can be empty. Think of a decision problem where there are only two different strategies and they induce the same payoffs over all horizons. Then neither of them can be called optimal according to this optimality criterion. Due to such instances, this particular notion is less appealing to us.

One could call the set of strategies in $G_{LHW}^{NB \succ} = G_{LHS}^{BA \succ}$ “strictly overtakingly optimal”. Thus, we denote them by G_{OTS} . From our previous results, it is obvious that a decision problem might have no strictly overtakingly optimal strategies. Moreover, if G_{OTS} is nonempty, then it contains exactly one element.

Fig. 3 illustrates the set inclusion relationships between the various optimality notions defined in this section.

Figure 3: The different optimality notions.



A solid boundary indicates that the set is (conjectured to be) non-empty.
A dashed or dotted boundary indicates that the set might be empty.
A dotted boundary also indicates that the set contains at most one element.

5. Pointwise limits

We now turn to examine the relationship between the pointwise limits of strategies optimal on finite horizons examined by Fudenberg and Levine (1983) and our optimality notions.

Definition 11. Suppose $\mathfrak{s} = (s_i)_{i \in \mathbb{N}}$ 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 an index m_h such that $s(h) = s_i(h)$ for every $i \in I$ with $i \geq m_h$.

It is easy to provide an iterative method for finding pointwise limits of a sequence 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 strategies that do not take this action at that history. 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.

For a set of strategies $G \subseteq S$, denote by G_{FL} the set of all strategies s for which there exists a sequence $\mathfrak{s} = (s_T)$ with $s_T \in G_T$ for each $T \in \{1, 2, \dots\}$ such that s is a pointwise limit of \mathfrak{s} . The following theorem follows from the results of Fudenberg and Levine (1983). Since they use a game-theoretic approach, and their formalism is somewhat different, we attach a short proof.

Theorem 3. *Suppose $G \subseteq S$ is nonempty and closed for pointwise limits. Then, G_{FL} is a nonempty subset of G_{CH} .*

Proof. Since G is nonempty and closed for pointwise limits, the nonemptiness of G_{FL} follows from the discussion above on the iterative method. We will now show that $G_{\text{FL}} \subseteq G_{\text{CH}}$. Take an arbitrary $s \in G_{\text{FL}}$. Then, there exists a sequence $(s_T)_{T \in \mathbb{N}}$, with $s_T \in G_T$ for each $T \in \mathbb{N}$, and an infinite index set $I \subseteq \mathbb{N}$ such that s is the pointwise limit of the subsequence $(s_T)_{T \in I}$. Let $s' \in G$. Because $s_T \in G_T$, we have $U_T(s_T) \geq U_T(s')$ for every $T \in \mathbb{N}$. Let R be the maximum of the absolute values of all payoffs, and for every $T \in \mathbb{N}$ let

$$w_T = \sum_{t=T+1}^{\infty} \delta^{t-1} R = \frac{\delta^T}{1-\delta} R.$$

Then, $U_T(s_T) \leq U_{T'}(s_T) + w_{T'}$ for every $T, T' \in \mathbb{N}$ with $T' \leq T$, and $U_T(s') \geq U_{\infty}(s') - w_T$ for every $T \in \mathbb{N}$. Consequently, for every $T, T' \in \mathbb{N}$ with $T' \leq T$ we have

$$U_{T'}(s_T) \geq U_T(s_T) - w_{T'} \geq U_T(s') - w_{T'} \geq U_{\infty}(s') - w_T - w_{T'}.$$

For a fixed T' , taking the limit when T tends to infinity within the set I yields

$$U_{T'}(s) \geq U_{\infty}(s') - w_{T'}.$$

Next, taking the limit when $T' \rightarrow \infty$ yields

$$U_{\infty}(s) \geq U_{\infty}(s').$$

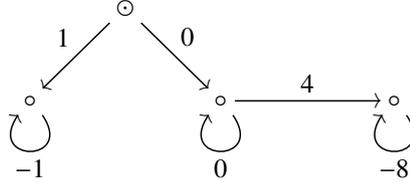
Because s' was arbitrary in G , we have shown that $s \in G_{\text{CH}}$. □

Remark. *A strategy generated through pointwise limits is not necessarily repeatedly optimal, i.e. G_{FL} is not necessarily a subset of G_{R} , even if G is closed for pointwise limits.*

To see this, consider the decision problem in Fig. 4. Suppose that \odot is the starting state and the discount factor is $\delta = 0.5$. Let g_t be the following strategy for $t > 1$: go southeast, stay there until t , go east in that period, then stay there.

First we show that g_t is the unique optimal strategy on horizon t . On this horizon, it induces the payoff sequence $(0, \dots, 0, 4)$ and hence a utility of $4 \cdot 0.5^{t-1}$. There are only three types of alternative strategies to it: taking less time or more time to get the payoff of 4, or going southwest. Since the payoff of 4 is taken at the last possible moment, any strategy that takes more time will

Figure 4: Pointwise limit of strategies optimal on finite horizons is dominated on all horizons.



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 -8 on every round after taking the payoff of 4. Last, going southwest induces a utility of 0.5^{t-1} . Thus g_t is uniquely optimal on horizon t .

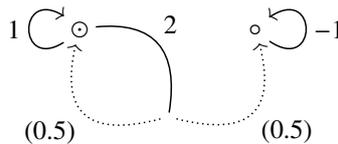
Since g_t is uniquely optimal on horizon t , it can easily be verified that G_{FL} is a singleton set, only consisting of the strategy s of going southeast, then staying at that point forever. This strategy earns 0 on all horizons. However, we have seen that going southwest earns 0.5^{t-1} on all finite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

6. Existence of repeatedly optimal strategies

Take $G \subseteq S$ to be non-empty and closed for pointwise limits. A natural attempt to prove the existence of a repeatedly optimal strategy within G would be by using Zorn's lemma, which is equivalent to the axiom of choice. For $s, s' \in G$, we write sRs' if and only if either $s = s'$, or $s >_{LHS} s'$. It follows from Thm. 1 that relation R is a partial order on G , i.e. R is reflexive, antisymmetric and transitive. Take a non-empty set $G' \subseteq G$. We call G' a chain in G if R is a total order on G' . A strategy $s \in G$ is called an upper bound for a chain G' if sRs' for every $s' \in G'$. By Zorn's lemma, if every chain G' in G has an upper bound, then there exists a strategy $s \in G$ for which there is no strategy $s' \in G$ such that $s' >_{LHS} s$. This would mean that $s \in G_{LHS}^{NB}$, i.e. that s is a repeatedly optimal strategy of G . So we would only have to prove that every chain G' in G has an upper bound. The following example demonstrates that this is not easy. In this example, we first construct a chain of countably infinite cardinality, and subsequently show that the unique pointwise limit of this chain is not an upper bound.

For the decision problem represented in Fig. 5, the decision maker has the options of going east (E) or staying in place (S) in the initial state \odot . By choosing to go east, he first gains a payoff of 2, and he has a 50/50 chance of either returning to the initial state or moving to the absorbing state. The behavior-relevant segment of a strategy is represented by a sequence of E -s and S -s: $(a_T)_{T \in \mathbb{N}}$, with $a_T \in \{E, S\}$ as the action for the decision maker if he is in state \odot in period T .

Figure 5: Pointwise limit of a sequence of strategies is dominated by each strategy in the sequence on all horizons.



Suppose the discount factor is $\delta = 0.5$. It is easy to see that any strategy is optimal on the complete horizon, earning an expected utility of 2. Moreover, it can be shown that for any finite horizon T , the number of E -s in the strategy up to and including T determine the expected utility up to T , their order being irrelevant. In general, if a strategy s contains k number of E -s up to T , then its expected utility up to T is

$$U_T(s) = 2 + \frac{1}{2^{T-1}} \left(1 - \frac{1}{2^{k-1}} \right),$$

which is increasing in k . Take the following sequence of strategies: $\mathfrak{s} = (s_T)_{T \in \mathbb{N}}$, where

$$s_T = \underbrace{SS \dots S}_{T-1} \underbrace{EE \dots E}_T SSSS \dots$$

Fix any horizon \bar{T} . We then have

$$s_1 \prec_T s_2 \prec_T \dots \prec_T s_{\bar{T}}$$

for all T greater or equal to $2\bar{T} - 1$. Therefore,

$$s_1 \prec_{\text{LHS}} s_2 \prec_{\text{LHS}} \dots \prec_{\text{LHS}} s_{\bar{T}}.$$

Obviously, \mathfrak{s} has a single pointwise limit, namely $s = SSS \dots$ that prescribes staying in \odot forever. However, s is not favored to *any* strategy in the sequence on *any* horizon, since $s \prec_T s_1$ for all T . Thus, the limit strategy, our most natural choice for an upper bound of this chain lacks the desired property.

7. Discussion

First, it should be noted that our examples concern non-generic decision problems. We do not regard this as a serious drawback, as many interesting decision problems are non-generic. In the generic case, all the sets of optimal strategies defined in the previous sections collapse into a single one.

The existence of repeatedly optimal strategies remains an intriguing open question, although we conjecture that for subsets of the strategy space closed for pointwise limits, repeatedly optimal strategies do exist. 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 an overtakingly optimal strategy can be found, there is even more reason to rejoice: compared with every other strategy, any regret for choosing the overtakingly optimal strategy will fully dissipate after just finitely many periods.

The assumption of exponential discounting can be relaxed for most of our results. We can use a more general, history-dependent discount function, as long as it is continuous at infinity in the sense of Fudenberg and Levine (1983).

While our model is presented in a decision-theoretic framework, the game-theoretic extension of the concepts introduced should be 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

- Brock, W. A. (1970): An axiomatic basis for the Ramsey-Weizsäcker overtaking criterion. *Econometrica* 38/6, 927–929.
- Brock, W. A. and Haurie, A. (1976): On existence of overtaking optimal trajectories over an infinite time horizon. *Mathematics of Operations Research* 1/4, 337–346.
- Fudenberg, D. and Levine, D. (1983): Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory* 31, 251–268.
- Gale, D. (1967): On optimal development in a multisector economy. *The Review of Economic Studies* 34/1, 1–18.
- Ramsey, F. P. (1928): A mathematical theory of savings. *Economic Journal* 38/152, 543–559.
- Rubinstein, A. (1979): Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory* 21, 1–9.
- von Weizsäcker, C. C. (1965): Existence of optimal programs of accumulation for an infinite time horizon. *The Review of Economic Studies* 32/2, 85–104.