Essays in infinite dynamic games

Citation for published version (APA):

Maes, J. (2020). *Essays in infinite dynamic games*. [Doctoral Thesis, Maastricht University]. Maastricht University. <https://doi.org/10.26481/dis.20200702jm>

Document status and date:

Published: 01/01/2020

DOI:

[10.26481/dis.20200702jm](https://doi.org/10.26481/dis.20200702jm)

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.

Essays in Infinite Dynamic Games

© Jasmine Maes, Maastricht 2020

All rights reserved. No part of this publication may be reproduced, stored in an automated data system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior permission of the author.

This book was typeset by the author using L^AT_EX.

Essays in Infinite Dynamic Games

Dissertation

To obtain the degree of Doctor at Maastricht University
on the authority of the Rector Magnificus
Prof. Dr. Rianne M. Letschert,
in accordance with the decision of the Board of Deans,
to be defended in public on July 2, 2020 at 12:00 hours.

by

Jasmine Maes

Supervisor: Prof. Dr. Jean-Jacques Herings

Co Supervisors: Dr. János Flesch
Dr. Arkadi Predtetchinski

Assessment Committee: Prof. Dr. Frank Thuijsman (Maastricht University)
Prof. Dr. Nicolas Vieille (HEC Paris)
Dr. Galit Ashkenazi-Golan (Tel-Aviv University)
Dr. Andrés Perea (Maastricht University)

Acknowledgements

The goal of doing a PhD is to grow into an independent researcher. As this process requires more than the effort and persistence from the PhD candidate some acknowledgements are in order.

First and most importantly, there is the technical aspect of learning how to do research. This includes acquiring knowledge of the field of game, learning to ask good question and learning to present research in an appealing manner, both in written form as well as orally.

For this, I would like to thank my supervisors dr. János Flesch, dr. Arkadi Predtetchinski and Prof. dr. Jean-Jacques Herings for the supervision they provided but also for convincing me to start the PhD in the first place. Over the years they listened to my many worries and always tried their best at answering my questions. I have no doubt that their supervision pushed me to become a better and more independent researcher. One the may research skills they taught me was how to fail well. They showed me how to transform the many failures which one encounters in research into opportunities. Furthermore I wish to thank Prof. dr. Nicolas Vieille for hosting my research visit at HEC Paris. The discussions with him helped tremendously in development of the research questions addressed in Chapter 3.

Doing research is a rewarding experience. Few things come close to the excitement and joy one can feel when completing a proof or discovering a new theorem. However every coin has two sides and the excitement of discovering something new is often preceded by lengthy periods of being stuck and not knowing how to proceed. The social support during these periods was as important as the technical support.

Therefore I would like to award a special place in these acknowledgements to Anastas Tenev, with whom I shared an office for 3 years. During this time he has observed me celebrate my success and grieve my failures. During the difficult times he always encouraged me to persist and try again or try differently. He proved to be an excellent sparring partner who was always willing to listen to my numerous research ideas and provide critical feedback. Furthermore Anastas also showed to be a wonderful teacher. Not only did he help me improve my tutoring skills, he also introduced me to new topics in economics and to the field of linguistics. I will be eternally grateful for all help he gave me and for all the interesting discussions we had free of judgment and political correctness.

I thank Henrik Zaunbrecher for teaching me various ways in which people can be irrational. Moreover, I also thank him for organizing the many PhD lunches and being the PhD-representative of our department. His efforts contributed greatly to a better working environment for all current and future PhD students.

Giang Tran's and Frauke Stehr's efforts to organise departmental lunches and create a birthday calendar contributed to a more social work environment which went beyond

the group of PhD-students. I greatly admire their optimistic realism and cherish all the moments which we shared together, especially those which involved eating ice cream.

I thank Benoît Duvocell for the many French conversations we had. His patience and kindness gave me the self-confidence to speak French when needed. A skill which has proven extremely useful in my current work environment. Other skills I acquired during my PhD was how to play pool. This was taught to me by Péter Bayer, an excellent neighbour, friend and colleague. Furthermore, I would like to express my gratitude to Riccardo Saulle, Nickolas Gagnon, Aditya Aradhye, Max Brüning and all the others with whom I shared many nice moments.

Lastly, I would like to thank my parents who facilitated my time in Maastricht by providing logistical support whenever I needed it. They, among other things, helped me move to Maastricht to start my PhD adventure and they allowed to move back home when my contract ended to look for a new adventure.

Zwijnaarde, December 22, 2019.

Contents

Introduction	1
1 Subgame maxmin strategies in zero-sum stochastic games with tolerance levels	6
1.1 Introduction	6
1.2 Two-player zero-sum stochastic games	9
1.3 Subgame ϕ -maxmin strategies	10
1.4 Conditions for strategies to be subgame ϕ -maxmin	14
1.4.1 n -Day maxmin strategies and equalizing strategies	14
1.4.2 The case of an upper semi-continuous payoff function	18
1.5 Existence of subgame ϕ -maxmin strategies	19
1.5.1 Definition of the switching strategy σ^ϕ	20
1.5.2 Definition and analysis of finite switching strategies	21
1.5.3 Properties of a switching strategy σ^ϕ	26
1.6 Subgame maxmin strategies	28
1.6.1 The proof of Theorem 1.6.3	29
1.6.2 The one-shot game Υ_h	30
1.6.3 Construction of the tolerance function ϕ^*	32
1.6.4 The proof of Theorem 1.6.1	33
1.7 Discussion	34
2 Individual upper semicontinuity and subgame perfect ϵ-equilibria in games with almost perfect information	37
2.1 Introduction	37
2.2 The model	39
2.3 Examples	41
2.4 Existence of a subgame perfect ϵ -equilibrium	44
2.4.1 The existence result	44
2.4.2 Optimal plays	45
2.4.3 The proof of Theorem 2.4.1	47
2.5 The topology induced by i -convergence	49
2.5.1 The topological space $(P, \mathcal{T}_i)$	49
2.5.2 The topological space $(P, \cap_{i \in I} \mathcal{T}_i)$	55
2.6 Discussion	56
2.6.1 Perfect information games	56

2.6.2	Necessity of the conditions in Theorem 2.4.1	57
2.6.3	Payoff functions that are i -lower semicontinuous	57
2.6.4	Stochastic transitions	58
3	Decreasing Competition in Competitive Search: Unequal Outcomes with Equal Chances	60
3.1	Introduction	60
3.2	Model	62
3.3	Expected payoff functions and Best-response strategies	64
3.3.1	Expected payoff functions	64
3.3.2	Best-response thresholds	65
3.4	Subgame perfect equilibria	66
3.4.1	The two player game	68
3.4.2	The monotone subgame perfect equilibrium	72
3.4.3	Examples	78
3.5	Welfare Analysis	81
3.5.1	The social optimum	82
3.5.2	Welfare with a priority ranking	82
3.5.3	Welfare under the equal chances policy	85
3.6	Conclusion and open problems	88
3.6.1	Other subgame perfect equilibria	88
3.6.2	Understanding the equity-efficiency trade-off	89
3.6.3	Understanding the effect of the underlying distribution	89
3.6.4	Generalizing the preferences and beliefs of the players	89
	Appendix	90
A	Universal measurability	90
B	The proof of inequality (1.6.7)	92
C	Symmetric equilibria in behavioral strategies.	95
	Bibliography	96
	Samenvatting (Dutch Summary)	101
	Valorisation	103
	Biography	105

Introduction

The three chapters in this thesis all study infinite dynamic games to some degree of abstraction. Because they each address different research questions, the relation to existing literature is explained in the beginning of each of these Chapters. The purpose of this general introduction is threefold. First, it introduces the main concepts used in this thesis in an informal manner. Second, it explains the commonalities between the different chapters illustrating the coherence of the thesis. And third, it lays out a road map for the rest of the thesis by giving an intuitive summary of each of the chapters.

Concepts in infinite dynamic games

In this thesis we study infinite dynamic games. Such games model situations where players interact with each other multiple times. We focus on dynamic games with an infinite horizon which can be used to model situations with strategic interactions which have an indefinite ending.

The contexts in which players interact with each other are modelled by states. The collection of these states is referred to as the state space. The dynamic games studied in this thesis are played according to the following process: at every state players can choose actions from an actions space. We will often allow players to chose mixed actions, which are probability distributions over the action space. After all players have chosen their mixed action, the actions are realized and they, possibly combined with an element of chance, determine the next state. This process repeats itself indefinitely, creating an infinite sequence of actions and states called a play. This realized play determines the payoff for each of the players involved. At any point, the actions and states which have been realised before constitute the history of the game. A subgame with history h is the “smaller” game that has to be played after the history h is realized.

When the transition between states does depend on an element of chance as well as the realized history, we call the game a stochastic game. We study such a game in Chapter 1. The games studied in Chapter 2 and Chapter 3 can be viewed as special cases of a stochastic game. In Chapter 2 we assume the transition between states is completely determined by the realized history, while in Chapter 3 the transition to the next state only depends on the current state and the actions chosen therein. The actions players chose at any point in the game are described by their strategy. A strategy is called stationary if it only depends on the current state. In this thesis we will often allow for behavioral strategies, i.e. strategies in which at each state game players are allowed to play mixed actions.

Players aim to maximize their own payoff function and will choose a strategy accordingly. Players objectives are intertwined because the actions of other players can influence

ones own payoff. A special case of this is the zero-sum setting which we study in Chapter 1. Here players have completely opposing objectives. As the name “zero-sum” suggests, an increase in the payoff of one player is always accompanied by an equal and opposite decrease in other players’ payoffs. When a zero-sum game has only two players, it is sufficient to focus on one payoff function u which one player aims to maximize and the other player aims to minimize. Intuitively, the lower value $\underline{v}$ of a two player zero-sum stochastic game is the highest value the maximizing player can guarantee he will receive no matter the strategy used by the minimizing player. Similarly, the upper value $\bar{v}$ is the lowest value the minimizing player can guarantee he will have to pay no matter the strategy used by the maximizing player. In any zero-sum stochastic game we have that $\underline{v} \leq \bar{v}$. When the payoff function is Borel-measurable and the state space is countable, then we have that $\underline{v} = \bar{v}$ (Martin (1998) and Maitra and Sudderth (1998)). This common value is referred to as “the value of the zero-sum game”. Every subgame can have a different upper and lower value. The upper and lower value are denoted with $\bar{v}(h)$ and $\underline{v}(h)$ respectively, where h is the history determining the subgame. If $\bar{v}(h) = \underline{v}(h)$, then $v(h)$ denotes the value of the subgame with history h . In dynamic two player zero-sum games one is often interested in finding the best strategy a player could use. Such strategies are called subgame optimal strategies and they are able to guarantee the value in each of the subgames.

In case the maximising behavior of one player is not necessarily detrimental to the other players’ payoffs, we need a different solution concept. A central solution concept in dynamic games is given by the notion of subgame perfect equilibrium, first introduced by Selten (1965) as a refinement of the well-known Nash equilibrium. A subgame perfect equilibrium is defined as a strategy profile such that in every subgame no player wants to change his strategy given that no other player changes his strategy. Expressed differently: there does not exist a subgame in which there is a player who has a profitable unilateral deviation. In this thesis we also use the more permissive notion of subgame perfect ϵ -equilibrium, which assumes that players do not deviate in any subgame if the corresponding increase in expected payoff is at most ϵ .

In finite horizon dynamic games a subgame perfect equilibrium can easily be constructed using backwards induction. This technique requires that one starts by first solving the terminal subgames of the larger game, this partial solution yields an expected payoff for each of the players if they were to reach this subgame. Using this solution one can then solve all penultimate subgames, continuing this process one will eventually construct an equilibrium solution to the entire game.

When dealing with infinite dynamic games it is not immediately clear how to construct a subgame perfect equilibrium as there may not be terminal subgames. Even the existence of a subgame perfect equilibrium or a subgame perfect ϵ -equilibrium may be problematic in these games. This is especially true when players have payoff functions which are discontinuous at infinity, meaning that actions in a very distant future can have a substantial impact on the payoffs player receive.

Infinite dynamic games can be studied to varying degrees of abstraction. In this thesis the chapters are organised in a decreasing order of abstraction. Each of these chapters illustrates an important trade-off present in all mathematical models: the more assumptions made in the model, the stronger the results one can derive, but also the narrower the scope of applicability. Given this trade-off it is therefore an important task of the researcher to select the appropriate level of abstraction for the given research question.

Zero-sum stochastic games and tolerance functions

In this chapter we study a two-player zero-sum dynamic game with finite action spaces and a countable state space. The assumptions made about the payoff-function are minimal, as we only require the function to be universally measurable. Consequently, we cannot guarantee the existence of the value and of subgame optimal strategies. Instead, we will take the perspective of the maximizing player and focus on creating strategies for this player that perform sufficiently well throughout the game. We call such a strategy a subgame ϕ -maxmin strategy. Here ϕ is a function which assigns a non-negative tolerance level to every subgame. More formally, a subgame ϕ maxmin strategy is a strategy of the maximizing player which at a subgame defined by a history h guarantees an expected payoff of $\underline{v}(h) - \phi(h)$.

We make the following contributions: First, we provide necessary and sufficient conditions for a strategy to be a subgame ϕ -maxmin strategy. As a special case we obtain a characterization for subgame maxmin strategies, i.e. strategies that exactly guarantee the lower value at every subgame. Secondly, we present sufficient conditions for the existence of a subgame ϕ -maxmin strategy. In particular we find that a subgame ϕ -maxmin strategy may not always exist for a fast decreasing tolerance function ϕ . Finally, we show the possibly surprising result that each game admits a strictly positive tolerance function ϕ^* with the following property: if a player has a subgame ϕ^* -maxmin strategy then he has a subgame maxmin strategy too. As a consequence, the existence of a subgame ϕ -maxmin strategy for every positive tolerance function ϕ is equivalent to the existence of a subgame maxmin strategy.

Individual upper semicontinuity and the existence of subgame perfect ϵ equilibrium

In Chapter 2 we study games with almost perfect information and an infinite time horizon. In such games, the players simultaneously choose actions from finite action sets at each stage, knowing the actions chosen at all previous stages. The payoff of each player is a function of all actions chosen during the game. We introduce the topological notion of individual upper semicontinuity and all players have an individually upper semi-continuous payoff function.

This notion of individual upper semicontinuity generalises the following more familiar notion of upper semicontinuity. A payoff function u_i for player i is called upper semi-continuous if for every sequence of plays $(p_t)_{t \in \mathbb{N}}$ that converges to the play p^1 we have that $\limsup_{t \rightarrow \infty} u_i(p_t) \leq u_i(p)$.

We can generalize this to the notion of individual upper semi-continuity by only considering sequences of plays in which player i is the only player who first deviates from the play p . To do this we introduce the notion of i -convergence. We say a sequence of plays $(p_t)_{t \in \mathbb{N}}$ i -converges to a play p if the length of the common history between the play p and the plays p_t goes to infinity the further we proceed in the sequence $(p_t)_{t \in \mathbb{N}}$ and if eventually player i is the first player who deviates from the play p . Hence there are fewer sequences

¹A sequence of plays $(p_t)_{t \in \mathbb{N}}$ is said to converge to the play p if the length of the common history between p_t and p grows to infinity when proceeding through the sequence of plays.

which i -converge to the play p than there are sequences that converge to it. Using this concept of i -convergence we can then define the notion of i -upper semi continuity. A payoff u_i of player i is i -upper semicontinuous if for any play p and for all the sequences $(p_t)_{t \in \mathbb{N}}$ that i -converge to the this play p , we have that $\limsup_{t \rightarrow \infty} u_i(p_t) \leq u_i(p)$.

From a game theoretical point of view such payoff functions are interesting because of the following property: if every player has an individually upper semicontinuous payoff function then for any play there exists a time after which all players can agree on following this play.

Using this property we can then apply the technique of backwards induction to construct a subgame perfect ϵ -equilibrium for any $\epsilon > 0$. The constructed subgame perfect ϵ -equilibrium has the property that eventually all players only play pure actions.

Decreasing competition in competitive search

In Chapter 3 we study a specific infinite dynamic game to better understand the effects of decreasing competition in a competitive search model.

More specifically we study the following infinite horizon game. There is a group of players who all want to get precisely one item. When the game starts an item is revealed and all player can either apply for this item or wait for a potentially better item. The item can be allocated to at most one player. If there is precisely one player who wants the item then the item is given to this player. If the item has multiple applicants then the item is allocated to each of the applicants with equal probability. If none of the players apply for the item then the item disappears. The player who got the item leaves the game and all other players continue to the next period where a new item is revealed and the process repeats itself. The game continues until all players have an item. The payoff of each player is the discounted value of the item they obtained. We assume that all players are identical, implying that they are equally impatient and have the same valuations for all items. Furthermore, we assume that the valuation of each item is independently drawn from the same continuous distribution. The fact that all players can acquire at most one item and leave the game as soon as they have it, leads to a shrinking pool of competing players. By analysing this game we obtain the following insights. Firstly, from a fairness point of view, even if all players are identical and are given equal chances, inequality in ex ante expected outcome may still occur in equilibrium. To illustrate this we construct a subgame perfect equilibrium in which all players have different expected equilibrium payoffs and maintain their relative rank in terms of expected payoffs throughout the game. Secondly, from an efficiency point of view, giving equal chances to all players always leads to a welfare loss. This contrasts with the socially optimal policy of imposing an exogenous priority ranking on the players.

Chapter 1

Subgame maxmin strategies in zero-sum stochastic games with tolerance levels

1.1 Introduction

Two-player zero-sum stochastic games model the repeated interaction between two agents with opposite objectives. The environment in which the interaction takes place is fully characterized by a state variable. The transition from one state variable to the next one is influenced by both players as well as an element of chance. Throughout the paper we take the perspective of the maximizing player. We are interested in strategies of the maximizing player that guarantee the lower value at every subgame and call such strategies subgame maxmin strategies. Under the assumptions as made in the paper, the value may not exist, which explains why we consider the lower value instead.

As the name subgame maxmin strategy suggests, this concept is closely related to the concept of a subgame perfect equilibrium as defined in Selten (1965). In two-player zero-sum games where the value exists, for conditions see Maitra and Sudderth (1998) and Martin (1998), the notions of a subgame maxmin strategy and a subgame minmax strategy coincide with the notion of a subgame optimal strategy. Moreover, in such games a strategy profile is a subgame perfect equilibrium if and only if it consists of a pair of subgame optimal strategies.

As illustrated by the Big Match, a game introduced in Gillette (1957) and analyzed in Blackwell and Ferguson (1968), even if the value exists, it is not guaranteed that optimal strategies exist, so a fortiori, subgame optimal strategies and subgame perfect equilibria may not exist. A large part of the literature therefore focuses on so-called subgame perfect ϵ -equilibria as defined in Radner (1980). This equilibrium concept is more permissive than a subgame perfect equilibrium and consists of a strategy pair such that every player obtains the value at each history up to a fixed error term of $\epsilon/2$.

Instead of having a fixed error term at each subgame, we allow the error term to vary across different subgames. This error term is expressed as a function ϕ of the histories and is called the tolerance function. The central topic of this paper is the concept of a subgame ϕ -maxmin strategy. This is a strategy of the maximizing player that guarantees the lower value at every subgame within the allowed tolerance level. Intuitively, a subgame ϕ -maxmin strategy performs sufficiently well across all subgames. This type of strategy is related to the concept of ϕ -tolerance equilibrium as defined in Flesch and Predtetchinski (2016). Indeed, if the value exists, then a strategy profile in which both players use a

subgame $(\phi/2)$ -optimal strategy is a ϕ -tolerance equilibrium.

One motivation for letting the tolerance level vary across subgames is given by Mailath, Postlewaite, and Samuelson (2005) when introducing the concept of a contemporaneous perfect ϵ -equilibrium. The authors focus on games in which the payoff function of the players is given by the discounted sum of periodic rewards. Due to this discounting, there exists a period after which the maximal total discounted reward a player can receive is smaller than ϵ . If the allowed tolerance level ϵ is fixed across all subgames, any strategy will be an ϵ -maxmin strategy of a subgame in such a period. Therefore, the concept of subgame ϵ -maxmin strategy does not impose any restrictions on the actions chosen at a very distant future. The issue here is that it would be more intuitive to discount not only the reward but also the allowed tolerance level.

Additional motivation for letting the tolerance level vary across subgames stems from the fact that the notion of what is considered sufficiently good might be relative. For instance, Tversky and Kahneman (1985) observe that people evaluate decisions with respect to a reference level. They find that significantly more people were willing to exert extra effort to save \$5 on a \$15 purchase than to save \$5 on a \$125 purchase. To apply this to the context of zero-sum games, consider the following game to which we will refer as the high stakes-low stakes game. In the first stage of the game, a chance move determines whether the player will engage in the high stakes or the low stakes variant of this game. The high stakes and the low stakes games are identical in terms of possible strategies. The only difference is that the payoffs in the high stakes game are a thousand fold the payoffs in the low stakes game. Furthermore, assume that in the high stakes subgame the payoff of player 1 ranges between 0 and 2000 and the value is 1000, while in the low stakes subgame the payoff of player 1 ranges between 0 and 2 and the value is 1. Players that evaluate decisions with respect to a reference level, may label a strategy which guarantees a payoff of 999 in the high stakes game as sufficiently good. However, in the low stakes game, 0 corresponds to the minimum payoff. Allowing the tolerance level to vary across subgames can therefore be used to accommodate such behavioral effects into the model of zero-sum stochastic games.

Another case where history dependent tolerance levels are natural is the following. In situations that commonly occur, a player may use a familiar strategy irrespective of the scale of the payoffs. To understand this better, imagine a player who is highly experienced in playing a certain zero-sum game. He has a trusted strategy which guarantees him the value of this game within some error. Now consider the high stakes-low stakes game again. The player might well use the trusted strategy in both the low stakes and the high stakes subgame. Therefore the error related to this strategy will be relative with respect to the value of the respective subgame.

Finally, in the class of stochastic games as identified in Flesch, Thuijsman and Vrieze (1998), the only way to obtain ϵ -optimality is to use strategies that are called improving. Improving strategies are closely related to subgame ϕ -maxmin strategies such that the tolerance level in some subgames is smaller than the tolerance level at the root.

With respect to the concept of subgame ϕ -maxmin strategies, this paper attempts to provide answers to the following two fundamental questions:

1. For positive tolerance functions ϕ , when does a subgame ϕ -maxmin strategy exist?
2. How is the existence of a subgame ϕ -maxmin strategy related to the existence of a subgame maxmin strategy?

To answer these questions we will first derive necessary and sufficient conditions for a strategy to be a subgame ϕ -maxmin strategy. This is done in Section 1.4. As a special case of these necessary and sufficient conditions, we obtain a characterization of subgame maxmin strategies. For the special class of positive and negative stochastic games, a related characterization of subgame maxmin strategies was obtained by Flesch, Predtetchinski and Sudderth (2018).

The necessary and sufficient conditions for strategies to be subgame ϕ -maxmin can be split into a local condition and an equalizing condition. Informally, the local condition states that the lower value one expects to get in the next subgame should always be at least the lower value of the current subgame. The equalizing condition requires that, for every strategy of the other player, a subgame ϕ -maxmin strategy almost surely results in a play with an eventually good enough payoff, where eventually good enough means being at least the lower value in very deep subgames up to the allowed tolerance level.

In Section 1.5 we consider existence of subgame ϕ -maxmin strategies for positive tolerance functions ϕ . The existence and construction of such strategies is important as they can serve as punishment strategies in multi-player games. This is illustrated in the paper of Mashiah-Yaakovi (2015).

We prove that for a positive tolerance function ϕ , a subgame ϕ -maxmin strategy exists if every play is either a point of upper semicontinuity of the payoff function or if the sequence of tolerance levels which occur along the play has a positive lower bound. We give a constructive proof of this statement using the sufficient conditions for a strategy to be subgame ϕ -maxmin.

A special case of our theorem, where the sequence of tolerance levels which occur along the play always has a positive lower bound, has been studied in Mashiah-Yaakovi (2015). In Proposition 11 of that paper, the existence of a subgame ϵ -optimal strategy in a two-player zero-sum stochastic game with Borel measurable payoff functions, finite action sets, and a countable state space has been shown. A subgame ϵ -optimal strategy corresponds to a constant tolerance function that is everywhere equal to ϵ .

In Section 1.6 we study the relation between the existence of subgame ϕ -maxmin strategies and subgame maxmin strategies. Our main result states that the existence of a particular subgame ϕ^* -maxmin strategy, with $\phi^* > 0$, is equivalent to the existence of a subgame maxmin strategy. Consequently, the existence of a subgame ϕ -maxmin strategies for all positive tolerance functions ϕ is equivalent to the existence of a subgame maxmin strategy. For upper semi-continuous payoff functions, our theorem in Section 1.5 guarantees the existence of a subgame ϕ -maxmin strategy for every positive tolerance function ϕ , so it follows that a subgame maxmin strategy exists if the payoff function is upper semi-continuous. The latter conclusion is related to a result in Laraki, Maitra and Sudderth (2013).

The connection between existence of subgame ϕ -maxmin strategies for every positive tolerance function ϕ and the existence of subgame maxmin strategies is not only useful to further understand the results obtained by Laraki, Maitra and Sudderth (2013) but also highlights an important and surprising difference between subgame ϕ -maxmin strategies and the closely related concept of subgame ϵ -maxmin strategies. Indeed, the existence of a subgame ϵ -maxmin strategy for every $\epsilon > 0$ does not imply the existence of a subgame maxmin strategy.

The rest of the paper is structured as follows. In Section 1.2 we formulate the model setup and in Section 1.3 we formally define the main concepts. Then in Section 1.4 we discuss the necessary and sufficient conditions for a strategy to be a subgame ϕ -maxmin strat-

egy and give a characterization for subgame maxmin strategies. We continue in Section 1.5 by providing sufficient conditions to guarantee the existence of a subgame ϕ -maxmin strategy. Then in Section 1.6 we explain why the existence of subgame ϕ -maxmin strategies for every positive tolerance function ϕ is equivalent to existence of a subgame maxmin strategy. Finally, in Section 1.7 we discuss the importance of the assumptions we made and whether they might be relaxed.

1.2 Two-player zero-sum stochastic games

We consider a two-player zero-sum stochastic game with finitely many actions and countably many states. The payoff function is required to be bounded and universally measurable. The model encompasses all two-player zero-sum games with perfect information and stochastic moves.

Actions, states, and histories: The action sets of players 1 and 2 are given by the finite sets $\mathcal{A}$ and $\mathcal{B}$, respectively. The state space is given by the countable set $\mathcal{X}$. Let x_0 denote the initial state and define the set $\mathcal{Z} = \mathcal{A} \times \mathcal{B} \times \mathcal{X}$. The game consists of an infinite sequence of one-shot games. At the initial state x_0 , the one-shot game $G(x_0)$ is played as follows: Players 1 and 2 simultaneously select an action from their respective action sets, denoted by a_1 and b_1 , respectively. Then the next state x_1 is selected according to the transition probability $q(\cdot | x_0, a_1, b_1)$. At the new state x_1 , this process repeats itself and both players play the one-shot game $G(x_1)$ by selecting actions a_2 and b_2 from their respective action sets. The next state x_2 is selected according to the transition probability $q(\cdot | x_0, a_1, b_1, x_1, a_2, b_2)$. This goes on indefinitely and creates a play $p = (x_0, a_1, b_1, x_1, a_2, b_2, \dots)$. Note that the transition probability can depend on the entire history.

For every $t \in \mathbb{N} = \{0, 1, 2, \dots\}$, let $\mathcal{H}^t = \{x_0\} \times \mathcal{Z}^t$ denote the set of all histories that are generated after t one-shot games. The set $\mathcal{H}^0$ consists of the single element x_0 . For $t \geq 1$, elements of $\mathcal{H}^t$ are of the form $(x_0, a_1, b_1, \dots, a_t, b_t, x_t)$. Let $\mathcal{H} = \cup_{t \in \mathbb{N}} \mathcal{H}^t$ denote the set of all histories. For all $h \in \mathcal{H}$, let $\|h\| = \|(x_0, a_1, b_1, \dots, a_t, b_t, x_t)\| = t$ denote the number of one-shot games that occurred during the history h . We refer to $\|h\|$ as the length of the history h . For all $t \leq \|h\|$, the restriction of the history h to the first t one-shot games is denoted by $h|_t = (x_0, a_1, b_1, \dots, a_t, b_t, x_t)$. We write $h' \leq h$ if there exists $t \leq \|h\|$ such that $h|_t = h'$, so if the history h extends the history h' .

The space of plays: Define $\mathcal{P} = \{x_0\} \times \mathcal{Z}^{\mathbb{N}}$ to be the set of plays. Elements of $\mathcal{P}$ are of the form $p = (x_0, a_1, b_1, x_1, a_2, b_2, \dots)$. For $t \in \mathbb{N}$, let $p|_t$ denote the prefix of p of length t : that is $p|_0 = x_0$ and $p|_t = (x_0, a_1, b_1, \dots, a_t, b_t, x_t)$ for $t \geq 1$. We write $h < p$ if a history h is a prefix of p . For every $h \in \mathcal{H}$, let $\mathcal{P}(h) = \{p \in \mathcal{P} | h < p\}$ denote the cylinder set consisting of all plays which extend history h .

We endow $\mathcal{Z}$ with the discrete topology and $\mathcal{P}$ with the product topology. The collection of all cylinder sets is a basis for the product topology on $\mathcal{P}$.

For $t \in \mathbb{N}$, let $\mathcal{F}^t$ be the sigma-algebra generated by the collection of cylinder sets $\{\mathcal{P}(h) | h \in \mathcal{H}^t\}$. Each set in $\mathcal{F}^t$ can be written as the union of sets in $\{\mathcal{P}(h) | h \in \mathcal{H}^t\}$. Let $\mathcal{F}^{\infty}$ be the sigma-algebra generated by $\cup_{t \in \mathbb{N}} \mathcal{F}^t$. This is exactly the Borel sigma-algebra generated by the product topology on $\mathcal{P}$. The sigma-algebra of universally measurable subsets of $\mathcal{P}$ is denoted by $\mathcal{F}$. Elements of $\mathcal{F}$ are sets that belong to the completion of each Borel probability measure on $\mathcal{P}$. For details of the definition of the sigma-algebra $\mathcal{F}$, the reader is referred

to Appendix A. It holds that $\mathcal{F}^t \subseteq \mathcal{F}^{t+1} \subseteq \dots \subseteq \mathcal{F}^\infty \subseteq \mathcal{F}$. The set of plays $\mathcal{P}$ together with the universally measurable sigma-algebra $\mathcal{F}$ determines a measurable space $(\mathcal{P}, \mathcal{F})$. A stochastic variable is a universally measurable function from $\mathcal{P}$ to $\mathbb{R}$.

Strategies: Let $\Delta(\mathcal{A})$ denote the set of probability measures over the action set of player 1 and $\Delta(\mathcal{B})$ the set of probability measures over the action set of player 2. A behavioral strategy for player 1 is a function $\sigma : \mathcal{H} \rightarrow \Delta(\mathcal{A})$. Hence, at each history player 1 chooses a mixed action. A pure strategy for player 1 is a function that assigns to every history an action, with a minor abuse of notation, $\sigma : \mathcal{H} \rightarrow \mathcal{A}$. Similarly, one can define a behavioral and a pure strategy τ for player 2. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ denote the sets of behavioral strategies of players 1 and 2, respectively.

It follows from the Ionescu Tulcea extension theorem that every history $h \in \mathcal{H}$ and strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$ determine a probability measure $\mathbb{P}_{h,\sigma,\tau}$ on the measurable space $(\mathcal{P}(h), \mathcal{F}_{\mathcal{P}(h)}^\infty)$, where $\mathcal{F}_{\mathcal{P}(h)}^\infty$ denotes the Borel sigma-algebra over the set of plays extending the history h . The measure $\mathbb{P}_{h,\sigma,\tau}$ can be extended to the measurable space $(\mathcal{P}, \mathcal{F}^\infty)$ in the obvious way. Taking the completion of the probability space $(\mathcal{P}, \mathcal{F}^\infty, \mathbb{P}_{h,\sigma,\tau})$ yields the probability space $(\mathcal{P}, \mathcal{F}, \mathbb{P}_{h,\sigma,\tau}^C)$. With a minor abuse of notation, we will omit the superscript C and write $\mathbb{P}_{h,\sigma,\tau}$ in the remainder of this paper.

Payoff function: We assume that the payoff function $u : \mathcal{P} \rightarrow \mathbb{R}$ of player 1 is bounded and universally measurable. We also assume the game to be zero-sum. The payoff function of player 2 is therefore given by $-u$. We denote the game as described above by $\Gamma_{x_0}(u)$. Throughout the paper we take the point of view of player 1. This is without loss of generality, since the situation of Player 2 in the game $\Gamma_{x_0}(u)$ is identical to that of Player 1 in the game $\Gamma_{x_0}(-u)$.

The expected payoff of player 1 corresponding to strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$ at history $h \in \mathcal{H}$ is given by $\mathbb{E}_{h,\sigma,\tau}[u]$, where the expectation is taken with respect to the probability measure $\mathbb{P}_{h,\sigma,\tau}$. The expected payoff of player 1 at the history x_0 is denoted by $\mathbb{E}_{\sigma,\tau}[u]$.

Unlike two-player zero-sum stochastic games with Borel measurable payoff functions, two-player zero-sum stochastic games with universally measurable payoff functions do not necessarily have a value, formally defined in Section 1.3. The core idea of this paper, the construction and existence of strategies that perform sufficiently well in every subgame, is independent of the problem of the existence of a value. The reader unfamiliar with universally measurable payoff functions may therefore imagine Borel measurable payoff functions throughout the paper.

1.3 Subgame ϕ -maxmin strategies

For every game $\Gamma_{x_0}(u)$, for every history $h \in \mathcal{H}$, we define the lower value $\underline{v}(h)$ and the upper value $\bar{v}(h)$ by

$$\underline{v}(h) = \sup_{\sigma \in \mathcal{S}_1} \inf_{\tau \in \mathcal{S}_2} \mathbb{E}_{h,\sigma,\tau}[u], \quad (1.3.1)$$

$$\bar{v}(h) = \inf_{\tau \in \mathcal{S}_2} \sup_{\sigma \in \mathcal{S}_1} \mathbb{E}_{h,\sigma,\tau}[u]. \quad (1.3.2)$$

Because the payoff function u is assumed to be bounded, we have that $\underline{v}(h), \bar{v}(h) \in \mathbb{R}$. Therefore, the lower and upper value exist in every subgame of $\Gamma_{x_0}(u)$. Furthermore, we have that

Figure 1.1: Characterization of ϕ such that pure subgame ϕ -maxmin strategies exist.

$\underline{v}(h) \leq \bar{v}(h)$. Whenever $\underline{v}(h) = \bar{v}(h)$ we say that the subgame at history h has a value and we denote it by $v(h)$. The lower value, the upper value, and the value at the initial state x_0 are denoted by $\underline{v}$, $\bar{v}$, and v , respectively. If u is Borel measurable, then the value exists by Maitra and Sudderth (1998) and Martin (1998). Since we do not assume u to be Borel measurable, we present our results in terms of the lower value.

Even if the value exists, player 1 may not have a strategy that guarantees the value in each subgame and the literature has therefore studied subgame ϵ -optimal strategies. These are strategies which guarantee the value in each subgame up to an allowed error term of ϵ . If the payoff function is bounded and Borel measurable, it has been shown by Mashiah-Yaakovi (2015) that for each $\epsilon > 0$ player 1 has a subgame ϵ -optimal strategy. The concept of a subgame ϵ -optimal strategy has a constant error term ϵ across all subgames. However, as argued in the introduction, there are instances in which it is more natural to consider a variable error term. Therefore, instead of considering a constant error term, we follow Flesch and Predtetchinski (2016), who allow the error term to vary across histories in their investigation of the ϕ -tolerance equilibrium. This leads us to the study of subgame ϕ -maxmin strategies, where $\phi : \mathcal{H} \rightarrow [0, \infty)$ is a tolerance function assigning an allowed tolerance level to each history.

Definition 1.3.1. Let $\phi : \mathcal{H} \rightarrow [0, \infty)$ be a tolerance function. A strategy $\sigma \in \mathcal{S}_1$ is a *subgame ϕ -maxmin strategy* in the game $\Gamma_{x_0}(u)$ if for every history $h \in \mathcal{H}$ it holds that

$$\forall \tau \in \mathcal{S}_2, \mathbb{E}_{h, \sigma, \tau} [u] \geq \underline{v}(h) - \phi(h). \quad (1.3.3)$$

In case ϕ is identically equal to zero, we omit it from the notation, and simply refer to a subgame maxmin strategy.

A subgame ϕ -maxmin strategy guarantees at each history h of the game the lower value up to the tolerance level $\phi(h)$. If the tolerance function is such that, for some $\epsilon \geq 0$, for every $h \in \mathcal{H}$, $\phi(h) = \epsilon$, then we refer to the strategy as a subgame ϵ -maxmin strategy. If the value exists, then the notion of subgame ϵ -maxmin strategy coincides with the notion of subgame ϵ -optimal strategy.

The following example illustrates that even for a strictly positive tolerance function ϕ player 1 may have no subgame ϕ -maxmin strategies. Interestingly, however, player 1 has a subgame ϵ -maxmin strategy for every positive $\epsilon > 0$.

Example 1.3.2. The decision problem depicted in Figure 1.1 corresponds to a two-player zero-sum stochastic game in which the state space is trivial and the second player is a dummy player. Whenever the state space or action sets are degenerate, the corresponding states and actions are omitted from the notation in examples. The set of actions of player 1 is $\mathcal{A} = \{c, q\}$, where c stands for continue and q for quit. The game stops as soon as player 1

chooses to quit. If player 1 decides to quit at period t , then his payoff is $t/(t+1)$. If player 1 never quits, his payoff is 0.

In this game, player 1 has a subgame ϵ -maxmin strategy for every positive $\epsilon > 0$. For example, the strategy which quits whenever quitting leads to a payoff of at least $1 - \epsilon$.

We now turn to the existence of a subgame ϕ -maxmin strategy. Clearly, there exist no subgame maxmin strategy. As we will see later, Theorem 1.6.1 then implies that there is some strictly positive tolerance function ϕ for which there does not exist a subgame ϕ -maxmin strategy. Intuitively, such a tolerance function forces player 1 to continue with such a large probability that the total probability of never quitting becomes nearly one.

In the remainder of this example we focus on pure strategies and we provide a characterization of tolerance functions ϕ for which there is a pure subgame ϕ -maxmin strategy.

CLAIM: *There exists a pure subgame ϕ -maxmin strategy if and only if*

1. *for every $t \in \mathbb{N}$, $\phi(c^t) > 0$,*
2. *for infinitely many $t \in \mathbb{N}$, $\phi'(c^t) = \min_{n \leq t} \phi(c^n) \geq \frac{1}{t+1}$.*

PROOF: For every $t \in \mathbb{N}$, the value $v(c^t)$ exists and is equal to 1. Hence any pure subgame ϕ -maxmin strategy σ has the property that, for every $t \in \mathbb{N}$, $u(\pi(\sigma, c^t)) \geq 1 - \phi(c^t)$, where $\pi(\sigma, c^t)$ denotes the play induced from history c^t when using strategy σ .

$\Rightarrow$ Because a payoff of exactly 1 can never be reached it is clear that $\phi(c^t) > 0$ for every $t \in \mathbb{N}$.

Let σ be a pure subgame ϕ -maxmin strategy. We distinguish three cases.

CASE 1: For every $t \in \mathbb{N}$, $\sigma(c^t) = c$. For every $t \in \mathbb{N}$ it holds that $u(\pi(\sigma, c^t)) = u(c^\infty) = 0$. Because σ is a subgame ϕ -maxmin strategy, we find that, for every $t \in \mathbb{N}$, $\phi(c^t) \geq 1$, so $\phi'(c^t) = \min_{n \leq t} \phi(c^n) \geq 1 \geq 1/(t+1)$.

CASE 2: The number of $t \in \mathbb{N}$ such that $\sigma(c^t) = q$ is positive and finite. Consider the increasing sequence of times $(t_k)_{k=0, \dots, k'}$ at which $\sigma(c^{t_k}) = q$. For $t \in \{0, \dots, t_0\}$, we have that

$$u(\pi(\sigma, c^t)) = u(c^{t_0}q) = \frac{t_0}{t_0+1},$$

so $\phi(c^t) \geq 1/(t_0+1)$ since σ is a subgame ϕ -maxmin strategy. We find that

$$\phi'(t_0) = \min_{t \in \{0, \dots, t_0\}} \phi(t) \geq \frac{1}{t_0+1}.$$

We argue next that if, for some $k \geq 1$, t_{k-1} and t_k are quitting times, then

$\min_{t \in \{t_{k-1}+1, \dots, t_k\}} \phi(c^t) \geq 1/(t_k+1)$. Indeed, for $t \in \{t_{k-1}+1, \dots, t_k\}$ we have that

$$u(\pi(\sigma, c^t)) = u(c^{t_k}q) = \frac{t_k}{t_k+1},$$

so $\phi(c^t) \geq 1/(t_k+1)$ since σ is a subgame ϕ -maxmin strategy. Using induction, we find for $k = 1, \dots, k'$ that

$$\phi'(c^{t_k}) \geq \min\{\phi'(c^{t_{k-1}}), \frac{1}{t_k+1}\} \geq \frac{1}{t_k+1}.$$

For every $t > t^{k'}$ it holds that $\sigma(c^t) = c$ and $u(\pi(\sigma, c^t)) = u(c^\infty) = 0$. For every $t > t^{k'}$, since σ is a subgame ϕ -maxmin strategy, we have $\phi(c^t) \geq 1$, so

$$\phi'(c^t) \geq \min\{\phi'(c^{t^{k'}}), 1\} \geq \frac{1}{t^{k'}+1} > \frac{1}{t+1},$$

which concludes this case.

CASE 3: The number of $t \in \mathbb{N}$ such that $\sigma(c^t) = q$ is infinite. Consider the increasing sequence of times $(t_k)_{k \in \mathbb{N}}$ at which $\sigma(c^{t_k}) = q$. As in Case 2 it can be shown that for every $k \in \mathbb{N}$ it holds that $\phi'(c^{t_k}) \geq 1/(t_k + 1)$.

$\Leftarrow$ Let the strategy σ be defined as follows. For $t \in \mathbb{N}$,

$$\sigma(c^t) = \begin{cases} q, & \text{if } \phi'(c^t) \geq \frac{1}{t+1}, \\ c, & \text{otherwise.} \end{cases} \quad (1.3.4)$$

We show first that σ is a subgame ϕ' -maxmin strategy. For every $t \in \mathbb{N}$, there exists $t' \geq t$ such that $\phi'(c^{t'}) \geq 1/(t'+1)$. Take the minimal t' with this property. We have that $u(\pi(\sigma, c^t)) = u(c^{t'}q) = t'/(t'+1)$. Because ϕ' is a non-increasing function, we have that

$$1 - \phi'(c^t) \leq 1 - \phi'(c^{t'}) \leq 1 - \frac{1}{t'+1} = \frac{t'}{t'+1} = u(\pi(\sigma, c^t)).$$

We conclude that σ is a subgame ϕ' -maxmin strategy. Because, for every $t \in \mathbb{N}$, $\phi'(c^t) \leq \phi(c^t)$, the strategy σ is a subgame ϕ -maxmin strategy as well. $\diamond$

To identify subgame ϕ -maxmin strategies, it is useful to define the function $\underline{u} : \mathcal{S}_1 \times \mathcal{H} \rightarrow \mathbb{R}$ by

$$\underline{u}(\sigma, h) = \inf_{\tau \in \mathcal{S}_2} \mathbb{E}_{h, \sigma, \tau} [u]. \quad (1.3.5)$$

The payoff $\underline{u}(\sigma, h)$ corresponds to the guarantee level that player 1 is expected to receive at history h when playing the strategy σ . A strategy $\sigma \in \mathcal{S}_1$ is called a $\phi(h)$ -maxmin strategy for the subgame at history h if $\underline{u}(\sigma, h) \geq \underline{v}(h) - \phi(h)$.

For every strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$, for every $t \in \mathbb{N}$, define the stochastic variables $U_{\sigma, \tau}^t$, $\underline{U}_\sigma^t$, and $\underline{V}^t$ by letting $U_{\sigma, \tau}^t(p) = \mathbb{E}_{p|t, \sigma, \tau} [u]$, $\underline{U}_\sigma^t(p) = \underline{u}(\sigma, p|t)$, and $\underline{V}^t(p) = \underline{v}(p|t)$, respectively, for each play $p \in \mathcal{P}$. All three stochastic variables are $\mathcal{F}^t$ -measurable. We have $\underline{U}_\sigma^t \leq U_{\sigma, \tau}^t$ and $\underline{U}_\sigma^t \leq \underline{V}^t$ everywhere on $\mathcal{P}$.

The next lemma states the submartingale property of guarantee levels. It says that the guarantee level that player 1 can expect to receive increases over time.

Lemma 1.3.3. (Submartingale property of guarantee levels) Let a strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$, $t \in \mathbb{N}$, and a history $h \in \mathcal{H}^t$ of length t be given.

[1] It holds that $\underline{u}(\sigma, h) \leq \mathbb{E}_{h, \sigma, \tau} [U_\sigma^{t+1}]$.

[2] The process $(\underline{U}_\sigma^{t+n})_{n \in \mathbb{N}}$ is a $\mathbb{P}_{h, \sigma, \tau}$ -submartingale.

Proof. Take $\delta > 0$. Let $\tau' \in \mathcal{S}_2$ be such that $\tau'(h) = \tau(h)$ and for each $(a, b, x) \in \mathcal{Z}$ it holds that

$$\mathbb{E}_{(h, a, b, x), \sigma, \tau'} [u] \leq \underline{u}(\sigma, (h, a, b, x)) + \delta.$$

We have that

$$\begin{aligned} \underline{u}(\sigma, h) &\leq \mathbb{E}_{h, \sigma, \tau'} [u] \\ &= \sum_{(a, b, x) \in \mathcal{Z}} \sigma(h)(a) \cdot \tau(h)(b) \cdot q(x|h, a, b) \cdot \mathbb{E}_{(h, a, b, x), \sigma, \tau'} [u] \\ &\leq \sum_{(a, b, x) \in \mathcal{Z}} \sigma(h)(a) \cdot \tau(h)(b) \cdot q(x|h, a, b) \cdot (\underline{u}(\sigma, (h, a, b, x)) + \delta) \\ &= \mathbb{E}_{h, \sigma, \tau} [\underline{U}_\sigma^{t+1}] + \delta. \end{aligned}$$

The first claim follows since $\delta > 0$ is arbitrary.

The second claim follows by Lemma A.1 in Appendix A. $\square$

1.4 Conditions for strategies to be subgame ϕ -maxmin

1.4.1 n -Day maxmin strategies and equalizing strategies

The goal of this section is to provide necessary and sufficient conditions for a strategy to be subgame ϕ -maxmin and to provide a characterization of subgame maxmin strategies.

Definition 1.4.1. A strategy $\sigma \in \mathcal{S}_1$ is an n -day ϕ -maxmin strategy in the game $\Gamma_{x_0}(u)$ if for every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$ of length t , and for every strategy $\tau \in \mathcal{S}_2$,

$$\mathbb{E}_{h,\sigma,\tau}[V^{t+n}] \geq \underline{v}(h) - \phi(h). \quad (1.4.1)$$

Definition 1.4.2. A strategy $\sigma \in \mathcal{S}_1$ is ϕ -equalizing in the game $\Gamma_{x_0}(u)$ if for every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$ of length t , and for every strategy $\tau \in \mathcal{S}_2$,

$$u \geq \limsup_{t \rightarrow \infty} \underline{V}^t - \phi(h), \quad \mathbb{P}_{h,\sigma,\tau}\text{-almost surely.} \quad (1.4.2)$$

When $\phi = 0$, we use the terms n -day maxmin and equalizing to mean n -day 0-maxmin and 0-equalizing, respectively.

The first definition is very intuitive. It says that player 1 should play in such a way that, on average, the lower value increases over time. The notion of 1-day maxmin strategies is particularly well known in dynamic programming and stochastic games, see Puterman (1994). A simple characterization of 1-day maxmin strategies is provided in following theorem.

Theorem 1.4.3. Consider a strategy $\sigma \in \mathcal{S}_1$ in the game $\Gamma_{x_0}(u)$. The following three conditions are equivalent:

1. For each $n \in \mathbb{N}$, σ is an n -day maxmin strategy.
2. σ is a 1-day maxmin strategy.
3. For each history $h \in \mathcal{H}^t$ of length t and each strategy $\tau \in \mathcal{S}_2$, the process $(\underline{V}^{t+n})_{n \in \mathbb{N}}$ is a $\mathbb{P}_{h,\sigma,\tau}$ -submartingale.

Proof. That [1] implies [2] is obvious. That [2] implies [3] follows by Lemma A.1 in Appendix A. Finally, that [3] implies [1] follows from the properties of a submartingale and Lemma A.1 in Appendix A. $\square$

A strategy is ϕ -equalizing if, roughly speaking, it almost surely results in a play with an eventually good enough payoff, where eventually good enough means being about as large as the lower value in very deep subgames.

The following example illustrates both the notion of an n -day maxmin strategy and an equalizing strategy.

Example 1.4.4. Consider the centipede game depicted in Figure 1.2. At every history the active player can choose to continue (c) or to quit (q). As soon as a player decides to quit the game ends and in that case the payoff is as given in Figure 1.2. If the game continues indefinitely then the payoff is 0. It is easily verified that, for every $t \in \mathbb{N}$, $\underline{v}(c^{2t}) = \underline{v}(c^{2t+1}) = (t+1)/(t+2)$.

Figure 1.2: Strategies that are n -day maxmin may not be subgame maxmin.

In what follows we focus our attention on pure strategies and characterize the pure strategies that are n -day maxmin and the pure strategies that are equalizing.

CLAIM 1: A pure strategy $\sigma \in \mathcal{S}_1$ is an n -day maxmin strategy for every $n \in \mathbb{N}$ if and only if $\sigma(c^{2t}) = c$ for every $t \in \mathbb{N}$.

PROOF: If the active history is a history of player 1, i.e. $h = c^{2t}$, and player 1 continues everywhere then for any strategy $\tau \in \mathcal{S}_2$ of the second player we have that $\mathbb{E}_{h,\sigma,\tau}[V^{t+1}] = \underline{v}(c^{2t+1})$. If the active history is a history of player 2, i.e. $h = c^{2t+1}$, then for any strategy $\tau \in \mathcal{S}_2$ of the second player we have that $\mathbb{E}_{h,\sigma,\tau}[V^{t+1}]$ is either $u(c^{2t+1}q)$ or $\underline{v}(c^{2t+2})$. Because in this game the lower value function is non-decreasing and $u(c^{2t+1}q) = \underline{v}(c^{2t})$ we have that $\mathbb{E}_{h,\sigma,\tau}[V^{t+1}] \geq \underline{v}(h)$ for every history $h \in \mathcal{H}$ and every strategy $\tau \in \mathcal{S}_2$. Hence σ is a 1-day maxmin strategy. Using Theorem 1.4.3 we conclude that σ is an n -day maxmin strategy for every $n \in \mathbb{N}$.

Conversely, assume there exists a history at which according to the strategy σ player 1 quits. Let c^{2t} denote this history. Then we have for every $\tau \in \mathcal{S}_2$ that $\mathbb{E}_{c^{2t},\sigma,\tau}[V^{t+1}] = u(c^{2t}q) < \underline{v}(c^{2t})$. We conclude that such a strategy σ cannot be a 1-day maxmin strategy.

CLAIM 2: A pure strategy $\sigma \in \mathcal{S}_1$ is equalizing if and only if for infinitely many $t \in \mathbb{N}$ it holds that $\sigma(c^{2t}) = q$.

PROOF: If for infinitely many $t \in \mathbb{N}$ it holds that $\sigma(c^{2t}) = q$, then for every strategy $\tau \in \mathcal{S}_2$ and every history $h \in \mathcal{H}$ there exists an $n \in \mathbb{N}$ such that the play p generated from the history h under the strategy profile (σ, τ) is $c^n q$. In this case it is clear that $\limsup_{t \rightarrow \infty} \underline{v}(c^n q|_t) = u(c^n q)$, which proves that the strategy σ is equalizing.

Conversely, assume $\sigma(c^{2t}) = q$ for at most finitely many $t \in \mathbb{N}$. Then there exists a history after which player 1 always plays continue. Let h denote this history and let $\tau \in \mathcal{S}_2$ denote the strategy of the second player in which he always continues. Then the play generated from the history h under the strategy profile (σ, τ) is c^∞ . Observing that $0 = u(c^\infty) < \limsup_{t \rightarrow \infty} \underline{v}(c^t) = 1$ concludes the proof that any strategy σ in which player 1 quits at most finitely many times cannot be equalizing.

CONSEQUENCE: In the centipede game depicted in Figure 1.2, player 1 does not have a pure strategy which is both n -day maxmin and equalizing. $\diamond$

The following theorem states sufficient conditions under which a strategy σ of player 1 is a subgame ϕ -maxmin strategy.

Theorem 1.4.5. (Sufficient condition) Let $\phi : \mathcal{H} \rightarrow [0, \infty)$ be a tolerance function. The strategy $\sigma \in \mathcal{S}_1$ is a subgame ϕ -maxmin strategy in the game $\Gamma_{x_0}(u)$ if there exist tolerance functions $\phi_1 : \mathcal{H} \rightarrow [0, \infty)$ and $\phi_2 : \mathcal{H} \rightarrow [0, \infty)$ such that $\phi_1 + \phi_2 \leq \phi$ and

1. for every $n \in \mathbb{N}$, σ is n -day ϕ_1 -maxmin,

2. σ is ϕ_2 -equalizing.

Proof. Let ϕ_1 , ϕ_2 , and σ be such that the conditions in the theorem are satisfied. We show that σ is a subgame ϕ -maxmin strategy.

Fix a history $h \in \mathcal{H}^t$ and a strategy $\tau \in \mathcal{S}_2$. Then we have that

$$\begin{aligned} \underline{v}(h) - \phi(h) &\leq \underline{v}(h) - \phi_1(h) - \phi_2(h) \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{E}_{h, \sigma, \tau}[\underline{V}^{t+n}] - \phi_2(h) \\ &\leq \mathbb{E}_{h, \sigma, \tau}[\limsup_{n \rightarrow \infty} \underline{V}^{t+n}] - \phi_2(h) \\ &\leq \mathbb{E}_{h, \sigma, \tau}[u], \end{aligned}$$

where the second inequality holds since σ is an n -day ϕ_1 -maxmin strategy, the third inequality is by Fatou lemma, and the last inequality holds since σ is ϕ_2 -equalizing. $\square$

According to Theorem 1.4.5, to conclude that σ is a subgame ϕ -maxmin strategy, we should find tolerance functions ϕ_1 and ϕ_2 such that at each history their sum does not exceed ϕ and the strategy σ is both n -day ϕ_1 -maxmin and ϕ_2 -equalizing.

Particularly natural are situations where the tolerance level does not increase as time progresses. More formally, the tolerance function ϕ is said to be non-increasing if $\phi(h) \geq \phi(h')$ whenever $h < h'$. The following result states necessary conditions for a strategy to be subgame ϕ -maxmin.

Theorem 1.4.6. (Necessary condition) Let $\sigma \in \mathcal{S}_1$ be a subgame ϕ -maxmin strategy in the game $\Gamma_{x_0}(u)$. Then it holds that:

1. For every $n \in \mathbb{N}$, σ is an n -day ϕ -maxmin strategy.
2. If the tolerance function ϕ is non-increasing, then σ is ϕ -equalizing.

Proof. Let $\sigma \in \mathcal{S}_1$ be a subgame ϕ -maxmin strategy in the game $\Gamma_{x_0}(u)$. Take a history $h \in \mathcal{H}^t$ of length t and a strategy $\tau \in \mathcal{S}_2$.

We prove condition 1. We have

$$\mathbb{E}_{h, \sigma, \tau}[\underline{V}^{t+n}] \geq \mathbb{E}_{h, \sigma, \tau}[\underline{U}_\sigma^{t+n}] \geq \mathbb{E}_{h, \sigma, \tau}[\underline{U}_\sigma^t] = \underline{u}(\sigma, h) \geq \underline{v}(h) - \phi(h),$$

where the first inequality holds since for each play $p \in \mathcal{P}(h)$ we have $\underline{V}^{t+n}(p) = \underline{v}(p_{|t+n}) \geq \underline{u}(\sigma, p_{|t+n}) = \underline{U}_\sigma^{t+n}(p)$. The second inequality holds by Lemma 1.3.3. The next equation holds since t is the length of history h , and the final inequality holds since σ is a subgame ϕ -maxmin strategy.

We prove condition 2. We have, $\mathbb{P}_{h, \sigma, \tau}$ -almost surely,

$$u(p) = \lim_{t \rightarrow \infty} \mathbb{E}_{p_{|t}, \sigma, \tau}[u] \geq \limsup_{t \rightarrow \infty} (\underline{v}(p_{|t}) - \phi(p_{|t})) \geq \limsup_{t \rightarrow \infty} \underline{V}^t(p) - \phi(h),$$

where the equality is by Levy's zero-one law (Lemma A.2 in Appendix A), the first inequality follows since σ is a subgame ϕ -maxmin strategy, and the second inequality holds since ϕ is non-increasing. $\square$

Notice that in the case of a non-zero tolerance function, the necessary and sufficient conditions do not coincide and we do not obtain a characterization of subgame ϕ -maxmin strategies. We now turn to the case where the tolerance function ϕ is identically equal to 0.

Figure 1.3: A 1-day ϕ -maxmin strategy may not be n -day ϕ -maxmin.

Corollary 1.4.7. A strategy $\sigma \in \mathcal{S}_1$ is a subgame maxmin strategy in the game $\Gamma_{x_0}(u)$ if and only if it is 1-day maxmin and equalizing.

When we compare the sufficient conditions of Theorem 1.4.5 to the sufficient conditions of Corollary 1.4.7, we notice that in the case of a non-zero tolerance function we require the strategy to be n -day ϕ -maxmin. In the case of a zero tolerance function the corresponding sufficient conditions only require the strategy to be 1-day maxmin. The reason for this difference is that we should avoid that strategies accumulate the allowed tolerance levels, causing them to become too permissive over time. The following example illustrates this issue.

Example 1.4.8. Consider the decision problem depicted in Figure 1.3. Each period the decision maker can choose to continue (c) or to quit (q). Notice that $v(c^t) = 1/(t+1)$ and hence $\lim_{t \rightarrow \infty} v(c^t) = 0$. Any strategy of player 1 is therefore equalizing. Furthermore, it is clear that in this decision problem the subgame maxmin strategy is unique and requires the decision maker to quit immediately. Now suppose the decision maker has the following tolerance function:

$$\phi(c^t) = v(c^t) - v(c^{t+1}) = \frac{1}{t+1} - \frac{1}{t+2}, \quad t \in \mathbb{N}.$$

Consider the strategy σ where the decision maker always chooses to continue. From the definition of the tolerance function it follows that the strategy σ is a 1-day ϕ -maxmin strategy. Indeed, it holds that $v(c^{t+1}) \geq v(c^t) - \phi(c^t)$. It is also easily seen that the strategy σ is equalizing. Nevertheless, it is clear that σ is not a subgame ϕ -maxmin strategy.

The underlying problem with the strategy σ is that every time the decision maker chooses to continue this causes an acceptable loss in the value, but over time these losses add up to an unacceptable loss. The requirement that for every $n \in \mathbb{N}$ a strategy is n -day ϕ -maxmin guarantees that the accumulated losses over any finite period of time remain acceptable.

If we require a strategy to be subgame maxmin, then we do not tolerate any losses. Therefore, the accumulation problem mentioned above will never occur and it will be sufficient to only require that the strategy is 1-day maxmin. $\diamond$

Example 1.4.9 is such that player 1 has a maxmin strategy in every subgame, but has no subgame maxmin strategy. From Corollary 1.4.7 it follows that any subgame maxmin strategy needs to be both 1-day maxmin and equalizing. The example presents a game where all the strategies are 1-day maxmin but none of them is equalizing and therefore subgame maxmin strategies do not exist.

Example 1.4.9. Consider the following perfect information game. Both players have two actions, left (L) and right (R), so $\mathcal{A} = \mathcal{B} = \{L, R\}$. The players take turns playing an action, which generates a play $p \in \{L, R\}^{\mathbb{N}}$. Let $r_1(p)$ and $r_2(p)$ denote the number of times

that player 1 and player 2, respectively, play action R during the play p and define $r(p) = \min\{r_1(p), r_2(p)\}$. Player 1 obtains a payoff of 0 if both players choose R infinitely often. When at least one of them chooses R only a finite number of times, then player 1 receives a payoff of $r(p)/(r(p) + 1)$. The payoff function u is therefore obtained by defining, for $p \in \{L, R\}^{\mathbb{N}}$,

$$u(p) = \begin{cases} \frac{r(p)}{r(p)+1}, & \text{if } r_1(p) \neq \infty \text{ or } r_2(p) \neq \infty, \\ 0, & \text{if } r_1(p) = r_2(p) = \infty. \end{cases} \quad (1.4.3)$$

At each history $h \in \mathcal{H}$ the value of the game exists and is given by $v(h) = r_2(h)/(r_2(h)+1)$, where $r_2(h)$ denotes the number of times player 2 has chosen R in the history h . Indeed, player 1 can guarantee this payoff by choosing the action R $\max\{r_2(h) - r_1(h), 0\}$ times after history h . Player 2 can guarantee to lose at most this amount by playing only left after history h . Hence at every history $h \in \mathcal{H}$ player 1 has a maxmin strategy.

For every history $h \in \mathcal{H}$ where player 1 takes an action and for every action $a \in \mathcal{A}$ we have that $r_2(ha) = r_2(h)$ and $v(ha) = v(h)$. Therefore, all strategies of player 1 are 1-day maxmin.

On the other hand, no equalizing strategy exists. To see this, take any strategy $\sigma \in \mathcal{S}_1$ for player 1 and consider the strategy $\tau \in \mathcal{S}_2$ in which player 2 always chooses R . Let p be the play generated by the strategy profile (σ, τ) . Then we have that $u(p) < 1$ and $\lim_{t \rightarrow \infty} v(p|_t) = 1$. It follows that σ is not equalizing. Using Corollary 1.4.7 we conclude that player 1 does not have a subgame maxmin strategy. $\diamond$

Example 1.4.10. (Non-leavable decision problems) We consider a stochastic decision problem for player 1, so player 2 is a dummy player. Let $r : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function that associates a reward to every state. The payoff function $u : \mathcal{P} \rightarrow \mathbb{R}$ is defined by

$$u(x_0, a_1, x_1, a_2, \dots) = \limsup_{t \rightarrow \infty} r(x_t), \quad (x_0, a_1, x_1, a_2, \dots) \in \mathcal{P}.$$

Maitra and Sudderth (1996) call decision problems with this structure non-leavable gambling problems and provide a characterization of optimal strategies in terms of thrifty and equalizing strategies. The 1-day maxmin strategies of Theorem 1.4.3 are the strategies that are thrifty after each history in Theorem 7.3 of Chapter 4 in Maitra and Sudderth (1996). Thus 1-day maxmin strategies can be seen as the “subgame” counterpart of thrifty strategies.

A strategy σ is equalizing at a history $h \in \mathcal{H}$ if and only if for each $\epsilon > 0$

$$\{t \in \mathbb{N} \mid r(x_t) \geq \underline{v}(p|_t) - \epsilon\} \text{ is infinite, } \mathbb{P}_{h,\sigma}\text{-almost surely.}$$

This follows from the fact that in a decision problem the process of lower values $\underline{V}^t$ is a $\mathbb{P}_{h,\sigma}$ -supermartingale, and hence is a convergent sequence $\mathbb{P}_{h,\sigma}$ -almost surely.

Using Theorem 7.7 of Chapter 4 in Maitra and Sudderth (1996), it follows that σ is equalizing at h according to their definition. $\diamond$

1.4.2 The case of an upper semi-continuous payoff function

The remainder of this section is devoted to the case where the payoff function is upper semi-continuous. We argue that in this case, any strategy of player 1 is equalizing. Because all upper semi-continuous functions are Borel measurable and because we assumed finite

action sets and a countable state space, the value exists, see Maitra and Sudderth (1998) and Martin (1998), and the lower value equals the value.

The function u is upper semi-continuous at a play $p \in \mathcal{P}$ if for every sequence $\{p_t\}_{t \in \mathbb{N}}$ of plays converging to p it holds that

$$\limsup_{t \rightarrow \infty} u(p_t) \leq u(p).$$

Lemma 1.4.11. Let the payoff function u be upper semi-continuous at the play p . Then we have that

$$u(p) \geq \limsup_{t \rightarrow \infty} \underline{v}(p_t). \quad (1.4.4)$$

Proof. Fix $\epsilon > 0$. For every $t \in \mathbb{N}$, define $h_t = p|_t$ and let $p_t \in \mathcal{P}(h_t)$ be such that $u(p_t) \geq \underline{v}(h_t) - \epsilon$. Such a play p_t exists as player 1 can guarantee a payoff of at least $\underline{v}(h_t) - \epsilon$ at history h_t . Since the sequence $\{p_t\}_{t \in \mathbb{N}}$ converges to p , we have

$$u(p) \geq \limsup_{t \rightarrow \infty} u(p_t) \geq \limsup_{t \rightarrow \infty} \underline{v}(h_t) - \epsilon.$$

Since this holds for every $\epsilon > 0$, the lemma follows. $\square$

In view of Lemma 1.4.11, we obtain the following corollary to Theorems 1.4.5, 1.4.6, and 1.4.7.

Corollary 1.4.12. Let $\Gamma_{x_0}(u)$ be such that u is upper semi-continuous. Then each strategy of player 1 is equalizing. The strategy $\sigma \in \mathcal{S}_1$ is a subgame ϕ -maxmin strategy if and only if for every $n \in \mathbb{N}$ it is an n -day ϕ -maxmin strategy. The strategy $\sigma \in \mathcal{S}_1$ is a subgame maxmin strategy if and only if it is a 1-day maxmin strategy.

Example 1.4.13. (Staying in the set game) Let some subset $\mathcal{X}^*$ of $\mathcal{X}$ be given. For a play $p = (x_0, a_1, b_1, x_1, a_2, b_2, \dots)$ we define $u(p)$ to be 1 if $x_t \in \mathcal{X}^*$ for every $t \in \mathbb{N}$ and to be 0 otherwise. Maitra and Sudderth (1996) refer to such a payoff function as “staying in a set” and in the computer science literature it goes under the name of “safety objective,” see Bruyère (2017). Since u is upper semi-continuous, any strategy $\sigma \in \mathcal{S}_1$ is equalizing. $\diamond$

Example 1.4.14. Consider again the centipede game depicted in Figure 1.2, but with one slight modification. If the game continues indefinitely, then player 1 receives a payoff of 2 instead of 0. The payoff function is now upper semi-continuous. As argued in Example 1.4.4, the strategy σ in which player 1 continues at each history is 1-day maxmin. Because of Corollary 1.4.12, we can conclude that the strategy σ is a subgame maxmin strategy. $\diamond$

Note that the limit average payoff is generally not upper semi-continuous. Hence, in games where the payoff function is given by the limit average payoff, there may exist strategies of the maximizing player which are not equalizing.

1.5 Existence of subgame ϕ -maxmin strategies

The goal of this section is to give sufficient conditions for the existence of a subgame ϕ -maxmin strategy if the tolerance function ϕ is positive. We will prove the following theorem.

Theorem. Let a game $\Gamma_{x_0}(u)$ and a tolerance function $\phi > 0$ be given such that for every $p \in \mathcal{P}$ at least one of the following two conditions holds:

1. (point of upper semicontinuity) The function u is upper semi-continuous at p .
2. (positive limit inferior) $\liminf_{t \rightarrow \infty} \phi(p|_t) > 0$.

Then there exists a subgame ϕ -maxmin strategy in the game $\Gamma_{x_0}(u)$.

The proof of this theorem will be constructive. A crucial element of this proof are switching strategies which are formally defined in Subsection 1.5.1. Intuitively, a switching strategy σ^ϕ checks at every history h whether the strategy is a $\phi(h)$ -maxmin strategy, and if this is not the case it switches to a new strategy which is $\phi(h)/2$ -maxmin. Such a switching strategy will always be an n -day ϕ -maxmin strategy for any $n \in \mathbb{N}$ as will be shown in Theorem 1.5.8. However it is not guaranteed that a switching strategy is a subgame ϕ -maxmin strategy, as the possibility that infinitely many switches occur, may jeopardize the ϕ -equalizing property. To guarantee that the ϕ -equalizing property is satisfied we put additional restrictions on the tolerance function or payoff function. Using the results developed in Subsection 1.4.2 we have that along plays which are points of upper semi-continuity of the payoff-function, the ϕ -equalizing condition is always fulfilled. Then using the boundedness of the payoff-function, we find that if along a play the sequence of tolerance levels is bounded away from 0, the switching strategy will only need to switch finitely often. This finite switching property guarantees that the ϕ -equalizing condition is fulfilled.

1.5.1 Definition of the switching strategy σ^ϕ

The construction of the subgame ϕ -maxmin strategy in Theorem 1.5.9 is as follows. Player 1 starts by playing a $(\phi(x_0)/2)$ -maxmin strategy. Next, at every history $h \in \mathcal{H}$ player 1 checks whether the strategy he is currently using is a $\phi(h)$ -maxmin strategy for the subgame at history h . If this is the case he keeps using the strategy. If not, he switches to a $(\phi(h)/2)$ -maxmin strategy for the subgame at history h . We then use Theorem 1.4.5 to show that this construction leads to a subgame ϕ -maxmin strategy. This type of construction is not new, and similar constructions were used for example in Rosenberg, Solan, and Vieille (2001), Solan and Vieille (2002), and Mashiah-Yaakovi (2015).

Fix a tolerance function $\phi > 0$. For every history $h \in \mathcal{H}$, player 1 has a $(\phi(h)/2)$ -maxmin strategy for the subgame at history h , denoted by σ^h . The function $\psi : \mathcal{H} \rightarrow \mathcal{H}$ maps histories into histories and is such that player 1 is going to use strategy $\sigma^{\psi(h)}$ at subgame $h \in \mathcal{H}$. The function ψ is used to describe when player 1 switches strategies and is recursively defined by setting $\psi(x_0) = x_0$ and, for every $h \in \mathcal{H}$, for every $z \in \mathcal{Z}$,

$$\psi(hz) = \begin{cases} \psi(h), & \text{if } \underline{u}(\sigma^{\psi(h)}, hz) \geq \underline{v}(hz) - \phi(hz), \\ hz, & \text{otherwise.} \end{cases}$$

The condition $\underline{u}(\sigma^{\psi(h)}, hz) \geq \underline{v}(hz) - \phi(hz)$ verifies whether the strategy to which player 1 switched last, $\sigma^{\psi(h)}$, is a $\phi(hz)$ -maxmin strategy for the subgame at history hz . If this is the case, then there is no need to switch and $\psi(hz) = \psi(h)$. Otherwise, player 1 switches to σ^{hz} , which is achieved by setting $\psi(hz) = hz$. Formally, we define the switching strategy $\sigma^\phi : \mathcal{H} \rightarrow \Delta(\mathcal{A})$ by

$$\sigma^\phi(h) = \sigma^{\psi(h)}(h), \quad h \in \mathcal{H}. \quad (1.5.1)$$

The following example illustrates the construction of the switching strategy σ^ϕ and shows that it may not be subgame ϕ -maxmin.

Example 1.5.1. Consider again the centipede game depicted in Figure 1.2. We recall that, for every $t \in \mathbb{N}$, $\underline{v}(c^{2t}) = \underline{v}(c^{2t+1}) = (t+1)/(t+2)$. Take a tolerance function ϕ with the property that, for every $t \in \mathbb{N}$, $\phi(c^{2t}) < 1/((t+1)(t+2))$.

Let $h \in \mathcal{H}$ be an active history for player 1 or player 2 and let $k \in \mathbb{N}$ be such that $h = c^{2k}$ or $h = c^{2k+1}$. The strategy σ^h in which player 1 chooses continue at periods $0, 2, \dots, 2k$ and chooses quit at every later period, i.e.

$$\sigma^h(c^{2t}) = \begin{cases} c, & \text{if } 2t \leq 2k, \\ q, & \text{otherwise,} \end{cases}$$

is a maxmin strategy at subgame h .

We now consider the switching strategy σ^ϕ . We show by induction that, for every $h \in H$, for every $z \in Z$, it holds that $\psi(hz) = hz$ if hz is an active history of player 1 and $\psi(hz) = h$ if hz is an active history of player 2. The statement trivially holds for the initial history. Let h be an active history of player 2 and let $t \in \mathbb{N}$ be such that $h = c^{2t+1}$. Since $\sigma^h = \sigma^{c^{2t}}$ is a maxmin strategy at subgame h , it holds that $\psi(c^{2t+1}) = c^{2t}$. Let h be an active history of player 1 and let $t \in \mathbb{N} \setminus \{0\}$ be such that $h = c^{2t}$. We have that

$$\underline{u}(\sigma^{\psi(c^{2t-1})}, c^{2t}) = \underline{u}(\sigma^{c^{2t-2}}, c^{2t}) = \frac{t}{t+1} = \underline{v}(c^{2t}) - \frac{1}{(t+1)(t+2)} < \underline{v}(c^{2t}) - \phi(c^{2t}),$$

so $\psi(c^{2t}) = c^{2t}$. Since the tolerance function ϕ is so stringent, it forces player 1 to switch at each of his active histories. For every $t \in \mathbb{N}$, it holds that $\sigma^\phi(c^{2t}) = \sigma^{c^{2t}}(c^{2t}) = c$, so under σ^ϕ player 1 chooses c at each of his active histories. The strategy σ^ϕ is not a subgame ϕ -maxmin strategy as it fails to be ϕ -equalizing, see Example 1.4.4. $\diamond$

1.5.2 Definition and analysis of finite switching strategies

Given the switching strategy σ^ϕ , for every $k \in \mathbb{N}$ we define the strategy $\sigma^k : \mathcal{H} \rightarrow \Delta(\mathcal{A})$ such that it coincides with σ^ϕ as long as at most k switches have been made and stops switching thereafter. Formally, we recursively define the function $\kappa : \mathcal{H} \rightarrow \mathbb{N}$ which counts the number of switches along a history h by setting $\kappa(x_0) = 0$ and, for all histories $h, hz \in \mathcal{H}$,

$$\kappa(hz) = \begin{cases} \kappa(h), & \text{if } \underline{u}(\sigma^{\psi(h)}, hz) \geq \underline{v}(hz) - \phi(hz), \\ \kappa(h) + 1, & \text{otherwise.} \end{cases}$$

For every $k \in \mathbb{N}$, we define the stopping time $T_k : \mathcal{P} \rightarrow \mathbb{N} \cup \{\infty\}$ by

$$T_k(p) = \inf\{t \in \mathbb{N} \mid \kappa(p|_t) = k\}, \quad p \in \mathcal{P}. \quad (1.5.2)$$

The stopping time T_k indicates the time at which switch k occurred. Since, for $t < \infty$, the expression $T_k(p) \leq t$ only depends on the history up to period t , it holds that the set $\{p \in \mathcal{P} \mid T_k(p) \leq t\}$ is $\mathcal{F}^t$ -measurable and T_k is a stopping time indeed.

We now formally define $\sigma^k : \mathcal{H} \rightarrow \Delta(\mathcal{A})$. Take any $p \in \mathcal{P}(h)$ and let

$$\sigma^k(h) = \begin{cases} \sigma^\phi(h), & \text{if } \kappa(h) \leq k, \\ \sigma^{h|_{T_k(p)}}(h), & \text{otherwise.} \end{cases} \quad (1.5.3)$$

If $\kappa(h) > k$, then the time at which switch k has occurred is the same for every $p \in \mathcal{P}(h)$, so it holds that σ^k is well defined.

For every $k \in \mathbb{N}$, let $\mathcal{R}_k \subseteq \mathcal{P}$ be the set of plays along which at least k switches occur,

$$\mathcal{R}_k = \{p \in \mathcal{P} \mid T_k(p) < \infty\}, \quad (1.5.4)$$

so $\mathcal{R}_1 \supseteq \mathcal{R}_2 \supseteq \mathcal{R}_3 \supseteq \dots$. Furthermore, let $\mathcal{R}_\infty \subseteq \mathcal{P}$ denote the set of plays along which infinitely many switches occur,

$$\mathcal{R}_\infty = \bigcap_{k=1}^{\infty} \mathcal{R}_k. \quad (1.5.5)$$

The next result is very intuitive. Consider the strategies $\sigma^k, \sigma^{k+1}, \dots$. All these strategies agree with σ^ϕ for as long as σ^ϕ does not require more than k switches. Consequently, the measures that these strategies induce on $\mathcal{P}$ assign the same probability to any event that is “determined” before switch $k + 1$ occurs, i.e. to any event in the sigma-algebra $\mathcal{F}^{T_{k+1}}$.

Lemma 1.5.2. Let a strategy $\tau \in \mathcal{S}_2$, a history $h \in \mathcal{H}$, and some $k \in \mathbb{N}$ be given. For $\sigma = \sigma^k, \sigma^{k+1}, \dots, \sigma^\phi$, the probability measures $\mathbb{P}_{h,\sigma,\tau}$ all coincide on the sigma-algebra $\mathcal{F}^{T_{k+1}}$. Furthermore, these probability measures all agree on each universally measurable subset of $\mathcal{P} \setminus \mathcal{R}_{k+1}$.

Proof. A set A of the universally measurable sigma-algebra $\mathcal{F}$ is called *agreeable* if for $\sigma = \sigma^k, \sigma^{k+1}, \dots, \sigma^\phi$ the measures $\mathbb{P}_{h,\sigma,\tau}$ all assign the same probability to A .

We argue first that each cylinder set in $\mathcal{F}^{T_{k+1}}$ is agreeable. A cylinder set $\mathcal{P}(h)$ is a member of $\mathcal{F}^{T_{k+1}}$ if and only if $\kappa(h) \leq k + 1$. Let a cylinder set $\mathcal{P}(h')$ in $\mathcal{F}^{T_{k+1}}$ be given. Since $\kappa(h') \leq k + 1$, we know that $\kappa(h'') \leq k$ for each history h'' preceding h' . It follows that $\sigma^k(h'') = \sigma^{k+1}(h'') = \dots = \sigma^\phi(h'')$. Since this applies to each history h'' that precedes h' , the set $\mathcal{P}(h')$ is agreeable.

Now take any $E \in \mathcal{F}^{T_{k+1}}$. For $t \in \mathbb{N}$, let $E_t = E \cap \{p \in \mathcal{P} \mid T_{k+1}(p) = t\}$ and let $E_\infty = E \cap \{p \in \mathcal{P} \mid T_{k+1}(p) = \infty\}$. To show that E is agreeable, it suffices to show that the sets E_t and E_∞ are.

Let some $t \in \mathbb{N}$ be given. We know that E_t is a member of $\mathcal{F}^t$. Consequently, E_t can be written as a disjoint union of cylinder sets in $\mathcal{F}^t$, say $E_t = \cup \{C_n \mid n \in \mathbb{N}\}$, with each C_n a member of $\mathcal{F}^t$. Since each C_n is a subset of the set $\{p \in \mathcal{P} \mid T_{k+1}(p) = t\}$, it is a member of $\mathcal{F}^{T_{k+1}}$, so is agreeable by the result of the second paragraph in the proof. It now follows that E_t is agreeable.

To show that E_∞ is agreeable, we make use of the fact that E_∞ is a Borel set and of the regularity of σ on the Borel sigma-algebra. Let O be any open subset of $\mathcal{P}$ containing E_∞ . The set O can be written as a disjoint union of cylinder sets, say $O = \cup \{\mathcal{P}(h_n) \mid n \in \mathbb{N}\}$. Without loss of generality it holds that, for every $n \in \mathbb{N}$ the set $\mathcal{P}(h_n)$ has a point in common with E_∞ . Thus in particular there is $p \in \mathcal{P}(h_n)$ with $T_{k+1}(p) = \infty$. This implies that $\kappa(h_n) \leq k$ and hence that $\mathcal{P}(h_n)$ is a member of $\mathcal{F}^{T_{k+1}}$. We conclude that each $\mathcal{P}(h_n)$ is agreeable by the result of the second paragraph in the proof. It follows that O is an agreeable set.

To prove the second claim, we notice that all Borel subsets of $\mathcal{P} \setminus \mathcal{R}_{k+1} = \{p \in \mathcal{P} \mid T_{k+1}(p) = \infty\}$ are members of $\mathcal{F}^{T_{k+1}}$, so are agreeable. The result for universally measurable subsets of $\{p \in \mathcal{P} \mid T_{k+1}(p) = \infty\}$ follows since each such set differs from a Borel set by a negligible set. $\square$

The following lemma is a special case of the optional sampling theorem with unbounded stopping times as presented in Yeh (1995, p. 139). Assume we have specified an $\mathcal{F}^\infty$ -measurable stochastic variable $\underline{U}_\sigma^\infty$. Let T be a stopping time. We define the stochastic

variable $\underline{U}_\sigma^T$ by letting it agree with $\underline{U}_\sigma^t$ on the set $\{p \in \mathcal{P} : T(p) = t\}$ for each $t \in \mathbb{N}$ and by letting it agree with $\underline{U}_\sigma^\infty$ on the set $\{p \in \mathcal{P} : T(p) = \infty\}$. The stochastic variable $\underline{U}_\sigma^T$ is then $\mathcal{F}^T$ -measurable (Yeh, 1995, Theorem 3.28).

Exactly how the stochastic variable $\underline{U}_\sigma^\infty$ must be chosen is a rather subtle matter. In Lemmas 1.5.3, 1.5.5, and 1.5.6, $\underline{U}_\sigma^\infty$ is taken equal to some Borel measurable function that agrees with u almost surely for the measure that is specified by the respective lemma. We cannot use u itself, since u is only assumed to be universally measurable. Neither can we fix $\underline{U}_\sigma^\infty$ in advance, because there is no function that would agree with u almost surely with respect to all the measures that arise henceforth.

Lemma 1.5.3. (Optional sampling for guarantee level) Let a strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$, $t \in \mathbb{N}$, and a history $h \in \mathcal{H}^t$ of length t be given. Let $\underline{U}_\sigma^\infty$ be an $\mathcal{F}^\infty$ -measurable stochastic variable that $\mathbb{P}_{h,\sigma,\tau}$ -almost surely coincides with u . Consider the stopping times S and T such that, for each $p \in \mathcal{P}(h)$, $t \leq S(p) \leq T(p)$. Then

$$\underline{U}_\sigma^T \leq \mathbb{E}_{h,\sigma,\tau}[u|\mathcal{F}^T], \quad \mathbb{P}_{h,\sigma,\tau}\text{-almost surely}, \quad (1.5.6)$$

$$\underline{U}_\sigma^S \leq \mathbb{E}_{h,\sigma,\tau}[\underline{U}_\sigma^T|\mathcal{F}^S], \quad \mathbb{P}_{h,\sigma,\tau}\text{-almost surely}. \quad (1.5.7)$$

Proof. The result follows by Theorem 8.16 in Yeh (1995, p. 139), applied to the process $(\underline{U}_\sigma^n)_{n \geq t}$ on the measurable space $(\mathcal{P}, \mathcal{F}, \mathbb{P}_{h,\sigma,\tau})$ with a filtration $(\mathcal{F}^n)_{n \geq t}$.

We verify that the conditions of Theorem 8.16 in Yeh (1995) are satisfied. The process $(\underline{U}_\sigma^n)_{n \geq t}$ is a $\mathbb{P}_{h,\sigma,\tau}$ -submartingale with respect to the filtration $(\mathcal{F}^n)_{n \geq t}$ by Lemma 1.3.3. Take ξ of Theorem 8.16 in Yeh (1995) equal to u . Since $\underline{U}_\sigma^\infty$ is an $\mathcal{F}^\infty$ -measurable stochastic variable that $\mathbb{P}_{h,\sigma,\tau}$ -almost surely coincides with u , it is a version of $\mathbb{E}_{h,\sigma,\tau}[u|\mathcal{F}^\infty]$, as is required by the theorem.

Lastly, we verify that condition (1) of Theorem 8.16 in Yeh (1995) is satisfied. Notice that, for every $n \geq t$, for every play $p \in \mathcal{P}$, we have $\underline{U}_\sigma^n(p) = \underline{u}(\sigma, p|_n) \leq \mathbb{E}_{p|_n,\sigma,\tau}[u]$. The right-hand side of this inequality is a version of $\mathbb{E}_{h,\sigma,\tau}[u|\mathcal{F}^n]$. Consequently, $\underline{U}_\sigma^n \leq \mathbb{E}_{h,\sigma,\tau}[u|\mathcal{F}^n]$ holds $\mathbb{P}_{h,\sigma,\tau}$ -almost surely, as desired. $\square$

The next lemma relates the guaranteed expected payoffs of strategies σ^k and σ^{k+1} at the moment of switch $k+1$. For this result, we choose $\underline{U}_{\sigma^k}^\infty = \underline{U}_{\sigma^{k+1}}^\infty$ to be any $\mathcal{F}^\infty$ -measurable stochastic variable. How this stochastic variable is related to u is unimportant. We write Φ^t to denote the $\mathcal{F}^t$ -measurable stochastic variable given by $\Phi^t(p) = \phi(p|_t)$.

Lemma 1.5.4. Let $k \in \mathbb{N}$ and $\mathcal{F}^\infty$ -measurable stochastic variables $\underline{U}_{\sigma^k}^\infty, \underline{U}_{\sigma^{k+1}}^\infty$ such that $\underline{U}_{\sigma^k}^\infty = \underline{U}_{\sigma^{k+1}}^\infty$ be given. Then it holds that

$$\underline{U}_{\sigma^k}^{T_{k+1}} \leq \underline{U}_{\sigma^{k+1}}^{T_{k+1}} - \frac{1}{2}\Phi^{T_{k+1}} \cdot I(T_{k+1} < \infty). \quad (1.5.8)$$

Proof. Let some $p \in \mathcal{P}$ be given. We distinguish the following two cases.

CASE 1: $T_{k+1}(p) < \infty$. In this case at least $k+1$ switches occur along the play p . For $h = p|_{T_{k+1}(p)}$ we have the following inequalities

$$\underline{u}(\sigma^k, h) < \underline{v}(h) - \phi(h) = \underline{v}(h) - \frac{1}{2}\phi(h) - \frac{1}{2}\phi(h) \leq \underline{u}(\sigma^{k+1}, h) - \frac{1}{2}\phi(h),$$

where the first inequality holds since σ^k is not a $\phi(h)$ -maxmin strategy for the subgame at history h and the second inequality holds because the strategy σ^{k+1} is a $(\phi(h)/2)$ -maxmin strategy for the subgame at history h . Since $I(T_{k+1}(p) < \infty) = 1$, (1.5.8) holds.

CASE 2: $T_{k+1}(p) = \infty$. In this case we have

$$\underline{U}_{\sigma^k}^{T_{k+1}}(p) = \underline{U}_{\sigma^k}^{\infty}(p) = \underline{U}_{\sigma^{k+1}}^{\infty}(p) = \underline{U}_{\sigma^{k+1}}^{T_{k+1}}(p).$$

Thus (1.5.8) holds as an equality. $\square$

The following lemma states the intuitive property that for histories with less than $k + 1$ switches or histories at which switch $k + 1$ occurs, the strategy σ^{k+1} guarantees at least the same payoff to player 1 than strategy σ^k .

Lemma 1.5.5. Let $t \in \mathbb{N}$ and a history $h \in \mathcal{H}^t$ of length t be given. Let $k \in \mathbb{N}$ be such that $T_{k+1}(p) \geq t$ for every $p \in \mathcal{P}(h)$. Then it holds that $\underline{u}(\sigma^k, h) \leq \underline{u}(\sigma^{k+1}, h)$.

Proof. Fix strategy $\tau \in \mathcal{S}_2$ of player 2.

We first define $\underline{U}_{\sigma^k}^{\infty}$. Consider the probability measure $\mathbb{Q}$ on the measurable space $(\mathcal{P}, \mathcal{F})$ given by $\mathbb{Q}(A) = \sum_{k \in \mathbb{N}} 2^{-k-1} \mathbb{P}_{h, \sigma^k, \tau}(A)$ for each $A \in \mathcal{F}$. Let $\bar{u}$ be an $\mathcal{F}^{\infty}$ -measurable stochastic variable with the property that $\bar{u} = u$, $\mathbb{Q}$ -almost surely. Since $\mathbb{P}_{h, \sigma^k, \tau}$ is absolutely continuous with respect to $\mathbb{Q}$ it holds that $\bar{u} = u$, $\mathbb{P}_{h, \sigma^k, \tau}$ -almost surely, for every $k \in \mathbb{N}$. We define $\underline{U}_{\sigma^k}^{\infty}$ to be equal to $\bar{u}$, for every $k \in \mathbb{N}$.

We now obtain the following inequalities. First, we have

$$\underline{u}(\sigma^k, h) \leq \mathbb{E}_{h, \sigma^k, \tau}[\underline{U}_{\sigma^k}^{T_{k+1}}]$$

as an instance of inequality (1.5.7) of Lemma 1.5.3 with $S = t$ and $T = T_{k+1}$. Secondly, from the fact that $\underline{U}_{\sigma^k}^{T_{k+1}}$ is an $\mathcal{F}^{T_{k+1}}$ -measurable stochastic variable, we obtain by Lemma 1.5.2 that

$$\mathbb{E}_{h, \sigma^k, \tau}[\underline{U}_{\sigma^k}^{T_{k+1}}] = \mathbb{E}_{h, \sigma^{k+1}, \tau}[\underline{U}_{\sigma^k}^{T_{k+1}}].$$

Thirdly, we know that

$$\underline{U}_{\sigma^k}^{T_{k+1}} \leq \underline{U}_{\sigma^{k+1}}^{T_{k+1}} \leq \mathbb{E}_{h, \sigma^{k+1}, \tau}[u | \mathcal{F}^{T_{k+1}}], \quad \mathbb{P}_{h, \sigma^{k+1}, \tau}\text{-almost surely,}$$

where the first of these inequalities follows from Lemma 1.5.4 and the second one follows by inequality (1.5.6) of Lemma 1.5.3. Taking the expectation of the last array of inequalities with respect to $\mathbb{P}_{h, \sigma^{k+1}, \tau}$ and making use of the law of iterated expectation yields

$$\mathbb{E}_{h, \sigma^{k+1}, \tau}[\underline{U}_{\sigma^k}^{T_{k+1}}] \leq \mathbb{E}_{h, \sigma^{k+1}, \tau}[u].$$

Combining these facts yields $\underline{u}(\sigma^k, h) \leq \mathbb{E}_{h, \sigma^{k+1}, \tau}[u]$. Taking the infimum over all strategies τ of player 2 completes the proof. $\square$

Notice that a switch occurring at history h increases the guarantee level of player 1 at h by at least $\phi(h)/2$. Although it is possible that player 1 switches infinitely many times along a play p , and therefore incurs infinitely many increases in his guarantee level along this play, the total overall increase in his guarantee level is bounded, since the payoff function itself is a bounded function. The next lemma provides the formal statement. We define

$$M = \sup_{p \in \mathcal{P}} |u(p)|.$$

Lemma 1.5.6. Let $t \in \mathbb{N}$, a history $h \in \mathcal{H}^t$ of length t , and a strategy $\tau \in \mathcal{S}_2$ of player 2 be given. Then

$$\sum_{k=\kappa(h)+1}^{\infty} \mathbb{E}_{h,\sigma^\phi,\tau} \left[\frac{1}{2} \Phi^{T_k} \cdot I(T_k < \infty) \right] \leq 2M. \quad (1.5.9)$$

Proof. As in the proof of Lemma 1.5.5, let $\bar{u}$ be any $\mathcal{F}^\infty$ -measurable stochastic variable with the property that, for every $k \in \mathbb{N}$, $\bar{u} = u$, $\mathbb{P}_{h,\sigma^k,\tau}$ -almost surely. For every $k \in \mathbb{N}$ we define $\underline{U}_{\sigma^k}^\infty = \bar{u}$.

Let some $k > \kappa(h)$ be given. For every $p \in \mathcal{P}(h)$, it holds that $t \leq T_k(p) \leq T_{k+1}(p)$, so Lemma 1.5.3 implies

$$\underline{U}_{\sigma^k}^{T_k} \leq \mathbb{E}_{h,\sigma^k,\tau} [\underline{U}_{\sigma^k}^{T_{k+1}} | \mathcal{F}^{T_k}], \quad \mathbb{P}_{h,\sigma^k,\tau}\text{-almost surely.} \quad (1.5.10)$$

We now have

$$\mathbb{E}_{h,\sigma^\phi,\tau} [\underline{U}_{\sigma^k}^{T_k}] = \mathbb{E}_{h,\sigma^k,\tau} [\underline{U}_{\sigma^k}^{T_k}] \leq \mathbb{E}_{h,\sigma^k,\tau} [\underline{U}_{\sigma^k}^{T_{k+1}}] = \mathbb{E}_{h,\sigma^\phi,\tau} [\underline{U}_{\sigma^k}^{T_{k+1}}],$$

where the two equalities follow from Lemma 1.5.2 and the fact that both $\underline{U}_{\sigma^k}^{T_k}$ and $\underline{U}_{\sigma^k}^{T_{k+1}}$ are $\mathcal{F}^{T_{k+1}}$ -measurable stochastic variables, and the inequality follows by taking the expectation on both sides of inequality (1.5.10) and the law of iterated expectation.

Using Lemma 1.5.4 we can conclude that

$$\begin{aligned} \mathbb{E}_{h,\sigma^\phi,\tau} \left[\frac{1}{2} \Phi^{T_{k+1}} \cdot I(T_{k+1} < \infty) \right] &\leq \mathbb{E}_{h,\sigma^\phi,\tau} \left[\underline{U}_{\sigma^{k+1}}^{T_{k+1}} \right] - \mathbb{E}_{h,\sigma^\phi,\tau} \left[\underline{U}_{\sigma^k}^{T_{k+1}} \right] \\ &\leq \mathbb{E}_{h,\sigma^\phi,\tau} \left[\underline{U}_{\sigma^{k+1}}^{T_{k+1}} \right] - \mathbb{E}_{h,\sigma^\phi,\tau} \left[\underline{U}_{\sigma^k}^{T_k} \right]. \end{aligned}$$

Summing the preceding inequality over $k = \kappa(h) + 1, \dots, K$, we obtain

$$\sum_{k=\kappa(h)+1}^K \mathbb{E}_{h,\sigma^\phi,\tau} \left[\frac{1}{2} \Phi^{T_k} \cdot I(T_k < \infty) \right] \leq \mathbb{E}_{h,\sigma^\phi,\tau} [\underline{U}_{\sigma^K}^{T_K}] - \mathbb{E}_{h,\sigma^\phi,\tau} [\underline{U}_{\sigma^{\kappa(h)+1}}^{T_{\kappa(h)+1}}] \leq 2M.$$

The result follows by taking the limit as $K \rightarrow \infty$. $\square$

The following lemma plays a crucial role in the proof of Theorem 1.5.9. Essentially it says that along almost any play $p \in \mathcal{P}$ only finitely many switches occur or the tolerance level goes to zero.

Lemma 1.5.7. For every history $h \in \mathcal{H}$, for every strategy $\tau \in \mathcal{S}_2$ of player 2, it holds that

$$\lim_{k \rightarrow \infty} \Phi^{T_k} \cdot I(T_k < \infty) = 0, \quad \mathbb{P}_{h,\sigma^\phi,\tau}\text{-almost surely.}$$

Proof. Let us write $X_k = \Phi^{T_k} \cdot I(T_k < \infty)$ and $k' = \kappa(h) + 1$. Since $X_k \geq 0$, the monotone convergence theorem implies that $\mathbb{E}_{h,\sigma^\phi,\tau} [\sum_{k=k'}^{\infty} X_k] = \sum_{k=k'}^{\infty} \mathbb{E}_{h,\sigma^\phi,\tau} [X_k]$. The latter expression is finite by Lemma 1.5.6. Thus $\sum_{k=k'}^{\infty} X_k$ has a finite mean with respect to the probability measure $\mathbb{P}_{h,\sigma^\phi,\tau}$. Hence $\sum_{k=k'}^{\infty} X_k < \infty$ holds $\mathbb{P}_{h,\sigma^\phi,\tau}$ -almost surely. This implies that $X_k \rightarrow 0$ holds $\mathbb{P}_{h,\sigma^\phi,\tau}$ -almost surely. $\square$

1.5.3 Properties of a switching strategy σ^ϕ .

Before moving to the main theorem of this section we show that the switching strategy σ^ϕ is always n -day ϕ -maxmin for every $n \in \mathbb{N}$.

Theorem 1.5.8. Let a game $\Gamma_{x_0}(u)$ and a tolerance function $\phi > 0$ be given. For every $n \in \mathbb{N}$, the switching strategy σ^ϕ is n -day ϕ -maxmin.

Proof. Let a history $h \in \mathcal{H}$ with length t be given. Let $k = \kappa(h)$ denote the number of switches that has occurred along the history h . We obtain the following chain of inequalities

$$\begin{aligned} \underline{v}(h) - \phi(h) &\leq \underline{u}(\sigma^k, h) \leq \underline{u}(\sigma^{t+n}, h) \leq \mathbb{E}_{h, \sigma^{t+n}, \tau}[\underline{U}_{\sigma^{t+n}}^{t+n}] \leq \mathbb{E}_{h, \sigma^{t+n}, \tau}[\underline{V}^{t+n}] \\ &= \mathbb{E}_{h, \sigma^\phi, \tau}[\underline{V}^{t+n}], \end{aligned}$$

where the first inequality holds since σ^k is a $\phi(h)$ -maxmin strategy for the subgame at history h , the second inequality holds by Lemma 1.5.5 since $k \leq t \leq t+n$, the third inequality holds by Lemma 1.3.3, and the fourth one follows since $\underline{U}_{\sigma^{t+n}}^{t+n} \leq \underline{V}^{t+n}$. The final equality follows from the fact that the stochastic variable $\underline{V}^{t+n}$ is $\mathcal{F}^{t+n}$ -measurable and the fact that the strategy σ^{t+n} coincides with σ^ϕ at least until time $t+n$. $\square$

We are now in a position to state the main result of this section, which gives sufficient conditions for the existence of subgame ϕ -maxmin strategies.

Theorem 1.5.9. Let a game $\Gamma_{x_0}(u)$ and a tolerance function $\phi > 0$ be given such that for every $p \in \mathcal{P}$ at least one of the following two conditions holds:

1. (point of upper semicontinuity) The function u is upper semi-continuous at p .
2. (positive limit inferior) $\liminf_{t \rightarrow \infty} \phi(p|_t) > 0$.

Then there exists a subgame ϕ -maxmin strategy in the game $\Gamma_{x_0}(u)$.

Proof. We define the tolerance function ϕ' by

$$\phi'(h) = \frac{1}{2} \min\{\phi(h') \mid h' \leq h\}, \quad h \in \mathcal{H}.$$

Thus ϕ' is a non-increasing tolerance function with $\phi' \leq \frac{1}{2}\phi$.

We show that $\sigma^{\phi'}$ is ϕ' -equalizing. Since $2\phi' \leq \phi$ and $\sigma^{\phi'}$ is an n -day ϕ' -maxmin strategy for every $n \in \mathbb{N}$ by Theorem 1.5.8, it then follows from Theorem 1.4.5 that $\sigma^{\phi'}$ is a subgame ϕ -maxmin strategy.

Let $\mathcal{U}$ denote the set of plays $p \in \mathcal{P}$ at which u is upper semi-continuous and let $\mathcal{I}$ denote the set of plays $p \in \mathcal{P}$ such that $\liminf_{t \rightarrow \infty} \phi(p|_t) > 0$. By the assumption of the theorem it holds that $\mathcal{P} = \mathcal{U} \cup \mathcal{I}$. By the definition of ϕ' , we have $\liminf_{t \rightarrow \infty} \phi'(p|_t) > 0$ for each $p \in \mathcal{I}$.

Let some $t \in \mathbb{N}$, a history $h \in \mathcal{H}^t$, and a strategy $\tau \in \mathcal{S}_2$ be given.

STEP 1: For every $p \in \mathcal{U}$, $u(p) \geq \limsup_{n \rightarrow \infty} \underline{V}^n(p) - \phi'$.

This follows directly from Lemma 1.4.11.

For every $k \in \mathbb{N}$, we define $\mathcal{J}_k = \{p \in \mathcal{I} \cap (\mathcal{R}_k \setminus \mathcal{R}_{k+1}) \mid \lim_{n \rightarrow \infty} U_{\sigma^k, \tau}^n(p) = u(p)\}$ and $\mathcal{J} = \bigcup_{k \in \mathbb{N}} \mathcal{J}_k$.

STEP 2: For every $p \in \mathcal{J}$, $u(p) \geq \limsup_{n \rightarrow \infty} \underline{V}^n(p) - \phi'$.

Let $k \in \mathbb{N}$ and $p \in \mathcal{J}_k$ be given. Exactly k switches occur along the play p and the last switch occurs at time $T_k(p)$. By our construction of $\sigma^{\phi'}$, this means that for each time $n > T_k(p)$ the strategy σ^k is a $\phi'(p|_n)$ -maxmin strategy for the subgame at history $p|_n$, so $\underline{u}(\sigma^k, p|_n) \geq \underline{v}(p|_n) - \phi'(p|_n)$. Since $U_{\sigma^k, \tau}^n(p) \geq \underline{u}(\sigma^k, p|_n)$ and since for $n \geq t$ we have $\phi'(p|_n) \leq \phi'(h)$, we conclude that for $n \geq t$

$$U_{\sigma^k, \tau}^n(p) \geq \underline{V}^n(p) - \phi'(h).$$

Taking the limit as n goes to infinity, and making use of the fact that $p \in \mathcal{J}_k$, we obtain

$$u(p) = \lim_{n \rightarrow \infty} U_{\sigma^k, \tau}^n(p) \geq \limsup_{n \rightarrow \infty} \underline{V}^n(p) - \phi'(h).$$

STEP 3: $\mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{I} \cap \mathcal{R}_\infty) = 0$.

Recall that by Lemma 1.5.7, $\Phi^{T_k} \cdot I(T_k < \infty)$ converges to 0 as k goes to infinity, $\mathbb{P}_{h, \sigma^{\phi'}, \tau}$ -almost surely. Also recall that $\mathcal{R}_\infty$ is the set of plays where infinitely many switches occur. Thus $I(T_k < \infty)$ is identically equal to 1 on $\mathcal{R}_\infty$. Furthermore, $\liminf_{k \rightarrow \infty} \Phi^{T_k} > 0$ everywhere on $\mathcal{I}$. We conclude that $\Phi^{T_k} \cdot I(T_k < \infty)$ does not converge to zero on $\mathcal{I} \cap \mathcal{R}_\infty$ and the result follows.

STEP 4: $\mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{U} \cup \mathcal{J}) = 1$.

For every $k \in \mathbb{N}$, it holds by Levy's zero-one law (Lemma A.2 in Appendix A) that

$$\mathbb{P}_{h, \sigma^k, \tau}(\mathcal{J}_k) = \mathbb{P}_{h, \sigma^k, \tau}(\mathcal{I} \cap (\mathcal{R}_k \setminus \mathcal{R}_{k+1})).$$

Using Lemma 1.5.2 twice yields

$$\mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{J}_k) = \mathbb{P}_{h, \sigma^k, \tau}(\mathcal{J}_k) = \mathbb{P}_{h, \sigma^k, \tau}(\mathcal{I} \cap (\mathcal{R}_k \setminus \mathcal{R}_{k+1})) = \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{I} \cap (\mathcal{R}_k \setminus \mathcal{R}_{k+1})).$$

We now have

$$\mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{J}) = \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{I} \setminus \mathcal{R}_\infty) = \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{I}),$$

where the last equality follows from Step 3. Finally, we obtain

$$\mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{U} \cup \mathcal{J}) \geq \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{U} \setminus \mathcal{I}) + \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{J}) = \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{U} \setminus \mathcal{I}) + \mathbb{P}_{h, \sigma^{\phi'}, \tau}(\mathcal{I}) = 1,$$

where the last equality follows from the fact that the sets $\mathcal{U}$ and $\mathcal{I}$ cover $\mathcal{P}(h)$. $\square$

Theorem 1.5.9 generalizes Proposition 11 in Mashiah-Yaakovi (2015), where the existence of a subgame ϵ -optimal strategy in a two-player zero-sum stochastic game with Borel measurable payoff functions is shown. The tolerance function ϕ there is given by $\phi(h) = \epsilon$ for every $h \in \mathcal{H}$, so for every play the tolerance function has a positive limit inferior. Theorem 1.5.9 yields the existence of a subgame ϵ -maxmin strategy. Because the Borel measurability of the payoff function guarantees the existence of the value (Maitra and Sudderth, 1998, and Martin, 1998), this is equivalent to proving the existence of a subgame ϵ -optimal strategy.

In Section 1.6, we argue that Theorem 1.5.9 also provides further insight into the main result of Laraki, Maitra, and Sudderth (2013). There the authors prove among other things that if the payoff function is bounded and upper semi-continuous, then the first player has a subgame optimal strategy.

1.6 Subgame maxmin strategies

The goal of this section is to explore the relationship between the concept of a subgame maxmin strategy and that of a subgame ϕ -maxmin strategy for $\phi > 0$. The main result of this section is the following theorem.

Theorem 1.6.1. For every game $\Gamma_{x_0}(u)$ one can construct a tolerance function $\phi^* > 0$ such that the following statements are equivalent:

1. The game $\Gamma_{x_0}(u)$ has a subgame maxmin strategy.
2. The game $\Gamma_{x_0}(u)$ has a subgame ϕ^* -maxmin strategy.

Because every subgame maxmin strategy is a subgame ϕ -maxmin strategy, statement 1 clearly implies statement 2. To prove the converse, we construct a tolerance function $\phi^* > 0$, the precise construction of which is discussed in Subsection 1.6.3. Intuitively, the construction goes as follows: we fix a positive sequence $(\delta_t)_{t \in \mathbb{N}}$ with the property that $\sum_{t=0}^{\infty} \delta_t < \infty$. Then choose ϕ^* such that (1) $\phi^*(p|_t) < \delta_t$ for every play $p \in \mathcal{P}$ and every $t \in \mathbb{N}$ and (2) every one-day ϕ^* -maxmin action is δ_t -close to the region of optimal mixed actions in the one-day game. The first condition implies that tolerance function ϕ^* decreases rapidly to zero, i.e. $\sum_{t=0}^{\infty} \phi^*(p|_t) < \infty$. By assumption player 1 has a subgame ϕ^* -maxmin strategy. We construct a subgame maxmin strategy as follows: at every history h_t we pick a mixed action which guarantees the lower value $\underline{v}(h_t)$ in the one-day game and is δ_t -close to the mixed action chosen by the subgame ϕ^* -maxmin strategy. This method creates a strategy which is clearly one-day maxmin. The condition that the tolerance function ϕ^* fastly decreases to 0 is then used to show that expected payoffs from a history h_t with length t obtained by the subgame ϕ^* -maxmin strategy and the constructed strategy are within a distance of $M \cdot \sum_{n=t}^{\infty} \delta_n$, where M is a constant. From this the equalizing property of the constructed strategy follows. Then Corollary 1.4.7 from Section 1.4 lets us conclude that the constructed strategy is indeed a subgame maxmin strategy.

Theorem 1.6.1 has the following corollary.

Corollary 1.6.2. The following statements are equivalent:

1. The game $\Gamma_{x_0}(u)$ has a subgame maxmin strategy.
2. For every $\phi > 0$, the game $\Gamma_{x_0}(u)$ has a subgame ϕ -maxmin strategy.

We would like to make two remarks. First, as the payoff function u need not be continuous, one cannot simply use a continuity argument to prove that statement 2 of Corollary 1.6.2 implies statement 1. Second, note that the existence of a subgame ϵ -maxmin strategy for every $\epsilon > 0$ is a weaker requirement than statements 1 and 2 of Corollary 1.6.2. Indeed, as Example 1.3.2 already illustrated, there exist games which admit a subgame ϵ -maxmin strategy for every $\epsilon > 0$ but do not admit a subgame maxmin strategy.

The special case where the payoff-function is upper semi-continuous deserves some additional attention. From Maitra and Sudderth (1998) and Martin (1998) it follows that every two-player zero-sum stochastic game with a countable state space, finite action sets, and a Borel measurable payoff function admits a value. Because every upper semi-continuous payoff function is Borel measurable, the existence of the value in our model is guaranteed. From Theorem 1.5.9 we obtain the existence of a subgame ϕ -optimal strategy for every $\phi > 0$.

Combining this with Corollary 1.6.2 shows that player 1 has a subgame optimal strategy. Hereby we obtain a special case of the result by Laraki, Maitra, and Sudderth (2013), where the authors allow the state space to be a Borel subset of a Polish space and transition probabilities to be determined by a Borel transition function.

When we are interested in pure strategies that guarantee the maxmin levels, we can strengthen the result of Theorem 1.6.1. In the context of simultaneous move games, pure strategies are of course rather restrictive. Still, there are important classes of games, such as perfect information games, in which they are natural and play a prominent role.

Theorem 1.6.3. For every game $\Gamma_{x_0}(u)$ there exists a tolerance function $\phi^* > 0$ such that the following statements are equivalent:

1. The pure strategy $\sigma \in \mathcal{S}_1$ is a subgame maxmin strategy.
2. The pure strategy $\sigma \in \mathcal{S}_1$ is a subgame ϕ^* -maxmin strategy.

This section is structured as follows. We start by proving Theorem 1.6.3, which is easier and helps us explain the main ideas. Then we turn to the proof of Theorem 1.6.1.

1.6.1 The proof of Theorem 1.6.3

In this subsection we prove Theorem 1.6.3. Since the statement of this theorem is about pure strategies, the proof is less technical and the intuition is more transparent.

We only need to prove that there is $\phi^* > 0$ such that statement 2 of Theorem 1.6.3 implies statement 1 of Theorem 1.6.3. From now on, for any $t \in \mathbb{N}$, history $h \in \mathcal{H}^t$ with some final state x , and mixed actions $m_1 \in \Delta(\mathcal{A})$ and $m_2 \in \Delta(\mathcal{B})$, we denote the expectation of the lower value at the next stage by

$$\mathbb{E}_{h,m_1,m_2} [\underline{V}^{t+1}] = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \sum_{x' \in \mathcal{X}} q(x'|h, a, b) \cdot \underline{v}(h, a, b, x'). \quad (1.6.1)$$

The proof consists of two steps.

STEP 1. Construction of $\phi^* > 0$.

For every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$, there exists a number $d(h) > 0$ such that for every action $a \in \mathcal{A}$ of player 1 either one of the following holds:

- Action a guarantees that the lower value does not drop in expectation: for every $m_2 \in \Delta(\mathcal{B})$, $\mathbb{E}_{h,a,m_2} [\underline{V}^{t+1}] \geq \underline{v}(h)$.
- There exists a mixed action $m_2 \in \Delta(\mathcal{B})$ for player 2 such that, if player 1 uses action a , the lower value drops in expectation by more than $d(h)$: $\mathbb{E}_{h,a,m_2} [\underline{V}^{t+1}] < \underline{v}(h) - d(h)$.

This statement is true because the action set $\mathcal{A}$ of player 1 is finite. We define $\phi^* > 0$ as follows. For every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$, let $\phi^*(h) = \min\{d(h), 2^{-t}\}$. The term 2^{-t} is included so that the tolerance levels tend to 0 along each play. For every $p \in \mathcal{P}$, it holds that $\lim_{t \rightarrow \infty} \phi^*(p_t) = 0$.

STEP 2. If the pure strategy $\sigma \in \mathcal{S}_1$ is a subgame ϕ^* -maxmin strategy, then σ is a subgame maxmin strategy.

Let the pure strategy $\sigma \in \mathcal{S}_1$ be subgame ϕ^* -maxmin. We verify that σ satisfies the conditions of Corollary 1.4.7.

First we show that σ is 1-day maxmin. Fix $t \in \mathbb{N}$ and $h \in \mathcal{H}^t$. Let $a = \sigma(h)$ denote the action that σ recommends at h . Then, for every mixed action $m_2 \in \Delta(\mathcal{B})$ and strategy $\tau \in \mathcal{S}_2$ with $m_2 = \tau(h)$, we have

$$\mathbb{E}_{h,a,m_2} [\underline{V}^{t+1}] = \mathbb{E}_{h,\sigma,\tau} [\underline{V}^{t+1}] \geq \underline{v}(h) - \phi^*(h) \geq \underline{v}(h) - d(h),$$

where the first inequality follows from the fact that σ is subgame ϕ^* -maxmin and by condition 1 of Theorem 1.4.6, and the second inequality follows from the definition of $\phi^*(h)$. Therefore, by the choice of $d(h)$ in Step 1, for every $m_2 \in \Delta(\mathcal{B})$ it holds that

$$\mathbb{E}_{h,a,m_2} [\underline{V}^{t+1}] \geq \underline{v}(h).$$

Hence, for every strategy $\tau \in \mathcal{S}_2$ it holds that

$$\mathbb{E}_{h,\sigma,\tau} [\underline{V}^{t+1}] = \mathbb{E}_{h,a,\tau(h)} [\underline{V}^{t+1}] \geq \underline{v}(h).$$

Thus, σ is 1-day maxmin.

We show that σ is equalizing. Take some $\tau \in \mathcal{S}_2$. Because σ is subgame ϕ^* -maxmin, for every $p \in \mathcal{P}(h)$, for every $n \geq t$, we have

$$U_{\sigma,\tau}^n(p) = \mathbb{E}_{p|n,\sigma,\tau} [u] \geq \underline{v}(p|n) - \phi^*(p|n) = \underline{V}^n(p) - \Phi^{*n}(p),$$

where $\Phi^{*n}(p) = \phi^*(p|n)$. We conclude that

$$\lim_{n \rightarrow \infty} U_{\sigma,\tau}^n \geq \limsup_{n \rightarrow \infty} (\underline{V}^n - \Phi^{*n}) = \limsup_{n \rightarrow \infty} \underline{V}^n - \lim_{n \rightarrow \infty} \Phi^{*n} = \limsup_{n \rightarrow \infty} \underline{V}^n,$$

where the last equality from the fact that for all $p \in \mathcal{P}$ we have $\lim_{n \rightarrow \infty} \phi^*(p|n) = 0$, so

$$u = \lim_{n \rightarrow \infty} U_{\sigma,\tau}^n \geq \limsup_{n \rightarrow \infty} \underline{V}^n, \quad \mathbb{P}_{h,\sigma,\tau}\text{-almost surely,}$$

where the equality follows from Lemma A.2 in Appendix A. We have shown that σ is equalizing.

1.6.2 The one-shot game Υ_h

To prove Theorem 1.6.1, we first analyze a one-shot zero-sum game in this subsection. For each history, the one-shot game is such that the payoff is given by the lower value at the next stage. In the next subsection, we use these one-shot games to construct the tolerance function $\phi^* > 0$.

For some $t \in \mathbb{N}$, let $h \in \mathcal{H}^t$ be a history in the game $\Gamma_{x_0}(u)$. The one-shot zero-sum game Υ_h is played as follows. Player 1 chooses an action $a \in \mathcal{A}$ and player 2 simultaneously chooses an action $b \in \mathcal{B}$. Then, player 1 receives from player 2 the amount $\mathbb{E}_{h,a,b} [\underline{V}^{t+1}]$. As the action sets $\mathcal{A}$ and $\mathcal{B}$ are finite, the game Υ_h has a value, which we denote by $w(h)$. Furthermore, both players have optimal mixed actions in the game Υ_h .

The following lemma states that the value $w(h)$ of the one-shot game Υ_h equals the lower value $\underline{v}(h)$ at the history h in the original game $\Gamma_{x_0}(u)$.

Lemma 1.6.4. For every history $h \in \mathcal{H}$, we have $w(h) = \underline{v}(h)$.

Proof. The proof is by contradiction. Fix $t \in \mathbb{N}$ and a history $h \in \mathcal{H}^t$. Now suppose that $w(h) \neq \underline{v}(h)$. Then we have either $w(h) > \underline{v}(h)$ or $w(h) < \underline{v}(h)$.

CASE 1: $w(h) > \underline{v}(h)$.

Let $\delta = w(h) - \underline{v}(h)$. We derive a contradiction by showing that, in the subgame of $\Gamma_{x_0}(u)$ at history h , player 1 can guarantee an expected payoff of at least $\underline{v}(h) + \delta/2$.

Let $m_1 \in \Delta(\mathcal{A})$ be an optimal mixed action for player 1 in the one-shot game Υ_h . Let $\sigma \in \mathcal{S}_1$ be such that $\sigma(h) = m_1$ and such that it induces a $(\delta/2)$ -maxmin strategy for the subgame at each history in period $t+1$, i.e., for every $h' \in \mathcal{H}^{t+1}$, for every $\tau \in \mathcal{S}_2$,

$$\mathbb{E}_{h',\sigma,\tau}[u] \geq \underline{v}(h') - \frac{\delta}{2}. \quad (1.6.2)$$

Then, for every $\tau \in \mathcal{S}_2$, it holds that

$$\begin{aligned} \mathbb{E}_{h,\sigma,\tau}[u] &= \mathbb{E}_{h,\sigma,\tau}[U_{\sigma,\tau}^t] = \mathbb{E}_{h,\sigma,\tau}[U_{\sigma,\tau}^{t+1}] \\ &\geq \mathbb{E}_{h,\sigma,\tau}[V_{\sigma,\tau}^{t+1}] - \frac{\delta}{2} = \mathbb{E}_{h,m_1,\tau(h)}[V^{t+1}] - \frac{\delta}{2} \geq w(h) - \frac{\delta}{2} = \underline{v}(h) + \frac{\delta}{2}, \end{aligned}$$

where the second equality follows from the fact that $(U_{\sigma,\tau}^n)_{n \geq t}$ is a $\mathbb{P}_{h,\sigma,\tau}$ -martingale, the first inequality follows from (1.6.2), and the second inequality follows from the choice of m_1 .

CASE 2: $w(h) < \underline{v}(h)$.¹

Let $\delta = \underline{v}(h) - w(h)$. We derive a contradiction by showing that, for every strategy of player 1, there is a strategy for player 2 such that the expected payoff is at most $\underline{v}(h) - \delta/2$ in the subgame of $\Gamma_{x_0}(u)$ at history h .

Fix $\sigma \in \mathcal{S}_1$ and let $m_1 = \sigma(h)$. Let $m_2 \in \Delta(\mathcal{B})$ be an optimal mixed action for player 2 in the one-shot game Υ_h . Let $\tau \in \mathcal{S}_2$ be such that $\tau(h) = m_2$ and the expected payoff under (σ, τ) in the subgame at each history h' at period $t+1$ is at most the lower value $\underline{v}(h') + \delta/2$, i.e., for every $h' \in \mathcal{H}^{t+1}$,

$$\mathbb{E}_{h',\sigma,\tau}[u] \leq \underline{v}(h') + \frac{\delta}{2}. \quad (1.6.3)$$

We have that

$$\begin{aligned} \underline{v}(h) - \frac{\delta}{2} &= w(h) + \frac{\delta}{2} \geq \mathbb{E}_{h,m_1,m_2}[V^{t+1}] + \frac{\delta}{2} \\ &\geq \mathbb{E}_{h,m_1,m_2}[U_{\sigma,\tau}^{t+1}] = \mathbb{E}_{h,\sigma,\tau}[U_{\sigma,\tau}^{t+1}] = \mathbb{E}_{h,\sigma,\tau}[U_{\sigma,\tau}^t] = \mathbb{E}_{h,\sigma,\tau}[u], \end{aligned}$$

where the first inequality follows from the choice of m_2 , the second inequality follows from (1.6.3), and the penultimate equality follows from the fact that $(U_{\sigma,\tau}^n)_{n \geq t}$ is a $\mathbb{P}_{h,\sigma,\tau}$ -martingale. Because σ is chosen arbitrarily, we have derived a contradiction with the definition of the lower value $\underline{v}(h)$. $\square$

The total variation distance between two mixed actions $m_1, n_1 \in \Delta(\mathcal{A})$ of player 1 is defined as

$$\|m_1 - n_1\|_{\text{TV}} = \sum_{a \in \mathcal{A}} |m_1(a) - n_1(a)|.$$

The total variation distance between two probability measures $\mathbb{P}$ and $\mathbb{P}'$ on $(\mathcal{P}, \mathcal{F})$ is defined as

$$\|\mathbb{P} - \mathbb{P}'\|_{\text{TV}} = \sup \left\{ \sum_{i=1}^n |\mathbb{P}(F_i) - \mathbb{P}'(F_i)| : F_1, \dots, F_n \in \mathcal{F} \text{ and } \{F_1, \dots, F_n\} \text{ is a partition of } \mathcal{P} \right\}.$$

¹The proof of this case is not symmetric to the proof of Case 1, because we consider the lower value. Imitating the proof of Case 1 for player 2 would yield results in terms of the upper value.

Let $t \in \mathbb{N}$ and a history $h \in \mathcal{H}^t$ be given. Let $O_h \subseteq \Delta(\mathcal{A})$ denote the set of optimal mixed actions of player 1 in the one-shot game Υ_h . By Lemma 1.6.4 it holds that

$$O_h = \{m_1 \in \Delta(\mathcal{A}) \mid \text{for every } m_2 \in \Delta(\mathcal{B}), \mathbb{E}_{h,m_1,m_2} [V^{t+1}] \geq \underline{v}(h)\}.$$

Note that O_h is a compact subset of $\Delta(\mathcal{A})$. For every $m_1 \in \Delta(\mathcal{A})$, the distance of m_1 to O_h is defined by

$$\|m_1 - O_h\|_{\text{TV}} = \min_{n_1 \in O_h} \|m_1 - n_1\|_{\text{TV}}.$$

Due to the compactness of O_h , the minimum is attained. For $\delta > 0$, let D_h^δ be the set of mixed actions of player 1 which have a distance of at least δ to the set O_h , so

$$D_h^\delta = \{m_1 \in \Delta(\mathcal{A}) \mid \|m_1 - O_h\|_{\text{TV}} \geq \delta\}. \quad (1.6.4)$$

The mixed actions in D_h^δ are not optimal in the one-shot game Υ_h . The following lemma says that the loss in utility caused by these mixed actions has a positive lower bound.

Lemma 1.6.5. Let $h \in \mathcal{H}$ and $\delta > 0$ be given. If D_h^δ is non-empty, then there is $\epsilon > 0$ such that for every $m_1 \in D_h^\delta$ there exists $b \in \mathcal{B}$ such that

$$\mathbb{E}_{h,m_1,b} [V^{t+1}] \leq \underline{v}(h) - \epsilon.$$

Proof. Assume D_h^δ is non-empty. Consider the function $e_h^\delta : D_h^\delta \rightarrow \mathbb{R}$ defined by

$$e_h^\delta(m_1) = \underline{v}(h) - \min_{b \in \mathcal{B}} \mathbb{E}_{h,m_1,b} [V^{t+1}], \quad m_1 \in D_h^\delta. \quad (1.6.5)$$

Since $\mathcal{B}$ is finite, the minimum exists. For every $m_1 \in D_h^\delta$, we have $m_1 \notin O_h$ and therefore there exists $m_2 \in \Delta(\mathcal{B})$ such that $\mathbb{E}_{h,m_1,m_2} [V^{t+1}] < \underline{v}(h)$. The function e_h^δ is therefore a positive and continuous function on a compact set, so has a positive minimum. $\square$

1.6.3 Construction of the tolerance function ϕ^*

In this subsection, we define a positive tolerance function ϕ^* . Fix a positive and decreasing sequence $(\delta_t)_{t \in \mathbb{N}}$ such that $\sum_{t=0}^\infty \delta_t < \infty$. Notice that this implies $\lim_{t \rightarrow \infty} \delta_t = 0$.

For every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$, we define the constant $c(h)$ as follows. If the set $D_h^{\delta_t}$ is non-empty, then $c(h)$ is equal to the positive number ϵ of Lemma 1.6.5 and $c(h) = \delta_t$ otherwise. We define

$$\phi^*(h) = \frac{\min\{c(h), \delta_t\}}{2}. \quad (1.6.6)$$

Notice that $0 < \phi^*(h) < \delta_t$.

We summarize the properties of the tolerance function ϕ^* :

1. For every history $h \in \mathcal{H}$, we have $\phi^*(h) > 0$.
2. For every play $p \in \mathcal{P}$, we have $\sum_{t=0}^\infty \phi^*(p|_t) \leq \sum_{t=0}^\infty \delta_t < \infty$.
3. For every play $p \in \mathcal{P}$, $\lim_{t \rightarrow \infty} \phi^*(p|_t) = 0$.
4. If the set $D_h^{\delta_t}$ is non-empty, then by the choice of $c(h)$, for every $m_1 \in D_h^{\delta_t}$ there exists $b \in \mathcal{B}$ such that

$$\mathbb{E}_{h,m_1,b} [V^{t+1}] \leq \underline{v}(h) - c(h) < \underline{v}(h) - \phi^*(h).$$

The importance of $\sum_{t=0}^{\infty} \delta_t < \infty$ is underlined by the following lemma. Recall that $M = \sup_{p \in \mathcal{P}} |u(p)|$.

Lemma 1.6.6. Let the strategies $\sigma, \sigma' \in \mathcal{S}_1$ be such that, for every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$, $\|\sigma(h) - \sigma'(h)\|_{\text{TV}} \leq \delta_t$. Then, for every strategy $\tau \in \mathcal{S}_2$, for every $t \in \mathbb{N}$, and for every history $h \in \mathcal{H}^t$,

$$|\mathbb{E}_{h,\sigma,\tau}[u] - \mathbb{E}_{h,\sigma',\tau}[u]| \leq M \cdot \sum_{n=t}^{\infty} \delta_n.$$

Proof. Let $\tau \in \mathcal{S}_2$ be given. It follows from a more general result in Abate, Redig, and Tkachev (2014, theorem 1) that, for every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$,

$$\|\mathbb{P}_{h,\sigma,\tau} - \mathbb{P}_{h,\sigma',\tau}\|_{\text{TV}} \leq \sum_{n=t}^{\infty} \delta_n. \quad (1.6.7)$$

For completeness, we provide a direct proof of this inequality in Lemma B.1 in Appendix B. The claim of Lemma 1.6.6 follows directly. $\square$

Lemma 1.6.6 says the following. Consider two arbitrary strategies $\sigma, \sigma' \in \mathcal{S}_1$ such that the total variation distance between the mixed actions at every history in period t is at most δ_t . Now consider $t \in \mathbb{N}$, a history $h \in \mathcal{H}^t$, and a strategy $\tau \in \mathcal{S}_2$ of player 2. Then the expected payoffs under (σ, τ) and (σ', τ) in the subgame at h differ at most the constant M times $\sum_{n=t}^{\infty} \delta_n$. Note that this bound does not depend on the strategy τ and it only depends on the history h through its period t . Moreover, these bounds tend to 0 as t goes to infinity.

The importance of property 4 of the tolerance function is shown by the following lemma. It says that if $\sigma \in \mathcal{S}_1$ is a subgame ϕ^* -maxmin strategy, then for every $h \in H$ the mixed action $\sigma(h)$ is close to the set of optimal mixed actions O_h in the one-shot game Υ_h .

Lemma 1.6.7. Let $\sigma \in \mathcal{S}_1$ be a subgame ϕ^* -maxmin strategy. Then, for every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$, we have $\sigma(h) \notin D_h^{\delta_t}$, so $\|\sigma(h) - O_h\|_{\text{TV}} < \delta_t$.

Proof. Because σ is a subgame ϕ^* -maxmin strategy, it follows from condition 1 of Theorem 1.4.6 that for every mixed action $m_2 \in \Delta(\mathcal{B})$ of player 2

$$\mathbb{E}_{h,\sigma(h),m_2} [\underline{V}^{t+1}] \geq \underline{v}(h) - \phi^*(h).$$

If $D_h^{\delta_t}$ is empty, then there is nothing to prove. If $D_h^{\delta_t}$ is non-empty, then property 4 of ϕ^* shows that $\sigma(h) \notin D_h^{\delta_t}$. $\square$

1.6.4 The proof of Theorem 1.6.1

Proof. Let $\Gamma_{x_0}(u)$ be a game and take the tolerance function ϕ^* as defined in Subsection 1.6.3. We only need to show that statement 2 implies statement 1. Let $\sigma \in \mathcal{S}_1$ be a subgame ϕ^* -maxmin strategy of $\Gamma_{x_0}(u)$.

With the help of σ , we define a strategy $\sigma^* \in \mathcal{S}_1$ in Step 1 of the proof. Then it is shown that σ^* is a subgame maxmin strategy of $\Gamma_{x_0}(u)$ in Steps 2 and 3 of the proof by verifying that σ^* satisfies the conditions of Corollary 1.4.7.

STEP 1: Definition of $\sigma^* \in \mathcal{S}_1$.

Take $t \in \mathbb{N}$ and a history $h \in \mathcal{H}^t$. In view of Lemma 1.6.7, it holds that $\|\sigma(h) - O_h\|_{TV} < \delta_t$. Therefore, there exists $m^*(h) \in O_h$ such that

$$\|\sigma(h) - m^*(h)\|_{TV} < \delta_t. \quad (1.6.8)$$

Now define $\sigma^*(h) = m^*(h)$.

STEP 2: σ^* is 1-day maxmin.

Consider some $\tau \in \mathcal{S}_2$. For every $t \in \mathbb{N}$, for every $h \in \mathcal{H}^t$, since $\sigma^*(h) \in O_h$ we have that

$$\mathbb{E}_{h, \sigma^*, \tau} [V^{t+1}] \geq \underline{v}(h).$$

STEP 3: σ^* is equalizing.

Consider some $\tau \in \mathcal{S}_2$. In view of (1.6.8) we can apply Lemma 1.6.6 to conclude that, for every $t \in \mathbb{N}$, for every $h \in \mathcal{H}^t$,

$$|\mathbb{E}_{h, \sigma, \tau}[u] - \mathbb{E}_{h, \sigma^*, \tau}[u]| \leq M \cdot \sum_{n=t}^{\infty} \delta_n.$$

Hence, for every $t \in \mathbb{N}$, for every history $h \in \mathcal{H}^t$, and for every $n \geq t$,

$$|U_{\sigma, \tau}^n[u] - U_{\sigma^*, \tau}^n[u]| \leq M \cdot \sum_{i=n}^{\infty} \delta_i, \quad \mathbb{P}_{h, \sigma^*, \tau}\text{-almost surely.} \quad (1.6.9)$$

Because σ is subgame ϕ^* -maxmin, for every history $h \in \mathcal{H}$, we have

$$\mathbb{E}_{h, \sigma, \tau}[u] \geq \underline{v}(h) - \phi^*(h). \quad (1.6.10)$$

Thus, for every history $h \in \mathcal{H}$ it holds that

$$u = \lim_{t \rightarrow \infty} U_{\sigma^*, \tau}^t = \lim_{t \rightarrow \infty} U_{\sigma, \tau}^t \geq \limsup_{t \rightarrow \infty} (\underline{V}^t - \Phi^{*t}) = \limsup_{t \rightarrow \infty} \underline{V}^t, \quad \mathbb{P}_{h, \sigma^*, \tau}\text{-almost surely,}$$

where the first equality is due to Lemma A.2 in Appendix A, the second equality follows from (1.6.9), the inequality is by (1.6.10), and the last equality is a consequence of property 3 of ϕ^* from Subsection 1.6.3. $\square$

Note that it is not necessary that the initial δ_t are small, or that for initial histories the tolerance level $\phi^*(h)$ is small. The conditions put on the sequence $(\delta_t)_{t \in \mathbb{N}}$ are there to ensure that the equalizing condition for subgame maxmin strategies is satisfied. The equalizing condition essentially only cares about what happens to the payoff at deep subgames. The requirement that the sequence $(\delta_t)_{t \in \mathbb{N}}$ decreases to zero very fast guarantees that $\lim_{t \rightarrow \infty} \|\mathbb{P}_{p|t, \sigma, \tau} - \mathbb{P}_{p|t, \sigma^*, \tau}\|_{TV} = 0$. This allows us to use the ϕ^* -equalizing property of the subgame ϕ^* -maxmin strategy to show that the strategy σ^* satisfies the equalizing condition.

1.7 Discussion

In the previous sections we have assumed that action sets are finite and the state space is countable. The goal of this section is to analyze these assumptions further and to pinpoint where and how they were used.

Throughout this paper the restrictions on the cardinalities of action sets and state space were used to ensure the following properties:

1. The measurability of the lower value.
2. The existence of 1-day optimal actions in the game Υ_h .

The measurability of the lower value is crucial for the sufficient (Theorem 1.4.5) and necessary (Theorem 1.4.6) conditions of subgame ϕ -maxmin strategies as well as for the characterization result (Corollary 1.4.7) for subgame maxmin strategies. Because the results in Sections 1.5 and 1.6 rely on this sufficient condition, the measurability of the lower value is indispensable throughout the paper. When working with infinite action sets and uncountable state spaces, Nowak (1985) and the references therein demonstrate that the (lower) value is not necessarily measurable.

Apart from ensuring that the lower value is measurable, the finiteness of action sets is used in Section 1.6 to guarantee the existence of 1-day optimal mixed actions in the game Υ_h . Indeed, if action sets are finite, then for every history $h \in \mathcal{H}$ the game Υ_h is a finite zero-sum game and hence both players have optimal strategies.

Chapter 2

Individual upper semicontinuity and subgame perfect ϵ -equilibria in games with almost perfect information

2.1 Introduction

Games with almost perfect information play a prominent role in the theory of dynamic games. In these games, at each stage, the players simultaneously choose actions, knowing the actions chosen at all previous stages. The payoff of each player is determined by the sequence of actions that have been chosen during the entire game. The notion of subgame perfect equilibrium and the weaker notion of subgame perfect ϵ -equilibrium, where $\epsilon > 0$ is an error term, are two main solution concepts in these games. We continue the discussion of these games under the assumptions that the number of players is finite and the payoffs are bounded.

A major body of the literature on games with almost perfect information examines the existence of subgame perfect equilibrium while assuming some kind of continuity of the payoff functions and compactness of the action spaces. Fudenberg and Levine (1983) showed that if the payoff functions are continuous and the action spaces are finite, then there exists a subgame perfect equilibrium. With infinite but compact action spaces the existence question becomes very subtle, see for instance Harris, Reny and Robson (1995), Mertens and Parthasarathy (2003), or Maitra and Sudderth (2007). For a detailed overview, we refer the reader to the recent survey by Jaśkiewicz and Nowak (2016).

In this paper we continue the investigation of the existence of subgame perfect ϵ -equilibria in games with almost perfect information. We assume that the set of available actions is finite. The sequence of actions chosen in the game is referred to as a play. The contribution of this paper is twofold. First, we introduce a new condition, called individual upper semicontinuity, on the payoff functions. Individual upper semicontinuity is reminiscent of upper semicontinuity, but it uses a new notion of convergence on the set of plays for each player separately.

Consider a player i . We say that a sequence of plays $(p_m)_{m \in \mathbb{N}}$ is i -convergent to a play p if (1) the length of the common history between the play p_m and p tends to infinity, and (2) for large $m \in \mathbb{N}$, either $p_m = p$ or the first deviation from p to p_m involves only player i . This

We thank Jérôme Renault, whose remarks inspired us to investigate this research question.

notion of convergence of plays is stronger than the conventional notion, which would only require property (1). That is, each i -convergent sequence of plays is also convergent in the conventional sense, but not vice versa. Based on the notion of i -convergence, we say that the payoff function of player i is i -upper semicontinuous if for each sequence of plays $(p_m)_{m \in \mathbb{N}}$ that is i -convergent to a play p , player i weakly prefers the payoff along p compared to the payoff along p_m for large m . If player i 's payoff function is upper semicontinuous in the conventional sense, then it is also i -upper semicontinuous, but not vice versa. We analyze the topological properties of i -convergence and the properties of i -upper semicontinuous payoff functions in detail.

Secondly, we show that if each player i 's payoff function is i -upper semicontinuous, then the game admits a subgame perfect ϵ -equilibrium for every $\epsilon > 0$. Moreover, these strategy profiles are eventually pure, meaning that in each subgame randomization is only used at finitely many stages with probability 1. The proof of our existence result is constructive. The key idea is that if each player i 's payoff function is i -upper semicontinuous, then each play eventually reaches a history after which no player has an incentive to deviate from this play. This enables us to essentially cut the time horizon of the game at some finite stages and to apply backward induction on the earlier stages. This existence result generalizes the result of Purves and Sudderth (2011), who study perfect information games with upper semicontinuous payoff functions, in two directions. We reduce the topological restrictions on the payoff functions of the players and we allow the players to move simultaneously.¹ Using the new concept of i -convergence we can weaken the topological structure on the game while still maintaining sufficient conditions that guarantee the existence of a subgame perfect ϵ -equilibrium. In normal form games with totally ordered compact strategy spaces, Prokopovych and Yannelis (2017) show that it is possible to ensure the existence of a pure Nash equilibrium, while significantly relaxing the requirements relating to the upper semicontinuity and single crossing properties of the payoff functions. Alós-Ferrer and Ritzberger (2017) study the problem of finding minimal topological conditions needed to guarantee the existence of a subgame perfect equilibrium in an extensive form game with continuous payoff functions.

Games with perfect information are an important special case of games with almost perfect information. Conditions related to semicontinuity of the payoff functions play an important role in several results for perfect information games. Not only do perfect information games where every player has an upper semicontinuous payoff function admit a subgame perfect ϵ -equilibrium for every $\epsilon > 0$ (Purves and Sudderth, 2011), so do perfect information games where every player has a lower semicontinuous payoff function as proven in Flesch et al. (2010). Moreover, these results were extended and unified in Flesch and Predtetchinski (2016). Questions related to semicontinuity are also studied in Le Roux and Pauly (2016) and Bruyère, Le Roux, Pauly and Raskin (2017). A counterexample by Flesch et al. (2014) shows that perfect information games with finite action sets and Borel measurable payoff functions do not always have subgame perfect ϵ -equilibria. This further illustrates the importance of topological properties of payoff functions in equilibrium analysis. For an overview on subgame perfect ϵ -equilibria in perfect information games, we refer to Jaśkiewicz and Nowak (2016) and Bruyère (2017). We remark that it follows from a result of Secchi and Sudderth (2001) that if the players have upper semicontinuous payoff functions in a stochastic game with a countable state space, then the game admits a Nash

¹Purves and Sudderth (2011) allow arbitrary action sets in perfect information games. In Section 2.6 we argue that for perfect information games, we can dispense with the assumption that the action sets are finite.

ϵ -equilibrium for every $\epsilon > 0$.

The remainder of this paper is structured as follows. In Section 2.2, we describe the model and introduce the notions of i -convergence for plays and i -upper semicontinuity for payoff functions of player i . In Section 2.3, we present four illustrating examples. In Section 2.4, we prove the above-mentioned existence result for subgame perfect ϵ -equilibria. In Section 2.5, we study the topology induced by i -convergence in detail. Finally, in Section 2.6, we provide some concluding remarks.

2.2 The model

In this section we describe the model and define the notion of individual upper semicontinuity of payoff functions. Let $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}^* = \{0, 1, 2, \dots\}$.

The game: We consider games with almost perfect information and an infinite time horizon. We assume that the set of players and the sets of available actions are finite. Such a game is given by a tuple $G = \{I, (A_i)_{i \in I}, H, (u_i)_{i \in I}\}$, where

1. I is a non-empty and finite set of players.
2. For each player $i \in I$, A_i is a non-empty finite set of actions. Let $A = \times_{i \in I} A_i$. The set A corresponds to the set of stage game outcomes.
3. $H \subseteq \cup_{t \in \mathbb{N}^*} A^t$ is a non-empty set of histories, where A^0 is the singleton $\{\emptyset\}$. At each history $h \in H$, let $A(h) = \{a \in A \mid ha \in H\}$ and for each player $i \in I$ let $A_i(h)$ denote the projection of $A(h)$ on A_i . The set $A(h)$ corresponds to the set of stage game outcomes at history h and $A_i(h)$ to the set of available actions for player i at history h . We assume that: (i) H is closed under initial segments: if for some $t \in \mathbb{N}^*$ and $ha \in A^{t+1}$ we have $ha \in H$, then we also have $h \in H$, (ii) for each history $h \in H$ and each player $i \in I$, the set $A_i(h)$ is non-empty, (iii) the set of stage game outcomes at history h is a product set: for each $h \in H$ it holds that $A(h) = \times_{i \in I} A_i(h)$.
4. Let P denote the set of all sequences $(a^1, a^2, \dots) \in A^{\mathbb{N}}$ such that for every $t \in \mathbb{N}$ we have $(a^1, \dots, a^t) \in H$. The set P is called the set of plays.² For each $i \in I$, $u_i : P \rightarrow \mathbb{R}$ is the payoff function of player i .

The game is played at stages in $\mathbb{N}$. At stage 1, each player $i \in I$ chooses an action a_i^1 from the set $A_i(\emptyset)$, independently of the other players. This yields a stage game outcome $a^1 = (a_i^1)_{i \in I}$. At a general stage $t \in \mathbb{N}$, knowing the previous stage game outcomes $a^1, \dots, a^{t-1}$, each player $i \in I$ chooses an action a_i^t from the set $A_i(a^1, \dots, a^{t-1})$, independently of the other players. This yields a stage game outcome $a^t = (a_i^t)_{i \in I}$. The payoff of each player $i \in I$ is given by $u_i(a^1, a^2, \dots)$.

Notations for histories and plays: For every $t \in \mathbb{N}^*$, for every history $h = (a^1, \dots, a^t) \in H$, let $\|h\| = t$ denote the number of stage games played during h . We refer to $\|h\|$ as the length of history h . For $k \in \mathbb{N}^*$ such that $k \leq t$, let $h|_k = (a^1, \dots, a^k)$ denote the truncated history

²Properties (3.i) and (3.ii) mean that H is a pruned tree on A . Thus, P can be seen as the set of all infinite branches of H .

available after stage game k . If $h, h' \in H$ are such that there exists $t \in \mathbb{N}^*$ for which $h|_t = h'$, then h is called an initial segment of h' , denoted by $h \leq h'$. We write $h < h'$ if $h \leq h'$ and $h \neq h'$. Furthermore if $h \leq h'$ then we can define the maximum of these histories as $\max(h, h') = h'$ and the minimum of these histories as $\min(h, h') = h$.

For every $t \in \mathbb{N}^*$, let $p|_t$ denote the history that arises by restricting p to the first t stages. Let $p|_{[t],i} = a_i^t$ denote the action that player i played at stage t along the play p . For a history $h \in H$ and a play $p \in P$ we write $h < p$ if p is an extension of h , i.e. for some $t \in \mathbb{N}^*$ we have $p|_t = h$. We call h a prefix of p . For any two distinct plays $p, q \in P$, let $\min(p, q)$ denote the longest common history shared by those plays and let $\ell(p, q)$ denote its length. Furthermore, define $\min(p, p) = p$ and $\ell(p, p) = \infty$. For any two distinct plays $p, q \in P$, let $I(p, q) \subseteq I$ denote the subset of players who deviated first from the play p to the play q , i.e. if $\ell(p, q) = t$ then $I(p, q) = \{i \in I | p|_{[t+1],i} \neq q|_{[t+1],i}\}$. We formally define $I(p, p) = \emptyset$. Note that I, ℓ , and $\min$ are all symmetric in their two arguments.

Strategies: A mixed action for player $i \in I$ at history $h \in H$ is a probability measure on $A_i(h)$. A strategy for player $i \in I$ is a mapping σ_i that assigns to each history $h \in H$ a mixed action $\sigma_i(h)$ for player i at history h . Let $\mathcal{S}_i$ denote the set of all strategies of player i . If for every history $h \in H$ the mixed action $\sigma_i(h)$ places probability 1 on one action, then σ_i is called a *pure strategy*.

A tuple of strategies $\sigma = (\sigma_i)_{i \in I}$ is called a strategy profile. The set of strategy profiles is denoted by $\mathcal{S} = \times_{i \in I} \mathcal{S}_i$. For each player $i \in I$, let $\sigma_{-i} = (\sigma_j)_{j \in I \setminus \{i\}}$ denote the profile of strategies of player i 's opponents. The strategy profile σ is called pure if each player i 's strategy σ_i is pure.

A pure strategy profile $\sigma \in \mathcal{S}$ induces a unique play from every history h , which we denote by $\pi(\sigma; h)$. For a general strategy profile $\sigma \in \mathcal{S}$, let H_σ denote the set of histories h with the following property: there is a play $p = (h, a^{\ell(h)+1}, a^{\ell(h)+2}, \dots)$ such that for each stage $k \geq \ell(h)$ and for each player $i \in I$, the mixed action $\sigma_i(h, a^{\ell(h)+1}, \dots, a^k)$ places probability 1 on action a_i^{k+1} . Intuitively, this means that at history h , the strategy profile σ induces the play p with probability 1. The strategy profile σ is called *eventually pure in every subgame* if for each history $h \in H$ and each play $p > h$ there is a history h' such that $h \leq h' < p$ and $h' \in H_\sigma$. This means that starting at any history, any continuation play reaches a history after which the strategy profile σ induces a unique play.

Topology on the set of plays: We now define the standard cylinder topology on the set of plays P .³ For each history $h \in H$, let $P(h)$ denote the set of all plays extending h , i.e. $P(h) = \{p \in P | p > h\}$. Let $\mathcal{T}$ denote the topology on the set of plays P induced by the collection $\{P(h) | h \in H\}$. It is easy to see that the collection $\{P(h) | h \in H\}$ forms a basis of the topology $\mathcal{T}$. That is, for a set $Q \subseteq P$, we have $Q \in \mathcal{T}$ exactly when Q can be written as a union of sets belonging to $\{P(h) | h \in H\}$. The topological space $(P, \mathcal{T})$ is metrizable. For example, the metric $d : P \times P \rightarrow [0, \infty)$ defined by $d(p, q) = 2^{-\ell(p, q)}$ induces the topology $\mathcal{T}$. Moreover, $\mathcal{T}$ coincides with the relative topology inherited from the product topology on $A^\mathbb{N}$. Since the latter space is compact and P is a closed subset of it, we conclude that $(P, \mathcal{T})$ is compact.

Let Σ denote the Borel sigma-algebra corresponding to the topology $\mathcal{T}$. Then, (P, Σ) is a measurable space. Similarly, for every history $h \in H$, we can construct the measurable

³For further reading, we refer to Kechris (1995).

space $(P(h), \Sigma_h)$, where Σ_h is the sigma-algebra generated by all sets $P(h')$ with $h' \geq h$. We have $\Sigma_h = \{Q \cap P(h) | Q \in \Sigma\}$.

As a consequence of the Ionescu Tulcea extension theorem, every strategy profile $\sigma \in \mathcal{S}$ induces for every history $h \in H$ a probability measure $\mathbb{P}_{h,\sigma}$ on the measurable space $(P(h), \Sigma_h)$. Note that σ is pure exactly when for each history $h \in H$, $\mathbb{P}_{h,\sigma}$ is a Dirac measure. Also, σ is eventually pure in every subgame exactly when for each history $h \in H$ and each play $p \succ h$ there exists a history h' such that $h \leq h' < p$ and $\mathbb{P}_{h',\sigma}$ is a Dirac measure.

Convergence and upper semicontinuity: In the topology $\mathcal{T}$, a sequence of plays $(p_m)_{m \in \mathbb{N}}$ is convergent to the play p if $\lim_{m \rightarrow \infty} \ell(p_m, p) = \infty$. The payoff function u_i of player i is called upper semicontinuous if for each play $p \in P$ and each sequence $(p_m)_{m \in \mathbb{N}}$ that is convergent to p , we have $\limsup_{m \rightarrow \infty} u_i(p_m) \leq u_i(p)$. If u_i is upper semicontinuous, then it is also Borel measurable.

We now strengthen the notion of convergence. Consider a player $i \in I$. We say that a sequence of plays $(p_m)_{m \in \mathbb{N}}$ is *i-convergent* to the play p and write $\lim_{m \rightarrow \infty}^{(i)} p_m = p$ if $\lim_{m \rightarrow \infty} \ell(p_m, p) = \infty$ and there exists $M \in \mathbb{N}$ such that for every $m \geq M$, $I(p_m, p) \subseteq \{i\}$. The notion of *i-convergence* thus strengthens the notion of convergence by additionally imposing that eventually, if the plays p_m and p differ, player i is the only player who causes the first difference between p_m and p . Clearly, if a sequence of plays is *i-convergent* then it is also convergent.

In turn, this leads to a weakening of upper semicontinuity that we refer to as individual upper semicontinuity. The payoff function u_i of player i is called *i-upper semicontinuous* if for each play $p \in P$ and each sequence $(p_m)_{m \in \mathbb{N}}$ that is *i-convergent* to p , i.e. $\lim_{m \rightarrow \infty}^{(i)} p_m = p$, we have $\limsup_{m \rightarrow \infty} u_i(p_m) \leq u_i(p)$. As each *i-convergent* sequence of plays is also convergent, it holds that each upper semicontinuous payoff function is also *i-upper semicontinuous*. The converse however is not true and there are even *i-upper semicontinuous* payoff functions that are not Borel measurable. These issues are illustrated by Example 2.3.1 in Section 2.3.

Subgame perfect ϵ -equilibrium: Assume that for every player $i \in I$ the payoff function u_i is bounded and Borel measurable. Then, for $\epsilon \geq 0$, a strategy profile $\sigma \in \mathcal{S}$ is called a *subgame perfect ϵ -equilibrium* if for each history $h \in H$, each player $i \in I$, and each strategy $\sigma'_i \in \mathcal{S}_i$, we have

$$\mathbb{E}_{h,\sigma} [u_i] \geq \mathbb{E}_{h,\sigma'_i,\sigma_{-i}} [u_i] - \epsilon.$$

In other words, σ is a subgame perfect ϵ -equilibrium if at each history $h \in H$ it induces a Nash ϵ -equilibrium. When $\epsilon = 0$, the concept of subgame perfect 0-equilibrium coincides with the usual concept of subgame perfect equilibrium.

2.3 Examples

In this section we discuss a few illustrative examples.

The first example demonstrates that *i-upper semicontinuity* of u_i does not imply that u_i is upper semicontinuous. In fact, it does not even imply that u_i is Borel measurable.

Example 2.3.1. Consider the following game with two players. At each stage, player 1 chooses an action from the set $\{1, 2\}$. Player 2 is a dummy player, who can only choose

action 0 at each stage. The set of plays is thus $P = (\{1, 2\} \times \{0\})^{\mathbb{N}}$. Let Q be a non-Borel set of P .⁴ Let the payoff function u_2 of player 2 be defined as $u_2(p) = 1$ if $p \in Q$ and $u_2(p) = 0$ if $p \in P \setminus Q$.

The payoff function u_2 is clearly not Borel measurable and hence not upper semicontinuous either. However, u_2 is 2-upper semicontinuous. Indeed, take any sequence of plays $(p_m)_{m \in \mathbb{N}}$ that is 2-converging to a play p . Since player 2 is a dummy player, we have $p_m = p$ for large m . Hence, $u_2(p_m) = u_2(p)$ for large m , proving that u_2 is 2-upper semicontinuous.

Of course, as long as the payoff function u_1 is bounded and Borel measurable, the game admits a subgame perfect ϵ -equilibrium for each $\epsilon > 0$. $\diamond$

The next example belongs to the class of so-called quitting games, see, for instance, Flesch et al. (1997) and Solan and Vieille (2001).

Example 2.3.2. Consider the following game with two players. The set of actions is $\{c_1, q_1\}$ for player 1 and $\{c_2, q_2\}$ for player 2, with c_1 and c_2 standing for “continue” and q_1 and q_2 for “quit”. The players choose actions simultaneously. For a play $p \in P$, let $K_1(p)$ denote the first stage at which player 1 chooses action q_1 . If player 1 always chooses action c_1 , then $K_1(p) = \infty$. We define $K_2(p)$ similarly for player 2. The payoffs $(u_1(p), u_2(p))$ for the players are defined as follows. They are equal to $(0, 0)$ if $K_1(p) = K_2(p) = \infty$, i.e. if the players always choose “continue”, $(-1, 1)$ if $K_1(p) < K_2(p)$, $(1, -1)$ if $K_1(p) > K_2(p)$, and $(2, 2)$ if $K_1(p) = K_2(p) < \infty$. Note that u_1 and u_2 are symmetric and continuous everywhere, except at the play $p^c = ((c_1, c_2), (c_1, c_2), \dots)$.

We focus on player 1. The payoff function u_1 is not upper semicontinuous. Indeed, for each $m \in \mathbb{N}$, consider the play p_m in which (c_1, c_2) is chosen at the first m stages and (q_1, q_2) at stage $m+1$. The payoff for player 1 is 2 at each p_m . However, the sequence $(p_m)_{m \in \mathbb{N}}$ converges to the play p^c , which only gives payoff 0 to player 1.

Nevertheless, u_1 is 1-upper semicontinuous. Indeed, consider a sequence of plays $(p_m)_{m \in \mathbb{N}}$ that is 1-convergent to the play p^c . Then for large m we have either $p_m = p^c$ or $K_1(p_m) < K_2(p_m)$. Hence, for large m we obtain $u_1(p_m) \leq 0 = u_1(p^c)$. So, u_1 is 1-upper semicontinuous.

A subgame perfect equilibrium in the game is obtained if player 1 always chooses action q_1 and player 2 always chooses action q_2 . Another subgame perfect equilibrium is for example to always choose actions c_1 and c_2 . $\diamond$

We now illustrate the usefulness of i -upper semicontinuity by examining a class of games that includes the game in Example 2.3.2 as a special case.

Example 2.3.3. Consider the two-player game of Example 2.3.2, where the action sets for player 1 and 2 is given by $\{c_1, q_1\}$ and $\{c_2, q_2\}$ respectively. At every stage both players pick an action from their action space. The game ends when at least one player picks quit. Let $u_{i,\{j\}}^t$ denote the payoff of player i when player j unilaterally decides to quit at stage $t \in \mathbb{N}$ and let $u_{i,\emptyset}$ denote the payoff player i receives when quitting never takes place. For the sake of simplicity, assume that all payoffs are integers. Note that a payoff function of this game is continuous everywhere, except possibly at the play $p^c = ((c_1, c_2), (c_1, c_2), \dots)$.

The payoff function u_i of player i is i -upper semicontinuous in this game if the payoff player i receives when being the only one who quits at a late stage is at most the payoff

⁴The set P is an uncountable Polish space (essentially the Cantor space $\{1, 2\}^{\mathbb{N}}$), so P contains a non-Borel set, see Corollary 6.7.11 in Bogachev (2007).

Figure 2.1: The non-existence of a subgame perfect equilibrium and the stopping times $\tau_{1/5}^m$

player i would receive if the game goes on indefinitely, i.e. if there exists a stage $K \in \mathbb{N}$ such that for each $k \geq K$ we have $u_{i,\{i\}}^k \leq u_{i,\emptyset}$. This follows from the fact that the play p^c is the only possible point of discontinuity of u_i and the assumption that all payoffs are integers.

We now explain why the game admits a subgame perfect equilibrium provided that for each player $i \in I$ the payoff function u_i is i -upper semicontinuous. A generalization of the main idea behind the construction constitutes the basis for the proof of Theorem 2.4.1. Assume that for each player $i \in I$ the payoff function u_i is i -upper semicontinuous, and define a strategy profile $\sigma \in \mathcal{S}$ as follows. Since for each player $i \in I$ the payoff function u_i is i -upper semicontinuous, there exists a stage K such that, for each stage $k \geq K$, for each player $i \in I$, we have $u_{i,\{i\}}^k \leq u_{i,\emptyset}$. Consequently, after stage K , there is a play that is most preferred by both players: the play p^c . Thus, after stage K , we define σ_1 and σ_2 to continue indefinitely with probability 1. This implies that if stage no one quits before stage K , σ gives player i a payoff of $u_{i,\emptyset}$. We complete the strategy profile σ by a backward induction argument on the stages $K-1, K-2, \dots, 1$. By construction, σ constitutes a subgame perfect equilibrium of the game. $\diamond$

The following example shows that games with bounded and upper semicontinuous payoff functions may have no subgame perfect equilibrium. This motivates the use of the concept of subgame perfect ϵ -equilibrium.

Example 2.3.4. Consider the two-player centipede game depicted in Figure 2.1. Before we proceed, we spell out some of our notational conventions.

First, whenever an action is taken from a singleton action set, it will be omitted from our notation. Notice that, the game in Figure 2.1 having perfect information, at any given history only one of the players has a non-singleton action set. In accordance with our convention, we will only be recording the action taken by the active player. Thus for each $t \in \mathbb{N}^*$ we write c^t to denote the history of length t obtained if the players, whenever it is their turn to move, take the action c . In particular c^0 is the empty history. The symbol $c^t q$ denotes the history of length $t+1$ obtained if the action q is chosen for the first time in period $t+1$.

Second, we assume that, once either player takes the action q , both players' action sets become singletons for the rest of the game. This implies that the game only has countably many plays, which we denote by $c^0 q, c^1 q, c^2 q, \dots$ and c^∞ . Here $c^t q$ denotes the play where action q is chosen for the first time in period $t+1$. The symbol c^∞ is the play such that no one ever plays q . With these conventions in place, we turn to the analysis of the game.

The payoff function of player 1 is upper semicontinuous and the payoff function of player 2 is continuous. We show that the game does not have a subgame perfect equilibrium.

Suppose that $\sigma = (\sigma_1, \sigma_2)$ is a subgame perfect equilibrium. For each $t \in \mathbb{N}^*$, let $\mathbb{P}_{c^t, \sigma}(c^\infty)$ denote the probability that, in the subgame at history c^t , under σ the play c^∞ is realized. We distinguish two cases and derive a contradiction in each of them.

Assume first that for each $t \in \mathbb{N}^*$ we have $\mathbb{P}_{c^t, \sigma}(c^\infty) = 0$. As σ is a subgame perfect equilibrium, it follows that at each history where player 1 is active, σ_1 must place probability 1 on action q . This implies that, at each history where player 2 is active, σ_2 places probability 1 on action c . But then player 1 would have a profitable deviation from σ_1 by always choosing the action c , which is a contradiction.

Assume now that for some $t \in \mathbb{N}^*$ we have $\mathbb{P}_{c^t, \sigma}(c^\infty) > 0$. Then, $\mathbb{P}_{c^k, \sigma}(c^\infty)$ converges to 1 as $k \rightarrow \infty$. So there exists $K \in \mathbb{N}^*$ such that, for each $k \geq K$, player 1's expected payoff at history c^k is strictly more than 1, $\mathbb{E}_{c^k, \sigma}[u_1] > 1$. As σ is a subgame perfect equilibrium, σ_1 places probability 1 on action c at each history where player 1 is active beyond stage K . This implies that σ_2 places probability 1 on action q at each history where player 2 is active beyond stage K . But this contradicts $\mathbb{P}_{c^t, \sigma}(c^\infty) > 0$. Consequently, the game does not have a subgame perfect equilibrium.

Note however that for each $\epsilon > 0$ there exists a pure subgame perfect ϵ -equilibrium. Indeed, fix $\epsilon > 0$. Then there is $K \in \mathbb{N}$ such that if neither player quits before stage K then all feasible payoffs for player 2 are at most ϵ . Now the following pure strategy profile constitutes a subgame perfect ϵ -equilibrium: player 1 quits at each history, whereas player 2 continues at each history before stage K and quits at each history after stage K . Note that this strategy profile can be found as follows: let player 2 quit at each history after stage K and, given this, use backward induction at the various parts of the game tree. This is in line with the construction that we provide in the proof of Theorem 2.4.1. $\diamond$

2.4 Existence of a subgame perfect ϵ -equilibrium

The goal of this section is to present and prove the existence result, Theorem 2.4.1, for subgame perfect ϵ -equilibria in games with almost perfect information as defined in Section 2.2.

Section 2.4 is structured as follows. In the first subsection, we present and discuss this existence result. In the second subsection, we examine some implications of individual upper semicontinuity in more detail. In the third subsection, we provide the proof.

2.4.1 The existence result

The main result of this section is the following theorem.

Theorem 2.4.1. Consider a game G with almost perfect information as defined in Section 2.2. If for every player $i \in I$ the payoff function u_i is bounded, Borel measurable, and i -upper semicontinuous, then, for each $\epsilon > 0$, the game admits a subgame perfect ϵ -equilibrium, which is eventually pure.

Theorem 2.4.1 generalizes the result of Purves and Sudderth (2011) on the existence of subgame perfect ϵ -equilibria in perfect information games with upper semicontinuous payoff functions in two directions. First, as each upper semicontinuous function u_i is also Borel measurable and i -upper semicontinuous, we reduce the topological restrictions on the payoff functions of the players. Second, we allow the players to move simultaneously.

We provide a constructive proof for Theorem 2.4.1 in the next subsections. We remark that our proof is different from the one in Purves and Sudderth (2011), who after discretization of the payoffs, make use of an induction argument on the cardinality of the available payoffs in the subgames.

We discuss extensions of Theorem 2.4.1 and the necessity of the conditions in Section 2.6. Note Example 2.3.4 demonstrates that the conditions in Theorem 2.4.1 do not guarantee the existence of a subgame perfect equilibrium.

2.4.2 Optimal plays

Consider a player $i \in I$. For every play $p \in P$ and stage $t \in \mathbb{N}^*$, we define the set $O_i(p, t)$ as the set consisting of the play p and all other plays $q \neq p$ such that q coincides with p until at least stage t and the first deviation from p to q is caused by player i alone. That is,

$$O_i(p, t) = \{p\} \cup \{q \in P \mid \ell(p, q) \geq t \text{ and } I(p, q) = \{i\}\}. \quad (2.4.1)$$

Note that $O_i(p, t)$ is non-empty as it contains the play p .

Now we formulate a condition such that a player does not wish to deviate from a play after a certain stage.

Definition 2.4.2. Consider a player $i \in I$ and some $\epsilon > 0$. A play $p \in P$ is (ϵ, i) -optimal after stage t if for each play $q \in O_i(p, t)$

$$u_i(p) \geq u_i(q) - \epsilon.$$

A play $p \in P$ is ϵ -optimal after stage t if it is (ϵ, i) -optimal after stage t for each player $i \in I$. A play $p \in P$ is ϵ -optimal after history h if $h < p$ and p is ϵ -optimal after stage $\|h\|$.

The concept of ϵ -optimality is defined as a property of a play and not as a property of a strategy profile. Further, if the play p is ϵ -optimal after stage t then, for each $k \geq t$, the play p is also ϵ -optimal after stage k .

Lemma 2.4.3 and Corollary 2.4.4 relate individual upper semicontinuity and ϵ -optimality.

Lemma 2.4.3. Consider a player $i \in I$ and assume that player i 's payoff function u_i is bounded. Then u_i is i -upper semicontinuous if and only if for each play $p \in P$ and for each $\epsilon > 0$, there exists a stage $t \in \mathbb{N}^*$ such that the play p is (ϵ, i) -optimal after stage t .

Proof. Assume that the payoff function u_i is i -upper semicontinuous. Take a play $p \in P$ and some $\epsilon > 0$.

Suppose that for each $t \in \mathbb{N}^*$ there exists a play $q_t \in O_i(p, t)$ such that

$$u_i(p) < u_i(q_t) - \epsilon.$$

By construction, the sequence of plays $(q_t)_{t \in \mathbb{N}^*}$ is i -convergent to p and $u_i(p) \leq \limsup_{t \rightarrow \infty} u_i(q_t) - \epsilon$. This contradicts the assumption that u_i is i -upper semicontinuous.

For the other direction, fix a play $p \in P$ and some $\epsilon > 0$. Then, by assumption there exists a stage $t \in \mathbb{N}^*$ such that for every play $q_t \in O_i(p, t)$, $u_i(p) \geq u_i(q_t) - \epsilon$. Take any sequence of plays $(p_m)_{m \in \mathbb{N}}$ that i -converges to the play p . Then there exists $M \in \mathbb{N}$ such that for every $m \geq M$, $p_m \in O_i(p, t)$. Consequently, for every $m \geq M$, $u_i(p) \geq u_i(p_m) - \epsilon$. We conclude that $u_i(p) \geq \limsup_{m \rightarrow \infty} u_i(p_m) - \epsilon$. Since this holds for any play $p \in P$ and any $\epsilon > 0$, it follows that u_i is i -upper semicontinuous. $\square$

Since the set of players is finite, we have the following corollary.

Corollary 2.4.4. Assume that, for each player $i \in I$, the payoff function u_i is bounded and i -upper semicontinuous. Then, for each play $p \in P$ and for each $\epsilon > 0$, there exists a stage $t \in \mathbb{N}^*$ such that the play p is ϵ -optimal after stage t .

The next definition presents the notion of an ϵ -optimal history.

Definition 2.4.5. Let $\epsilon > 0$. A history $h \in H$ is ϵ -optimal if there exists a play $p \in P$ which is ϵ -optimal after history h .

In our construction of a subgame perfect ϵ -equilibrium, strategy profiles will be such that in a subgame corresponding to an ϵ -optimal history h , an ϵ -optimal play after history h is followed.

For $\epsilon > 0$, the number of times a history $h \in H$ has an initial segment which is ϵ -optimal is denoted by $n_\epsilon(h)$, so

$$n_\epsilon(h) = |\{t \in \mathbb{N}^* \mid t \leq \ell(h) \text{ and } h|_t \text{ is } \epsilon\text{-optimal}\}|.$$

Next, for $m \in \mathbb{N}$, the stopping time τ_ϵ^m is defined by

$$\tau_\epsilon^m(p) = \min\{t \in \mathbb{N}^* \mid n_\epsilon(p|_t) = m\}, \quad p \in P.$$

In words, $\tau_\epsilon^m(p)$ is the stage at which it occurs for the m -th time that a prefix of p is ϵ -optimal. Note that this does not mean that the play p needs to be ϵ -optimal from this prefix.

The following lemma claims that all these stopping times are uniformly bounded.

Lemma 2.4.6. For every $\epsilon > 0$, for every $m \in \mathbb{N}$, there exists $K_\epsilon^m \in \mathbb{N}^*$ such that for each play $p \in P$ we have $\tau_\epsilon^m(p) \leq K_\epsilon^m$.

Proof. Let some $\epsilon > 0$ and some $m \in \mathbb{N}$ be given.

Suppose that $\sup_{p \in P} \tau_\epsilon^m(p) = \infty$. Then, by finiteness of the set $A(\emptyset)$ of possible stage game outcomes at stage 1, there is $a^1 \in A(\emptyset)$ such that $\sup_{p > (a^1)} \tau_\epsilon^m(p) = \infty$. Then, for the same reason, there is $a^2 \in A(a^1)$ such that $\sup_{p > (a^1, a^2)} \tau_\epsilon^m(p) = \infty$. Continuing this way, we obtain a play $p = (a^1, a^2, \dots)$ for which $\tau_\epsilon^m(p) = \infty$. This is however in contradiction with Corollary 2.4.4. $\square$

Example 2.4.7. Consider again the two-player centipede game depicted in Figure 2.1 and discussed in Example 2.3.4. For the remainder of the example take $\epsilon = 1/5$. We will now compute the stopping times $\tau_{1/5}^m(p)$ for all $m \in \mathbb{N}$ and all plays $p \in P$. To do this we start by finding all $1/5$ -optimal histories.

Part 1: All histories $h \notin \{\emptyset, c^2\}$ are $1/5$ -optimal.

First note that all histories which contain a quitting action are clearly $1/5$ -optimal. All other histories have the form c^t for some $t \in \mathbb{N}^*$.

The history $c^0 = \emptyset$ is not $1/5$ -optimal. Indeed, the only $(1/5, 1)$ -optimal play after $\emptyset$ is c^∞ . However, c^∞ is not an $(1/5, 2)$ -optimal play after $\emptyset$ because the play cq gives player 2 a payoff strictly higher than $1/5$. Likewise the history c^2 is not $1/5$ -optimal: c^∞ is the only $(1/5, 1)$ -optimal play after history c^2 , but it is not $(1/5, 2)$ -optimal after c^2 , because the play c^3q gives player 2 a payoff strictly higher than $1/5$.

The history c is $1/5$ -optimal because the play cq is $1/5$ -optimal after history c . To see this, notice that player 1 cannot profitably deviate from cq after history c since history c is

Figure 2.2: Construction of a subgame perfect ϵ -equilibrium

controlled by player 2. And any deviation from the play cq by Player 2 after history c can increase player 2's payoff by no more than $1/2 - 1/3 < 1/5$. Using an analogous argument one can easily show that history c^3 is $1/5$ -optimal as the play c^3q is $1/5$ -optimal after this history.

All histories c^t with $t \geq 4$ are $1/5$ -optimal because the play c^∞ becomes $1/5$ -optimal after history c^4 . Indeed, once history c^4 has occurred neither player can gain an additional payoff of more than $1/5$ by quitting.

Part 2: The stopping times $\tau^m(p)$.

From part 1 it now easily follows that $\tau_{1/5}^m(p) = m$ if $m = 1$ or if $p \in \{q, cq\}$. While $\tau_{1/5}^m(p) = m + 1$ if $m \geq 2$ and $p \notin \{q, cq\}$. $\diamond$

2.4.3 The proof of Theorem 2.4.1

In this subsection, we prove Theorem 2.4.1. Assume that for every player $i \in I$ the payoff function u_i is bounded, Borel measurable, and i -upper semicontinuous. Fix $\epsilon > 0$. Our goal is to construct a subgame perfect ϵ -equilibrium σ . The construction is illustrated graphically in Figure 2.2. The trees T_h^k are defined in the proof.

STEP 1: We start from the root of the game, at history $\emptyset$. Along each play $p \in P$, the first ϵ -optimal history is given by $p|_{\tau_\epsilon^1(p)}$. Let H^1 denote the set of all these ϵ -optimal histories:

$$H^1 = \{p|_{\tau_\epsilon^1(p)} \in H \mid p \in P\}.$$

For each history $h \in H^1$, choose a play p^h that is ϵ -optimal after h .

Let $T_\emptyset^1$ denote the subtree that consists of all the histories that can occur at or before this stopping time τ_ϵ^1 , so

$$T_\emptyset^1 = \{h \in H \mid \exists h^1 \in H^1 : h \leq h^1\}.$$

In view of Lemma 2.4.6, the tree $T_\emptyset^1$ is finite. That is, the set of non-terminal nodes of $T_\emptyset^1$ is

the finite set

$$Z^1 = \{h \in H \mid \exists h^1 \in H^1 : h < h^1\}$$

and the terminal histories of $T_\emptyset^1$ are exactly the histories in H^1 . At each terminal history $h^1 \in H^1$, for each player $i \in I$, we define the terminal payoff of $T_\emptyset^1$ to be $u_i(p^{h^1})$. Intuitively, if h^1 is reached, then player i receives the payoff corresponding to the ϵ -optimal play p^{h^1} extending h^1 . Given these terminal payoffs, we can find a subgame perfect equilibrium $\sigma_\emptyset^1$ for $T_\emptyset^1$ by backward induction. Note that $\sigma_\emptyset^1$ is not necessarily pure, due to the potential presence of simultaneous moves.

For each $h^1 \in H^1$, let $W^1(h^1)$ denote the prefixes of the play p^{h^1} after the subtree $T_\emptyset^1$:

$$W^1(h^1) = \{h \in H \mid h^1 \leq h < p^{h^1}\}.$$

Define

$$W^1 = \cup_{h^1 \in H^1} W^1(h^1).$$

We now define the strategy profile σ for histories belonging to Z^1 and W^1 . The strategy profile σ equals $\sigma_\emptyset^1$ at histories in Z^1 and follows the ϵ -optimal plays at histories in W^1 . Thus, for every $h \in Z^1$ we define $\sigma(h) = \sigma_\emptyset^1(h)$ and, for every $h^1 \in H^1$, for every $h \in W^1(h^1)$, $\sigma(h)$ puts probability 1 on $p_{[\|h\|+1]}^{h^1}$.

STEP 2: Now we proceed by considering the minimal histories outside $Z^1 \cup W^1$. These are the histories that arise outside $T_\emptyset^1$ when along a play p^{h^1} with $h^1 \in H^1$ a deviation occurs from p^{h^1} . That is, we are considering the set of histories

$$R^2 = \bigcup_{h^1 \in H^1} \{ha \in H \mid h \in W^1(h^1), a \in A(h)\} \setminus W^1(h^1).$$

For each history $h^\circ = ha \in R^2$, we execute the following. Similarly to step 1, by using the boundedness of stopping times, we let $H^2(h^\circ)$ be the set of minimal ϵ -optimal histories h^2 with initial segment h° such that $n_\epsilon(h^2) = n_\epsilon(h) + 1$, so

$$H^2(h^\circ) = \{p_{\tau_\epsilon^{n_\epsilon(h)+1}(p)} \in H \mid p \in P(h^\circ)\}.$$

Let $T_{h^\circ}^2$ denote the finite subtree with root h° that consists of all histories that can occur at or before the stopping time $\tau_\epsilon^{n_\epsilon(h)+1}$, so

$$T_{h^\circ}^2 = \{h' \in H \mid \exists h^2 \in H^2(h^\circ) : h^\circ \leq h' \leq h^2\}.$$

Let $Z^2(h^\circ)$ denote the set of non-terminal histories belonging to $T_{h^\circ}^2$. The terminal histories of $T_{h^\circ}^2$ are precisely the histories in $H^2(h^\circ)$. For each history $h^2 \in H^2(h^\circ)$, choose a play p^{h^2} that is ϵ -optimal after h^2 . At each terminal history $h^2 \in H^2(h^\circ)$, for each player $i \in I$, we define the terminal payoff of $T_{h^\circ}^2$ to be $u_i(p^{h^2})$. Given these terminal payoffs, we can find a subgame perfect equilibrium $\sigma_{h^\circ}^2$ for $T_{h^\circ}^2$ by backward induction. Once again, $\sigma_{h^\circ}^2$ is not necessarily pure.

Just like in step 1, for each $h^2 \in H^2(h^\circ)$, let $W^2(h^\circ, h^2)$ denote the histories along the play p^{h^2} after the subtree $T_{h^\circ}^2$, and let $W^2(h^\circ)$ be the union of these sets. We define the strategy profile σ at histories belonging to $Z^2(h^\circ)$ and $W^2(h^\circ)$: the strategy profile σ equals $\sigma_{h^\circ}^2$ at histories in $Z^2(h^\circ)$ and follows the ϵ -optimal plays at histories in $W^2(h^\circ)$.

ALL FURTHER STEPS, AND CONCLUSION: By repeating this construction with steps in $\mathbb{N}$, we eventually consider each history of the game: certain histories belong to a finite subtree and certain histories are ϵ -optimal. This yields a fully specified strategy profile σ . By construction, σ is a subgame perfect ϵ -equilibrium. Indeed, a player cannot gain more than ϵ by deviating along an ϵ -optimal play, and given this, it is never profitable to deviate at histories belonging to a finite subtree.

Moreover, as the prescription for σ along the ϵ -optimal plays does not use randomization, σ is eventually pure in every subgame. This completes the proof of Theorem 2.4.1.

2.5 The topology induced by i -convergence

In this section, we examine the topology induced by the notion of i -convergence. We give criteria for metrizability, compactness, and separability for this topology.

2.5.1 The topological space $(P, \mathcal{T}_i)$

In this subsection, we fix a player $i \in I$ and define a topology $\mathcal{T}_i$ on the set of plays and show that a sequence of plays converges to a play p with respect to this topology $\mathcal{T}_i$ exactly when this sequence of plays i -converges to p . We also examine the relationship between the topology $\mathcal{T}_i$ and the topology $\mathcal{T}$.

For every play $p \in P$ and stage $t \in \mathbb{N}^*$, the set $O_i(p, t)$ is defined in (2.4.1). Let $\mathcal{T}_i$ be the topology on P that is induced by the collection of sets $\mathcal{O}_i = \{O_i(p, t) | p \in P, t \in \mathbb{N}^*\}$. That is, $\mathcal{T}_i$ is the smallest topology on P that contains each set belonging to $\mathcal{O}_i$. As the next lemma shows, the collection $\mathcal{O}_i$ forms a basis of the topology $\mathcal{T}_i$. This means that for a set $O \subseteq P$, we have $O \in \mathcal{T}_i$ exactly when O can be written as a union of sets belonging to $\mathcal{O}_i$.

Lemma 2.5.1. The collection $\mathcal{O}_i$ is a basis for the topology $\mathcal{T}_i$.

Proof. We only need to show (cf. Aliprantis and Border, page 25) that, for any two sets $O_i(p, s) \in \mathcal{O}_i$ and $O_i(q, t) \in \mathcal{O}_i$, the intersection $O_i(p, s) \cap O_i(q, t)$ can be written as a union of sets in $\mathcal{O}_i$. Let $O = O_i(p, s) \cap O_i(q, t)$. We can assume without loss of generality that $s \leq t$. We distinguish a number of cases and show in each case that O is indeed such a union.

Case 1: $p = q$. In this case, $O = O_i(q, t)$, so we are done.

Case 2: $p \neq q$. Then $\ell(p, q) < \infty$. We divide this case into three subcases.

Subcase 2.1: $\ell(p, q) < s \leq t$. This case is trivial, as $O = \emptyset$.

Subcase 2.2: $s \leq \ell(p, q) < t$. If $I(p, q) = \{i\}$ then $O = O_i(q, t)$. If $I(p, q) \neq \{i\}$ then $O = \emptyset$.

Subcase 2.3: $s \leq t \leq \ell(p, q)$. If $I(p, q) = \{i\}$ then $O = O_i(q, t)$. If $I(p, q) \neq \{i\}$, then O is the union of the sets $P(ha)$ where the history $h \in H$ and the stage outcome $a \in A$ have the properties: (a) h is a common prefix of p and q with length $t \leq \|h\| \leq \ell(p, q) - 1$, (b) $ha \in H$, (c) at the history h , the stage game outcome $a \in A$ deviates from the common prefix of p and q , but the only difference is caused by player i . Note that each such set $P(ha)$ is a union of sets in $\mathcal{O}_i$, hence we are done. $\square$

The next lemma gives some basic properties of the topological space $(P, \mathcal{T}_i)$.

Lemma 2.5.2. The topological space $(P, \mathcal{T}_i)$ is:

1. Hausdorff.
2. First countable.

3. Sequential.

Proof.

1. If $p \neq q$ then $\ell(p, q) < \infty$ and hence $O_i(p, \ell(p, q) + 1)$ and $O_i(q, \ell(p, q) + 1)$ are disjoint open sets containing p and q , respectively.
2. Take a play $p \in P$. Then the sequence $O_i(p, 1) \supseteq O_i(p, 2) \supseteq \dots$ is a countable neighborhood basis for p .
3. Follows immediately from (2).

□

Because the topological space $(P, \mathcal{T}_i)$ is sequential, it is fully determined by its convergent sequences. In the following lemma we show that the convergent sequences of $(P, \mathcal{T}_i)$ are precisely the i -convergent sequences, thereby proving that $\mathcal{T}_i$ is indeed the topology induced by i -convergence.

Lemma 2.5.3. A sequence of plays $(p_m)_{m \in \mathbb{N}}$ converges to the play p in the topological space $(P, \mathcal{T}_i)$ if and only if the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to the play p .

Proof. Take a sequence of plays $(p_m)_{m \in \mathbb{N}}$ that converges to the play p in the topological space $(P, \mathcal{T}_i)$. For each $t \in \mathbb{N}^*$, the set $O_i(p, t)$ is an open neighborhood of the play p , i.e. $p \in O_i(p, t) \in \mathcal{T}_i$. Hence, for each $t \in \mathbb{N}^*$, there exists a constant $M_t \in \mathbb{N}$ such that for all $m \geq M_t$ we have $p_m \in O_i(p, M_t)$. Therefore, $\lim_{m \rightarrow \infty} \ell(p_m, p) = \infty$ and for all $m \geq M_1$ we have $I(p_m, p) \subseteq \{i\}$. This means that the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to the play p .

Conversely, take a sequence of plays $(p_m)_{m \in \mathbb{N}}$ that i -converges to the play p . Consider an open neighborhood O of the play p , i.e. $p \in O \in \mathcal{T}_i$. By Lemma 2.5.1, there exists $t \in \mathbb{N}^*$ such that $O_i(p, t) \subseteq O$. Since we assumed that $(p_m)_{m \in \mathbb{N}}$ is i -convergent to p , there exists $M \in \mathbb{N}$ such that for each $m \geq M$ we have (a) $\ell(p_m, p) \geq t$ and (b) $I(p_m, p) \subseteq \{i\}$. Thus, for each $m \geq M$ we have $p_m \in O_i(p, t)$ and in particular $p_m \in O$. Therefore, $(p_m)_{m \in \mathbb{N}}$ converges to p in the topological space $(P, \mathcal{T}_i)$. □

Lemma 2.5.4. The topology $\mathcal{T}_i$ is finer than the topology $\mathcal{T}$, that is, $\mathcal{T}_i \supseteq \mathcal{T}$.

Proof. For every history $h \in H$ and for each play $p \in P(h)$ it holds that $p \in O_i(p, \|h\|) \subseteq P(h)$, and hence we have that $P(h) = \cup_{p \in P(h)} O_i(p, \|h\|)$. Therefore, we have that $\mathcal{T}_i$ contains each set $P(h)$. As the topology $\mathcal{T}$ is induced by these sets, we obtain $\mathcal{T}_i \supseteq \mathcal{T}$. □

Note that $\mathcal{T}_i$ is strictly larger than $\mathcal{T}$ in certain games. Consider the game in Example 2.3.4. The set $O_1(c^\infty, 1)$ belongs to $\mathcal{T}_1$ and it contains the play c^∞ and all plays in which player 1 quits at stage 3 or later. However, $O_1(c^\infty, 1)$ does not belong to $\mathcal{T}$ for the following reason. There is no history $h \in H$ such that $c^\infty \in P(h) \subseteq O_1(c^\infty, 1)$. As the sets $P(h)$, where $h \in H$, form a basis of the topology $\mathcal{T}$, we conclude that $O_1(c^\infty, 1) \notin \mathcal{T}$.

Now we turn to the question when $\mathcal{T}_i$ coincides with $\mathcal{T}$. The answer, and also the answer to many other questions, depends on the size of the action sets along each play. For this purpose, we introduce the notions of i -finiteness and i -cofiniteness of a play.

Definition 2.5.5. A play $p \in P$ is i -finite if there exists a stage $K \in \mathbb{N}^*$ such that at every prefix of p with length at least K the set of available actions to player i is a singleton, so for every $k \geq K$, $|A_i(p|_k)| = 1$.

Figure 2.3: The topology $\mathcal{T}_1$. Figure 2.4: The topology $\mathcal{T}_2$. Figure 2.5: The topology $\mathcal{T}_1 \cap \mathcal{T}_2$. Figure 2.6: The topology $\mathcal{T}$.

Figure 2.7: Local basis of the play $p = (00)^\infty$ in the topology $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_1 \cap \mathcal{T}_2$ and $\mathcal{T}$, respectively, for a 2-player multistage game where, for every $h \in H$, for $i = 1, 2$, $A_i(h) = \{0, 1\}$. The Euclidean distance between any play and the play p is inversely related to the length of the longest common history, while the circle sector indicates the subset of players who first deviated

Thus, if a play p is i -finite then, after finitely many stages, it is not within the power of player i to deviate from p . Let $F_i = \{p \in P \mid p \text{ is } i\text{-finite}\}$ denote the subset of i -finite plays of the P .

Definition 2.5.6. A play $p \in P$ is i -cofinite if there exists a stage $K \in \mathbb{N}^*$ such that at every prefix of p with length at least K the set of available actions to every player $j \neq i$ is a singleton, so for every $k \geq K$, for every $j \neq i$, $|A_j(p|_k)| = 1$.

Thus, if a play p is i -cofinite then, after finitely many stages, player i has full control over the realization of the play p . Let $C_i = \{p \in P \mid p \text{ is } i\text{-cofinite}\}$ denote the subset of i -cofinite plays of the P .

Note that there can be plays that are neither i -finite nor i -cofinite and that a play can be both i -finite and i -cofinite.

As it turns out, $\mathcal{T}_i = \mathcal{T}$ holds if and only if the topological space $(P, \mathcal{T}_i)$ is compact if and only if each play is i -cofinite.

Proposition 2.5.7. The following statements are equivalent:

- (1) $\mathcal{T}_i = \mathcal{T}$.
- (2) The topological space $(P, \mathcal{T}_i)$ is compact.
- (3) Each play $p \in P$ is i -cofinite, i.e. $P = C_i$.

Proof. We prove that (1) and (2) are equivalent and that (1) and (3) are equivalent as well.

(1) implies (2): This is immediate because $(P, \mathcal{T})$ is compact.

(2) implies (1): As we know, $(P, \mathcal{T})$ is a compact Hausdorff space. By assumption, $(P, \mathcal{T}_i)$ is compact as well and by Lemma 2.5.2 it is Hausdorff too. Because $\mathcal{T}_i \supseteq \mathcal{T}$ due to Proposition 2.5.4, the maximality principle of compact Hausdorff spaces implies that $\mathcal{T}_i = \mathcal{T}$.⁵

(1) implies (3): Suppose there exists a play $p \in P$ which is not i -cofinite. Then there is an infinite set $T \subseteq \mathbb{N}^*$ of stages such that at each prefix of p with length $t \in T$ there is a player $j_t \neq i$ with an action set containing at least two alternatives, so $|A_{j_t}(p|_t)| \geq 2$. For each $t \in T$, let p_t be a play in which the first difference from p is caused by the action of player j_t at stage t . Then, the sequence of plays $(p_t)_{t \in T}$ is convergent to the play p , but not i -convergent. Because both $\mathcal{T}_i$ and $\mathcal{T}$ are fully determined by the convergent sequences, this contradicts the assumption that $\mathcal{T}_i = \mathcal{T}$. Consequently, it follows that every play $p \in P$ is i -cofinite.

(3) implies (1): Because both $\mathcal{T}_i$ and $\mathcal{T}$ are determined by their convergent sequences it is sufficient to show that if all plays are i -cofinite then every convergent sequence is also i -convergent. To this end, fix a play $p \in P$ and a sequence $(p_m)_{m \in \mathbb{N}}$ that converges to the play p , i.e. $\lim_{m \rightarrow \infty} p_m = p$. Because the play p is i -cofinite, there exists a time t^* such that for every player $j \neq i$ and every $t \geq t^*$ $|A_j(p|_t)| = 1$. Because the sequence $(p_m)_{m \in \mathbb{N}}$ is convergent to the play p we have that there exists an $M \in \mathbb{N}^*$ such that, for all $m \geq M$, $\ell(p_m, p) \geq t^*$. Furthermore we have that $I(p_m, p) \subseteq \{i\}$ for all $m \geq M$. We can conclude that $\lim_{m \rightarrow \infty}^{(i)} p_m = p$. $\square$

A topological space is called separable if it contains a countable dense subset. The separability of $(P, \mathcal{T}_i)$ depends on the cardinality of the set of plays which are i -finite. The reason is that each i -finite play, as a singleton, is open in $\mathcal{T}_i$.

Proposition 2.5.8. The topological space $(P, \mathcal{T}_i)$ is separable if and only if the subset of i -finite plays F_i is countable.

Proof. First notice that, for every $p \in F_i$, the singleton $\{p\}$ is open in $\mathcal{T}_i$. Indeed, as p is i -finite, there is a stage t such that at any prefix of p with length at least t player i 's action set is a singleton. It follows that $\{p\} = O_i(p, t) \in \mathcal{T}_i$.

Part 1: Assume that $(P, \mathcal{T}_i)$ is separable and let D be a countable dense subset of P . For each $p \in F_i$, the singleton $\{p\}$ is open in $\mathcal{T}_i$, so we have $F_i \subseteq D$. It follows that F_i is countable.

Part 2: Assume that F_i is countable. We show that $(P, \mathcal{T}_i)$ is separable by constructing a countable dense subset of P .

For every history $h \in H$, fix a play $p^h \in P(h)$. Because the set H of histories is countable and because F_i is countable by assumption, the set $D = \{p^h | h \in H\} \cup F_i$ is countable too. We claim that D is dense in P with respect to $\mathcal{T}_i$.

⁵If $(X, \mathcal{V})$ is a compact Hausdorff space and $\mathcal{W}$ is another topology on X such that $\mathcal{W}$ strictly includes $\mathcal{V}$, then the topological space $(X, \mathcal{W})$ is not compact. Indeed, take a set $U \in \mathcal{W} \setminus \mathcal{V}$. As U is not open in $\mathcal{V}$, the set $X \setminus U$ is not closed in $\mathcal{V}$. Because $(X, \mathcal{V})$ is Hausdorff, $X \setminus U$ is not compact in $\mathcal{V}$ and hence $X \setminus U$ is not compact in $\mathcal{W}$. However, by the choice of U , the set $X \setminus U$ is closed in $\mathcal{W}$. Hence, $(X, \mathcal{W})$ cannot be compact.

It suffices to show that, for every $p \notin D$, for every stage $t \in \mathbb{N}^*$, $O_i(p, t) \cap D \neq \emptyset$. Take a play $p \notin D$ and a stage $t \in \mathbb{N}^*$. As $p \notin D$ we also have $p \notin F_i$, and hence there exists a prefix h of p such that $\|h\| \geq t$ and $|A_i(h)| \geq 2$. Consider a history h' of length $\|h\| + 1$ such that (a) h' coincides with h at the first $\|h\|$ stages and (b) at stage $\|h\| + 1$, h' differs from the play p only by the action of player i . Then, $p^{h'} \in O_i(p, t) \cap D$. $\square$

The following proposition shows that $(P, \mathcal{T}_i)$ is not metrizable under mild conditions.

Proposition 2.5.9. If the set $F_i \cup C_i$ is finite and $P \setminus (F_i \cup C_i)$ is infinite, then the topological space $(P, \mathcal{T}_i)$ is not metrizable.

Proof. Suppose $d_i : P \times P \rightarrow [0, \infty)$ is a metric which induces the topology $\mathcal{T}_i$.

Step 1: Construction of two sequences of plays $(p_m)_{m \in \mathbb{N}}$ and $(q_m)_{m \in \mathbb{N}}$.

In this step we inductively construct two sequences of plays $(p_m)_{m \in \mathbb{N}}$ and $(q_m)_{m \in \mathbb{N}}$ as illustrated in Figure 2.8.

Start with any play p_1 which is not i -finite. Because the play p_1 is not i -finite and there are infinitely many plays which are not i -cofinite there exists a play q_1 with the following three properties (1) $I(p_1, q_1) = \{i\}$, (2) q_1 is not i -cofinite, and (3) $d_i(p_1, q_1) < 1/2$. Indeed, the fact that the play p_1 is not i -finite guarantees that there exists an eventually non-constant sequence of plays which i -converges to the play p_1 . This guarantees that this sequence of plays contains an element q_1 having all the desired properties.

Now assume that for some $m \in \mathbb{N}$ the plays p_m and q_m are defined such that p_m is not i -finite and q_m is not i -cofinite. Take any play p_{m+1} that is not i -finite such that p_{m+1} coincides with the play q_m on a longer history than the play p_m and such that the first time the play p_{m+1} differs from q_m is not solely due to player i , so

$$\ell(p_m, q_m) < \ell(p_{m+1}, q_m) \text{ and } I(p_{m+1}, q_m) \setminus \{i\} \neq \emptyset. \quad (2.5.1)$$

Because the play p_{m+1} is not i -finite it is possible to choose the play q_{m+1} in such a way that q_{m+1} is not i -cofinite and

$$\ell(p_{m+1}, q_m) < \ell(p_{m+1}, q_{m+1}), \quad I(p_{m+1}, q_{m+1}) = \{i\}, \text{ and } d_i(p_{m+1}, q_{m+1}) < 2^{-(m+1)}. \quad (2.5.2)$$

Step 2: Finding a contradiction

For $m \in \mathbb{N}$, let $h_m = \min(p_m, q_m)$ denote the longest common history between the play p_m and the play q_m . Notice that by construction we have $h_1 < h_2 < \dots$. Therefore there exists a unique play p extending all histories h_m . Note that for every $m \in \mathbb{N}$ we have that

$$\ell(p, p_m) = \ell(p_m, q_m) < \ell(p_{m+1}, q_m) < \ell(p_{m+1}, q_{m+1}) = \ell(p, p_{m+1}).$$

It follows that $\lim_{m \rightarrow \infty} \ell(p, p_m) = \infty$. Furthermore, we have that $I(p, p_m) = I(p_m, q_m) = \{i\}$.

Therefore, we can conclude that $\lim_{m \rightarrow \infty}^{(i)} p_m = p$ and hence $\lim_{m \rightarrow \infty} d_i(p, p_m) = 0$. Furthermore, we have by construction that $\lim_{m \rightarrow \infty} d_i(p_m, q_m) = 0$. It follows from the triangle inequality that, for every $m \in \mathbb{N}$, $d_i(p, q_m) \leq d_i(p, p_m) + d_i(p_m, q_m)$. This yields that $\lim_{m \rightarrow \infty} d_i(p, q_m) = 0$ and consequently $\lim_{m \rightarrow \infty}^{(i)} q_m = p$. However, by construction we have that, for every $m \in \mathbb{N}$, $I(p, q_m) = I(p_m, q_m) \not\subseteq \{i\}$, so we have obtained a contradiction. $\square$

Figure 2.8: The non-metrizability of the topology of i -convergence

Example 2.5.10. Proposition 2.5.9 implies that the topological space $(P, \mathcal{T}_i)$ induced by an infinitely repeated stage game in which all players have at least two actions is not metrizable. Indeed, every play in such a game is neither i -finite nor i -cofinite and there are infinitely many plays.

Even though in general the topological space $(P, \mathcal{T}_i)$ is not metrizable, there are games for which the associated topological space is metrizable.

Lemma 2.5.11. If $P = F_i \cup C_i$, then the topological space $(P, \mathcal{T}_i)$ is metrizable.

Proof. Let $F_{i,k}$ denote the set of i -finite plays such that $A_i(p_t)$ is a singleton for each $t \geq k$. Thus $F_{i,1}, F_{i,2}, \dots$ is a non-decreasing sequence of sets converging to the set of i -finite plays F_i . Let $\delta(p) = 0$ for each $p \notin F_i$ and let $\delta(p) = 2^{-k}$ for each $p \in F_{i,k} \setminus F_{i,k-1}$.

Define $d_i(p, p) = 0$ for each $p \in P$, and for two distinct plays $p, q \in P$ we let $d_i(p, q) = \max\{d(p, q), \delta(p), \delta(q)\}$, where d is the usual ultrametric defined by $d(p, q) = 2^{-\ell(p, q)}$. We now show that because d is an ultrametric, d_i is as well. By definition we have that $d_i(p, p) = 0$ for all $p \in P$, furthermore it is trivial to see that d_i is non-negative and symmetric. It remains to show that $d_i(p, q) \leq \max\{d_i(p, r), d_i(r, q)\}$ for any $p, q, r \in P$. We have

$$\begin{aligned} d_i(p, q) &\leq \max\{d(p, r), d(r, q), \delta(p), \delta(q)\} \\ &\leq \max\{d(p, r), d(r, q), \delta(p), \delta(q), \delta(r)\} \\ &= \max\{d_i(p, r), d_i(r, q)\}. \end{aligned}$$

We now show that the metric d_i induces the topology $\mathcal{T}_i$ of i -convergence.

Part 1: If $\lim_{m \rightarrow \infty} d_i(p_m, p) = 0$ then the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to p .

Suppose the sequence $(p_m)_{m \in \mathbb{N}}$ converges to p under the metric d_i . First assume that $p \in F_i$ then because p is i -finite there exists $M \in \mathbb{N}$ such that $p_m = p$ for all $m \geq M$, which implies that $I(p_m, p) = \emptyset$ and $\ell(p_m, p) = \infty$ for all $m \geq M$. We can conclude that $(p_m)_{m \in \mathbb{N}}$ i -converges to p . Now assume that $p \in C_i$ and note that because $d \leq d_i$ we have that the sequence

$(p_m)_{m \in \mathbb{N}}$ converges to p under the metric d . Furthermore observe that any sequence of plays that converges to $p \in C_i$ under the metric d also i -converges to p .

Part 2: If the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to p then $\lim_{m \rightarrow \infty} d_i(p_m, p) = 0$

If the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to p , then $\lim_{k \rightarrow \infty} d(p_k, p) = 0$ and there exists an $M \in \mathbb{N}$ such that $I(p_m, p) \subseteq \{i\}$ for every $m \geq M$. Let $n_m = \ell(p_m, p)$. Then for each $m \geq M$ we have that either $p_m = p$ or the set $A_i(p_m|n_m) = A_i(p|n_m)$ is not a singleton. From this we can conclude that p and p_m are not elements of F_{i, n_m} . Therefore $\delta(p) \leq 2^{-n_m}$ and $\delta(p_m) \leq 2^{-n_m}$. Because $\lim_{m \rightarrow \infty} n_m = \infty$ and $\lim_{m \rightarrow \infty} d(p_m, p) = 0$ we conclude that p_m converges to p under the metric d_i . $\square$

2.5.2 The topological space $(P, \cap_{i \in I} \mathcal{T}_i)$

In this section, we take a closer look at the collection $\cap_{i \in I} \mathcal{T}_i$ of sets which are open in every topological space $(P, \mathcal{T}_i)$. By Proposition 2.5.4 we have $\mathcal{T}_i \supseteq \mathcal{T}$ and hence $\cap_{i \in I} \mathcal{T}_i \supseteq \mathcal{T}$. It is now a natural question to ask whether in fact $\cap_{i \in I} \mathcal{T}_i = \mathcal{T}$. As we will see, this is the case for perfect information games, but not in general.

For every play $p \in P$ and stage $t \in \mathbb{N}^*$, let

$$\begin{aligned} O^*(p, t) &= \cup_{i \in I} O_i(p, t) \\ &= \{q \in P \mid \ell(p, q) \geq t \text{ and } |I(p, q)| \leq 1\}. \end{aligned}$$

Let $\mathcal{T}^*$ be the topology on P that is induced by the collection $\mathcal{B}^* = \{O^*(p, t) \mid p \in P, t \in \mathbb{N}^*\}$. That is, $\mathcal{T}^*$ is the smallest topology on P that contains each set belonging to $\mathcal{B}^*$. As the next lemma shows, the collection $\mathcal{B}^*$ forms a basis of the topology $\mathcal{T}^*$. This means that for a set $O \subseteq P$, we have $O \in \mathcal{T}^*$ exactly when O can be written as a union of sets in $\mathcal{B}^*$.

Lemma 2.5.12. The collection $\mathcal{B}^*$ is a basis for the topology $\mathcal{T}^*$.

Proof. The proof is similar to that of Lemma 2.5.1. We only need to show (cf. Aliprantis and Border, page 25) that, for any two sets $O^*(p, s) \in \mathcal{B}^*$ and $O^*(q, t) \in \mathcal{B}^*$, the intersection $O^*(p, s) \cap O^*(q, t)$ can be written as a union of sets in $\mathcal{B}^*$. Let $O^* = O^*(p, s) \cap O^*(q, t)$. We can assume without loss of generality that $s \leq t$. We distinguish a number of cases and prove in each case that O^* is indeed such a union.

Case 1: $p = q$. In this case $O^* = O^*(q, t)$ and we are done.

Case 2: $p \neq q$. It follows that $\ell(p, q) < \infty$. We divide this case into three subcases.

Subcase 2.1: $\ell(p, q) < s \leq t$. In this case it holds that $O^* = \emptyset$.

Subcase 2.2: $s \leq \ell(p, q) < t$. If $I(p, q)$ is a singleton, then $O^* = O^*(q, t)$. If $I(p, q)$ is not a singleton, so multiple players cause the first difference between p and q , then $O^* = \emptyset$.

Subcase 2.3: $s \leq t \leq \ell(p, q)$. If $I(p, q)$ is a singleton, then $O^* = O^*(p, s) = O^*(q, t)$. If $I(p, q)$ is not a singleton, then O^* is the union of the sets $P(ha)$ where the history $h \in H$ and the stage outcome $a \in A$ have the properties: (a) h is a common prefix of p and q with length $t \leq \|h\| \leq \ell(p, q) - 1$, (b) $ha \in H$, (c) at the history h , the stage game outcome $a \in A$ deviates from the common prefix of p and q and the difference is caused by exactly one player. Because for every history $h \in H$ we have $P(h) = \cup_{p \in P(h)} O_i(p, \|h\|)$, each such set $P(ha)$ is a union of sets in $\mathcal{B}^*$, and hence we are done. $\square$

We now show that the topologies $\mathcal{T}^*$ and $\cap_{i \in I} \mathcal{T}_i$ coincide. Hence, $\mathcal{B}^*$ is a basis for the topology $\cap_{i \in I} \mathcal{T}_i$.

Lemma 2.5.13. $\mathcal{T}^* = \cap_{i \in I} \mathcal{T}_i$.

Proof.

Part 1: $\mathcal{T}^* \supseteq \cap_{i \in I} \mathcal{T}_i$.

Let $O \in \cap_{i \in I} \mathcal{T}_i$ be given. It is sufficient to show that for every $p \in O$ there exists $t \in \mathbb{N}^*$ such that $O^*(p, t) \subseteq O$. Let some $p \in O$ be given. For every $i \in I$ there exists $t_i \in \mathbb{N}^*$ such that $O_i(p, t_i) \subseteq O$. For $t = \max_{i \in I} t_i$ it holds that, for every $i \in I$, $O_i(p, t) \subseteq O$. It follows that $O^*(p, t) = \cup_{i \in I} O_i(p, t) \subseteq O$.

Part 2: $\mathcal{T}^* \subseteq \cap_{i \in I} \mathcal{T}_i$.

Fix a play $p \in P$ and a time $t \in \mathbb{N}^*$. It is sufficient to prove that, for every $i \in I$, $O^*(p, t) \in \mathcal{T}_i$. Fix some player $i \in I$. Let some $q \in O^*(p, t)$ be given. It is sufficient to show that there exists $k \in \mathbb{N}^*$ such that $O_i(q, k) \subseteq O^*(p, t)$. If $q = p$, then let $k = t$, so we have that $O_i(p, t) \subseteq \cup_{j \in I} O_j(p, t) = O^*(p, t)$ and we are done. If $q \neq p$, then there exists $k \in \mathbb{N}$ such that $k > \ell(p, q) \geq t$. It is immediate that $O_i(q, k) \subseteq O^*(p, t)$. $\square$

Because in perfect information games at every stage only one player can deviate from a given play, we obtain the following result.

Proposition 2.5.14. Consider a game G as defined in Section 2.2 with perfect information. It holds that $\mathcal{T}^* = \mathcal{T}$.

Proof. Because G has perfect information, we have for every play $p \in P$ and stage $t \in \mathbb{N}^*$ that $O^*(p, t) = \{q \in P \mid \ell(p, q) \geq t\} = P(p_t)$. Hence $\mathcal{T}^*$ and $\mathcal{T}$ have the same basis and the statement follows. $\square$

Corollary 2.5.15. If a sequence of plays is convergent in $\mathcal{T}^*$, then it is convergent in $\mathcal{T}$ as well. The converse holds for perfect information games.

2.6 Discussion

2.6.1 Perfect information games

In Theorem 2.4.1, we assumed that the set of available actions is always finite. This has two consequences. First, when applying backward induction in the proof of Theorem 2.4.1, we were guaranteed to have a subgame perfect equilibrium in each finite tree. Second, each stopping time τ_ϵ^k is bounded.

The second consequence is not essential. Even with unbounded, but finite, stopping times we would obtain subtrees in the proof of Theorem 5 that have no infinite branches. Hence, it would still be possible to apply backward induction by the following idea. Start at the root of a subtree. If there is a child history that is non-terminal, then take this history, and repeat this process as long as it is possible to choose a non-terminal child history. Since there is no infinite branch, this process will eventually stop at a history such that all of its children are terminal histories. Then, apply a step of backward induction at this history. By means of a transfinite procedure, applying a step of backward induction at each iteration step, we finally end up with a complete strategy profile for this subtree, and this strategy profile is thus a subgame perfect equilibrium for this subtree.

The first consequence is important though. Still, it would be enough to have a one-shot ϵ -equilibrium in each possible stage game, for every $\epsilon > 0$. In particular, this would be the

case if the game under consideration has perfect information, even if the set of available actions is not finite. To be more precise, this leads to the following statement.

Consider a perfect information game that satisfies the assumptions of Section 2.2, without the requirement that the set of available actions of each player is finite. If the payoff function u_i of every player $i \in I$ is bounded and i -upper semicontinuous, then for each $\epsilon > 0$, the game admits a pure strategy profile such that in any subgame, no player can gain more than ϵ by unilaterally deviating to another pure strategy.

This statement generalizes the corresponding result in Purves and Sudderth (2011), by relaxing the topological condition on the payoff functions.

2.6.2 Necessity of the conditions in Theorem 2.4.1

In this subsection, we discuss to which extent the conditions assumed in Theorem 2.4.1 are necessary for the existence of a subgame perfect ϵ -equilibrium, for every $\epsilon > 0$.

The assumption that the set of available actions is always finite was already discussed in the previous subsection.

We also assumed in Theorem 2.4.1 that the payoffs are bounded. This is a standard assumption, and without this assumption even very simple 1-player games fail to have ϵ -optimal solutions.

The assumption of individual upper semicontinuity is of course the main condition on the payoffs and is heavily used.

We also assumed that the payoffs are Borel measurable. The subgame perfect ϵ -equilibria that we constructed in Theorem 2.4.1 are eventually pure in every subgame. Hence, the calculation of the expected payoffs for these strategy profiles, or even if a player deviates to a pure strategy, does not require the payoffs to be Borel measurable. Rather, Borel measurability is needed to be able to calculate the expected payoffs when a player deviates to a non-pure strategy.

The assumption that the number of players is finite was used to obtain the crucial Corollary 2.4.4.

2.6.3 Payoff functions that are i -lower semicontinuous

It is not known whether all games as defined in Section 2, where each player's payoff function is bounded and lower semicontinuous admits an ϵ -equilibrium, for every $\epsilon > 0$.⁶

Nevertheless, using i -convergent sequences we can define the notion of i -lower semicontinuity. A payoff function u_i is i -lower semicontinuous if for every play p and every sequence $(p_m)_{m \in \mathbb{N}}$ of plays that is i -convergent to p we have $\liminf_{m \rightarrow \infty} u_i(p_m) \geq u_i(p)$. With this concept, for perfect information games we can conclude the following statement. Flesch et al. (2010) showed that in every perfect information game, with arbitrary action spaces, if each player's payoff function is bounded and lower semicontinuous then there exists a pure strategy profile such that in any subgame, no player can gain more than ϵ by unilaterally deviating to another pure strategy. After a careful look at the proof, the only place where the lower semicontinuity of the payoff functions is used is on page 750 in order to prove that Equation (18) follows from Equation (19). Since the sequence of plays considered there

⁶It is not even known in the context of quitting games, see for instance Solan and Solan (2017) and the references therein.

is not only convergent but also i -convergent, the result of Flesch et al. (2010) remains valid even if lower semicontinuity is relaxed to i -lower semicontinuity.

2.6.4 Stochastic transitions

It is a natural question whether or not Theorem 2.4.1 can be extended to games with stochastic transitions. Even if this is possible, the proof would be substantially more complex, as we would not be able to work with ϵ -optimal plays any more.

Chapter 3

Decreasing Competition in Competitive Search: Unequal Outcomes with Equal Chances

3.1 Introduction

Imagine a group of new graduates in a high demand trade all wanting to secure a high paying job. For simplicity assume they have identical skills and identical preferences. Every period a job appears to which they can apply. If a job has multiple applicants, all applicants are given equal chances of securing the job. Once a graduate accepts a job he cannot apply to any future jobs. Graduates that do not yet have a job, continue searching, however waiting for a potentially better job opportunity is costly. This paper aims to show that even though all players have identical preferences and have equal chances of obtaining a job they applied for, the ex ante expected equilibrium payoffs of the different players may be different. This expected outcome inequality is solely driven by the fact that the pool of players actively looking for a job is expected to shrink over time.

From a fairness point of view, it may be desirable to give every player equal chance. However this paper shows the possibly surprising fact that this is insufficient to guarantee that all players have an ex ante equal expected outcomes. Alternatively, this paper illustrates that one cannot conclude that differences in ex ante expected outcomes are necessarily the result either of discrimination or differences in preferences.

This paper also studies the welfare aspects of this equal chances policy for a decreasing group of competitors. We find that the equal chances policy always leads to a utilitarian welfare loss. This contrasts to the first best policy where players are initially given a priority ranking. The found efficiency loss stems from the fact that when players are given equal chances, they are too selective and this leads them to wait inefficiently long. While if players are given a priority ranking, the lowest ranked player knows he has no chance for winning high values and therefore will be less selective in applying for low value items.

We study the following infinite horizon dynamic game. There is a fixed set of homogeneous players who want to get an item. Every player can obtain at most one item and every item can be allocated to at most one player. The items are revealed over time and the players have the same valuations for each of the items. Every period $t \in \mathbb{N}$ an item appears,

We thank Nicolas Vieille for his comments and remarks during the development of this chapter

with valuation x_t drawn independently from a known distribution F . In the same period all players who do not yet have an item can apply it. If there are multiple applicants, all applicants will obtain the item with equal chance. If an item with value x_t is won in period t , the payoff to the winning applicant is $\delta^t x_t$, where δ is the discount factor which models the cost of waiting. All players are identical and hence are equally impatient. There is no cost of applying for an item, but an allocated item cannot be refused. Hence the only cost associated with winning an item is the opportunity cost of not being able to apply for possibly better items in future periods. If no player applies for an item then the item disappears and the next period starts.

The game described above is a multiplayer optimal stopping problem with the property that the set of players who do not yet have an item will decrease over time. The literature on multiplayer stopping problems is vast as problems can vary on a lot of different levels. For example one can differentiate between zero-sum and non zero-sum games, games with an infinite or finite horizon, games where the distribution of the values is known or unknown, or games where players have only partial or no information about the valuations of their competitors (see for example Whitmeyer (2018)). A survey by Abdelaziz and Krichen (2007) gives an organized overview of multiplayer optimal stopping problems.

Competitive secretary problems are a well-known subclass of multiplayer stopping problems. In a competitive secretary problem, several competing players sequentially observe items, called secretaries, with an associated value. Every period a secretary is revealed, players can apply for this secretary or wait and hope for a better one. If multiple players apply for the same secretary, the secretary is allocated according to a tie breaking rule. Sakaguchi (1980), Fushimi (1981) Chen, Rosenberg and Shepp (1997), Immorlica, Kleinberg and Mahdian (2006), Cownden and Steinsaltz (2014) Karlin and Lei (2015), all study variations of the competitive secretary problem where the information stream of the players, the objective functions of the players or the tie breaking rules differ. This paper is close to the contribution by Immorlica et al. (2006) where the authors assign equal probability to all applicants. However, in contrast to our paper Immorlica et al. consider an finite horizon game where the distribution of the valuations of the items is unknown.

Apart from the relation to multiplayer stopping problems, the game in this paper can be seen as a multiplayer version of the undirected job search problem. Interpreting the item as a job whose value is given by the total discounted earnings of the job, the current game models a job search problem where competition is expected to decrease over time. In fact, the one player version of the game described in this paper reduces to the undirected job search problem studied by McCall (1970) and Lippman and McCall (1979). This broad literature, among other things, tries to explain when jobseekers should accept a job or when they should wait for better opportunities. See Lippman and McCall (1979) and Rogerson, Shimer and Wright (2005) for an overview of job search models. More recently, Mazalov and Konovalchikova (2015) studied a two player competitive job search problem with incomplete information, where each player wants to select a better item than his opponent.

We show that when restricting the strategy space to pure threshold strategies, players may use different thresholds in equilibrium, making some of them more selective than others, despite having the same discount factor. Because all players have the same preferences, less competitive periods are intuitively more desirable to every single player. This incen-

tivizes individual players to wait for the periods with less competition. However, if all players decide to wait, the competition remains the same. Therefore, players differentiate their strategies: some players act less selective, increasing their chances of getting an item early on, while other players are more selective, increasing their expected waiting time but also the expected value of their allocated item.

Apart from studying the asymmetry of equilibrium payoffs, this paper also investigates the welfare consequences of imposing an equal chances policy. We contrast this policy to the socially optimal policy of giving the players a priority ranking, where in case an item has multiple applicants, the item is allocated to the player with the highest priority. The welfare consequences of various types of priority rankings are studied in the more recent literature on dynamic one sided matching where a stream of items, such as social housing or kidneys, has to be allocated to a queue of waiting players. Bloch and Catalana (2017) study the dynamic assignment of items to queuing agents considering both a setup with homogeneous and heterogeneous players keeping the size of the queue constant throughout the game. Hence whenever a player leaves the queue, a new player enters it. They investigate the concept of probabilistic queuing disciplines in which a probability is assigned to each of the possible priority orderings of players in the waiting list. This encompasses the special case where all priority rankings are equally likely. Su and Zenios (2004) study the dynamic allocations of kidneys to homogeneous patients on a waiting list who are able to refuse a proposed kidney to wait for a better one. Players enter the waiting list and leave it unmatched according to Poisson processes. The authors investigate the welfare aspect of different policies that assign priorities to newly entering players. An important difference between our game and these models of one sided dynamic matching is that in our game there is no growing or constant queue of applicants. Instead, because every player can at most receive one item, and no new players enter, the set of active players decreases over time. Furthermore, we assume all applicants for an item to receive the item with equal probability.

This paper continues as follows: First, in Section 3.2 we introduce the model. In Section 3.3 we explain how to derive the expected payoff functions and the best-response function. In Section 3.4 we construct a subgame perfect equilibrium in threshold strategies and illustrate that inequality in outcomes may persist in equilibrium. Then in Section 3.5 we study the welfare effects of the equal chances policy. We conclude the paper in Section 3.6 where we pose some open questions.

3.2 Model

The Game: Let P be a finite set of players. There is an interval of available values $X = [0, \bar{v}]$ with $\bar{v} > 0$. The players play the following infinite horizon game. At every period $t \in \mathbb{N}$ an item with value x_t is drawn from X according to a cumulative distribution F , independently of $x_1, \dots, x_{t-1}$. We assume that F has a density f with $f(x) > 0$ for every $x \in X$. After the value x_t is revealed, all the players can either apply for the item with this value or wait. If precisely one player applies for this item, this player receives the item. If none of the players apply for the item, the item disappears and cannot be chosen at a later stage. If multiple players apply for the item, the item is allocated to each of the applicants with equal probability. Once a player is allocated an item, he leaves the game. If a player is allocated an item with value x_t at period t , then his payoff is given by $\delta^t x_t$, where $0 < \delta < 1$ denotes a

discount factor. Note that we assume that players are homogeneous, hence all players have the same discount factor δ and observe the same common value x_t . We will refer to players who are not yet allocated an item as the *active players*. Because each player can have at most one item, the set of active players decreases throughout the game.

Strategy: A tuple (x, J) consisting of a value x and a non-empty subset of active players J is called a state. Let $S = \{(x, J) | x \in X, \emptyset \neq J \subseteq P\}$ denote the set of states, and let $S^p = \{(x, J) | x \in X, p \in J \subseteq P\}$ be the set of states where player p is active. For every player $p \in P$ a pure stationary strategy is a function $\sigma^p : S^p \rightarrow \{a, r\}$. Here the action a stands for “apply” and the action r stands for “reject”. Let Σ^p denote the set of all pure stationary strategies of player p . A pure stationary threshold strategy is a pure stationary strategy which has the following property: for every subset of active players J containing p , there exists a threshold τ_J^p such that $\sigma^p(x, J) = a$ if $x \geq \tau_J^p$ and $\sigma^p(x, J) = r$ if $x < \tau_J^p$. Since we will only deal with pure stationary threshold strategies, we will simply refer to them as “strategies”.

A collection of strategies $\sigma = (\sigma^p)_{p \in P}$ is called a strategy profile. Furthermore, $\sigma^{-p} = (\sigma^{p'})_{p' \in P \setminus \{p\}}$ denotes the strategy profile σ excluding player p . A strategy profile σ is called symmetric if for every $p, p' \in P$, for every $x \in X$ and for every J including p and p' we have $\sigma^p(x, J) = \sigma^{p'}(x, J)$.

Subgame: Consider a non-empty set of players J . A subgame corresponding to J is the continuation of the game when the set of active players is J . Note that there may exist different subgames where J is the set of active players, however all of these subgames will be treated in an identical manner. A strategy profile σ induces an expected payoff for every active player $p \in J$ in this subgame, denoted by $\pi_J^p(\sigma)$. The precise computation of this expected payoff will be thoroughly discussed in Section 3.3. Given a player $p \in J$ and a strategy profile σ^{-p} , a strategy σ^p for player $p \in J$ is called a best-response in the subgame at J if $\pi_J^p(\sigma) \geq \pi_J^p(\tilde{\sigma}^p, \sigma^{-p})$ for all strategies $\tilde{\sigma}^p$ of player p .

Subgame perfect equilibrium: A strategy profile σ is called a subgame perfect equilibrium if for every set of active players J , for every $p \in J$ and for every pure stationary¹ strategy $\tilde{\sigma}^p \in \Sigma^p$ of player p :

$$\pi_J^p(\sigma) \geq \pi_J^p(\tilde{\sigma}^p, \sigma^{-p}).$$

When all strategies in a subgame perfect equilibrium are threshold strategies, we refer to it as a subgame perfect equilibrium in threshold strategies. A subgame perfect equilibrium σ is called *symmetric* if the strategy profile σ is symmetric. A subgame perfect equilibrium σ is called *monotone* if for any set of active players J and any two players $p, p' \in J$:

$$\pi_P^p(\sigma) \geq \pi_P^{p'}(\sigma) \Rightarrow \pi_J^p(\sigma) \geq \pi_J^{p'}(\sigma).$$

This means that if player p has a higher expected payoff than player p' initially (when the set of active players is P), then this will also hold in the rest of the game as long both players are active.

Ordering the players: In order to make the presentation of the paper more transparent we order the players according to their threshold values. Since these threshold values may change throughout the game, so can the order of the players.

¹In this game the restriction to only consider deviations to stationary strategies is without loss of generality. This is because only the length of the history h influences the expected payoff of the active players at the history h , all other information contained in the realized history is irrelevant for the expected payoff. This is because all players are treated equally and the value of the presented item is independent from the history.

For every strategy profile σ and every set of active players J , fix a bijection $\phi_J^\sigma : J \rightarrow \{1, \dots, |J|\}$ such that whenever $\phi_J^\sigma(p) < \phi_J^\sigma(p')$ then $\tau_J^p \geq \tau_J^{p'}$. That is, ϕ_J^σ orders the active players according to decreasing threshold values. If $\phi_J^\sigma(p) = i$, then we will say that player p has rank i when the set of active players is J , and we will consistently use τ_J^i to denote his threshold. We will also refer to the player p as player (i, J) , as his rank i within the set of active players J uniquely identifies him. Thus, for example, player $(1, J)$ uses the highest threshold (but possibly not alone).

Continuation payoff: Consider a strategy profile σ , a set of active players J and two indices $\ell, i \in \{1, \dots, |J|\}$ with $\ell \neq i$. Let p_ℓ be the player with rank $\phi_J^\sigma(p_\ell) = \ell$, and similarly, let p_i be the player with rank $\phi_J^\sigma(p_i) = i$. We define $\pi_{J,-\ell}^i(\sigma)$ as the expected *continuation payoff* $\pi_{J \setminus \{p_\ell\}}^{p_i}(\sigma)$. That is, $\pi_{J,-\ell}^i(\sigma)$ is the expected payoff of player p_i when player p_ℓ leaves the game. This notation has the advantage that for example $\pi_{J,-1}^2(\sigma)$ is the expected payoff to the player with second highest threshold, thus player $(2, J)$, when the player with the highest threshold, player $(1, J)$, obtains the item and leaves the game. A *continuation strategy* $\sigma_{J,-\ell}^i$ is the strategy player (i, J) will follow when player (ℓ, J) leaves, yielding an expected payoff of $\pi_{J,-\ell}^i$. A *continuation plan* for player (i, J) is a vector $\sigma_{J,-}^i = (\sigma_{J,-1}^i, \dots, \sigma_{J,-(i-1)}^i, \sigma_{J,-(i+1)}^i, \dots, \sigma_{J,-k}^i)$ of continuation strategies. A continuation plan fully determines how player (i, J) will play when each of the other players leave the game. Therefore the combination of a threshold τ_J^i and continuation plan $\sigma_{J,-}^i$ gives the strategy of player (i, J) in the subgame J .

Notational conventions: For the ease of the presentation we will use the following notational guidelines. (1) Throughout the document, the superscript i is used only as a rank. (2) Whenever the strategy profile σ is clear from the context, we will omit it as an argument. (3) Whenever a strategy profile σ has the property that the strategy of each active player only considers the current value and the number of active players $k = |J|$, but not their identities, then we will use the more transparent notation k instead of J .

3.3 Expected payoff functions and Best-response strategies

3.3.1 Expected payoff functions

The expected payoff of the player (i, J) conditional on J being the set of active players before a value x is drawn consists of the sum of the following two parts: A_J^i which denotes the expected payoff when the player applies for the item, and R_J^i which denotes the expected payoff when the player rejects an item. Note that player (i, J) will apply for an item with value x iff $x \geq \tau_J^i$. However, the probability that he wins this specific item depends on the number of competing players that also apply for it. For notational convenience, define $\tau_J^0 = \bar{v}$, where $\bar{v}$ is the highest possible value an item can have. Let $k = |J|$ denote the number of active players in the subgame J .

Part A_J^i : Suppose that $x \geq \tau_J^i$, so player (i, J) will apply. Note that we can ignore the case where $x = \bar{v}$ as this happens with zero probability. Therefore there exists a $j \in \{0, 1, \dots, |J|\}$ such that $\tau_J^j > x \geq \tau_J^{j+1} \geq \tau_J^i$. Hence all the players $(j+1, J), \dots, (k, J)$, including player (i, J) , will apply for the item with value x . There are two cases:

Case 1: Player (i, J) wins the item. This happens with probability $1/(k-j)$, in which case player (i, J) receives the payoff x .

Case 2: One of the other applicants wins the item and leaves the game. This has probability $(k-j-1)/(k-j)$. As all applicants are equally likely to receive the item, the conditional expected payoff of player (i, J) in this case is equal to $1/(k-j-1) \sum_{\ell=j+1, \ell \neq i}^k \delta \pi_{J,-\ell}^i$.

Combining the two cases gives:

$$A_J^i = \sum_{j=0}^{i-1} \left[\frac{1}{k-j} \mathbb{E}[xI(\tau_J^j \geq x \geq \tau_J^{j+1})] + \frac{k-j-1}{k-j} \mathbb{P}(\tau_J^j \geq x \geq \tau_J^{j+1}) \left(\frac{1}{k-j-1} \sum_{\ell=j+1, \ell \neq i}^k \delta \pi_{J,-\ell}^i \right) \right].$$

Note that for this formula we have used the fact that we can ignore the case where x equals one of the thresholds, as this is a zero probability event.

Part R_J^i : If $x < \tau_J^i$, then player (i, J) will not apply for the item with value x . However, if $\tau_J^j > x \geq \tau_J^{j+1}$ for some $j \leq k-1$ then at least one player will apply. In this case player (i, J) will transition to a subgame with precisely $k-1$ active players. Because all of the $k-j$ applying players $(j+1, J), \dots, (k, J)$ are equally likely to receive the item, the expected payoff of player (i, J) when rejecting the item with value x is given by $1/(k-j) \sum_{\ell=j+1}^k \delta \pi_{J,-\ell}^i$. On the other hand, if $x < \tau_J^k$, then no player applies for the item with value x and the game will repeat itself. Using the stationarity assumption player (i, J) expects to receive $\delta \pi_J^i$ in this case. Therefore, the expected value for the range where player τ_J^i rejects the item is:

$$R_J^i = \underbrace{\sum_{j=i}^{k-1} \mathbb{P}(\tau_J^j \geq x \geq \tau_J^{j+1}) \left(\frac{1}{k-j} \sum_{\ell=j+1}^k \delta \pi_{J,-\ell}^i \right)}_{T_J^i} + \delta F(\tau_J^k) \pi_J^i.$$

Here the term T_J^i indicates the conditional expected payoff of player (i, J) in case he rejects and the game transitions to a subgame with fewer active players.

Combining A_J^i and R_J^i : Based on the previous two parts, the expected payoff of player (i, J) is equal to $\pi_J^i = A_J^i + R_J^i$. To simplify the expression for π_J^i , we introduce the function $\eta : X \rightarrow \mathbb{R}$ defined by:

$$\eta(\tau) = \frac{1}{1 - \delta F(\tau)}.$$

Then we obtain:

$$\begin{aligned} \pi_J^i &= \eta(\tau_J^k) (A_J^i + T_J^i) \\ &= \eta(\tau_J^k) \left(\sum_{j=0}^{i-1} \left[\frac{1}{k-j} \mathbb{E}[xI(\tau_J^j \geq x \geq \tau_J^{j+1})] + \frac{k-j-1}{k-j} \mathbb{P}(\tau_J^j \geq x \geq \tau_J^{j+1}) \left(\frac{1}{k-j-1} \sum_{\ell=j+1, \ell \neq i}^k \delta \pi_{J,-\ell}^i \right) \right] \right) \\ &\quad + \eta(\tau_J^k) \left(\sum_{j=i}^{k-1} \mathbb{P}(\tau_J^j \geq x \geq \tau_J^{j+1}) \left(\frac{1}{k-j} \sum_{\ell=j+1}^k \delta \pi_{J,-\ell}^i \right) \right). \end{aligned}$$

3.3.2 Best-response thresholds

The notion of best-response thresholds will be useful when applying backwards induction on the number of active players.

Fix a set of active players J and a partial strategy profile σ_J . Let $|J| = k$. Then the thresholds $\tau_J^1 \geq \dots \geq \tau_J^k$ are defined and so are the expected continuation payoffs $\pi_{J,-\ell}^i$ for every $i, \ell \leq k$ with $i \neq \ell$. For any player (i, J) we can now compute the *best-response threshold* $BR_J^i(\tau_J^1, \dots, \tau_J^{i-1}, \tau_J^{i+1}, \dots, \tau_J^k, (\pi_{J,-\ell}^i)_{\substack{\ell \leq k \\ i \neq \ell}})$ as follows

Step 1: Splitting up the interval of values.

Use the $k - 1$ thresholds $\tau_J^1, \dots, \tau_J^{i-1}, \tau_J^{i+1}, \dots, \tau_J^k$ to split up the interval X in k adjacent subintervals. For each $h \in \{1, \dots, k\}$ let

$$X_h^i = \begin{cases} [\tau_J^h, \tau_J^{h-1}] & h < i, \\ [\tau_J^{h+1}, \tau_J^{h-1}] & h = i, \\ [\tau_J^{h+1}, \tau_J^h] & h > i, \end{cases}$$

denote one of these subintervals.

Step 2: Finding the locally maximizing thresholds $\tau_h^{i,*} \in X_h^i$.

Figure 3.1: Splitting up the interval $[0, \bar{v}]$ for the construction of the best-response function BR_J^i .

For each interval X_h^i there is a fixed number of competitors who apply for an item with the value $x \in X_h^i$. Because the expected continuation payoffs $\pi_{J,-\ell}^i$ are known for all $\ell \neq i$ and all the thresholds $\tau_J^1, \dots, \tau_J^{i-1}, \tau_J^{i+1}, \dots, \tau_J^k$ are fixed, it is possible to derive the expected payoff function of a player using a threshold $\tau \in X_h^i$ and receiving the expected continuation payoffs $(\pi_{J,-\ell}^i)_{\substack{\ell \leq k \\ \ell \neq i}}$. As this function will be continuous and the interval X_h^i is bounded, there will exist a threshold $\tau_h^{i,*} \in X_h^i$ that maximizes this function.

Step 3: Finding a globally maximizing threshold.

As there are only finitely many intervals X_h^i and each of these intervals has a locally maximizing threshold $\tau_h^{i,*}$ it is now trivial to select a globally maximizing threshold. It is this globally maximizing threshold to which we will refer to as the best-response threshold $BR_J^i(\tau_J^1, \dots, \tau_J^{i-1}, \tau_J^{i+1}, \dots, \tau_J^k, (\pi_{J,-\ell}^i)_{\substack{\ell \leq k \\ i \neq \ell}})$.

3.4 Subgame perfect equilibria

In this section we show that there does not exist a pure symmetric subgame perfect equilibrium in threshold strategies and we construct a pure asymmetric subgame perfect equilibrium. The constructed equilibrium will be monotone, i.e. it will have that property that players maintain their relative rankings in terms of expected payoffs throughout the game. Furthermore, the equilibrium payoff vector in every subgame will only depend on the number of active players.

The key intuition behind the non-existence of a pure symmetric equilibrium is due to the fact that the number of active players goes down over time. As all players observe

an item with a common value x , players may prefer a subgame in which there are fewer competitors present. So the players might be inclined to wait to enter a less competitive period. However, as the players are impatient, waiting is costly and if all players wait the competition will not decrease. Therefore, some players might act less picky, increasing their chances of winning an item quickly and given this, other players might act more picky, expecting to wait longer for a higher value item.

We start with a technical lemma which will be useful in the rest of the paper.

Lemma 3.4.1. For any $\tau^a, \tau^b \in [0, \bar{v}]$ with $\tau^a > \tau^b$ we have:

$$\tau^b(F(\tau^a) - F(\tau^b)) < \mathbb{E}[xI(\tau^a > x \geq \tau^b)] < \tau^a(F(\tau^a) - F(\tau^b)). \quad (3.4.1)$$

Proof. From the assumptions made on the distribution F we have that $F(\tau^a) - F(\tau^b) > 0$. We have that there exists an $\epsilon > 0$ such that $\tau^a > \tau^a - \epsilon > \tau^b$. We have that:

$$\begin{aligned} \mathbb{E}[xI(\tau^a > x \geq \tau^b)] &= \int_{\tau^b}^{\tau^a} x dF = \int_{\tau^b}^{\tau^a - \epsilon} x dF + \int_{\tau^a - \epsilon}^{\tau^a} x dF \\ &\leq (\tau^a - \epsilon)(F(\tau^a - \epsilon) - F(\tau^b)) + \tau^a(F(\tau^a) - F(\tau^a - \epsilon)) \\ &< \tau^a(F(\tau^a) - F(\tau^b)). \end{aligned}$$

Because $\tau^a > \tau^b$ there exists an $\epsilon > 0$ such that $\tau^a > \tau^b + \epsilon > \tau^b$.

$$\begin{aligned} \mathbb{E}[xI(\tau^a > x \geq \tau^b)] &= \int_{\tau^b}^{\tau^a} x dF = \int_{\tau^b + \epsilon}^{\tau^a} x dF + \int_{\tau^b}^{\tau^b + \epsilon} x dF \\ &\geq (\tau^b + \epsilon)(F(\tau^a) - F(\tau^b + \epsilon)) + \tau^b(F(\tau^b + \epsilon) - F(\tau^b)) \\ &> \tau^b(F(\tau^a) - F(\tau^b)). \end{aligned}$$

□

Lemma 3.4.2. In each subgame with only one player, the optimal strategy for the active player $p \in P$ is given by

$$\sigma^p(x, \{p\}) = \begin{cases} a & \text{if } x \geq \tau_1^1, \\ r & \text{otherwise,} \end{cases}$$

where τ_1^1 is the unique value which solves

$$\tau_1^1 = \delta \frac{\mathbb{E}[xI(x \geq \tau_1^1)]}{1 - \delta F(\tau_1^1)} = \delta \eta(\tau_1^1) \mathbb{E}[xI(x \geq \tau_1^1)]. \quad (3.4.2)$$

Using this threshold strategy will yield an optimal expected payoff $\pi_1^1 = \frac{\tau_1^1}{\delta}$.

Proof. Note that in any subgame where there is only one player this problem reduces to an unconstrained optimal stopping problem which is known to have a unique threshold solution. See for example Lippman and McCall (1979) or Rogerson et al. (2005). □

3.4.1 The two player game

We start by analysing the two player game as this analysis will be illustrative for the k player case.

We first explain intuitively why there cannot exist a pure symmetric subgame perfect equilibrium in threshold strategies. Assume by way of contradiction there exists a common threshold $\tilde{\tau}_2$ which both players use in equilibrium and let $\tilde{\pi}_2$ denote the corresponding equilibrium payoff. Because $\tilde{\tau}_2$ is a threshold used in the subgame perfect equilibrium it should be a best-response for both players to apply for an item with the threshold value $\tilde{\tau}_2$. Hence $\tilde{\tau}_2 \geq \delta\pi_1^1$, where $\delta\pi_1^1$ is the optimal expected payoff for rejecting the item and moving to the one player subgame. Using the condition of a subgame equilibrium a second time and incorporating the stationarity assumption, we find that it should be a best-response for both players to reject any item with a value slightly below the threshold $\tilde{\tau}_2$. Doing so will give them both an expected payoff of $\delta\tilde{\pi}_2$. Combining those equilibrium conditions it follows that the expected equilibrium payoff of the two player subgame is at least the optimal payoff in the one-player subgame, i.e. $\tilde{\pi}_2 \geq \pi_1^1$. This will lead to a contradiction. Because with two players present, the players are not sure that when applying for an item it will be allocated to them, which suggests that the discounted expected payoff $\delta\tilde{\pi}_2$ of both players when using the threshold $\tilde{\tau}_2$ is strictly lower than the optimal expected payoff in a subgame where they are alone. Because compared to a one player subgame, in a two player subgame the active players either have to wait longer for a good item or be less picky.

The proof of the following theorem formalizes this intuition. Note that it is sufficient to show that there is no pure symmetric subgame perfect equilibrium in a subgame with two players as games with more players always contain subgames with precisely two players.

Theorem 3.4.3. There does not exist a pure symmetric subgame perfect equilibrium in threshold strategies.

Proof. Proof by contradiction. Suppose there is a pure symmetric subgame perfect equilibrium in threshold strategies and denote this equilibrium with $\tilde{\sigma}$. Note that as soon as a player enters a subgame where he is the only active player, he will use the unique optimal threshold τ_1^1 yielding an expected payoff of π_1^1 as derived in Lemma 3.4.2. Now consider a subgame with two active players and let $\tilde{\tau}_2$ denote the common threshold they use in this two player subgame. Then the resulting expected payoff $\tilde{\pi}_2$ to both players, given that they use the optimal threshold τ_1^1 in the one player subgame, is given by:

$$\tilde{\pi}_2 = \frac{\frac{1}{2}\mathbb{E}[xI(x \geq \tilde{\tau}_2)] + \frac{1}{2}(1 - F(\tilde{\tau}_2))\delta\pi_1^1}{1 - \delta F(\tilde{\tau}_2)} = \eta(\tilde{\tau}_2) \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tilde{\tau}_2)] + \frac{1}{2}(1 - F(\tilde{\tau}_2))\delta\pi_1^1 \right).$$

Note that $\tilde{\tau}_2 < \bar{v}$ because waiting for the highest value $\bar{v}$ is a dominated strategy yielding an expected payoff of 0. Because $\tilde{\sigma}$ is a subgame perfect equilibrium, we have that in any state with two active players it is a best-response for any player to use a threshold $\tilde{\tau}_2$ given that the other player uses the threshold $\tilde{\tau}_2$. Hence for all $x \geq \tilde{\tau}_2$ it is better to apply for x , than to reject and progress to the subgame where the player is alone. Therefore $1/2x + 1/2\delta\pi_1^1 \geq \delta\pi_1^1$ which simplifies to $x \geq \delta\pi_1^1$. Because this has to hold for all $x \geq \tilde{\tau}_2$ we have that $\tilde{\tau}_2 \geq \delta\pi_1^1$.

Similarly, for all $x < \tilde{\tau}_2$ it is better to wait and receive a discounted expected payoff of $\delta\tilde{\pi}_2$ than to pick x . Hence $x \leq \delta\tilde{\pi}_2$. Because this holds for all $x < \tilde{\tau}_2$ we have that for every $\epsilon > 0$, $\tilde{\tau}_2 - \epsilon \leq \delta\tilde{\pi}_2$. Combining these equilibrium conditions we find that for all $\epsilon > 0$:

$$\delta\tilde{\pi}_2 + \epsilon \geq \tilde{\tau}_2 \geq \delta\pi_1^1$$

Taking the limit for ϵ going to 0 and the fact that $\delta > 0$ we conclude that:

$$\tilde{\pi}_2 \geq \pi_1^1$$

However observe that:

$$\begin{aligned} \tilde{\pi}_2 &= \eta(\tilde{\tau}_2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tilde{\tau}_2)] + \frac{1}{2}(1 - F(\tilde{\tau}_2))\delta\pi_1^1 \right) \\ &\leq \eta(\tilde{\tau}_2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tilde{\tau}_2)] + \frac{1}{2}(1 - F(\tilde{\tau}_2))\tilde{\tau}_2 \right) \\ &< \eta(\tilde{\tau}_2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tilde{\tau}_2)] + \frac{1}{2} \mathbb{E}[xI(x \geq \tau_1^1)] \right) \\ &= \eta(\tilde{\tau}_2) \mathbb{E}[xI(x \geq \tilde{\tau}_2)] \\ &\leq \eta(\tilde{\tau}_1^1) \mathbb{E}[xI(x \geq \tau_1^1)] \\ &= \pi_1^1. \end{aligned}$$

Here the first inequality follows from the fact that $\delta\pi_1^1 \leq \tilde{\tau}_2$, the second inequality follows from Lemma 3.4.1 and the fact that $\tilde{\tau}_2 < \bar{v}$. The last inequality follows from the fact that τ_1^1 maximizes $\eta(\tau)\mathbb{E}[xI(x \geq \tau)]$ as shown in Lemma 3.4.2. We found that $\tilde{\pi}_2 < \pi_1^1$, a contradiction. $\square$

Given this non-existence of a pure symmetric subgame perfect equilibrium in threshold strategies, we are interested to find another pure subgame perfect equilibrium in threshold strategies. We start with the two player game. Note that in any pure subgame perfect equilibrium if a player enters a subgame where he is the only active player he will use the unique threshold τ_1^1 and receive an expected payoff of π_1^1 . Let $T = \{(\tau_2^1, \tau_2^2) \in [0, \bar{v}]^2 | \tau_2^1 \geq \tau_2^2\}$ denote the set of possible thresholds used in the two player game. Then the associated expected payoff functions $\pi_2^1 : T \rightarrow \mathbb{R}$ and $\pi_2^2 : T \rightarrow \mathbb{R}$ are given by:

$$\pi_2^1(\tau_2^1, \tau_2^2) = \eta(\tau_2^2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1) \delta\pi_1^1 + \mathbb{P}(\tau_2^1 > x \geq \tau_2^2) \delta\pi_1^1 \right), \quad (3.4.3)$$

$$\pi_2^2(\tau_2^1, \tau_2^2) = \eta(\tau_2^2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1) \delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)] \right). \quad (3.4.4)$$

We will show that in equilibrium the highest threshold used τ_2^1 is the same as the optimal threshold in the one player subgame, i.e. $\tau_2^1 = \tau_1^1 = \delta\pi_1^1$, and the lowest threshold τ_2^2 equals the associated discounted expected payoff when both players remain, i.e. $\tau_2^2 = \delta\pi_2^2$. The intuition behind this equilibrium is straightforward. If both players use strictly different thresholds, then the player that uses the highest threshold will only apply to items with values that give at least the discounted expected value he would receive when he enters a one player subgame. He knows that the other player will apply for items with values just beneath this threshold level, guaranteeing that he moves to a subgame without competition and receives an expected payoff of $\delta\pi_1^1$. The other player, using a lower threshold, will apply for an item as long as its value is at least the value of waiting another period and repeating the same game again. Note here the threshold τ_2^2 is defined as a fixed point of the equation $\tau_2^2 = \delta\pi_2^2(\tau_2^1, \tau_2^2)$. The following lemma shows that this fixpoint exists and is unique.

Lemma 3.4.4. Let $\tau_2^1 = \delta\pi_1^1$. Then there exists a unique solution to the fixpoint equation $\tau = \delta\pi_2^2(\tau_2^1, \tau)$ with $\tau \in [0, \tau_2^1]$.

Proof. Note that $\pi_2^2(\tau_2^1, 0) > 0$ and

$$\begin{aligned}\delta\pi_2^2(\tau_2^1, \tau_2^1) &= \delta\eta(\tau_2^1) \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 \right) \\ &< \delta\eta(\tau_2^1) (\mathbb{E}[xI(x \geq \tau_2^1)]) = \delta\pi_1^1 = \tau_2^1.\end{aligned}$$

Because the function $\pi_2^2(\tau_2^1, \cdot)$ is continuous there exists a fixpoint $\tau_2^2 \in [0, \tau_2^1]$ such that $\tau_2^2 = \delta\pi_2^2(\tau_2^1, \tau_2^2)$. To show the uniqueness of this fixpoint observe that $\frac{\partial\pi_2^2}{\partial\tau_2^2}(\tau_2^1, \tau_2^2) = 0 \Leftrightarrow \tau_2^2 = \delta\pi_2^2(\tau_2^1, \tau_2^2)$. Indeed, we have that $f(\tau_2^2)\eta(\tau_2^2)^2 > 0$ for all $\tau_2^2 \in [0, \bar{v}]$ and

$$\begin{aligned}\frac{\partial\pi_2^2}{\partial\tau_2^2}(\tau_2^1, \tau_2^2) &= f(\tau_2^2)\eta(\tau_2^2)^2 \left(-\tau_2^2(1 - \delta F(\tau_2^2)) \right. \\ &\quad \left. + \delta \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)] \right) \right).\end{aligned}$$

Furthermore the function $\pi_2^2(\tau_2^1, \cdot)$ is strictly increasing if $\tau_2^2 < \delta\pi_2^2(\tau_2^1, \tau_2^2)$ and strictly decreasing if $\tau_2^2 > \delta\pi_2^2(\tau_2^1, \tau_2^2)$. Therefore any fixpoint $\tau_2^2 = \delta\pi_2^2(\tau_2^1, \tau_2^2)$ is a local maximum. Hence there can only be one fixpoint. $\square$

The following theorem constructs a pure subgame perfect equilibrium for a game with 2 players. This construction is illustrative for the one used in the general k player case.

Theorem 3.4.5. If $n = 2$ then there exists a pure subgame perfect equilibrium in threshold strategies which is unique up to the permutation of the players, such that $\tau_2^1 = \delta\pi_1^1$ and $\tau_2^2 = \delta\pi_2^2$ and $\pi_1^1 > \pi_2^1 > \pi_2^2$.

Proof. Step 1: Finding the best-response thresholds and proving uniqueness.

Because of Lemma 3.4.3 we have that if both players use a threshold strategy then both players will use a different threshold. Let τ_2^1, τ_2^2 denote the highest, resp. lowest threshold, then the expected payoff functions are given by Equations 3.4.3 and 3.4.4. First assume $\tau_2^2 < \delta\pi_1^1$. Then observe that for all thresholds τ_2^1 with $\tau_2^1 > \tau_2^2$, the threshold strategy where $\tau_2^1 = \delta\pi_1^1$ maximizes π_2^1 . Indeed, we have that:

$$\left. \frac{\partial\pi_2^1(\tau, \tau_2^2)}{\partial\tau} \right|_{\tau=\tau_2^1} = \eta(\tau_2^2)f(\tau_2^1) \left[-\frac{1}{2}\tau_2^1 - \frac{1}{2}\delta\pi_1^1 + \delta\pi_1^1 \right] = 0$$

and $\frac{\partial\pi_2^1(\tau, \tau_2^2)}{\partial\tau} > 0$ for all $\tau_2^2 < \tau < \tau_2^1$ and $\frac{\partial\pi_2^1(\tau, \tau_2^2)}{\partial\tau} < 0$ for all $\tau > \tau_2^1$. Therefore it is clear that τ_2^1 is the unique maximizing value of $\pi_2^1(\cdot, \tau_2^2)$. Let $\tau_2^2 = BR_2^2(\tau_2^1, \pi_1^1) = \arg \max_{\tau \in [0, \tau_2^1]} \pi_2^2(\tau_2^1, \tau)$. We have that $BR_2^2(\tau_2^1, \pi_1^1) = \delta\pi_2^2(\tau_2^1, \tau_2^2)$. Indeed,

$$\begin{aligned}\left. \frac{\partial\pi_2^2(\tau_2^1, \tau)}{\partial\tau} \right|_{\tau=\tau_2^2} &= \eta(\tau_2^2)f(\tau_2^2) \left[-\tau_2^2 + \delta\eta(\tau_2^2) \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)] \right) \right]\end{aligned}$$

From this it follows that $\left. \frac{\partial\pi_2^2(\tau_2^1, \tau)}{\partial\tau} \right|_{\tau=\tau_2^2} = 0$ if $\tau_2^2 = \delta\pi_2^2(\tau_2^1, \tau_2^2)$. Note that in the interval $[0, \tau_2^1]$ the fixpoint $\tau_2^2 = \delta\pi_2^2(\tau_2^1, \tau_2^2)$ exists and is unique because of Lemma 3.4.4. It can be verified

that $\tau_2^2 = \delta\pi_2^2$ is indeed a maximum. To prove that $\tau_2^1 = BR_2^1(\tau_2^2, \pi_1^1)$ we still have to verify that $\tau_2^1 > \tau_2^2$. Note that for this it is sufficient to show that $\pi_1^1 > \pi_2^1(\tau_2^1, \tau_2^2) > \pi_2^2(\tau_2^1, \tau_2^2)$.

Step 2: $\pi_2^1(\tau_2^1, \tau_2^2) > \pi_2^2(\tau_2^1, \tau_2^2)$.

We have that $\pi_2^1(\tau_2^1, \tau_2^2) > \pi_2^2(\tau_2^1, \tau_2^2) \Leftrightarrow \mathbb{P}(\tau_2^1 > x \geq \tau_2^2)\delta\pi_1^1 > \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)]$. Using the fact that $\tau_2^1 = \delta\pi_1^1$ and Lemma 3.4.1 we conclude that indeed $\mathbb{P}(\tau_2^1 > x \geq \tau_2^2)\delta\pi_1^1 = \mathbb{P}(\tau_2^1 > x \geq \tau_2^2)\tau_2^1 > \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)]$ and hence $\pi_2^1(\tau_2^1, \tau_2^2) > \pi_2^2(\tau_2^1, \tau_2^2)$.

Step 3: $\pi_1^1 > \pi_2^1(\tau_2^1, \tau_2^2)$.

Proof by contradiction. Assume that $\pi_1^1 \leq \pi_2^1(\tau_2^1, \tau_2^2)$, then using Equation 3.4.3 have that:

$$\begin{aligned} (1 - \delta F(\tau_2^2))\pi_1^1 &\leq (1 - \delta F(\tau_2^2))\pi_2^1(\tau_2^1, \tau_2^2) \\ &= \frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 + \mathbb{P}(\tau_2^1 > x \geq \tau_2^2)\delta\pi_1^1 \\ &\Rightarrow (1 - \delta F(\tau_2^2))\pi_1^1 \leq \frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1. \end{aligned}$$

Hence we have that:

$$\begin{aligned} \pi_1^1 &\leq \eta(\tau_2^1) \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 \right) \\ &= \eta(\tau_2^1) \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 \right) \\ &< \eta(\tau_2^1) \left(\frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2}\mathbb{E}[xI(x \geq \tau_2^1)] \right) \\ &= \eta(\tau_2^1)\mathbb{E}[xI(x \geq \tau_2^1)] \\ &= \eta(\tau_1^1)\mathbb{E}[xI(x \geq \tau_1^1)] \\ &= \pi_1^1. \end{aligned}$$

Here, for the strict inequality we have used the fact that $\tau_2^1 = \delta\pi_1^1 < \bar{v}$. We found a contradiction. We can conclude that $\tau_2^1 > \tau_2^2$ and hence $\tau_2^1 = BR_2^1(\tau_2^2, \pi_1^1)$. Furthermore we can conclude that $\pi_1^1 > \pi_2^1(\tau_2^1, \tau_2^2) > \pi_2^2(\tau_2^1, \tau_2^2)$.

Step 4: Verification of a subgame perfect equilibrium.

It is still necessary to verify that neither player will deviate to another (possibly non-threshold) strategy. It is sufficient to check no deviation in the two player subgame. We distinguish the following three cases:

- If $x \geq \delta\pi_1^1$, then it is clear that the best-response is to apply given that the other player applies as this gives an expected payoff of $1/2x + 1/2\delta\pi_1^1$, while rejecting only yields an expected payoff of $\delta\pi_1^1$.
- If $\delta\pi_1^1 > x \geq \delta\pi_2^2(\tau_2^1, \tau_2^2)$, then given that player τ_2^2 applies it is a best-response for player τ_2^1 to reject and receive an expected payoff of $\delta\pi_1^1$. Given that player τ_2^1 rejects, it is a best response for player τ_2^2 to apply as rejecting the item as well would only yield an expected payoff of $\delta\pi_2^2(\tau_2^1, \tau_2^2)$.
- If $x < \delta\pi_2^2$, then because $\pi_1^1 > \pi_2^1(\tau_2^1, \tau_2^2) > \pi_2^2(\tau_2^1, \tau_2^2)$ it is a dominant strategy for all players to reject x .

□

Corollary 3.4.6. When restricting ourselves to subgame perfect equilibria in threshold strategies, we find that in any two-player subgame, the player who uses the highest threshold will use the same threshold that he would use if he is alone.

Figure 3.2: A two-player subgame with $X \sim U[0, 1]$

The computation of the equilibrium thresholds as well as the resulting expected equilibrium payoffs can easily be done with mathematical software. We computed the equilibrium thresholds τ_2^1 and τ_2^2 and the corresponding equilibrium payoffs for varying levels of δ when $X \sim U[0, 1]$. Our results are plotted in Figure 3.2. In Figure 3.2a we plotted the thresholds τ_2^1 and τ_2^2 as a function of δ . The fact that both functions are strictly increasing in δ is intuitive as the higher the δ , the more patient both players are, hence the pickier they can be. In Figure 3.2b we plotted the difference in expected equilibrium payoffs. We see that indeed $\pi_2^1 > \pi_2^2$ for all $\delta \in (0, 1)$.

We see that small values of δ model impatient players. In this case both players will apply for the item even if the item has a very low value, as it is too costly to wait for a potentially higher value item. If δ is close to one, both players can afford to wait to obtain an item with a very high value. Because waiting is almost free and the distribution of values for the items is the same every period, the influence of the competing player is almost negligible in this case.

In this subsection we showed that when restricting ourselves to pure threshold strategies, there does not exist a symmetric subgame perfect equilibrium. It is a natural question to ask whether there exists a symmetric equilibrium if one allows for behavioral strategies, i.e. strategies which for any value x and set of competitors J specify the probability with which one applies for the item with value x . It turns out that there may exist a symmetric subgame perfect equilibrium in behavioral strategies. For the interested reader we provide an example of the construction for a symmetric behavioral equilibrium in the Appendix C.

3.4.2 The monotone subgame perfect equilibrium

We now extend the two player construction to the multi-player setting and show that there exists a monotone subgame perfect equilibrium in threshold strategies, i.e. there exists a subgame perfect equilibrium in which the players are endogenously ranked in terms of expected payoff throughout the game. The construction of this equilibrium is intuitive and

generalizes the construction of the two player case. The thresholds the players use are determined recursively and the underlying recursive procedure is illustrated in Figure 3.3. Once an equilibrium has been constructed for a subgame with $k - 1$ active players, we will use the same set of equilibrium strategies in all subgames with $k - 1$ players.

To extend the subgame perfect equilibrium from a state with $k - 1$ active players to a state with k players one must define all the thresholds used in the k -player subgame as well as all continuation plans for the smaller subgames. These thresholds and continuation plans are defined as follows. Let J denote a set of players with k active players. Then player $(1, J)$ will always use the strategy σ_{k-1}^1 yielding the highest payoff π_{k-1}^1 in any of the $k - 1$ active player subgames. Player $(2, J)$ will use the strategy σ_{k-1}^2 yielding the second highest payoff in any of the $k - 1$ active player subgames, except when player $(1, J)$ was the player who left, in which case he will use σ_{k-1}^1 . Similarly (i, J) will use the strategy σ_{k-1}^i yielding the i -th highest payoff in any of the $k - 1$ active player subgames, except when any of the players $(1, J), \dots, (i - 1, J)$ leaves, in which case he will use σ_{k-1}^{i-1} . Finally, player (k, J) will always use the strategy σ_{k-1}^{k-1} yielding the lowest expected payoff in the subgame with $k - 1$ active players. We show that this player (k, J) will use the lowest thresholds which will at most be the lowest expected payoff of the subgame with $k - 1$ active players. Hence if an item with the value $\delta\pi_{k-1}^i$ appears, player (k, J) will apply. This will guarantee that the game moves to a subgame with only $k - 1$ players. This will prompt the other players $(i, J), \dots, (k - 1, J)$ to also all apply for the item with value $\delta\pi_{k-1}^i$. Under such a construction the thresholds τ_k^i for $i \neq k$ will be recursively determined and given by $\tau_k^i = \delta\pi_{k-1}^i$. Then the last threshold τ_k^k can be defined as the best-response threshold given the other thresholds.

Figure 3.3: Construction of a monotone subgame perfect equilibrium for a subgame with k active players.

Theorem 3.4.7. There exists a monotone subgame perfect equilibrium in threshold strategies in which at every subgame all active players have distinct expected equilibrium payoffs.

Proof. We will construct a monotone subgame perfect equilibrium using backwards induction on the number of active players. The construction is illustrated in Figure 3.3.

PART 1: STRATEGY WHEN THERE IS ONLY 1 ACTIVE PLAYER

Because of Lemma 3.4.2 there is a unique optimal strategy in the one player subgame, i.e. to apply for x if $x \geq \tau_1^1$ and to reject otherwise, where τ_1^1 is the unique value which solves $\tau_1^1 = \delta \eta(\tau_1^1) \mathbb{E}[xI(x \geq \tau_1^1)]$. Using this threshold strategy yields an optimal expected payoff $\pi_1^1 = \frac{\tau_1^1}{\delta}$.

PART 2: DEFINITION OF THE THRESHOLDS AND CONTINUATION STRATEGIES

Because of Theorem 3.4.5 for $k = 2$ there exists a monotone subgame perfect equilibrium for the two player subgame with $\pi_2^1 > \pi_2^2$. Assume there exists a monotone subgame perfect equilibrium for a subgame with $k - 1$ active players, where $k \geq 3$. Fix such an equilibrium. Let $\pi_{k-1}^1 > \dots > \pi_{k-1}^{k-1}$ denote the ranked equilibrium payoffs and let $\sigma_{k-1}^1, \dots, \sigma_{k-1}^{k-1}$ denote the corresponding equilibrium strategies. We will use the equilibrium strategies of this subgame perfect equilibrium in all subgames with $k - 1$ active players. Now consider the subgame J with k active players. Let the thresholds $\tau_k^1, \dots, \tau_k^k$ be defined as follows:

$$\tau_k^i = \delta \pi_{k-1}^i \quad \forall i \leq k - 1 \quad (3.4.5)$$

$$\tau_k^k = BR_k^k(\tau_k^1, \dots, \tau_k^{k-1}, (\pi_{k-1}^{k-1}, \dots, \pi_{k-1}^1)). \quad (3.4.6)$$

Let the set of strategies used by the active players in any of the subgames with $k - 1$ active players be equal to $\{\sigma_{k-1}^i | i \leq k - 1\}$. Note that $\tau_k^1 > \tau_k^{k-1}$, so it is possible to uniquely distinguish $k - 1$ players based on the threshold they use. For each of the thresholds τ_k^i with $i \leq k - 1$ fix a player that uses this threshold and rename him player i . Doing so has renamed precisely $k - 1$ from the k players. Rename the remaining player, i.e. the player using τ_k^k , player k . Now for each of these players let the continuation strategy be defined as follows: if player ℓ leaves player i uses the strategy σ_{k-1}^i if $i < \ell$ and the strategy σ_{k-1}^{i-1} if $i > \ell$.

PART 3: EXISTENCE OF EQUILIBRIUM

Following the induction hypothesis, once the game enters a state with $k - 1$ players, none of the players will have an incentive to deviate at any state. It remains to show that in the subgame with k players, no player has an incentive to deviate. The subgame with k active players is defined by the thresholds $\tau_k^1, \dots, \tau_k^k$ and the continuation strategies.

From the definition it is clear that $\tau_k^1 > \dots > \tau_k^{k-1}$. We now verify that $\tau_k^{k-1} \geq \tau_k^k$. Note that from the best-response threshold definition of τ_k^k it is clear that the expected continuation payoff of this player in a smaller subgame is always π_{k-1}^{k-1} .

Step 1²: $\tau_k^{k-1} \geq \tau_k^k$

Proof by contradiction. Assume that $\tau_k^k > \tau_k^{k-1} = \delta \pi_{k-1}^{k-1}$. Because $\tau_k^k > \tau_k^{k-1}$ there exists a threshold $\tilde{\tau}_k^k > \tau_k^{k-1} > \tau_k^k$. Fix any such threshold $\tilde{\tau}_k^k$. We now argue that a player who always faces a continuation payoff of $\delta \pi_{k-1}^{k-1}$ will have a strictly higher expected payoff when using the threshold $\tilde{\tau}_k^k$ then when he would use τ_k^k . Indeed, the difference between both expected payoffs can only be due to the actions chosen for items with values in the range $[\tilde{\tau}_k^k, \tau_k^k)$. For these item, using a threshold $\tilde{\tau}_k^k$ implies applying for the item, while using

²We conjecture this inequality is actually strict but cannot prove this.

τ_k^k implies rejecting the item. By construction each of these items has a value x strictly above $\delta\pi_{k-1}^{k-1}$. Applying for such an item, yields an expected payoff of $\alpha x + (1 - \alpha)\delta\pi_{k-1}^{k-1} > \delta\pi_{k-1}^{k-1}$, for some strictly positive weight α which depends on the number of competitors. while rejecting such an item always yields an expected payoff of $\delta\pi_{k-1}^{k-1}$, as at least one other player will apply. Because $\mathbb{P}([\tilde{\tau}_k^k, \tau_k^k]) > 0$ we can conclude that the expected payoff of using a threshold $\tilde{\tau}_k^k$ is strictly higher than the expected payoff of using the threshold τ_k^k . A contradiction with the definition of τ_k^k as best-response threshold.

Step 2: Determining the payoff function π_k^i for every $i \leq k$

For ease of notation we will use the convention that $\pi_{k-1}^0 = \pi_{k-1}^1$ and $\pi_{k-1}^k = \pi_{k-1}^{k-1}$. Then for every $i \leq k$ and $0 \leq j \leq i - 1$ let

$$a_j^i = \frac{1}{k-j} \mathbb{E}[xI(\tau_k^j \geq x \geq \tau_k^{j+1})] + \frac{k-j-1}{k-j} \mathbb{P}(\tau_k^j \geq x \geq \tau_k^{j+1}) \left(\frac{i-j-1}{k-j-1} \delta\pi_{k-1}^{i-1} + \frac{k-i}{k-j-1} \delta\pi_{k-1}^i \right)$$

denote the expected payoff of player i when applying for an item with a value in $[\tau_k^j, \tau_k^{j+1}]$. From Step 1 we have that τ_k^k is the lowest threshold used hence the expected payoff of player i in the k -player subgame is given by:

$$\pi_k^i(\tau_k^1, \dots, \tau_k^k) = \eta(\tau_k^k) \left(\sum_{j=0}^{i-1} a_j^i + \mathbb{P}(\tau_k^i \geq x \geq \tau_k^k) \delta\pi_{k-1}^i \right) \quad (3.4.7)$$

Step 3: $\pi_k^1 > \pi_k^2 > \dots > \pi_k^{k-1} > \pi_k^k$

From the induction hypothesis we have that $\pi_{k-1}^1 > \dots > \pi_{k-1}^{k-1}$ and $k \geq 3$. Because of the definition of τ_k^j we have that $\mathbb{P}(\tau_k^j > x \geq \tau_k^{j+1}) > 0$ for all $0 \leq j \leq k-2$. Furthermore we have that

$$\frac{i-1}{k-1} \delta\pi_{k-1}^{i-1} + \frac{k-i}{k-j-1} \delta\pi_{k-1}^i > \frac{i}{k-1} \delta\pi_{k-1}^i + \frac{k-i-1}{k-1} \delta\pi_{k-1}^{i+1} \quad \text{for } 1 \leq i \leq k-1. \quad (3.4.8)$$

Indeed, for $2 \leq i \leq k-2$ this is trivial as then $\pi_{k-1}^{i-1} > \pi_{k-1}^i > \pi_{k-1}^{i+1}$. When $i = 1$ Equation 3.4.8 reduces to $\delta\pi_{k-1}^1 > 1/(k-1)\delta\pi_{k-1}^1 + (k-2)/(k-1)\delta\pi_{k-1}^2$ which holds because $k \geq 3$. When $i = k-1$ Equation 3.4.8 reduces to $(k-2)/(k-1)\delta\pi_{k-1}^{k-2} + 1/(k-1)\pi_{k-1}^{k-1} > \delta\pi_{k-1}^{k-1}$ which also holds because $k \geq 3$. Using Equation 3.4.8 we can conclude that $a_0^{i+1} < a_0^i$ for all $i \in \{1, \dots, k-1\}$. Now fix an $i \in \{1, \dots, k-1\}$ then:

$$\begin{aligned} & (1 - \delta F(\tau_k^k)) \pi_k^{i+1} \\ &= \sum_{j=0}^i a_j^{i+1} + \mathbb{P}(\tau_k^{i+1} \geq x \geq \tau_k^k) \delta\pi_{k-1}^{i+1} \\ &= \sum_{j=0}^{i-1} a_j^{i+1} + \frac{1}{k-i} \mathbb{E}[xI(\tau_k^i \geq x \geq \tau_k^{i+1})] + \frac{k-i-1}{k-i} \mathbb{P}(\tau_k^i \geq x \geq \tau_k^{i+1}) \delta\pi_{k-1}^{i+1} + \mathbb{P}(\tau_k^{i+1} \geq x \geq \tau_k^k) \delta\pi_{k-1}^{i+1} \\ &< \sum_{j=0}^{i-1} a_j^i + \frac{1}{k-i} \mathbb{P}(\tau_k^i \geq x \geq \tau_k^{i+1}) \delta\pi_{k-1}^i + \frac{k-i-1}{k-i} \mathbb{P}(\tau_k^i \geq x \geq \tau_k^{i+1}) \delta\pi_{k-1}^{i+1} + \mathbb{P}(\tau_k^{i+1} \geq x \geq \tau_k^k) \delta\pi_{k-1}^{i+1} \\ &\leq \sum_{j=0}^{i-1} a_j^i + \delta \mathbb{P}(\tau_k^i \geq x \geq \tau_k^k) \pi_{k-1}^i \\ &= (1 - \delta F(\tau_k^k)) \pi_k^i \end{aligned}$$

Note that the last inequality is strict precisely when $i \neq k - 1$. We can conclude that:

$$\pi_k^1 > \pi_k^2 > \dots > \pi_k^k.$$

Step 4: $\tau_k^k \leq \delta \pi_k^k(\tau_k^1, \dots, \tau_k^k)$

From the definition of τ_k^k and Step 1 we have that:

$$\tau_k^k = BR_k^k(\tau_k^1, \dots, \tau_k^{k-1}, (\pi_k^{k-1}, \dots, \pi_k^{k-1})) = \arg \max_{\tau \in [0, \tau_k^{k-1}]} \pi_k^k(\tau_k^1, \tau_k^2, \dots, \tau_k^{k-1}, \tau) \quad (3.4.9)$$

From Equation 3.4.7 we have that:

$$\frac{\partial \pi_k^k}{\partial \tau} = \eta(\tau)^2 f(\tau) [-(1 - \delta F(\tau))\tau + \delta(1 - \delta F(\tau))\pi_k^k]$$

Because the function π_k^k is continuous in τ_k^k it attains a maximum in the interval $[0, \tau_k^{k-1}]$. The maximum will not be attained at $\tau = 0$. Indeed the right partial derivative evaluated in 0, $\left. \frac{\partial}{\partial \tau} \pi_k^k \right|_{\tau=0} > 0$, because $\pi_k^k(\tau_k^1, \dots, \tau_k^{k-1}, 0) > 0$. Hence either the maximum will be attained in a internal point of the interval or the maximum will be attained at the cornerpoint, i.e. $\tau = \tau_k^{k-1}$. Note that any fixpoint $\tau = \delta \pi_k^k(\tau_k^1, \dots, \tau_k^{k-1}, \tau)$ is a local maximum. Consequently if such a fixpoint exists, it will be unique and it will be a global maximum. If such a fixpoint does not exist, then the function $\delta \pi_k^k(\tau_k^1, \dots, \tau_k^{k-1}, \tau) > \tau$ as $\delta \pi_k^k(\tau_k^1, \dots, \tau_k^{k-1}, 0) > 0$. Hence we have that

$$\tau_k^k \leq \delta \pi_k^k(\tau_k^1, \dots, \tau_k^{k-1}, \tau_k^k).$$

Step 5: For every player the best-response is given by the threshold $\tau_k^i = \delta \pi_{k-1}^i$

Fix any player $i \leq k - 1$ and assume that all other players $j \neq i$ use the threshold τ_k^j and that the continuation strategies are defined such as in Part 2. We now show that given the continuation strategy of player i , in the $k - 1$ -player subgames, the best-response in the k -player subgame is indeed given by a threshold and this threshold is equal to $\delta \pi_{k-1}^i$. Distinguish the following two cases

Case 1: player $i \leq k - 1$ uses threshold τ_k^i .

For this observe the following:

1. If $x \geq \tau_{k-1}^{i-1}$, then by construction, if the player τ_k^i does not receive the item with value x , he would move to a subgame where his expected equilibrium payoff is either $\delta \pi_{k-1}^i$ or $\delta \pi_{k-1}^{i-1}$. As $x \geq \tau_{k-1}^{i-1} = \delta \pi_{k-1}^{i-1} \geq \delta \pi_{k-1}^i$ it is always a best-response to apply for x .
2. If $\tau_{k-1}^{i-1} > x \geq \tau_{k-1}^i$ then only the players $\tau_k^i, \dots, \tau_k^k$ will apply. By construction of the equilibrium, if the player τ_k^i does not receive the item with value x would then move to a state where his expected equilibrium payoff is given by $\delta \pi_{k-1}^i$. Because $x \geq \tau_{k-1}^i = \delta \pi_{k-1}^i$, it is always a best-response for player τ_k^i to apply for x .
3. If $\tau_k^i > x \geq \tau_k^k$, then player τ_k^i has no reason to apply for an item with this value. Waiting gives a guaranteed payoff of $\delta \pi_{k-1}^i$ and applying for the value does not yield more.
4. If $\tau_k^k > x$ then not applying for the value x gives $\delta \pi_k^i$ as no player will apply for the value and the equilibrium strategies are stationary. From Step 2 and Step 3 we have that $\delta \pi_k^i > \delta \pi_k^k \geq \tau_k^k > x$. Hence not applying for x is a best-response given that no other player applies for x .

$k = 5$	$k = 4$	$k = 3$	$k = 2$	$k = 1$
$\pi_5^1 = 0.59564 \times \delta$ $\tau_5^1 = 0.59670 \swarrow$	$\pi_4^1 = 0.62810 \times \delta$ $\tau_4^1 = 0.63226 \swarrow$	$\pi_3^1 = 0.66554 \times \delta$ $\tau_3^1 = 0.67411 \swarrow$	$\pi_2^1 = 0.70959 \times \delta$ $\tau_2^1 = 0.72395 \swarrow$	$\pi_1^1 = 0.76205 \times \delta$ $\tau_1^1 = 0.72395 \swarrow$
$\pi_5^2 = 0.59525 \times \delta$ $\tau_5^2 = 0.59603 \swarrow$	$\pi_4^2 = 0.62740 \times \delta$ $\tau_4^2 = 0.63084 \swarrow$	$\pi_3^2 = 0.66410 \times \delta$ $\tau_3^2 = 0.67036 \swarrow$	$\pi_2^2 = 0.70564 \times \delta$ $\tau_2^2 = 0.67036 \swarrow$	
$\pi_5^3 = 0.59465 \times \delta$ $\tau_5^3 = 0.59489 \swarrow$	$\pi_4^3 = 0.62618 \times \delta$ $\tau_4^3 = 0.62781 \swarrow$	$\pi_3^3 = 0.66085 \times \delta$ $\tau_3^3 = 0.62781 \swarrow$		
$\pi_5^4 = 0.59362 \times \delta$ $\tau_5^4 = 0.59228 \swarrow$	$\pi_4^4 = 0.62345 \times \delta$ $\tau_4^4 = 0.59228 \swarrow$			
$\pi_5^5 = 0.59126 \times \delta$ $\tau_5^5 = 0.56170 \swarrow$				

Table 3.1: Payoffs and thresholds (rounded) of the monotone equilibrium constructed in Theorem 3.4.7 when $X \sim U[0, 1]$ and $\delta = 0.95$.

We conclude that it is a best-response to apply for an item with value x if and only if $x \geq \tau_k^i$.

Case 2: player k uses threshold τ_k^k .

For this observe the following:

1. If $x \geq \tau_k^{k-1}$, then by construction, at least one other player will apply for the item with value x . Hence if player k does not receive the item with value x he moves to a state where his expected equilibrium payoff is $\delta\pi_{k-1}^{k-1}$. As $x \geq \tau_k^{k-1} = \delta\pi_{k-1}^{k-1}$. It is always a best-response to apply to the item with value x .
2. Now take 2 items with respective values x and y such that $x \leq y < \tau_k^{k-1}$. Then player k is the only player potentially applying for those items, which implies that if he applies for an item he will get it. Therefore, if it is a best-response to apply to an item with a value of x , it must be a best-response to apply for the item with value y .

From those two points it follows that the best-response of player k is indeed given by a threshold. It is now trivial to conclude that $\tau_k^k = BR_k^k(\tau_k^1, \dots, \tau_k^{k-1}, (\pi_{k-1}^{k-1}, \dots, \pi_{k-1}^{k-1}))$.

PART 4: THE EQUILIBRIUM IS MONOTONIC

Consider the player i and the player j with $i > j$. Then we have that player i has a larger expected payoff than player j because $\pi_k^i > \pi_k^j$. It follows from the definition of the continuation strategies that if both players continue to a subgame with $k - 1$ players then $\pi_{k-1,\ell}^i > \pi_{k-1,\ell}^j$ for all $\ell \notin \{i, j\}$. This verifies that the equilibrium is monotonic. $\square$

The recursive computations to determine the thresholds and expected equilibrium payoffs of the monotone equilibrium described above can be easily done by mathematical software. Table 3.1 illustrates the results of the monotone equilibrium when $X \sim U[0, 1]$ and $\delta = 0.95$.

The constructed monotone subgame perfect equilibrium has the property that in the subgame with k players $\pi_k^1 > \pi_k^2 > \dots > \pi_k^k$. This fact is also illustrated by the computations in Table 3.1. It is now a very natural question to ask whether the i -th highest expected equilibrium payoff in a subgame with k players is lower than in the subgame with $k - 1$ players, i.e. is $\pi_k^i < \pi_{k-1}^i$. Table 3.1 suggests that this might be the case. Unfortunately we were only able to prove this for the highest expected payoffs π_k^1 . That is, the highest expected equilibrium payoff is lower in a subgame with more active players.

Theorem 3.4.8. The monotone subgame perfect equilibrium constructed in Theorem 3.4.7 has the property that for all $1 \leq k \leq n$

$$\pi_k^1 < \pi_{k-1}^1.$$

Proof. We prove the theorem using induction on k . From Lemma 3.4.5 it follows that $\pi_1^1 > \pi_2^1$. Furthermore assume that $\pi_{k-2}^1 > \pi_{k-1}^1$. We now show that $\pi_{k-1}^1 > \pi_k^1$. Assume this is not the case and we have that $\pi_{k-1}^1 \leq \pi_k^1$. Then:

$$\begin{aligned} (1 - \delta F(\tau_k^k))\pi_{k-1}^1 &\leq (1 - \delta F(\tau_k^k))\pi_{k-1}^1 \\ &= \frac{1}{k}\mathbb{E}[xI(x \geq \tau_k^1)] + \frac{k-1}{k}\mathbb{P}(x \geq \tau_k^1)\delta\pi_{k-1}^1 + \mathbb{P}(\tau_k^1 > x \geq \tau_k^k)\delta\pi_{k-1}^1 \\ \Rightarrow (1 - \delta F(\tau_{k-1}^{k-1}))\pi_{k-1}^1 &\leq \frac{1}{k}\mathbb{E}[xI(x \geq \tau_k^1)] + \frac{k-1}{k}\delta\mathbb{P}(x \geq \tau_k^1)\pi_{k-1}^1 + \delta\mathbb{P}(\tau_k^1 > x \geq \tau_{k-1}^{k-1})\pi_{k-1}^1 \end{aligned}$$

Now applying the induction hypothesis that $\pi_{k-2}^1 > \pi_{k-1}^1$ and hence $\tau_{k-1}^1 > \tau_k^1$ we find:

$$\begin{aligned} &(1 - \delta F(\tau_{k-1}^{k-1}))\pi_{k-1}^1 \\ &\leq \frac{1}{k}\mathbb{E}[xI(x \geq \tau_k^1)] + \frac{k-1}{k}\mathbb{P}(x \geq \tau_k^1)\delta\pi_{k-1}^1 + \mathbb{P}(\tau_k^1 > x \geq \tau_{k-1}^{k-1})\delta\pi_{k-1}^1 \\ &= \frac{1}{k}\mathbb{E}[xI(x \geq \tau_{k-1}^1)] + \frac{k-1}{k}\mathbb{P}(x \geq \tau_{k-1}^1)\delta\pi_{k-1}^1 \\ &\quad + \frac{1}{k}\mathbb{E}[xI(\tau_{k-1}^1 > x \geq \tau_k^1)] + \frac{k-1}{k}\mathbb{P}(\tau_{k-1}^1 > x \geq \tau_k^1)\delta\pi_{k-1}^1 + \mathbb{P}(\tau_k^1 > x \geq \tau_{k-1}^{k-1})\delta\pi_{k-1}^1 \\ &< \frac{1}{k-1}\mathbb{E}[xI(x \geq \tau_{k-1}^1)] + \frac{k-2}{k-1}\mathbb{P}(x \geq \tau_{k-1}^1)\delta\pi_{k-2}^1 \\ &\quad + \frac{1}{k}\mathbb{E}[xI(\tau_{k-1}^1 > x \geq \tau_k^1)] + \frac{k-1}{k}\mathbb{P}(\tau_{k-1}^1 > x \geq \tau_k^1)\delta\pi_{k-1}^1 + \mathbb{P}(\tau_k^1 > x \geq \tau_{k-1}^{k-1})\delta\pi_{k-1}^1 \\ &< \frac{1}{k-1}\mathbb{E}[xI(x \geq \tau_{k-1}^1)] + \frac{k-2}{k-1}\mathbb{P}(x \geq \tau_{k-1}^1)\delta\pi_{k-2}^1 \\ &\quad + \frac{1}{k}\mathbb{P}(\tau_{k-1}^1 > x \geq \tau_k^1)\delta\pi_{k-2}^1 + \frac{k-1}{k}\mathbb{P}(\tau_{k-1}^1 > x \geq \tau_k^1)\delta\pi_{k-2}^1 + \mathbb{P}(\tau_k^1 > x \geq \tau_{k-1}^{k-1})\delta\pi_{k-2}^1 \\ &< \frac{1}{k-1}\mathbb{E}[xI(x \geq \tau_{k-1}^1)] + \frac{k-2}{k-1}\mathbb{P}(x \geq \tau_{k-1}^1)\delta\pi_{k-2}^1 + \mathbb{P}(\tau_{k-1}^1 > x \geq \tau_{k-1}^{k-1})\delta\pi_{k-2}^1 \\ &= (1 - \delta F(\tau_{k-1}^{k-1}))\pi_{k-1}^1 \end{aligned}$$

A contradiction. So we have that $\pi_k^1 < \pi_{k-1}^1$. □

3.4.3 Examples

The construction of the monotone subgame perfect equilibrium raises the natural question whether all subgame perfect equilibria are monotone. In the two player subgame the equilibrium was unique, up to a permutation of players. However as soon as there are more players there may exist multiple subgame perfect equilibria. The purpose of the following example is to illustrate this.

Example 3.4.9. [Non-uniqueness of subgame perfect equilibria]

Consider the 3 player game with $X \sim U[0, 1]$ and $\delta = 0.95$. Applying Theorem 3.4.5 yields

Figure 3.4: Expected continuation plan: a non-monotone subgame perfect equilibrium

that in the two player subgame the remaining players will receive an expected payoff of $\pi_2^1 = 0.70916$ and $\pi_2^2 = 0.70567$. Let σ_2^1 and σ_2^2 denote the continuation strategies yielding these respective expected payoffs. Consider the following strategy profile, described in terms of thresholds and continuation plans.

$$\begin{aligned}\tau_3^1 &= \delta\pi_2^1 \text{ and } \sigma_{3,-}^1 = (\sigma_2^1, \sigma_2^1) \\ \tau_3^2 &= \delta\pi_2^2 \text{ and } \sigma_{3,-}^2 = (\sigma_2^2, \sigma_2^2) \\ \tau_3^3 &= \delta \cdot 0.661753 \text{ and } \sigma_{3,-}^3 = (\sigma_2^1, \sigma_2^2).\end{aligned}$$

Note that $\tau_3^1 > \tau_3^2 > \tau_3^3$. The continuation plans are illustrated in Figure 3.4. The strategy profile described above yields the following expected payoffs: $\pi_3^1 = 0.665181$, $\pi_3^2 = 0.663006$ and $\pi_3^3 = 0.661753 = \tau_3^3/\delta$. We now verify that this is indeed a subgame perfect equilibrium.

- Let $x \geq \delta\pi_2^1$. Then it is a best-response for all players to apply, given that the others apply. Indeed any continuation payoff they can receive, can never exceed x .
- Let $\delta\pi_2^1 > x \geq \delta\pi_2^2$. It is a best-response for player τ_3^1 to reject because if the others apply he has a guaranteed continuation payoff of $\delta\pi_2^1$. It is a best-response for player τ_2^2 to apply as his expected payoff in this case is $1/2x + 1/2\delta\pi_2^2$ which exceeds the expected payoff of rejecting $\delta\pi_2^2$. A similar argument holds for player π_3^3 .
- Let $\delta\pi_2^2 > x \geq \delta\pi_3^3$. Then it is a best-response for player τ_3^3 to apply and receive x , as rejecting would only yield an expected payoff of $\delta\pi_3^3$. Given this, it is clear that it is a best-response for the other players to reject.
- $x < \delta\pi_3^3$ Because $\pi_3^1 > \pi_3^2 > \pi_3^3 > x$ it is a best-response for any player to reject x given that the others reject x as well.

Note that this subgame perfect equilibrium is not monotone. Indeed the expected payoff of the player τ_3^2 is higher than the expected payoff of player τ_3^3 . However in the subgame where player τ_3^1 leaves, the expected payoff of player τ_3^2 is lower than that of player τ_3^3 .

Furthermore, observe the possibly surprising fact that even though the expected continuation payoff of the player using τ_3^2 in the two player subgame is always the lowest, it is still possible that he obtains a higher expected payoff and is more picky than a player with a more favourable continuation payoff in the three player subgame.

Subgame perfect equilibria can be easily constructed using backwards induction. Given that in all subgames with $k-1$ active players the equilibrium strategies are given by $\sigma_{k-1}^1, \dots, \sigma_{k-1}^{k-1}$ then for the construction of the k player subgame perfect equilibrium we need to find thresholds $\tau_k^1, \dots, \tau_k^k$ and for each threshold a continuation plan, that is, a list of continuation

strategies this player will follow for each of the leaving players. Suppose that for each leaving player, each of the remaining players plays one of the equilibrium strategies of the smaller subgame, such that in each of the smaller subgames the strategies player are a subgame perfect equilibrium. Can we always find thresholds for all the players such that these thresholds combined with the continuation plans constitute a subgame perfect equilibrium in the larger subgame? The next example illustrates that this is not always possible.

Figure 3.5: Not all continuation plans can be supported in equilibrium.

Example 3.4.10. [Not all continuation plans can be supported in a subgame perfect equilibrium with threshold strategies.]

Let π_2^1 and π_2^2 denote the unique expected equilibrium payoffs of the two player subgame as derived in Theorem 3.4.5 and let σ_2^1 and σ_2^2 denote the corresponding equilibrium strategies. Consider the 3 player subgame with $\tau_3^1 \geq \tau_3^2 \geq \tau_3^3$. Let the continuation plans be given by

$$\begin{aligned}\sigma_{3,-}^1 &= (\sigma_2^2, \sigma_2^1) \\ \sigma_{3,-}^2 &= (\sigma_2^1, \sigma_2^2) \\ \sigma_{3,-}^3 &= (\sigma_2^2, \sigma_2^1)\end{aligned}$$

These continuation plans are illustrated in Figure 3.5 and have the property that every player will in one of the two player subgames use the strategy σ_2^1 and in the other two player subgame use the strategy σ_2^2 . We now show that in the three player subgame there do not exist thresholds τ_3^1, τ_3^2 and τ_3^3 that together with these continuation plans constitute a subgame perfect equilibrium in threshold strategies. We can distinguish four cases:

Case 1: $\tau_3^1 = \tau_3^2 = \tau_3^3$

Assume all players use the same threshold denoted by τ . Because of the definition of the continuation plan any player who will continue to the two player subgame is expected to receive $\delta\pi_2^a = \delta\frac{1}{2}(\pi_2^1 + \pi_2^2)$. Therefore, the expected payoff to all three players using a common threshold τ is given by $\pi_3(\tau) = \eta(\tau) \left(\frac{1}{3}\mathbb{E}[xI(x \geq \tau)] + \frac{2}{3}\mathbb{P}(x \geq \tau)\delta\pi_2^a \right)$. To be an equilibrium we need to have that for all $x \geq \tau$, it is a best-response to apply for x given that all other players apply for x . From this it follows that $\tau \geq \delta\pi_2^a$. Moreover, for all $x < \tau$ it is a best-response to reject given that all other players reject. Hence $\delta\pi_3(\tau) \geq \tau$. Then we have that

$$\delta\pi_3(\tau) \geq \tau \geq \delta\pi_2^a > \delta\pi_2^2.$$

Furthermore observe that:

$$\begin{aligned}\pi_3(\tau) &= \eta(\tau) \left(\frac{1}{3}\mathbb{E}[xI(x \geq \tau)] + \frac{2}{3}\mathbb{P}(x \geq \tau)\delta\pi_2^a \right) \\ &\leq \eta(\tau) \left(\frac{1}{3}\mathbb{E}[xI(x \geq \tau)] + \frac{2}{3}\mathbb{P}(x \geq \tau)\tau \right) \leq \eta(\tau)\mathbb{E}[xI(x \geq \tau)] \leq \pi_1^1.\end{aligned}$$

From Theorem 3.4.5 we have that $\tau_2^1 = \delta\pi_1^1$. Hence $\tau_2^1 \geq \delta\pi_3(\tau) \geq \tau \geq \delta\pi_2^a$. Using this we find that

$$\begin{aligned}
\pi_3(\tau) &= \eta(\tau) \left(\frac{1}{3} \mathbb{E}[xI(x \geq \tau)] + \frac{2}{3} \mathbb{P}(x \geq \tau) \delta\pi_2^a \right) \\
&\leq \eta(\tau) \left(\frac{1}{3} \mathbb{E}[xI(x \geq \tau)] + \frac{2}{3} \mathbb{P}(x \geq \tau) \tau \right) \\
&\leq \eta(\tau) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau)] + \frac{1}{2} \mathbb{P}(x \geq \tau) \tau \right) \\
&= \eta(\tau) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{E}[xI(\tau_2^1 > x \geq \tau)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1) \tau + \frac{1}{2} \mathbb{P}(\tau_2^1 \geq x \geq \tau) \tau \right) \\
&\leq \eta(\tau) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1) \delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau)] \right) \leq \pi_2^2.
\end{aligned}$$

A contradiction with the fact that $\pi_3(\tau) > \pi_2^2$. Therefore, we can conclude that there is no equilibrium in which all players use the same threshold.

Case 2: $\tau_3^1 = \tau_3^2 > \tau_3^3$.

Then, for all $\tau_3^3 \leq x < \tau_3^2$ it is a best-response for the player τ_3^3 to pick, while the other players will reject and move to the next stage to receive an expected payoff of $\delta\pi_2^1$ and $\delta\pi_2^2$ respectively. Hence for all $\tau_3^3 \leq x < \tau_3^2$ we have that $x \leq \delta\pi_2^1$ and $x \leq \delta\pi_2^2$ from which it follows that $\tau_3^h \leq \delta\pi_2^2$. Therefore for $x = \delta\pi_2^2$ all players apply. However, given that the other players apply for $\delta\pi_2^2$ it is a best-response for any player to wait and receive an expected payoff of $\delta\frac{1}{2}(\pi_2^1 + \pi_2^2)$.

Case 3: $\tau_3^1 > \tau_3^2 = \tau_3^3$.

Notice that for any $i \neq j$ we have that $\pi_{3,j}^i \neq \pi_{3,i}^j$. Now consider an item with the value x such that $\tau_3^1 > x > \tau_3^2$ then it is a best-response for the players using τ_3^2 and τ_3^3 to apply for x . Then because of the precise definition of the continuation plan we need to have that in equilibrium $x \geq \delta\pi_2^1$ and $x \geq \delta\pi_2^2$. Furthermore, for player τ_3^1 it is a best-response to wait and receive $\delta\frac{1}{2}(\pi_2^1 + \pi_2^2)$. Hence $x \leq \delta\frac{1}{2}(\pi_2^1 + \pi_2^2)$. A contradiction with the fact that $x \geq \delta\pi_2^1$.

Case 4: $\tau_3^1 > \tau_3^2 > \tau_3^3$.

Let $\tau_3^1 > \tau_3^2 > \tau_3^3$. Take x such that $\tau_3^1 > x > \tau_3^2$. Then the best-response of the player using τ_3^1 is to reject, while the other players both apply, hence $x \leq \delta\frac{1}{2}(\pi_2^1 + \pi_2^2)$. Among the two applying players there always exists a player who, given that the other applies and he waits, he will transition to a state with an expected payoff of $\delta\pi_1^1$. Therefore it is only a best-response to apply for x if $x \geq \delta\pi_1^1$. Hence, in equilibrium we need have that $\delta\pi_1^1 \leq \delta\frac{1}{2}(\pi_2^1 + \pi_2^2)$. A contradiction.

3.5 Welfare Analysis

A policy which gives all applicants equal chances may be desirable because of fairness reasons. However, when the number of active players is expected to decrease, this will come at a cost. The goal of this section is to analyze the welfare consequences of the equal chances policy. In this section we show that any pure subgame perfect equilibrium is associated with a strictly positive welfare loss. Consider a strategy profile σ and a subgame J with $|J| = k$ such that the vector of equilibrium payoffs is given by $\pi_J^1, \dots, \pi_J^k$. Then the utilitar-

ian welfare in the subgame J associated with the strategy profile σ is given by

$$W_J^\sigma = \sum_{i=1}^k \pi_J^i.$$

3.5.1 The social optimum

To find the social optimum in the subgame J with k active players, we need to analyze the following decision problem. Each period an item with a value distributed according to a cumulative distribution F appears. Then the decision maker, hereafter referred to as the social planner, can either decide to pick this item or reject it. If the social planner picks this item, the item is given at random to any of the active players and the game moves to a subgame with only $k - 1$ active players. If the social planner decides to reject the item the game repeats itself. This process goes on until there are no more active players. If in period t_i the social planner picked the item with value x_{t_i} then the utilitarian welfare W_k^* is given by $\sum_{i=1}^k \delta^{t_i} x_{t_i}$. We can assume that in every subgame the social planner uses a threshold strategy. Note that because all players have identical payoff functions, the optimal threshold used by a social planner in a certain subgame J will only depend on the number of active players and not on their identity.

Let $|J| = k$ and let τ_k denote the threshold the social planner uses in a subgame with precisely k active players. Using the stationarity assumption we find that the expected welfare W_k^* of the social planner in such a subgame is given by the following recursive equation:

$$W_k^*(\tau_k) = \mathbb{E}[xI(x \geq \tau_k)] + \delta \mathbb{P}(x \geq \tau_k) W_{k-1}^*(\tau_{k-1}) + \mathbb{P}(x < \tau_k) W_k^*(\tau_k).$$

Hence:

$$W_k^*(\tau_k) = \eta(\tau_k) \left(\mathbb{E}[xI(x \geq \tau_k)] + \delta \mathbb{P}(x \geq \tau_k) W_{k-1}^*(\tau_{k-1}) \right).$$

The optimal thresholds τ_k^* can be computed recursively.

It is clear that $\tau_1^* = \eta(\tau_1^*) \delta \mathbb{E}[xI(x \geq \tau_1^*)] = \tau_1^1$. Now assume the values $\tau_1^*, \dots, \tau_{k-1}^*$ are all determined. Then $\tau_k^* \in [0, \bar{v}]$ should maximise W_k^* . This maximum exists as the function $W_k^* : [0, \bar{v}] \rightarrow \mathbb{R}$ is continuous over a compact interval. Note that the maximum will be obtained at an internal point. Indeed, $W_k^*(\bar{v}) = 0$ for all k , which implies that $k-1$ allocations are expected to only take a finite waiting time. To see that it is never optimal for the social planner to pick an item with value 0, note that picking an item with value 0 and moving to a smaller subgame is worse than first optimally playing the smaller subgame and then when there is only one player left picking any item with a strictly positive value. We have that:

$$\frac{dW_k^*}{d\tau_k} = \eta(\tau_k) f(\tau_k) \left(-\tau_k - \delta W_{k-1}^*(\tau_{k-1}^*) + \delta W_k^*(\tau_k) \right).$$

Therefore the optimal threshold τ_k^* needs to satisfy

$$\tau_k^* = \delta \left(W_k^*(\tau_k^*) - W_{k-1}^*(\tau_{k-1}^*) \right). \quad (3.5.1)$$

3.5.2 Welfare with a priority ranking

In this subsection we consider an alternative policy where from the start of the game players are given a priority ordering. When an item has multiple applicants, the item will be

allocated to the applicant with the highest priority. Even though the absolute ranking of a player may decrease during the game as some players will leave the game, the relative ranking of players compared to each other does not change. This setup is a special case of the dynamic assignment model studied by Su and Zenios (2004), where new players enter the game according to a Poisson process and unmatched players may leave the game according to a different Poisson process. Their results imply that in the special case where there is no entry, the threshold used by the player with the lowest priority equals the socially optimal threshold, a fact we will use in the welfare analysis when all applicants are given equal chances. For the ease of the reader we give a short analysis of the case when players are ranked. The threshold a player uses in a certain subgame will only depend on his priority ranking in that subgame. The thresholds used by the k -th ranked player can be determined by inductively solving optimal stopping problems.

The first-ranked player faces a decision problem. The presence of lower-ranked players is irrelevant to his optimal strategy. Indeed, if he wants an item with a specific value he can guarantee that he gets it. The optimal threshold τ_1^r used by the first ranked player is given by the fixpoint $\tau_1^r = \eta(\tau_1^r)\delta\mathbb{E}[xI(x \geq \tau_1)] = \tau_1^* = \tau_1^1$. We can now inductively define the threshold used by the k -th ranked player. Assume that $\tau_1^r \geq \tau_2^r \geq \dots \geq \tau_{k-1}^r$ are the thresholds used by the first ranked, second ranked, etc. and let $\pi_1^r, \dots, \pi_{k-1}^r$ denote the corresponding equilibrium payoffs. Then we have that the k -th ranked player can either use a threshold strictly below τ_{k-1}^r or equal to τ_{k-1}^r . Indeed, a threshold strictly higher than the one used by τ_{k-1}^r serves no purpose and is equivalent to using τ_{k-1}^r as the player with the higher priority ranking will get the item. The expected payoff of the k -th ranked player given that he uses a threshold $\tau_k \leq \tau_{k-1}^r$ is given by

$$\pi_k^r(\tau_k) = \eta(\tau_k^r) \left(\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k)] + (1 - F(\tau_{k-1}^r))\delta\pi_{k-1}^r \right).$$

The following lemma states that it is always optimal for the k -th ranked player to use a threshold strictly lower than the optimal threshold of the $(k-1)$ -st ranked player.

Lemma 3.5.1. When all players are given a priority ranking, we have that $\tau_{k-1}^r > \tau_k^r$ and the thresholds τ_k^r used by the k -th ranked player is the unique fixpoint of $\tau_k^r = \delta\pi_k^r(\tau_k^r) > 0$.

Proof. The k -th ranked player solves the following constraint optimization problem:

$$\max_{\tau_k \in [0, \tau_{k-1}^r]} \eta(\tau_k) \left(\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k)] + (1 - F(\tau_{k-1}^r))\delta\pi_{k-1}^r \right)$$

We have that $\tau_1^r > 0$ and $\tau_1^r = \delta\pi_1^r$ and we assume by induction that $\tau_{k-1}^r > 0$ and $\tau_{k-1}^r = \delta\pi_{k-1}^r$. The Lagrangian is given by $\mathcal{L}(\tau_k, \mu) = \eta(\tau_k) \left(\mathbb{E}_X[xI(\tau_{k-1}^r \geq x \geq \tau_k)] + \delta(1 - F(\tau_{k-1}^r))\pi_{k-1}^r \right) - \mu_1(\tau_k - \tau_{k-1}^r) + \mu_2\tau_k$.

According to the KKT conditions we have that $\mu_2, \mu_1 \geq 0$, $\mu_1(\tau_k - \tau_{k-1}^r) = 0$, $\mu_2\tau_k = 0$ and

$$\frac{\partial \mathcal{L}}{\partial \tau_k} = 0 \Rightarrow \mu_1 - \mu_2 = \eta(\tau_k)f(\tau_k)(-\tau_k + \delta\pi_k^r) \quad (3.5.2)$$

First, assume $\mu_1 > 0$ then because $\tau_{k-1}^r > 0$ we have $\mu_2 = 0$. Furthermore $\mu_1 > 0$ implies that $\tau_k = \tau_{k-1}^r$. Then from Equation 3.5.2 it follows that

$$\tau_{k-1}^r = \tau_k < \delta\pi_k^r(\tau_k) = \delta\pi_k^r(\tau_{k-1}^r) = \eta(\tau_{k-1}^r)(1 - F(\tau_{k-1}^r))\delta\pi_{k-1}^r < \delta\pi_{k-1}^r. \quad (3.5.3)$$

A contradiction with the induction hypothesis $\tau_{k-1}^r = \delta\pi_{k-1}^r$. If $\mu_2 > 0$ then $\tau_k = 0$ and Equation 3.5.2 would imply that $\pi_k^r(0) < 0$, a contradiction because $\pi_k^r(0) > 0$. We can conclude that $0 < \tau_k^r < \tau_{k-1}^r$. Therefore Equation 3.5.2 implies that $\tau_k^r = \delta\pi_k^r(\tau_k^r)$. The uniqueness of this fixpoint follows using a similar argument as in Lemma 3.4.4. Indeed the function π_k^r is strictly increasing if $\tau_k < \pi_k^r(\tau_k^r)$ and strictly decreasing if $\tau_k > \pi_k^r(\tau_k^r)$. Therefore any fixpoint $\tau_k = \delta\pi_k^r(\tau_k)$ is a local maximum.

If there are 2 distinct local maxima attained in τ_k^{+1} and τ_k^{+2} , then because of the continuity of π_k^r there must exist a threshold τ_k^- which lies in between the thresholds τ_k^{+1} and τ_k^{+2} and in which π_k^r attains a local minimum. However then τ_k^- must be a fixpoint of the $\tau_k = \delta\pi_k^r(\tau_k)$. A contradiction with the fact that all fixpoints of this equation are local maxima and the function π_k^r obtains a global maximum in an internal point of $[0, \tau_{k-1}^r]$. $\square$

The following theorem is a special case of a result in Su and Zenios (2004, proposition 3). For completeness we provide a proof for this special case.

Theorem 3.5.2. Utilitarian welfare is maximised when players are given a priority ranking. Moreover, the optimal threshold τ_k^* used by the social planner when there are k active players is equal to the threshold used by the k -th ranked player, i.e. $\tau_k^* = \tau_k^r$.

Proof. The proof is by induction on the number of active players. Let $W_k^r = \sum_{i=1}^k \pi_i^r$ denote the welfare when all players are given a priority ranking. Note that for $k = 1$ the statement follows trivially from the fact that both the social planner and the first ranked player face the same decision problem. Now assume that $\tau_{k-1}^* = \tau_{k-1}^r$ and $W_{k-1}^* = W_{k-1}^r$. Then we have that:

$$\begin{aligned}\tau_k^* &= \delta[W_k^*(\tau_k^*) - W_{k-1}^*] \\ \tau_k^r &= \delta\eta(\tau_k^r) [\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^r)] + \mathbb{P}(x \geq \tau_{k-1}^r)\delta\pi_{k-1}^r]\end{aligned}$$

Using the induction hypothesis $\tau_{k-1}^* = \tau_{k-1}^r$ it follows that $W_l^* = W_l^r$ for all $l \leq k-1$ and

$$\begin{aligned}(1 - \delta F(\tau_{k-1}^r)) \frac{\tau_{k-1}^r}{\delta} &= ((1 - \delta F(\tau_{k-1}^r))(W_{k-1}^*(\tau_{k-1}^r) - W_{k-2}^*)) \\ &= \mathbb{E}[xI(x \geq \tau_{k-1}^r)] + \delta\mathbb{P}(x \geq \tau_{k-1}^r)W_{k-2}^* - W_{k-2}^* + \delta F(\tau_{k-1}^r)W_{k-2}^* \\ &= \mathbb{E}[xI(x \geq \tau_{k-1}^r)] + (\delta - 1)W_{k-2}^r.\end{aligned}$$

We have that:

$$\begin{aligned}(1 - \delta F(\tau_k^*))\tau_k^* &= (1 - \delta F(\tau_k^*))\delta[W_k^*(\tau_k^*) - W_{k-1}^*] \\ &= \delta[\mathbb{E}[xI(x \geq \tau_k^*)] + \delta\mathbb{P}(x \geq \tau_k^*)W_{k-1}^* - (1 - \delta F(\tau_k^*))W_{k-1}^*] \\ &= \delta[\mathbb{E}[xI(x \geq \tau_k^*)] + (\delta - 1)W_{k-1}^*] \\ &= \delta[\mathbb{E}[xI(\tau_{k-1}^* \geq x \geq \tau_k^*)] + \mathbb{E}[xI(x \geq \tau_{k-1}^*)] + (\delta - 1)W_{k-1}^*] \\ &= \delta[\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^*)] + \mathbb{E}[xI(x \geq \tau_{k-1}^r)] + (\delta - 1)W_{k-1}^r] \\ &= \delta[\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^*)] + (\delta - 1)\pi_{k-1}^r + \mathbb{E}[xI(x \geq \tau_{k-1}^r)] + (\delta - 1)W_{k-2}^r] \\ &= \delta[\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^*)] + (\delta - 1)\pi_{k-1}^r + (1 - \delta F(\tau_{k-1}^r))\frac{\tau_{k-1}^r}{\delta}] \\ &= \delta[\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^*)] + \delta\pi_{k-1}^r - \delta F(\tau_{k-1}^r)\pi_{k-1}^r] \\ &= \delta[\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^*)] + \mathbb{P}(x \geq \tau_{k-1}^r)\delta\pi_{k-1}^r].\end{aligned}$$

From this we have that:

$$\tau_k^* = \delta \frac{\mathbb{E}[xI(\tau_{k-1}^r \geq x \geq \tau_k^*)] + \delta \mathbb{P}(x \geq \tau_{k-1}^r) \pi_{k-1}^r}{(1 - \delta F(\tau_k^*))} = \delta \pi_k^r(\tau_k^*).$$

Because the fixpoint of this equation is unique, we can conclude that $\tau_k^* = \tau_k^r$. Moreover, because $\tau_k^r = \delta \pi_k^r$ and $\tau_k^* = \delta(W_k^* - W_{k-1}^*)$, we have that:

$$W_k^* = \frac{\tau_k^*}{\delta} + W_{k-1}^* = \frac{\tau_k^r}{\delta} + W_{k-1}^r = \pi_k^r + W_{k-1}^r = W_k^r.$$

□

3.5.3 Welfare under the equal chances policy

We now show that any pure subgame perfect equilibrium in threshold strategies that can occur under the equal chances policy induces a strictly positive welfare loss. The proof is based on induction on the number of active players. We first show that in the two-player subgame there is a positive welfare loss. This welfare loss is caused by the fact that the players wait inefficiently long. This is because under the equal chances policy, the least selective player still has the opportunity to win high value items, while in the case of an exogenous priority ranking, the lowest ranked player has less incentive to wait longer as he knows he has no chance of receiving the high value items. After we establish the welfare loss in the two-player subgames, we can easily show that this welfare loss carries through to subgames with a higher number of active players.

Theorem 3.5.3. Any pure subgame perfect equilibrium in threshold strategies σ induces a strictly positive welfare loss in all subgames J with $|J| = k > 1$. Hence

$$W_J^\sigma < W_J^*.$$

Proof. We prove this theorem by induction on the number of active players. Fix a subgame perfect equilibrium σ and let W_J^σ denote the expected utilitarian welfare in a subgame with root J .

STEP 1: $W_J^\sigma < W_2^*$ where $|J| = 2$

Note that the subgame perfect equilibrium in threshold strategies for a subgame with two players is uniquely defined by Lemma 3.4.5. Let the expected equilibrium payoffs in this two-player game be denoted with π_2^1 and π_2^2 and let the optimal expected payoff in the one-player game be denoted by π_1^1 . Then we have that $\tau_2^1 = \delta \pi_1^1$, $\tau_2^2 = \delta \pi_2^2$ and

$$\begin{aligned} W_J^\sigma &= \pi_2^1 + \pi_2^2 \\ &= \eta(\tau_2^2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1) \delta \pi_1^1 + \mathbb{P}(\tau_2^1 > x \geq \tau_2^2) \delta \pi_1^1 \right) \\ &\quad + \eta(\tau_2^2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1) \delta \pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)] \right) \\ &= \eta(\tau_2^2) (\mathbb{E}[xI(x \geq \tau_2^1)] + \mathbb{P}(x \geq \tau_2^1) \delta \pi_1^1 + \mathbb{P}(\tau_2^1 > x \geq \tau_2^2) \delta \pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)]) \\ &= \eta(\tau_2^2) (\mathbb{E}[xI(x \geq \tau_2^2)] + \mathbb{P}(x \geq \tau_2^2) \delta \pi_1^1). \end{aligned}$$

Furthermore, we have that

$$W_2^* = \eta(\tau_2^2) (\mathbb{E}[xI(x \geq \tau_2^*)] + \mathbb{P}(x \geq \tau_2^*)\delta\pi_1^1).$$

We now need to compare τ_2^2 to τ_2^* . From Theorem 3.5.2 we have that the socially optimal threshold τ_2^* used by the social planner equals the threshold τ_2^r used by the second ranked player. Using Lemma 3.5.1 we have that $\tau_2^* = \tau_2^r = \delta\pi_2^r$, where π_2^r denotes the maximal expected payoff the second ranked player obtains using this threshold. Furthermore, we have that $\tau_1^r = \delta\pi_1^1 = \tau_2^1$. We now show that $\tau_2^2 > \tau_2^r$. Suppose the contrary, then $\tau_2^2 \leq \tau_2^r$ and

$$\begin{aligned} (1 - \delta F(\tau_2^2))\tau_2^2 &= \delta \left(\frac{1}{2} \mathbb{E}[xI(x \geq \tau_2^1)] + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)] \right) \\ &> \delta \left(\frac{1}{2} \mathbb{P}(x \geq \tau_2^1)\tau_2^1 + \frac{1}{2} \mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)] \right) \\ &= \delta (\mathbb{P}(x \geq \tau_2^1)\delta\pi_1^1 + \mathbb{E}[xI(\tau_2^1 > x \geq \tau_2^2)]) \\ &= \delta (\mathbb{P}(x \geq \tau_1^r)\delta\pi_1^1 + \mathbb{E}[xI(\tau_1^r > x \geq \tau_2^r)] + \mathbb{E}[xI(\tau_2^r \geq x \geq \tau_2^2)]) \\ &\geq \delta (\mathbb{P}(x \geq \tau_1^r)\delta\pi_1^1 + \mathbb{E}[xI(\tau_1^r > x \geq \tau_2^r)] + \mathbb{P}(\tau_2^r \geq x \geq \tau_2^2)\tau_2^2) \\ &= \delta (\mathbb{P}(x \geq \tau_1^r)\delta\pi_1^1 + \mathbb{E}[xI(\tau_1^r > x \geq \tau_2^r)]) + \delta F(\tau_2^r)\tau_2^2 - \delta F(\tau_2^2)\tau_2^2, \end{aligned}$$

where the strict inequality follows from the fact that $\tau_2^1 < \bar{v}$. Then we have

$$\tau_2^2 > \delta \frac{\mathbb{P}(x \geq \tau_1^r)\pi_1^1 + \mathbb{E}[xI(\tau_1^r > x \geq \tau_2^r)]}{1 - \delta F(\tau_2^r)} = \tau_2^r.$$

A contradiction with the assumption that $\tau_2^2 \leq \tau_2^r$. From Theorem 3.5.2 we have that the optimal threshold used by the social planner is such that $\tau_2^r = \tau_2^*$. Hence we can conclude that

$$\tau_2^2 > \tau_2^r = \tau_2^*.$$

Now consider the function $w : [0, \bar{v}] \rightarrow \mathbb{R}$ defined by $w(\tau) = \frac{\mathbb{E}[xI(x \geq \tau)] + \mathbb{P}(x \geq \tau)\delta\pi_1^1}{1 - \delta F(\tau)}$. Then we have that $W_2^* = w(\tau_2^*)$ and $W_J^\sigma = w(\tau_2^2)$. We have the function w attains it a maximum at τ_2^* and for any $\tau > \tau_2^*$ the function is strictly decreasing. From this we can conclude that if $|J| = 2$ then

$$W_J^\sigma < W_2^*. \quad (3.5.4)$$

STEP 2: $W_J^\sigma = \eta(\tau_J^k) \left(\mathbb{E}[xI(x \geq \tau_J^k)] + \sum_{j=0}^{k-1} \mathbb{P}(\tau_J^j \geq x \geq \tau_J^{j+1}) \frac{1}{k-j} \sum_{l=j+1}^k \delta W_{J,l}^\sigma \right)$ where $|J| = k$.

Take a subgame perfect equilibrium in threshold strategies σ and a subgame J such that $|J| = k$ and let $\tau_J^1 \geq \dots \geq \tau_J^k$ denote the thresholds used. We have that the sum of all expected equilibrium payoff in any subgame equals the expected welfare, i.e.

$$\sum_{i=1, i \neq l}^k \pi_{J,-l}^i = W_{J,-l}^\sigma.$$

We have that:

$$W_J^\sigma = \sum_{i=1}^k \pi_J^i$$

$$\begin{aligned}
&= \sum_{i=1}^k \sum_{j=0}^{i-1} \eta(\tau_j^k) \left[\frac{1}{k-j} \mathbb{E}[xI(\tau_j^j \geq x \geq \tau_j^{j+1})] \right. \\
&\quad + \frac{k-j-1}{k-j} \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \left(\frac{1}{k-j-1} \sum_{l=j+1, l \neq i}^k \delta \pi_{j,-l}^i \right) \\
&\quad \left. + \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \left(\frac{1}{k-j} \sum_{l=j+1}^k \delta \pi_{j,-l}^i \right) \right]
\end{aligned}$$

Notice that the term $\frac{1}{k-j} \mathbb{E}[xI(\tau_j^j \geq x \geq \tau_j^{j+1})]$ appears in the payoff of every player $\tau_j^{j+1}, \dots, \tau_j^k$ as they use a threshold at most τ_j^{j+1} , hence this term appears precisely $k-j$ times in the sum $\sum_{i=1}^k \pi_j^i$. Hence we have that

$$\sum_{i=1}^k \sum_{j=0}^{i-1} \frac{1}{k-j} \mathbb{E}[xI(\tau_j^j \geq x \geq \tau_j^{j+1})] = \sum_{j=0}^{k-1} \mathbb{E}[xI(\tau_j^j \geq x \geq \tau_j^{j+1})] = \mathbb{E}[xI(x \geq \tau_j^k)].$$

Now, fix a factor $\mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1})$, with $j < k-1$. Then for every player τ_j^i there is precisely one term in the expression of π_j^i involving the factor $\mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1})$. We can distinguish 2 cases:

Case 1: $i \leq j$

Then player i will not pick the item with value x . Therefore he will transition to the subgame with $k-1$ players and receive the expected payoff $\frac{1}{k-j} \sum_{l=j+1}^k \pi_{j,-l}^i$. So the term belonging to the factor $\mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1})$ equals

$$\mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \left(\frac{1}{k-j} \sum_{l=j+1}^k \delta \pi_{j,-l}^i \right).$$

Case 2: $i > j$

Then player i will apply for the item and with probability $\frac{k-j-1}{k-j}$ will transition to a subgame with $k-1$ players, yielding an expected payoff of $\frac{1}{k-j-1} \sum_{l=j+1, l \neq i}^k \delta \pi_{j,-l}^i$. So the term belonging to the factor $\mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1})$ equals

$$\begin{aligned}
&\frac{k-j-1}{k-j} \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \left(\frac{1}{k-j-1} \sum_{l=j+1, l \neq i}^k \delta \pi_{j,-l}^i \right) \\
&= \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \left(\frac{1}{k-j} \sum_{l=j+1, l \neq i}^k \delta \pi_{j,-l}^i \right).
\end{aligned}$$

Taking the sum over all i yields

$$\begin{aligned}
\sum_{i=1}^k \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \left(\frac{1}{k-j} \sum_{l=j+1, l \neq i}^k \delta \pi_{j,-l}^i \right) &= \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \frac{1}{k-j} \sum_{i=1}^k \sum_{l=j+1, l \neq i}^k \delta \pi_{j,-l}^i \\
&= \mathbb{P}(\tau_j^j \geq x \geq \tau_j^{j+1}) \frac{1}{k-j} \sum_{l=j+1}^k \delta W_{j,-l}^\sigma.
\end{aligned}$$

We find that:

$$W_J^\sigma = \eta(\tau_J^k) \left(\mathbb{E}[xI(x \geq \tau_J^k)] + \sum_{j=0}^{k-1} \mathbb{P}(\tau_J^j \geq x \geq \tau_J^{j+1}) \frac{1}{k-j} \sum_{l=j+1}^k \delta W_{J,-l}^\sigma \right).$$

STEP 3: $W_J^\sigma < W_k^*$

Using the fact that $\frac{1}{k-j} \sum_{l=j+1}^k W_{J,-l}^\sigma \leq \max_{l \in J} W_{J,-l}^\sigma$ we can conclude that

$$W_J^\sigma \leq \eta(\tau_J^k) \left(\mathbb{E}[xI(x \geq \tau_J^k)] + \mathbb{P}(x \geq \tau_J^k) \max_{l \in J} \delta W_{J,-l}^\sigma \right).$$

We now use induction on the number of active players. The induction basis follows from Step 1. Fix J and let $|J| = k$. Then the induction hypothesis states that $W_{J,-l}^\sigma < W_{k-1}^*$ for all $l \leq k$, which using the finiteness of J implies that $\max_{l \in J} \delta W_{J,-l}^\sigma < \delta W_{k-1}^*$. Then we find

$$\begin{aligned} W_J^\sigma &\leq \eta(\tau_J^k) \left(\mathbb{E}[xI(x \geq \tau_J^k)] + \mathbb{P}(x \geq \tau_J^k) \max_{l \in J} \delta W_{J,-l}^\sigma \right) \\ &< \eta(\tau_J^k) \left(\mathbb{E}[xI(x \geq \tau_J^k)] + \mathbb{P}(x \geq \tau_J^k) \delta W_{k-1}^* \right) \\ &\leq \eta(\tau_k^*) \left(\mathbb{E}[xI(x \geq \tau_k^*)] + \mathbb{P}(x \geq \tau_k^*) \delta W_{k-1}^* \right) \\ &= W_k^*. \end{aligned}$$

Where we have used that $\tau_J^k < \bar{v}$ for all subgame perfect equilibria σ and where the last inequality follows from the fact that τ_k^* maximizes $\eta(\tau) (\mathbb{E}[xI(x \geq \tau)] + \delta \mathbb{P}(x \geq \tau) W_{k-1}^*)$. $\square$

3.6 Conclusion and open problems

This paper investigates the policy of giving equal chances to players in a competitive search problem with decreasing competition. The fact that competition is expected to decrease brings interesting elements into this game. When focusing on pure threshold strategies, we found that players may have ex ante unequal expected outcomes. To illustrate this we constructed a monotone subgame perfect equilibrium, i.e. a subgame perfect equilibrium in which players maintain the same relative endogenous ranking in terms of expected payoff throughout the game. Furthermore, the policy of giving equal chances to all applicants will always cause a utilitarian welfare loss in equilibrium. Yet a lot of interesting questions remain unanswered. We see four directions for possible further research.

3.6.1 Other subgame perfect equilibria

In Theorem 3.4.7 we show that there exists a pure subgame perfect equilibrium in threshold strategies which has the property that players maintain their relative ranking in terms of expected payoffs throughout the game. A very fundamental and yet unanswered question about this specific equilibrium is how the set of expected payoffs evolves over time. This is particularly interesting because it would shed more light on the effects of competition on expected payoffs. Apart from the constructed monotone subgame perfect equilibrium, there may exist other equilibria as illustrated by Example 3.4.9. It is a natural question to ask whether it is possible to characterise all equilibria.

3.6.2 Understanding the equity-efficiency trade-off

We saw in Section 3.5 that any pure subgame perfect equilibrium results in a welfare loss. This contrasts to the socially optimal policy where players are given an exogenous priority ranking. Inspired by this it would be natural to study the equity-efficiency trade-off of all policies that lie in between the equal chances policy and the exogenous priority ranking policy. More precisely, we wonder whether it is true that the closer a policy is to the priority ranking, the higher the expected welfare, but also the higher the expected outcome inequality. Alternatively, one could restrict oneself to the subgame perfect equilibria which occur under the equal chances policy. In particular we wonder whether the monotone equilibrium maximizes both the utilitarian welfare and the range of expected equilibrium payoffs. To understand this last question, it seems crucial that one is able to develop a better understanding of how the lowest threshold is related to the opportunities of the players using this threshold in deeper subgames. Related to this, one could also attempt to find the worst subgame perfect equilibrium in terms of expected welfare or more generally to provide a bound on the price of anarchy for this game.

3.6.3 Understanding the effect of the underlying distribution

The attentive reader may have noticed that the majority of the results relied on the fact that the possible valuations for the items were assumed to be distributed according to a continuous distribution on a compact interval. This continuity property was crucial as it allowed players to slightly undercut each other. It is hence a natural question whether similar results hold in a game with a discrete distribution. It seems that in such a setting the relationship between the discount factor δ the cumulative distribution F and the set of values X becomes crucial. Even in games with a continuous distribution it would be interesting to know which distribution causes the highest or lowest welfare loss and which distribution induces the highest or lowest outcome inequality.

3.6.4 Generalizing the preferences and beliefs of the players

In the current model all players value everything in the same way and have full information about the preferences of their competitors. It would be natural to generalize both these assumptions and see to which extent our findings hold in this context. Firstly, one could consider generalizing this common value setting to an affiliated value setting. Secondly, one could try to incorporate an element of incomplete information in the game. This could include for example an element where players have to learn the distribution F or where players have incomplete information about the preferences of their competitors.

Appendix

A Universal measurability

This appendix contains a review of the definitions of universally measurable sets, integrals of universally measurable functions, stopping times, and conditional expectations, as well as two technical lemmas that are used in the paper.

Universally measurable sets: Let $(\mathcal{P}, \mathcal{F}^\infty)$ be a measurable space, where $\mathcal{F}^\infty$ denotes the Borel sigma-algebra and let $\mathcal{M}$ denote the collection of all probability measures over this measurable space. For each probability measure $\mathbb{P} \in \mathcal{M}$, we can extend the probability space $(\mathcal{P}, \mathcal{F}^\infty, \mathbb{P})$ to a complete probability space $(\mathcal{P}, \mathcal{F}_\mathbb{P}, \mathbb{P}^c)$ by including all $\mathbb{P}$ -negligible sets. To be more precise, let

$$\mathcal{F}_\mathbb{P}^0 = \{Q \subset \mathcal{P} \mid \exists Q' \in \mathcal{F}^\infty \text{ such that } \mathbb{P}(Q') = 0 \text{ and } Q \subseteq Q'\}$$

be the set of all subsets of $\mathbb{P}$ -negligible sets of $\mathcal{F}^\infty$. We define

$$\mathcal{F}_\mathbb{P} = \{Q \cup Q^0 \subseteq \mathcal{P} \mid Q \in \mathcal{F}^\infty \text{ and } Q^0 \in \mathcal{F}_\mathbb{P}^0\}$$

and we define $\mathbb{P}^c : \mathcal{F}_\mathbb{P} \rightarrow [0, 1]$ by

$$\mathbb{P}^c(Q \cup Q^0) = \mathbb{P}(Q), \quad Q \in \mathcal{F}^\infty, \quad Q^0 \in \mathcal{F}_\mathbb{P}^0.$$

It can be shown that $(\mathcal{P}, \mathcal{F}_\mathbb{P}, \mathbb{P}^c)$ is a probability space. Now define

$$\mathcal{F} = \bigcap_{\mathbb{P} \in \mathcal{M}} \mathcal{F}_\mathbb{P}.$$

It can be shown that $\mathcal{F}$ is a sigma-algebra that contains the Borel sigma-algebra $\mathcal{F}^\infty$ as a proper subset. The collection $\mathcal{F}$ is the universally measurable sigma-algebra and the elements of $\mathcal{F}$ are called universally measurable sets.

Integrals of universally measurable functions: A function $u : \mathcal{P} \rightarrow \mathbb{R}$ is called universally measurable if $u^{-1}[a, b] \in \mathcal{F}$ for every $[a, b] \subseteq \mathbb{R}$. The class of universally measurable functions contains the class of Borel measurable functions. Furthermore, for every universally measurable function there exists a Borel measurable function that coincides with it almost everywhere.

A function $g : \mathcal{P} \rightarrow \mathbb{R}$ is called a simple universally measurable function if it is of the form $g(p) = \sum_{i=1}^n c_i I(p \in Z_i)$, where $\{Z_1, \dots, Z_n\}$ is a partition of $\mathcal{P}$ with $Z_i \in \mathcal{F}$ for all $i = 1, \dots, n$. The expected value of a simple universally measurable payoff with respect to a probability measure $\mathbb{P}$ is defined as

$$\int_{p \in \mathcal{P}} g(p) \mathbb{P}(dp) = \sum_{i=1}^n c_i \mathbb{P}^c(Z_i). \quad (\text{A.1})$$

Let $\mathcal{G}$ denote the set of simple universally measurable functions. The expected value of a bounded universally measurable function $u : \mathcal{P} \rightarrow \mathbb{R}$ is then given by

$$\int_{p \in \mathcal{P}} u(p) \mathbb{P}(dp) = \sup_{\substack{g \in \mathcal{G} \\ g \leq u}} \int_{p \in \mathcal{P}} g(p) \mathbb{P}(dp) = \inf_{\substack{g \in \mathcal{G} \\ g \geq u}} \int_{p \in \mathcal{P}} g(p) \mathbb{P}(dp). \quad (\text{A.2})$$

Stopping times: A stopping time is a function $T : \mathcal{P} \rightarrow \mathbb{N} \cup \{\infty\}$ such that for each $t \in \mathbb{N}$ the set $\{p \in \mathcal{P} \mid T(p) = t\}$ is an element of $\mathcal{F}^t$. Given a stopping time T , let $\mathcal{F}^T$ denote the sigma-algebra of sets $A \in \mathcal{F}^\infty$ such that $A \cap \{p \in \mathcal{P} \mid T(p) = t\} \in \mathcal{F}^t$ for each $t \in \mathbb{N}$.

For every $t \in \mathbb{N}$, let X^t be an $\mathcal{F}^t$ measurable stochastic variable, and let X^∞ be an $\mathcal{F}^\infty$ measurable stochastic variable. The stochastic variable X^T is defined by letting it coincide with X^t on $\{p \in \mathcal{P} \mid T(p) = t\}$ for each $t \in \mathbb{N}$ and with X^∞ on $\{p \in \mathcal{P} \mid T(p) = \infty\}$. Following Yeh (1995, Theorem 3.28), it holds that X^T is $\mathcal{F}^T$ measurable.

Conditional expectations: Consider a bounded stochastic variable $F : \mathcal{P} \rightarrow \mathbb{R}$, a strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$, and a history $h \in \mathcal{H}^\ell$ of length ℓ . Let some $t \geq \ell$ be given. The conditional expectation of F with respect to the sigma-algebra $\mathcal{F}^t$ and the measurable space $(\mathcal{P}, \mathcal{F}, \mathbb{P}_{h, \sigma, \tau})$ is denoted by $\mathbb{E}_{h, \sigma, \tau}[F | \mathcal{F}^t]$. The conditional expectation $\mathbb{E}_{h, \sigma, \tau}[F | \mathcal{F}^t]$ can be identified with the stochastic variable $p \mapsto \mathbb{E}_{p|t, \sigma, \tau}[F]$.³

In what follows, we make repeatedly use of the following construction: Given a bounded function $f : \mathcal{H} \rightarrow \mathbb{R}$, we define for each $t \in \mathbb{N}$ the stochastic variable F^t by letting $F^t(p) = f(p|_t)$ for each play $p \in \mathcal{P}$. Notice that F^t is $\mathcal{F}^t$ measurable.

Lemma A.1. Let $f : \mathcal{H} \rightarrow \mathbb{R}$ be a bounded function and let $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$ be a strategy profile. The following two statements are equivalent:

- [1] For each $t \in \mathbb{N}$ and each history $h \in \mathcal{H}^t$ of length t it holds that $\mathbb{E}_{h, \sigma, \tau}[F^{t+1}] \geq f(h)$.
- [2] For each each history $h \in \mathcal{H}^\ell$ of length ℓ , the process $(F^t)_{t \geq \ell}$ is a $\mathbb{P}_{h, \sigma, \tau}$ -submartingale with respect to the filtration $(\mathcal{F}^t)_{t \geq \ell}$: for each $t \geq \ell$

$$\mathbb{E}_{h, \sigma, \tau}[F^{t+1} | \mathcal{F}^t] \geq F^t, \quad \mathbb{P}_{h, \sigma, \tau}\text{-almost surely.} \quad (\text{A.3})$$

Proof. To see that [1] implies [2], take a history $h \in \mathcal{H}^\ell$ of length ℓ and time $t \geq \ell$. Take a play $p \in \mathcal{P}(h)$. Evaluating the left-hand side of (A.3) at p yields $\mathbb{E}_{p|t, \sigma, \tau}[F^{t+1}]$, which is at least $f(p|_t) = F^t(p)$ by condition [1].

To see that condition [2] implies [1] take a $t \in \mathbb{N}$ and history $h \in \mathcal{H}^t$. Take any play $p \in \mathcal{P}(h)$. The left-hand side of (A.3) is simply $\mathbb{E}_{h, \sigma, \tau}[F^{t+1}]$. The right-hand side of (A.3), evaluated at p , is $f(h)$. Condition [1] follows as $\mathbb{P}_{h, \sigma, \tau}$ is carried by the set $\mathcal{P}(h)$. $\square$

Lemma A.2 states a version of Levy's zero-one law for universally measurable functions. It relies on the fact that a universally measurable function can be approximated by a Borel measurable function.

Lemma A.2. (Levy's zero-one law for universally measurable functions) For every strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$, for every history $h \in \mathcal{H}$, we have

$$\lim_{t \rightarrow \infty} U_{\sigma, \tau}^t = u, \quad \mathbb{P}_{h, \sigma, \tau}\text{-almost surely.} \quad (\text{A.4})$$

³As usual, a conditional expectation is not defined uniquely, since some histories might not be reached with positive probability under the strategy profile (σ, τ) . Our particular choice is both convenient and inconsequential, since any two conditional probability systems coincide $\mathbb{P}_{h, \sigma, \tau}$ -almost surely. We refer to Bogachev (2007, p. 350) for a careful discussion of conditional expectations.

Proof. Fix a strategy profile $(\sigma, \tau) \in \mathcal{S}_1 \times \mathcal{S}_2$ and a history $h \in \mathcal{H}$. Since u is universally measurable, there exists a Borel measurable function $\bar{u}$ such that $u = \bar{u}$, $\mathbb{P}_{h,\sigma,\tau}$ -almost surely. Then

$$\lim_{t \rightarrow \infty} U_{\sigma,\tau}^t = \lim_{t \rightarrow \infty} \mathbb{E}_{h,\sigma,\tau} [u | \mathcal{F}^t] = \lim_{t \rightarrow \infty} \mathbb{E}_{h,\sigma,\tau} [\bar{u} | \mathcal{F}^t] = \mathbb{E}_{h,\sigma,\tau} [\bar{u} | \mathcal{F}^\infty] = \bar{u},$$

$\mathbb{P}_{h,\sigma,\tau}$ -almost surely. In the first equality we use the definition of the stochastic variable $U_{\sigma,\tau}^t$. The second equality follows from the fact that $u = \bar{u}$ $\mathbb{P}_{h,\sigma,\tau}$ -almost surely. The third equality follows from Levy's zero-one law (see e.g. Bogachev, 2007, Example 10.3.15). The last equality follows from the fact that $\bar{u}$ is $\mathcal{F}^\infty$ measurable. This is because $\bar{u}$ is Borel measurable and $\mathcal{F}^\infty$ is the Borel sigma-algebra on $\mathcal{P}$. The fact that $u = \bar{u}$ $\mathbb{P}_{h,\sigma,\tau}$ -almost surely concludes the proof. $\square$

B The proof of inequality (1.6.7)

The following statement follows from the more general result of Theorem 1 in Abate, Redig, and Tkachev (2014). For completeness, we provide a direct proof in this section.

Lemma B.1. Let $\sigma, \sigma' \in \mathcal{S}_1$ be such that, for every $t \in \mathbb{N}$, for every $h \in \mathcal{H}^t$, $\|\sigma(h) - \sigma'(h)\|_{\text{TV}} \leq \delta_t$. Then, for every strategy $\tau \in \mathcal{S}_2$, for every $t \in \mathbb{N}$, and for every history $h \in \mathcal{H}^t$,

$$\|\mathbb{P}_{h,\sigma,\tau} - \mathbb{P}_{h,\sigma',\tau}\|_{\text{TV}} \leq \sum_{i=t}^{\infty} \delta_i.$$

The class $\mathcal{V} \subseteq \mathcal{F}$ is called an inner (outer) approximating class for the class $\mathcal{W} \subseteq \mathcal{F}$ if $\mathcal{V}$ is closed under unions (intersections) and if for every $\epsilon > 0$, and for every probability measure $\mathbb{P} \in \mathcal{M}$ and every set $W \in \mathcal{W}$ there exists a set $V \in \mathcal{V}$ such that $V \subseteq W$ ($V \supseteq W$) and $|\mathbb{P}(W) - \mathbb{P}(V)| \leq \epsilon$.

Lemma B.2. If $\mathcal{V}$ is an inner (outer) approximating class for the class $\mathcal{W}$ and $\mathcal{V} \subseteq \mathcal{W}$, then it holds for every two probability measures $\mathbb{P} \in \mathcal{M}$ and $\mathbb{P}' \in \mathcal{M}$ that

$$\sup_{W \in \mathcal{W}} |\mathbb{P}(W) - \mathbb{P}'(W)| = \sup_{V \in \mathcal{V}} |\mathbb{P}(V) - \mathbb{P}'(V)|.$$

Proof. Fix $\epsilon > 0$. Assume that the class $\mathcal{V}$ is an inner approximating class for $\mathcal{W}$. The proof for an outer approximating class is similar. Fix two probability measures $\mathbb{P}$ and $\mathbb{P}'$ and a set $W \in \mathcal{W}$. Then there exist sets $V \in \mathcal{V}$ and $V' \in \mathcal{V}$ such that $V \subseteq W$, $V' \subseteq W$, $|\mathbb{P}(W) - \mathbb{P}(V)| \leq \epsilon$, and $|\mathbb{P}'(W) - \mathbb{P}'(V')| \leq \epsilon$. Let $\tilde{V} = V \cup V'$. Because $\mathcal{V}$ is closed under unions we have that $\tilde{V} \in \mathcal{V}$. Furthermore, it follows trivially that $\tilde{V} \subseteq W$, $|\mathbb{P}(W) - \mathbb{P}(\tilde{V})| \leq \epsilon$, and $|\mathbb{P}'(W) - \mathbb{P}'(\tilde{V})| \leq \epsilon$. We find that

$$\begin{aligned} |\mathbb{P}(W) - \mathbb{P}'(W)| &= |\mathbb{P}(W) - \mathbb{P}(\tilde{V}) + \mathbb{P}(\tilde{V}) - \mathbb{P}'(\tilde{V}) + \mathbb{P}'(\tilde{V}) - \mathbb{P}'(W)| \\ &\leq |\mathbb{P}(W) - \mathbb{P}(\tilde{V})| + |\mathbb{P}(\tilde{V}) - \mathbb{P}'(\tilde{V})| + |\mathbb{P}'(\tilde{V}) - \mathbb{P}'(W)| \\ &\leq |\mathbb{P}(\tilde{V}) - \mathbb{P}'(\tilde{V})| + 2\epsilon. \end{aligned}$$

Because this holds for any $\epsilon > 0$ and any set $W \in \mathcal{W}$, we can conclude that

$$\sup_{W \in \mathcal{W}} |\mathbb{P}(W) - \mathbb{P}'(W)| \leq \sup_{V \in \mathcal{V}} |\mathbb{P}(V) - \mathbb{P}'(V)|.$$

Because $\mathcal{V} \subseteq \mathcal{W}$ it is clear that:

$$\sup_{W \in \mathcal{W}} |\mathbb{P}(W) - \mathbb{P}'(W)| \geq \sup_{V \in \mathcal{V}} |\mathbb{P}(V) - \mathbb{P}'(V)|.$$

□

In the following lemma we use Lemma B.2 to simplify the computation of the total variation norm. Instead of having to compute the supremum over all sets of the universally measurable sigma-algebra it will be sufficient to compute the supremum over the subclass of open sets $\mathcal{O}^t$, $t \in \mathbb{N}$, defined by

$$\mathcal{O}^t = \{\cup_{h \in H} \mathcal{P}(h) | H \subseteq \mathcal{H}^t\} \quad (\text{B.1})$$

The set $\mathcal{O}^t$ is the class of open sets such that all the plays sharing a common history at time t are such that either all or none of them are contained in a specific open set.

Lemma B.3. For every $h \in \mathcal{H}$ it holds that

$$\|\mathbb{P}_{h,\sigma,\tau} - \mathbb{P}_{h,\sigma',\tau}\|_{\text{TV}} = \sup_{t \in \mathbb{N}, O \in \mathcal{O}^t} |\mathbb{P}_{h,\sigma,\tau}(O) - \mathbb{P}_{h,\sigma',\tau}(O)|.$$

Proof. We define $\mathcal{O} = \cup_{t \in \mathbb{N}} \mathcal{O}^t$. Since any set in $\mathcal{O}$ is a union of open sets, it holds that $\mathcal{O}$ is a subset of the class of all open sets, $\mathcal{O}^*$. We now show that $\mathcal{O}$ is an inner approximating class for $\mathcal{O}^*$. Fix an open set $O^* \in \mathcal{O}^*$. For every $t \in \mathbb{N}$, we define $O^t = \{p \in \mathcal{P} | \mathcal{P}(p|_t) \subseteq O^*\}$. It is clear that for every $t \in \mathbb{N}$ we have $O^t \in \mathcal{O}$ and $O^t \subseteq O^*$. Furthermore, we have that $O^1 \subseteq O^2 \subseteq \dots$ and $O^* = \cup_{t \in \mathbb{N}} O^t$. For every probability measure $\mathbb{P}$ it therefore holds that $\lim_{t \rightarrow \infty} \mathbb{P}(O^t) = \mathbb{P}(O^*)$, so for every $\epsilon > 0$ there exists $t \in \mathbb{N}$ and $O^t \in \mathcal{O}^t$ such that $|\mathbb{P}(O^*) - \mathbb{P}(O^t)| < \epsilon$. Hence $\mathcal{O}$ is an inner approximating class of $\mathcal{O}^*$.

Because of the outer regularity of Borel measures on metric spaces, we have that the class of open sets is an outer approximating class for the class of Borel sets. In addition we have that the class of Borel sets is an inner approximating class for the class of universally measurable sets. Indeed, for any probability measure and any universally measurable set, there exists a Borel set which is contained in the universally measurable set and has the same probability. Repeated application of Lemma B.2 concludes the proof. □

Proof of Lemma B.1: Fix $t \in \mathbb{N}$ and a history $h_t \in \mathcal{H}^t$. We prove by induction that for every $n \geq t$, for every $O \in \mathcal{O}^{n+1}$,

$$|\mathbb{P}_{h_t,\sigma,\tau}(O) - \mathbb{P}_{h_t,\sigma',\tau}(O)| \leq \sum_{i=t}^n \delta_i. \quad (\text{B.2})$$

Induction basis ($n = t$).

We prove that for every $O \in \mathcal{O}^{t+1}$, $|\mathbb{P}_{h_t,\sigma,\tau}(O) - \mathbb{P}_{h_t,\sigma',\tau}(O)| \leq \delta_t$. Fix a set $O \in \mathcal{O}^{t+1}$. We define

$$\mathcal{Z}_{h_t} = \{(a, b, x) \in \mathcal{A} \times \mathcal{B} \times \mathcal{X} | \mathcal{P}(h_t a b x) \subseteq O\}.$$

Let $\mathcal{A}_{h_t} = \{a \in \mathcal{A} | \exists (b, x) \in \mathcal{B} \times \mathcal{X} : (a, b, x) \in \mathcal{Z}_{h_t}\}$ be the projection of the set $\mathcal{Z}_{h_t}$ on the set $\mathcal{A}$. Let x_t denote the state at the history h_t . We have that

$$\mathbb{P}_{h_t,\sigma,\tau}(O) = \sum_{(a,b,x) \in \mathcal{Z}_{h_t}} \sigma(h_t)(a) \cdot \tau(h_t)(b) \cdot q(x|a, b, x_t),$$

$$\mathbb{P}_{h_t, \sigma', \tau}(O) = \sum_{(a, b, x) \in \mathcal{Z}_{h_t}} \sigma'(h_t)(a) \cdot \tau(h_t)(b) \cdot q(x|a, b, x_t).$$

We find that

$$\begin{aligned} & |\mathbb{P}_{h_t, \sigma, \tau}(O) - \mathbb{P}_{h_t, \sigma', \tau}(O)| \\ &= \left| \sum_{(a, b, x) \in \mathcal{Z}_{h_t}} (\sigma(h_t)(a) - \sigma'(h_t)(a)) \cdot \tau(h_t)(b) \cdot q(x|a, b, x_t) \right| \\ &\leq \sum_{(a, b, x) \in \mathcal{Z}_{h_t}} |\sigma(h_t)(a) - \sigma'(h_t)(a)| \cdot \tau(h_t)(b) \cdot q(x|a, b, x_t) \\ &\leq \sum_{(a, b, x) \in \mathcal{A}_{h_t} \times \mathcal{B} \times \mathcal{X}} |\sigma(h_t)(a) - \sigma'(h_t)(a)| \cdot \tau(h_t)(b) \cdot q(x|a, b, x_t) \\ &= \left[\sum_{a \in \mathcal{A}_{h_t}} |\sigma(h_t)(a) - \sigma'(h_t)(a)| \right] \cdot \left[\sum_{(b, x) \in \mathcal{B} \times \mathcal{X}} \tau(h_t)(b) \cdot q(x|a, b, x_t) \right] \\ &\leq \sum_{a \in \mathcal{A}} |\sigma(h_t)(a) - \sigma'(h_t)(a)| \\ &= \|\sigma(h_t) - \sigma'(h_t)\|_{\text{TV}} \leq \delta_t. \end{aligned}$$

Induction step.

From the induction hypotheses it follows that for every open n -level set $O^n \in \mathcal{O}^n$ with $n - 1 \geq t$,

$$|\mathbb{P}_{h_t, \sigma, \tau}(O^n) - \mathbb{P}_{h_t, \sigma', \tau}(O^n)| \leq \sum_{i=t}^{n-1} \delta_i. \quad (\text{B.3})$$

Fix a set $O \in \mathcal{O}^{n+1}$. We can assume without loss of generality that $\mathbb{P}_{h_t, \sigma, \tau}(O) - \mathbb{P}_{h_t, \sigma', \tau}(O) \geq 0$. Define

$$\mathcal{H}_+^n = \{h_n \in \mathcal{H}^n | \mathbb{P}_{h_t, \sigma, \tau}(\mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \geq 0\}. \quad (\text{B.4})$$

We have that

$$\begin{aligned} & \mathbb{P}_{h_t, \sigma, \tau}(O) - \mathbb{P}_{h_t, \sigma', \tau}(O) \\ &= \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) \mathbb{P}_{h_t, \sigma, \tau}(\mathcal{P}(h_n)) - \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma', \tau}(O | \mathcal{P}(h_n)) \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \\ &= \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) \mathbb{P}_{h_t, \sigma, \tau}(\mathcal{P}(h_n)) - \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \\ &\quad + \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) - \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma', \tau}(O | \mathcal{P}(h_n)) \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \\ &= \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) \cdot (\mathbb{P}_{h_t, \sigma, \tau}(\mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n))) \\ &\quad + \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \cdot (\mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(O | \mathcal{P}(h_n))) \\ &\leq \sum_{h_n \in \mathcal{H}_+^n} \mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) \cdot (\mathbb{P}_{h_t, \sigma, \tau}(\mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n))) \\ &\quad + \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \cdot |(\mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(O | \mathcal{P}(h_n)))| \\ &\leq \sum_{h_n \in \mathcal{H}_+^n} (\mathbb{P}_{h_t, \sigma, \tau}(\mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n))) + \sum_{h_n \in \mathcal{H}^n} \mathbb{P}_{h_t, \sigma', \tau}(\mathcal{P}(h_n)) \cdot \delta_n \\ &\leq |\mathbb{P}_{h_t, \sigma, \tau}(\cup_{h_n \in \mathcal{H}_+^n} \mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(\cup_{h_n \in \mathcal{H}_+^n} \mathcal{P}(h_n))| + \delta_n \\ &\leq \sum_{i=t}^{n-1} \delta_i + \delta_n = \sum_{i=t}^n \delta_i, \end{aligned}$$

where the fact that $|(\mathbb{P}_{h_t, \sigma, \tau}(O | \mathcal{P}(h_n)) - \mathbb{P}_{h_t, \sigma', \tau}(O | \mathcal{P}(h_n)))| \leq \delta_n$ follows by assumption.

Using Lemma B.3 we can conclude that

$$\begin{aligned} \|\mathbb{P}_{h_t, \sigma, \tau} - \mathbb{P}_{h_t, \sigma', \tau}\|_{\text{TV}} &= \sup_{n \in \mathbb{N}, O \in \mathcal{O}^n} |\mathbb{P}_{h_t, \sigma, \tau}(O) - \mathbb{P}_{h_t, \sigma', \tau}(O)| \\ &= \sup_{n \geq t, O \in \mathcal{O}^{n+1}} |\mathbb{P}_{h_t, \sigma, \tau}(O) - \mathbb{P}_{h_t, \sigma', \tau}(O)| \\ &\leq \sup_{n \geq t} \sum_{i=t}^n \delta_i = \sum_{i=t}^{\infty} \delta_i. \end{aligned}$$

C Symmetric equilibria in behavioral strategies.

Definition C.1. A function $\Delta : S^p \rightarrow [0, 1]$ is called a behavioral strategy of player p . Here $\Delta(x, J)$ denotes the probability of applying to an item with value x when the set of active players is J .

Assume that once a player enters a subgame where he is the only active player, he will use the optimal threshold τ_1^1 securing an expected payoff of π_1^1 . Let Δ_2^1 and Δ_2^2 denote the behavioral strategies used by both players in the two player subgame then the associated expected payoff function is given by

$$\begin{aligned} \pi_2^1(\Delta_2^1, \Delta_2^2) = \int_X \left(\Delta_2^1(x, J_2) \left(\Delta_2^2(x, J_2) \left(\frac{1}{2}x + \frac{1}{2}\delta\pi_1^1 \right) + (1 - \Delta_2^2(x, J_2))x \right) \right. \\ \left. + (1 - \Delta_2^1(x, J_2))(\Delta_2^2(x, J_2)\delta\pi_1^1 + (1 - \Delta_2^2(x, J_2))\delta\pi_2^2) \right) f(x)dx \end{aligned}$$

$$\begin{aligned} \pi_2^2(\Delta_2^1, \Delta_2^2) = \int_X \left(\Delta_2^2(x, J_2) \left(\Delta_2^1(x, J_2) \left(\frac{1}{2}x + \frac{1}{2}\delta\pi_1^1 \right) + (1 - \Delta_2^1(x, J_2))x \right) \right. \\ \left. + (1 - \Delta_2^2(x, J_2))(\Delta_2^1(x, J_2)\delta\pi_1^1 + (1 - \Delta_2^1(x, J_2))\delta\pi_2^2) \right) f(x)dx \end{aligned}$$

The following example illustrates the existence of a symmetric subgame perfect equilibrium in behavioral strategies.

Example C.2. [symmetric subgame perfect equilibrium in behavioral strategies]

Consider the two-player game where $X \sim U[0, 1]$. Let π_1^1 denote the optimal payoff at a one-player subgame. Fix $\delta \in (0, 1)$. Let J_1 resp. J_2 denote sets with respectively one or two active players.

Part 1: Assume there exists a symmetric subgame perfect equilibrium (Δ, Δ) . Then the equilibrium strategy Δ is given by:

$$\Delta(x, J_2) = \begin{cases} 1 & \text{if } x \geq \pi_1^1, \\ \frac{2x - 2\delta\pi_2}{\delta\pi_1^1 - x - 2\delta\pi_2} & \text{if } \pi_2 < x < \pi_1^1, \\ 0 & \text{otherwise,} \end{cases} \quad \Delta(x, J_1) = I(x \geq \tau_1^1)$$

where π_2 denotes the expected equilibrium payoff of the two-player subgame.

Note that in the one player subgame, the optimal strategy is to always apply for items with a value at least τ_1^1 resulting in an expected payoff of π_1^1 .

Step 1: $\Delta(x, J_2) = 1 \Leftrightarrow x \geq \delta\pi_1^1$

If $\Delta(x, J_2) = 1$ then both players apply for the item with value x . This implies that $\frac{1}{2}x + \frac{1}{2}\delta\pi_1^1 \geq \delta\pi_1^1$. Which reduces to $x \geq \delta\pi_1^1$. Conversely, if $x \geq \delta\pi_1^1$ it is a dominant strategy for each of the players to apply.

Step 2: $\Delta(x, J_2) = 0 \Leftrightarrow x \geq \delta\pi_2$

If $\Delta(x) = 0$ then both players wait. Hence $\delta\pi_2 \geq x$. Furthermore if $\delta\pi_2 > x$ it is a dominant strategy to wait.

Step 3: Determining $\Delta(x, J_2)$ for $x \in (\delta\pi_2, \delta\pi_1^1)$

From Step 1 and 2 it follows that $\Delta(x, J_2) \in (0, 1)$ for all $x \in (\delta\pi_2, \delta\pi_1^1)$. Given that the other

player uses $\Delta(x, J)$ the expected payoff of applying is $1/2\Delta(x, J_2)(x + \delta\pi_1^1) + (1 - \Delta(x, J_2))x$. While the payoff of waiting is $\Delta(x, J_2)\delta\pi_1^1 + (1 - \Delta(x, J_2))\delta\pi_2$. As $\Delta(x, J_2) \in (0, 1)$ and (Δ, Δ) is an equilibrium we need that:

$$\frac{1}{2}\Delta(x, J_2)(x + \delta\pi_1^1) + (1 - \Delta(x, J_2))x = \Delta(x, J_2)\delta\pi_1^1 + (1 - \Delta(x, J_2))\delta\pi_2.$$

Indeed, if this was not the case the best-response would be pure, which contradicts the fact that $\Delta(x, J_2) \in (0, 1)$. Solving for $\Delta(x, J_2)$ we have that if $x \in (\delta\pi_2, \delta\pi_1^1)$

$$\Delta(x, J_2) = \frac{2x - 2\delta\pi_2}{\delta\pi_1^1 - x - 2\delta\pi_2}.$$

Note that in this equation the value π_2 is actually unknown.

Part 2: Determining a value for π_2

If π_2 denotes the expected payoff of the strategy profile (Δ, Δ) then π_2 needs to satisfy $\pi_2 = \phi(\pi_2)$ where the function $\phi : [0, \pi_1^1] \rightarrow \mathbb{R}$ defined by:

$$\begin{aligned} \phi(\pi_2) = & \eta(\delta\pi_2) \int_{\delta\pi_2}^{\delta\pi_1^1} \left[\left(\frac{2x - 2\delta\pi_2}{\delta\pi_1^1 - x - 2\delta\pi_2} \right) x + \left(\frac{\delta\pi_1^1 - x}{\delta\pi_1^1 - x - 2\delta\pi_2} \right) \delta\pi_1^1 \right] dx \\ & + \eta(\delta\pi_2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \delta\pi_1^1)] + \frac{1}{2} \mathbb{P}(x \geq \delta\pi_1^1) \delta\pi_1^1 \right). \end{aligned}$$

Observe that this function is continuous and that $\phi(0) > 0$ and $\phi(\pi_1^1) < \pi_1^1$. Indeed,

$$\phi(\pi_1^1) = \eta(\delta\pi_2) \left(\frac{1}{2} \mathbb{E}[xI(x \geq \delta\pi_1^1)] + \frac{1}{2} \mathbb{P}(x \geq \delta\pi_1^1) \delta\pi_1^1 \right) < \eta(\delta\pi_1^1) \mathbb{E}[xI(x \geq \delta\pi_1^1)] = \pi_1^1.$$

Hence there exists a fixpoint $\tilde{\pi}_2 = \phi(\tilde{\pi}_2)$. It is now trivial to verify that the strategy profile (Δ, Δ) with $\pi_2 = \tilde{\pi}_2$ is a subgame perfect equilibrium.

The previous example can be easily generalized to construct a symmetric subgame perfect equilibrium of any continuous distribution F . Similarly one can use an inductive construction to find a symmetric subgame perfect equilibrium in larger subgames.

Bibliography

- [1] Abate, A., Redig, F., and Tkachev, I. (2014). On the effect of perturbation of conditional probabilities in total variation. *Statistics and Probability Letters*, 88, 1-8.
- [2] Abdelaziz, F. B., & Krichen, S. (2007). Optimal stopping problems by two or more decision makers: a survey. *Computational Management Science*, 4(2), 89.
- [3] Alós-Ferrer, C., & Ritzberger, K.: Characterizing existence of equilibrium for large extensive form games: a necessity result. *Economic Theory*, 63(2), 407-430 (2017)
- [4] Blackwell, D., Ferguson, T. (1968). The big match. *Annals of Mathematical Statistics*, 39, 159-163.
- [5] Bloch, F., & Cantala, D. (2017). Dynamic assignment of objects to queuing agents. *American Economic Journal: Microeconomics*, 9(1), 88-122.
- [6] Bogachev V.I. (2007). *Measure Theory: Volume II*, Springer-Verlag, Berlin Heidelberg.
- [7] Bruyère, V.: Computer Aided Synthesis: A Game-Theoretic Approach. In: *Developments in Language Theory* pp. 3-35. Springer International Publishing (2017)
- [8] Bruyère, V., Le Roux, S., Pauly, A., & Raskin, J.F.: On the existence of weak subgame perfect equilibria. In *International Conference on Foundations of Software Science and Computation Structures*, pp. 145-161. Springer, Berlin, Heidelberg (2017)
- [9] Chen, R. W., Rosenberg, B., & Shepp, L. A. (1997). A secretary problem with two decision makers. *Journal of Applied Probability*, 34(4), 1068-1074.
- [10] Cownden, D., & Steinsaltz, D. (2014). Effects of competition in a secretary problem. *Operations Research*, 62(1), 104-113.
- [11] Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Shmaya, E., Solan, E. & Vrieze, K.: Non-existence of subgame-perfect ε -equilibrium in perfect information games with infinite horizon. *International Journal of Game Theory*, 43(4), 945-951 (2014)
- [12] Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Solan, E. & Vrieze, K.: Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operations Research*, 35(4), 742-755 (2010)
- [13] Flesch, J., Predtetchinski, A. and Sudderth, W. (2018). Characterization and simplification of optimal strategies in positive stochastic games. Forthcoming in *Journal of Applied Probability*.

- [14] Flesch, J., Predtetchinski, A. (2016). On refinements of subgame perfect ϵ -equilibrium. *International Journal of Game Theory*, 45, 523-542.
- [15] Flesch, J., & Predtetchinski, A.: Subgame-perfect ϵ -equilibria in perfect information games with common preferences at the limit. *Mathematics of Operations Research*, 41(4), 1208-1221 (2016).
- [16] Flesch, J., Thuijsman, F., & Vrieze, K.: Cyclic Markov equilibria in stochastic games. *International Journal of Game Theory*, 26(3), 303-314 (1997)
- [17] Flesch, J., Thuijsman, F., and Vrieze, O.J. (1998). Improving strategies in stochastic games. *Proceedings of the 37th IEEE Conference on Decision and Control*, Volume 3, IEEE, pp. 2674-2679.
- [18] Fudenberg, D., Levine, D. (1983). Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31, 251-268.
- [19] Fushimi, M. (1981). The secretary problem in a competitive situation. *Journal of the Operations Research Society of Japan*, 24(4), 350-359.
- [20] Gillette, D. (1957). Stochastic games with zero stop probabilities. In Drescher, M., Tucker, A.W., and Wolfe, P. (eds), *Contributions to the Theory of Games, Vol III, Annals of Mathematics Studies, Volume 39*, Princeton University Press, Princeton, New Jersey, pp. 179-187.
- [21] Harris C., Reny P., & Robson, A.: The existence of subgame-perfect equilibrium in continuous games with almost perfect information: a case for public randomization. *Econometrica*, 63(3), 507-544 (1995)
- [22] Immorlica, N., Kleinberg, R., & Mahdian, M. (2006, December). Secretary problems with competing employers. In *International Workshop on Internet and Network Economics* (pp. 389-400). Springer, Berlin, Heidelberg.
- [23] Jaśkiewicz, A., & Nowak, A.: Non-Zero-Sum Stochastic Games. In: *Handbook of Dynamic Game Theory* (pp. 1-64). Springer International Publishing (2016)
- [24] Karlin, A., & Lei, E. (2015, February). On a competitive secretary problem. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- [25] Kechris, A.: *Classical Descriptive Set Theory*. Springer-Verlag (1995)
- [26] Laraki, R., Maitra, A., and Sudderth, W. (2013). Two-person zero-sum stochastic games with semicontinuous payoff. *Dynamic Games and Applications*, 3, 162-171.
- [27] Le Roux, S., & Pauly, A.: Equilibria in multi-player multi-outcome infinite sequential games. Preprint: Retrieved from arXiv:1401.3325 (2016)
- [28] Lippman, S. A., & McCall, J. J. (1976). The economics of job search: A survey. *Economic inquiry*, 14(2), 155-189.
- [29] Mailath, G., Postlewaite, A., and Samuelson, L. (2005). Contemporaneous perfect epsilon-equilibria. *Games and Economic Behavior*, 53, 126-140.

- [30] Maitra, A., Sudderth, W. (1996). *Discrete Gambling and Stochastic Games*, Springer-Verlag, New York.
- [31] Maitra, A., Sudderth, W. (1998). Finitely additive stochastic games with Borel measurable payoffs. *International Journal of Game Theory*, 27, 257-267.
- [32] Maitra, A. P., & Sudderth, W. D.: Subgame-perfect equilibria for stochastic games. *Mathematics of Operations Research*, 32(3), 711-722 (2016)
- [33] Martin, D.A. (1998). The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63, 1565-1581.
- [34] Mashiah-Yaakovi, A. (2015). Correlated equilibria in stochastic games with Borel measurable payoffs. *Dynamic Games and Applications*, 5, 120-135.
- [35] Mazalov, V., & Konovalchikova, E. (2015). The job-search problem with incomplete information. *Procedia Computer Science*, 55, 159-164.
- [36] McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 113-126.
- [37] Mertens, J. F., & Parthasarathy, T.: Equilibria for discounted stochastic games. In: *Stochastic Games and Applications* pp. 131-172. Springer (2003)
- [38] Munkres, J. R.: *Topology*. Prentis Hall (2000)
- [39] Nowak, A. (1985). Universally measurable strategies in zero-sum stochastic games. *Annals of Probability*, 13, 269-287.
- [40] Prokopovych, P., & Yannelis, N.C.: On strategic complementarities in discontinuous games with totally ordered strategies. *Journal of Mathematical Economics*, 70, 147-153 (2017)
- [41] Purves, R. A., & Sudderth, W.D.: Perfect information games with upper semicontinuous payoffs. *Mathematics of Operations Research*, 36(3), 468-473. (2011)
- [42] Puterman, M.L. (1994). *Markov decision processes, discrete stochastic dynamic programming*, John Wiley and Sons, Hoboken, New Jersey.
- [43] Radner, R. (1980). Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22, 136-154.
- [44] Rogerson, R., Shimer, R., & Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of economic literature*, 43(4), 959-988.
- [45] Rosenberg, D., Solan, E., and Vieille, N. (2001). Stopping games with randomized strategies. *Probability Theory and Related Fields*, 119, 433-451.
- [46] Sakaguchi, M. (1980). Non-zero-sum games related to the secretary problem. *Journal of the Operations Research Society of Japan*, 23(3), 287-293.

- [47] Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträchtigkeit: Teil I: Bestimmung des dynamischen Preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301-324.
- [48] Secchi, P., & Sudderth, W. D.: N -Person stochastic games with upper semi-continuous payoffs. *International Journal of Game Theory*, 30(4), 491-502 (2001)
- [49] Shapley, L. S.: Stochastic games. *Proceedings of the national academy of sciences*, 39(10), 1095-1100 (1953)
- [50] Solan, E., & Solan, O.: Quitting games and linear complementarity problems. Preprint: Retrieved from arXiv:1707.02598 (2017)
- [51] Solan, E., & Vieille, N.: Quitting games. *Mathematics of Operations Research*, 26(2), 265-285 (2001)
- [52] Solan, E., Vieille, N. (2002). Uniform value in recursive games. *Annals of Applied Probability*, 12, 1185-1201.
- [53] Su, X., & Zenios, S. (2004). Patient choice in kidney allocation: The role of the queueing discipline. *Manufacturing & Service Operations Management*, 6(4), 280-301.
- [54] Tversky, A., Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453-458.
- [55] Whitmeyer, M. (2018). A Competitive Optimal Stopping Game. *The BE Journal of Theoretical Economics*, 18(1).
- [56] Yeh, J., (1995). *Martingales and Stochastic Analysis: Volume 1*, World Scientific.

Samenvatting

Deze doctoraatsthesis levert nieuwe inzichten in het domein van oneindige dynamische spelen. Zulke spelen kunnen gebruikt worden om situaties te modelleren waarbij spelers elkaar meermaals. Elk van de spelers beschikt over een strategy die op elk moment in het spel zegt welke actie de desbrennende speler dient te nemen.

Een centraal concept in de studie van deze spelen is het subgame perfect evenwicht. Intutief gezien is een subgame perfect evenwicht een collectie van strategiën, voor elke speler één, waarbij geen enkele speler van strategie wil veranderen als geen van de andere spelers van strategie verandert, en dit op elk mogelijk moment tijdens het spel. In deze thesis gebruiken we ook een veralgemening van het subgame perfect evenwicht, namelijk het subgame perfect ϵ -evenwicht. Hierbij veronderstellen we dat spelers niet gevoelig zijn voor zeer kleine verschillen ϵ in hun payoff.

De studie van oneindig dynamische spelen heeft toepassingen in de economische wetenschappen alsook in de computerwetenschappen. In deze thesis dragen we bij tot beide toepassingsgebieden. De modellen in Hoofdstukken 1 en 2 zijn eerder gerelateerd aan toepassingen in de computerwetenschappen, terwijl het model in Hoofdstuk 3 duidelijke toepassingen kent in de economische wetenschappen.

In hoofdstuk 1 bestuderen we een oneindig stochastisch nul-som spel met twee spelers. Meer bepaald zijn we geïntreseed in subgame ϕ -maxmin strategiën. Dit zijn strategiën van de maximiserende speler die in elk deelspel, ook wel subgame genoemd, goed genoeg presteren. Wat precies verstaan wordt onder “goed genoeg” wordt gemodelleerd door middel van een tolerantiefunctie. We geven voldoende voorwaarden onder dewelke een subgame ϕ -maxmin strategie bestaat. Verder tonen we aan dan een subgame ϕ -maxmin niet noodzakelijk bestaat in elk twee speler stochastisch nul-som spel. Problematisch hierbij zijn spelen waarbij de payoff functie niet upper semicontinu is en waarbij de tolerantiefunctie zeer snel naar nul convergeert. Echter wanneer een subgame ϕ -maxmin strategie bestaat voor elke strikt positieve tolerantiefunctie ϕ , dan bestaat er een subgame maxmin strategie. Dit laatste resultaat is eerder verrassend en illustreert een belangrijk verschil met de couranter gebruikte notie van subgame ϵ -optimale strategiën, waarbij de toegestane fout in elk deelspel dezelfde is.

In hoofdstuk 2 bestuderen we dynamische spelen met bijna perfecte informatie en een oneindige horizon. Bij deze spelen ontmoeten spelers elkaar meermaals in verschillende situaties. Tijdens zo’n ontmoeting kiezen alle spelers gelijktijdig hun acties, rekening houdend met de eerder gespeelde acties door alle spelers. De eerder gekozen acties bepalen op een deterministische manier bepaald de nieuwe toestand. Deze spelen zijn een speciaal geval van stochastische spelen, waarbij de overgang tussen verschillende toestanden niet alleen bepaald wordt door eerder gekozen acties, maar ook afhankelijk is van een toevalselement.

In de context van dynamische spelen met een oneindige horizon introduceren we een nieuwe notie van continuïteit die we “individuele upper semicontinuiteit” noemen. Deze vorm van continuïteit veralgemeent de bekendere notie van upper semicontinuiteit. We merken op dat als in een dynamisch spel met bijna perfecte informatie en een oneindige horizon, alle spelers een payoff functie hebben die individueel upper semicontinuu is, er een moment komt wanneer alle spelers het eens zijn over de acties die nog genomen moeten worden. Dit cruciaal inzicht staat ons toe het bestaan te bewijzen van een subgame perfect ϵ -evenwicht, voor elke $\epsilon > 0$ in deze spelen. Dit resultaat veralgemeent een eerder bekomen resultaat van Secchi and Sudderth (2001).

In hoofdstuk 3 bestuderen we een concreet dynamisch spel met als doel een beter inzicht te verkrijgen in het effect van een afnemende competitiviteit op een gelijkekansenbeleid. Meer bepaald bestuderen we het volgende spel. Elke dag wordt er een object gepresenteerd aan een groep identieke spelers. De waarde van het object volgt een continue verdeling F . Na de presentatie van het object mag iedere speler beslissen of hij het object al dan niet wil hebben. Indien er precies één speler het object wil hebben, krijgt deze speler het desbetreffende object. Als er meerdere spelers het object wensen, krijgt iedere geïnteresseerde het object met een gelijke kans. Als er niemand het object wil hebben, verdwijnt het. Zodra een speler een object bemachtigd heeft verdwijnt deze speler uit het spel. We gaan ervan uit dat de spelers ongeduldig zijn, zodat de uiteindelijke payoff gegeven wordt door de verdisconteerde waarde van het object dat ze verkregen hebben.

We vinden de volgende resultaten. Ondanks het feit dat alle spelers identiek worden verondersteld en gelijke kansen krijgen bij verkrijgen van een object, zijn er subgame perfect evenwichten waarbij de verwachte uitkomst van de spelers verschillend is. Meer bepaald bestaat er een subgame perfect evenwicht waarbij spelers kunnen worden geranscht aan de hand van hun verwachte payoff, en waarbij deze onderlinge rangschikking dezelfde blijft de gehele duurtijd van het spel. We noemen zo subgame perfect evenwicht monotoon. Deze uitkomstongelijkheid is een direct gevolg van het feit dat de competitiviteit afneemt tijdens het spel. Vanuit een welwaarts perspectief vinden we het gelijkekansenbeleid steeds leidt tot een utilitair welvaartsverlies. Dit staat in contrast met het perfect discriminerende beleid dat geen utilitair welvaartsverlies veroorzaakt.

Valorisation

The purpose of this valorisation is to provide an answer to the following question: “What value can society derive from the research presented in this thesis?”

To answer this question we first need to realize that no research exists in a vacuum and neither does the research discussed in this thesis. This observation is especially important given the theoretical and abstract nature of this thesis. Because of this the immediate opportunities for concrete applications are limited. The findings of this thesis contribute to the existing academic literature on dynamic games. How precisely this research contributes to this existing academic literature was already answered elaborately in each of the chapters, there is therefore no need to repeat it here. Instead, I will attempt to sketch the value of the study of dynamic games.

The mathematical framework of dynamic games can be used to model interactions between players who encounter each other in different environments. These players often face a trade-off between choosing a good action now and trying to ensure a favourable future game. Such games are of interest to economists as they provide a natural framework to study the interactions of competing players over time. In Chapter 3 we studied a concrete dynamic game which had a clear economic motivation. The mathematical framework we developed there, provided insight in how decreasing competition can lead to expected outcome inequality, even when all competitors are identical and get equal opportunities. Apart from this, the economic applications of dynamic games are numerous. Other examples include the study of bargaining procedures, reputation building and the (over)exploitation of a common resource.

Dynamic games are also of interest to computer scientists. One application is the study of reactive systems. Such systems must be able to continuously react to uncontrollable events in the environment in which they interact. An example mentioned by Bruyère (2017) is the autopilot of a plane which controls the speed. This system must adjust the speed continuously depending on the weather conditions. This can be modelled by a two player zero-sum game. The two players in question are the system and the environment. The environment is assumed to be hostile and its objective is to do whatever it can to make the system fail. If the system can always succeed, no matter what the environment does, then the system has a winning strategy. Also non-zero sum dynamic games are of interest to computer science as they can be used to model complex systems, i.e. systems which consist of multiple components where each component has its own objective function.

In the applications of infinite dynamic games in computer science the payoff-function is often not “continuous at infinity”. Which implies that events that happen in a distant future can have a substantial impact on the outcome of the game. One example of such a payoff function is the reachability condition. Here the system gets a payoff of one if it

reaches a certain set of states and zero otherwise. The effect of the non-continuity of the payoff-functions on the complexity of the mathematical analysis of the game should not be underestimated.

To analyse an infinite horizon dynamic game where the payoff function is continuous at infinity, the game can essentially be analyzed as a finite horizon game, given that one is indifferent to small changes in the payoff. The analysis of finite horizon dynamic games can be easily done using the technique of backward induction. Unfortunately, this technique will not necessarily be successful in infinite games with discontinuous payoff-functions. Hence different techniques need to be developed for the analysis of such games.

It is in this context that Chapters 1 and 2 aim to contribute. In both chapters the discontinuity of the payoff functions plays a central role. Chapter 1 studies two player zero-sum stochastic games where players have universally measurable payoff functions. The class of universally measurable payoff functions is very broad and essentially includes all possible payoff functions for which an expected payoff can be computed for any given strategy profile. Chapter 2 defines a new type of continuity which we call individually upper-semi continuity, which extends the notation of upper semi-continuity. We showed, by means of construction, that a subgame perfect ϵ equilibrium always exists in dynamic games with almost perfect information in which all players have individually upper semicontinuous payoff functions.

Biography

Jasmine Maes was born in Ghent, Belgium in 1992. She obtained a masters degree in Mathematics for the University of Ghent in July 2016, specializing in Applied Mathematics. In September 2016 she joined the department of Microeconomics and Public Economics at Maastricht University to pursue a PhD in game theory under the supervision of Prof. dr. Jean-Jacques Herings, dr. János Flesch and dr. Arkadi Predtetchinski. During her time as a PhD-student she visited Prof. dr. Nicolas Vieille at HEC in Paris.