

Four Essays on Decision and Game Theory

Inauguraldissertation
zur
Erlangung des Doktorgrades
der
Wirtschafts- und Sozialwissenschaftlichen Fakultät
der
Universität zu Köln

2014

vorgelegt

von

Dipl. Math. Johannes Kern

aus

Villingen-Schwenningen

Referent: Prof. Dr. Carlos Alós-Ferrer

Korreferent: Prof. Dr. Klaus Ritzberger

Tag der Promotion: 10.7.2014

Danksagung

Ich möchte allen Leuten danken, die mir direkt und indirekt bei der Fertigstellung dieser Arbeit geholfen haben. Insbesondere

- Carlos Alós-Ferrer für die erstklassige Betreuung,
- Klaus Ritzberger für die Erstellung des Zweitgutachtens,
- Đura-Georg Granić für mehrfaches Korrekturlesen, gute Kommentare und moralische Unterstützung,
- Johannes Buckenmaier, Jaume García-Segarra, Sabine Hügelschäfer, Ana Križić, Jiahui Li, Alexander Ritschel, Fei Shi und Alexander Wagner für die tolle Zusammenarbeit während der letzten Jahre,
- meiner Mutter Hildegard Kern sowie Verena Kempf für die Unterstützung während der letzten Jahre.

TABLE OF CONTENTS

List of Tables	v
List of Figures	vi
Introduction and Summary	1
1 Repeated Games in Continuous Time as Extensive Form Games	5
1.1 Introduction	5
1.2 Repeated Games in Continuous Time	9
1.2.1 Extensive Form Games Without Discreteness Assump- tions	9
1.2.2 Existing Approaches to Extensive Form Games in Con- tinuous Time	13
1.2.3 The Action-Reaction Framework	15
1.3 A Possibility Result	22
1.3.1 Strategies and Outcomes in General Extensive Form Games	23
1.3.2 Strategies and Outcomes in the Action-Reaction Frame- work	25
1.4 An Alternative Approach: Strategy Constraints	26
1.5 An Equivalence Result	31
1.5.1 Outcome-Equivalence and Equivalence Classes	31
1.5.2 Equivalence of CRM and Action-Reaction Framework	34
1.6 Relation to the Literature	36
1.6.1 Maximal Strategy Sets	36
1.6.2 Staying Quiet	38

1.6.3	Conditioning on Counterfactuals	38
1.7	Conclusion	41
	Appendix 1.A: Proofs from Sections 1.2 and 1.3	43
	Appendix 1.B: Proofs from Section 1.4	51
	Appendix 1.C: Proofs from Section 1.5	58
	Appendix 1.D: Proposition 5	76
	References Chapter 1	77
2	Comment on “Trees and Extensive Forms”	81
2.1	Introduction	81
2.2	Preliminaries	81
2.3	Corrected Formulation and Changes	82
2.4	Example	86
	References Chapter 2	88
3	Circulant Games	89
3.1	Introduction	89
3.2	Circulant Games	92
3.3	Main Results	97
3.3.1	Preliminaries	97
3.3.2	The Number of Nash Equilibria	99
3.3.3	The Structure of Nash Equilibria	100
3.4	Generalizations	103
3.5	Conclusion	104
	Appendix 3.A: Transformation of Games	106
	Appendix 3.B: Proofs of Main Results	107
	Appendix 3.C: Tables	119
	References Chapter 3	121
4	Preference Reversals: Time and Again	123
4.1	Introduction	123
4.2	A Simple Model of Preference Reversals and Decision Times	126
4.2.1	Model and Rationale	127
4.2.2	Predictions	129

4.2.3	Order Effects and Preference Reappraisal	132
4.2.4	Process Conflict and Decision Times	134
4.3	Experiment 1: Preference Reversals and Decision Times	135
4.3.1	Experimental Design and Procedures	136
4.3.2	Procedures	138
4.3.3	Results of Experiment 1	139
4.3.4	Regression Analysis for Experiment 1	142
4.3.5	Discussion of Experiment 1	146
4.4	Experiment 2: Eliminating Reversals	147
4.4.1	Motivation and Hypotheses	147
4.4.2	Design of Experiment 2	148
4.4.3	Procedures	150
4.4.4	Results of Experiment 2	150
4.4.5	Regression Analysis for Experiment 2	156
4.4.6	Discussion of Experiment 2	158
4.5	General Discussion and Conclusion	161
	Appendix 4.A: Proofs	163
	Appendix 4.B: Lotteries	168
	Appendix 4.C: Screenshots	169
	References Chapter 4	170

LIST OF TABLES

3.1	Examples of iso-circulant games.	119
3.2	Examples of counter-circulant games.	120
4.1	Preference reversal rates, Experiment 1.	139
4.2	Random effects panel regressions for decision times, Experiment 1.	144
4.3	Preference reversal rates, Experiment 2.	151
4.4	Random effects panel regressions for decision times, Experiment 2.	157
4.5	The lottery pairs.	168

LIST OF FIGURES

4.1 Average number of reversals per subject, Experiment 1. 140

4.2 Average decision time per individual in the choice task, Experiment 1.141

4.3 Average non-reversal decision time per individual in the choice task,
Experiment 1. 143

4.4 Average number of reversals per subject, Experiment 2. 152

4.5 Average decision time per individual in the choice task, Experiment 2.154

4.6 Average non-reversal decision time per individual in the choice task,
Experiment 2. 155

4.7 Screen displays. 169

INTRODUCTION AND SUMMARY

This dissertation consists of four research papers, covering topics from decision and game theory. Chapters 1 and 2 concern continuous-time games and extensive forms. Chapter 3 presents results on the number of Nash equilibria in a particular class of games called circulant games, while Chapter 4 covers the preference reversal phenomenon. In the following, I present a brief overview of the four chapters summarizing the main findings.

Chapter 1, entitled “Repeated Games in Continuous Time as Extensive Form Games”, is the result of joint work with Carlos Alós-Ferrer (University of Cologne). Continuous-time games suffer from a severe conceptual issue, namely that some strategy profiles induce multiple outcomes while other profiles induce no outcome at all. Since preferences are defined on the set of ultimate outcomes, neither profiles leading to a multiplicity of outcomes nor profiles that “evaporate” can be evaluated, hence making it impossible to analyze such games for example in terms of equilibria. The literature has proposed several ways to deal with this issue. The most common one requires players to stick to a chosen action for some strictly positive amount of time. Indeed, it can be shown that any profile of such strategies induces a unique outcome. This approach is, however, problematic from a game-theoretic point of view. Fixing the extensive form of the game (i.e. decision nodes and choices) determines the players’ strategies as these are mappings from the set of decision nodes to the set of choices. Placing exogenous restrictions on the set of strategies hence implicitly changes (and in the worst case destroys) the extensive form. Our paper presents a game-theoretically well-founded framework for modeling repeated games in continuous time. It further provides a clarification as to which restrictions on strategies can be allowed in the sense that the resulting strategies can be derived from a well-defined extensive form. Work on this paper was shared among authors as follows: Johannes Kern 50%, Carlos Alós-Ferrer 50%.

Chapter 2 is based on the paper “Comment on ‘Trees and Extensive Forms’”, which is joint work with Carlos Alós-Ferrer (University of Cologne) and Klaus Ritzberger (IHS Vienna) and has been published in the *Journal of Economic Theory*, Vol. 146, No. 5, September 2011, pp. 2165–2168. The paper comments on the definition of Extensive Form in Alós-Ferrer and Ritzberger (2008) and shows that one of the properties there needs to be adjusted. It provides counterexamples showing that with the original version of this property some results do not hold as stated and presents a corrected formulation of the property as well as the corrected statement of the results. It further provides proofs for these results under the new formulation. Work on this paper was shared among authors as follows: Johannes Kern $33\frac{1}{3}\%$, Carlos Alós-Ferrer $33\frac{1}{3}\%$, Klaus Ritzberger $33\frac{1}{3}\%$.

Chapter 3 entitled “Circulant Games” is joint work with Đura-Georg Granić (University of Cologne). Games with a cyclical structure are ubiquitous in game theory and are routinely used to generate popular examples, starting with Matching Pennies and Rock-Paper-Scissors. For these as well as larger games, the cyclical structure can be captured by circulant payoff matrices in which each row vector is rotated by one element relative to the preceding row vector. In our paper we study a class of two-player games in which both players payoffs are given by such circulant matrices. Given that these payoffs are ordered, we are able to determine the exact number of (pure and mixed) Nash equilibria. This number only depends on the number of strategies, the position of one of the player’s largest payoff in the first row of his payoff matrix, and whether the players’ payoff matrices “cycle” in the same or in different directions. Our results further allow us to describe the support of each Nash equilibrium strategy. Work on this paper was shared among authors as follows: Johannes Kern 50%, Đura Georg Granić 50%.

Chapter 4, “Preference Reversals: Time and Again”, is the result of joint work with Carlos Alós-Ferrer, Đura-Georg Granić, and Alexander K. Wagner (all at the University of Cologne). Experiments documenting the preference reversal phenomenon highlight that, contrary to the invariance assumption underlying most economic theories of choice, preferences may actually be influenced by the elicitation method employed. In the most basic setup of such

experiments, subjects are asked to choose from pairs of lotteries containing one lottery with a high chance of paying a moderate amount of money (P -bet) and one lottery with a moderate chance of paying a high amount of money ($\$$ -bet). They are then asked to state prices for each of the lotteries. A preference reversal occurs if either the $\$$ -bet receives a higher price in a pair where the P -bet is chosen (predicted reversal) or the P -bet receives a higher price in a pair where the $\$$ -bet is chosen (unpredicted reversal). The preference reversal phenomenon is characterized by a significantly higher rate of predicted reversals. We present a new, simple process-based model that explains the preference reversal phenomenon and makes novel predictions about the associated decision times in the choice phase. The phenomenon is jointly caused by noisy lottery evaluations and an overpricing phenomenon associated with the compatibility hypothesis. A laboratory experiment confirmed the model's predictions for both choice data and decision times. Choices associated with reversals take significantly longer than non-reversals, and non-reversal choices take longer whenever long-shot lotteries are selected. A second experiment showed that the overpricing phenomenon can be shut down, greatly reducing reversals, by using ranking-based, ordinally-framed evaluation tasks. This experiment also disentangled the two determinants of the preference reversal phenomenon since noisy evaluations still deliver testable predictions on decision times even in the absence of the overpricing phenomenon. Work on this paper was shared among authors as follows: Johannes Kern 25%, Carlos Alós-Ferrer 25%, Đura Georg Granić 25%, Alexander K. Wagner 25%.

References

ALÓS-FERRER, C., AND K. RITZBERGER (2008): "Trees and Extensive Forms," *Journal of Economic Theory*, 43(1), 216–250.

CHAPTER 1

REPEATED GAMES IN CONTINUOUS TIME AS
EXTENSIVE FORM GAMES

1.1 Introduction

Suppose two players play a continuous-time version of the infinitely repeated Prisoner's Dilemma, starting at time $t = 0$. A player is then free to choose a strategy conditioning on arbitrary events in the past. For instance, a player could specify the following grim-trigger strategy: cooperate as long as both players have always cooperated in the past, otherwise defect forever. Now suppose both players use this strategy. One is tempted to conclude that the outcome of the strategy profile is eternal cooperation. Indeed, this outcome is *compatible* with the strategy profile in the sense that, at every point in time, instantaneous cooperation is prescribed by the strategy profile given the past history contained in the outcome. However, if time is continuous, there are infinitely many other outcomes which are equally compatible with these grim-trigger strategies. Fix any arbitrary time T , and consider the outcome where both players cooperate up to and including time T , and defect at any later point in time. Since there is no first point in time where players defect, this outcome never contradicts the prescriptions of the grim-trigger strategy profile and hence is also compatible with it. We conclude that the strategy profile induces a continuum of different outcomes. As a consequence, even if every outcome has a well-defined payoff, the payoff of the considered strategy profile is not well-defined, and a game-theoretical analysis becomes impossible.

Outcome multiplicity is not the only problem in continuous-time repeated games. Consider a different strategy profile where each player starts cooperating and further decides to cooperate *unless* only cooperation has been

observed in the past. What is the outcome? Obviously, eternal cooperation cannot be the outcome. But, if a defection occurred at any strictly positive point in time, this must mean that no defection occurred before, and hence the strategies prescribe a defection at every previous, strictly positive point in time, a contradiction. Hence, this simple strategy profile induces no outcome at all.

These problems have been previously pointed out by Anderson (1984), Simon and Stinchcombe (1989), Stinchcombe (1992), and Alós-Ferrer and Ritzberger (2008), among others. As shown in Alós-Ferrer and Ritzberger (2008, 2013a), they are not exclusive of continuous-time settings: intuitively, it suffices for the time axis to have an accumulation point towards the past to generate such problems, as e.g. in the case of the time set $\{1/n\}_{n=1,2,\dots} \cup \{0\}$. We now have a good understanding of the underlying reasons for these problems. Alós-Ferrer and Ritzberger (2008) (see also Alós-Ferrer, Kern, and Ritzberger, 2011) formulated out a *characterization* of the set of extensive forms where every profile of pure strategies generates a unique outcome (and hence a normal-form game can be defined). This characterization can be argued to describe the domain of game theory, for games outside the characterized set cannot be “solved” in any sense of the word. Unsurprisingly, perfect-information continuous-time games are outside this domain; technically, they fail a condition called “up-discreteness” in Alós-Ferrer and Ritzberger (2008), which precludes accumulation points toward the past.

This state of affairs has not prevented economic theory from venturing into the realm of continuous-time games (the literature is of course too extensive to review it here). And neither should it. On the one hand, continuous time is often analytically convenient due to the possibility of employing techniques from differential calculus along the time dimension. On the other hand, discrete time sometimes creates artificial phenomena which vanish away in continuous time; and it is their vanishing in the latter framework which proves their artificiality in the former. However, the problems pointed out above create serious difficulties with the interpretation of continuous-time applications. For instance, if certain strategy combinations fall out of the framework by virtue of creating outcome existence or uniqueness prob-

lems, the meaning of any equilibrium concept becomes questionable, since some deviations might be excluded for merely technical reasons, and not the self-interest of the deviator. Further, if a proper extensive form game cannot be specified for a continuous-time model, notions of “time consistency” cannot rely on subgame perfection or other equilibrium refinements based on backward induction, since in the absence of a properly formulated extensive form, it is not possible to determine the full collection of subgames capturing the strategic, intertemporal structure of the problem.

One typical approach for developing a coherent framework in continuous time is to admit an exogenous restriction on the set of pure strategies and declare some of those inadmissible. In the case of differential games (Friedman, 1994), this approach often leads to the specification of a normal-form game, where strategies are required to be e.g. differentiable or integrable functions of some state variable. In other domains, the analysis has been restricted to strategies incorporating some Markov structure, as e.g. in the case of the literature on (individual) strategic experimentation (e.g. Keller and Rady, 1999; Keller, Rady, and Cripps, 2005). The approach was most effectively described by Stinchcombe (1992), who set out to identify a *maximal* set of strategies for a continuous-time game such that every strategy profile induces a unique outcome. The result incorporates elements of a framework introduced by Anderson (1984) and also studied by Bergin and MacLeod (1993) and Bergin (1992, 2006), and rests on the condition that a strategy must always identify the player’s *next* move.

Stinchcombe (1992) identifies the best that can be done through strategy constraints once one accepts the inconvenient fact that unconstrained continuous time games cannot be solved. From a game-theoretic point of view, however, restricting the strategy set is an unsatisfactory approach. On the one hand, since certain strategies are excluded on purely technical grounds, we face the problems with the interpretation of equilibria and time consistency pointed out above. On the other hand, there is a more fundamental, conceptual problem. An extensive form game incorporates a complete description of the possible choices of every player at every decision node. A behavioral strategy is merely a collection of possible “local” decisions at the

nodes, and any possible combination thereof is a feasible behavioral strategy. Once the game is specified, there can be no further freedom in the specification of the possible local decisions, since those have already been fixed in the extensive form. The set of possible behavioral strategies is thus automatically specified once the extensive form is given.

A restriction prohibiting a given combination of local decisions in order to preserve some property of the outcome, no matter how desirable, lacks any decision-theoretic justification. Worse, it is then unclear whether the extensive form structure survives the restriction, raising doubts as to whether the resulting formal object is simply a (constrained) normal-form game.

Here we propose a different approach to the study of continuous-time games. The basic idea is as follows. Continuous time is a convenient device; its modelization within an extensive form game, however, needs only go so far as it is useful for game-theoretic purposes. The formalizations analyzed until now might have “gone too far”, in the sense that the associated extensive forms become too large and restoring tractability requires restricting their strategy spaces. The literature has concentrated on providing ideas and rationales for restricting the strategy space in an ex-post way. In this paper, we prove that continuous-time decisions can be captured by applying those ideas to the very definition of the game. The resulting formal object can then still be considered a well-defined “continuous-time game”; it is, however, a fully solvable extensive form game, i.e. every strategy profile induces a unique profile, *without any restriction on the set of behavioral strategies*. The advantage is that the framework is an extensive form game without any caveat, and standard game-theoretic concepts and methods can then be applied. In other words, our message is a positive one: we show that continuous-time modeling is possible without giving up the benefits and the conceptual discipline resulting from well-defined extensive form games.

In this paper, we focus on the repeated-game framework with observable actions.¹ Specifically, we show how repeated games in continuous time can

¹This is the framework where the problems we mentioned above are the most severe. Continuous-time models are also customarily used in different frameworks, e.g. games with imperfect monitoring (Sannikov, 2007). Intuitively, the fact that players cannot condition on as many events as in the case of perfectly observable actions shrinks the strategy space

be formalized incorporating natural conditions from the onset. The construction is not trivial, and in order to describe it we must carefully detail the appropriate game tree and choice structure. Once this is in place, we show that, by virtue of fulfilling the appropriate conditions, the resulting game is well-behaved without any restrictions on the strategy sets. In order to link our construction to the literature, we then show that it is possible to retrace our steps and prove an equivalence result between the unrestricted behavioral strategies in our repeated game and a restricted class of strategies in a more naïvely specified (and hence, in our view, problematic) continuous-time repeated game.

The paper is structured as follows. Section 1.2 lies out the general framework for repeated games in continuous time, the Action-Reaction Framework. Section 1.3 presents our main result, showing that in our framework all strategy profiles induce unique outcomes. Section 1.4 presents the alternative approach through restricted strategies (Conditional Response Mappings) and Section 1.5 proves an equivalence result, which allows us to link our extensive form to the previous literature in Section 1.6. Section 1.7 concludes. The construction and the main arguments are detailed in the main text but specific proofs are relegated to the appendix.

1.2 Repeated Games in Continuous Time

1.2.1 Extensive Form Games Without Discreteness Assumptions

Working definitions of extensive form games frequently incorporate strong restrictions in the form of explicit finiteness or discreteness assumptions. Since we aim to view continuous-time repeated games as extensive form games, we need a more general approach. We will rely on a definition of extensive form games allowing for infinite time horizon, continuous time axis, and arbitrary action sets. This concept is the basis for a general framework developed in

and makes it easier to obtain a well-defined extensive form game. See Alós-Ferrer and Kern (2013) for a comment.

Alós-Ferrer and Ritzberger (2005, 2008, 2013a,b).

The definition comes in two parts. The first is a general concept of *game tree*, capturing the order and nature of decisions. The second is a definition of *extensive decision problem* (given the game tree) which incorporates all appropriate consistency conditions on the *choices* that players can make.

Let us start with game trees. Following Kuhn (1953), a game tree is just an ordered set of “decision points” or *nodes* which can be represented as an abstract graph. Alternatively, von Neumann and Morgenstern (1944) focus on ultimate outcomes as the primitive objects and consider nodes as sets of such outcomes, which become finer as decisions are taken. A result arising in the work quoted above is that there exists exactly one way of defining game trees such that both approaches are equivalent. As a consequence, there is no loss of generality in assuming a game tree where nodes are taken to be sets of ultimate outcomes, as in the following definition.

Definition 1. A (rooted) *game tree* $T = (N, \supseteq)$ is a collection of nonempty subsets $x \in N$ (called nodes) of a given set W partially ordered by set inclusion such that $W \in N$ (W is called the root) and

- (TI) “Trivial Intersection:” if $x, y \in N$ with $x \cap y \neq \emptyset$, then $x \subsetneq y$ or $y \subseteq x$.
- (IR) “Irreducibility:” if $w, w' \in W$ with $w \neq w'$, then there exist $x, x' \in N$ such that $w \in x \setminus x'$ and $w' \in x' \setminus x$.
- (BD) “Boundedness:” for every nonempty chain $h \subseteq N$ there exists $w \in W$ such that $w \in x$ for all $x \in h$.²

A *play* is a chain of nodes $h \subseteq N$ that is maximal in N , i.e. there is no $x \in N \setminus h$ such that $h \cup \{x\}$ is a chain. Plays are the natural objects on which preferences can be defined in a setting where the time horizon is not assumed to be finite. The advantage of game trees is that the underlying set W can also be identified with the set of plays. Specifically, Alós-Ferrer and Ritzberger (2005, Theorem 3(c)) show that an element $w \in W$ can be seen either as a possible outcome (element of some node) or as a play (maximal

²A chain is a subset of N that is *completely* ordered by set inclusion.

chain of nodes), and a node $x \in N$ can be identified with the set of plays passing through it.

For a game tree (N, \supseteq) with set of plays/outcomes W and an arbitrary subset $a \subseteq W$ (not necessarily a node), define the *up-set* $\uparrow a$ and the *down-set* $\downarrow a$ by

$$\uparrow a = \{y \in N \mid y \supseteq a\} \text{ and } \downarrow a = \{y \in N \mid a \supseteq y\}.$$

The key implication of (TI) is that $\uparrow x$ is a chain for all $x \in N$, which is contained in (can be “prolonged to”) the play $\uparrow\{w\}$ for any $w \in x$. Further, if h is a play, by (BD) there exists a unique outcome $w \in W$ such that $\bigcap_{x \in h} x = \{w\}$, or, equivalently, $h = \uparrow\{w\}$. This fact is the basis for the equivalence between outcomes and plays, which essentially reduces to the fact that, for $w \in W$ and $x \in N$, $w \in x$ if and only if $x \in \uparrow\{w\}$. When a distinction is called for, we write w for the outcome and $\uparrow\{w\}$ for the play (chain of nodes).

We now turn to the second part of the definition. In an extensive form game, players make decisions at nodes that are properly followed by other nodes, called *moves*. Let $X = \{x \in N \mid \downarrow x \setminus \{x\} \neq \emptyset\}$ be the set of all moves.³ In finite, perfect information examples, the possible actions or options available to a player at a given move can be identified with the nodes following that move (its *immediate successors*). In more general settings, we need a more general object. The possible alternatives faced by players are modeled through *choices*, which are subsets $c \subseteq W$ satisfying a number of consistency conditions.

Before we present those conditions, we need a notion of when a choice c is available at a move x . For an arbitrary set of outcomes/plays $a \subseteq W$, the set of *immediate predecessors* of a is defined by

$$P(a) = \{x \in N \mid \exists y \in \downarrow a : \uparrow x = \uparrow y \setminus \downarrow a\}.$$

Since nodes in a game tree are sets of plays, they too may, but need not, have immediate predecessors. Since choices are also sets of plays, the set of

³All other nodes are called *terminal*. It follows from (IR) that a node $x \in N$ is terminal if and only if there is $w \in W$ such that $x = \{w\}$.

immediate predecessors of a choice is well defined, and we will say that a choice c is *available* at a move $x \in X$ if $x \in P(c)$. This is the key element in the following definition.

Definition 2. An *extensive decision problem* (EDP) with player set I is a pair (T, C) , where $T = (N, \supseteq)$ is a game tree with set of plays W and $C = (C_i)_{i \in I}$ is a system consisting of collections C_i (the sets of players' choices) of nonempty unions of nodes (hence, sets of plays) for all $i \in I$ such that

(EDP.i) if $P(c) \cap P(c') \neq \emptyset$ and $c \neq c'$, then $P(c) = P(c')$ and $c \cap c' = \emptyset$, for all $c, c' \in C_i$ for all $i \in I$;

(EDP.ii) $x \cap [\bigcap_{i \in I(x)} C_i] \neq \emptyset$ for all $(c_i)_{i \in I(x)} \in A(x)$ and for all $x \in X$;

(EDP.iii) if $y, y' \in N$ with $y \cap y' = \emptyset$ then there are $c, c' \in C_i$ for some player $i \in I$ such that $y \subseteq c$, $y' \subseteq c'$, and $c \cap c' = \emptyset$;

(EDP.iv) if $x \supseteq y \in N$, then there is $c \in A_i(x)$ such that $y \subseteq c$ for all $i \in I(x)$, for all $x \in X$;

where $A(x) = \times_{i \in I(x)} A_i(x)$, $A_i(x) = \{c \in C_i \mid x \in P(c)\}$ are the choices available to $i \in I$ at $x \in X$, and $I(x) = \{i \in I \mid A_i(x) \neq \emptyset\}$ is the set of decision makers at x , which is required to be nonempty, for all $x \in X$.

An *extensive form game* is an extensive decision problem together with a specification of players' preferences on the set of plays.

The interpretation of the conditions above is as follows (see Alós-Ferrer and Ritzberger, 2005, Section 5 or Alós-Ferrer and Ritzberger, 2008, Section 3 for additional details). (EDP.i) stands in for information sets: if two distinct choices $c, c' \in C_i$ are ever simultaneously available, then they are disjoint and available at the same moves—at those in the *information set* $P(c) = P(c')$. (EDP.ii) requires that simultaneous decisions by different players at a common move do select some outcome. (EDP.iii) states that for any two disjoint nodes, there is a player who can eventually make a decision that selects among them. Finally, (EDP.iv) states that, if a player takes a decision at a given node, he must be able not to discard any given successor of the node. This excludes absent-mindedness (Piccione and Rubinstein, 1997), as in the original formulation of Kuhn (1953).

An important point about EDPs is that they allow several players to decide at the same move. This sometimes simplifies both the representation of a game and the equilibrium analysis (see Alós-Ferrer and Ritzberger, 2013a, for examples). This will also be important for our present purposes, for in repeated games players act simultaneously at every time point. If we adopted the convention that each move is assigned to one player only, we would be forced to incorporate artificial “cascading information sets” to accommodate this characteristic.

1.2.2 Existing Approaches to Extensive Form Games in Continuous Time

We now turn to the specific problem of modeling a repeated game in continuous time explicitly as an extensive form game. A first, direct approach to this task is to define strategies as mappings from the set of history-time pairs to the set of possible actions with the minimal requirement that at time t the same action is prescribed for two histories that agree on $[0, t[$. Indeed, this approach can be readily formalized as an EDP (Alós-Ferrer and Ritzberger, 2005, 2008).

Let W be the set of functions $f : \mathbb{R}_+ \rightarrow A$, where $A = \prod_{i \in I} A_i$ and each A_i is some fixed set of actions containing at least two elements. W is the set of all possible outcomes in the continuous-time repeated game. Let the set of nodes be $N = \{x_t(f) \mid t \in \mathbb{R}_+, f \in W\}$, where $x_t(f) = \{g \in W \mid g(\tau) = f(\tau) \forall \tau \in [0, t[\}$ for $f \in W$ and $t \in \mathbb{R}_+$. A node $x_t(f)$ contains all functions that agree with f on $[0, t[$ while all possibilities of values at t and afterwards are still open. (N, \supseteq) can be shown to be a game tree (Alós-Ferrer and Ritzberger, 2005).

A strategy in this framework is a mapping assigning a choice of the form $c_t^i(f, a_i) = \{g \in x_t(f) \mid g_i(t) = a_i\}$ (for some $a_i \in A_i$) to every move of the form $x_t(f)$. However, it is then possible to define strategies like the ones described in the introduction that induce no outcome or that induce a continuum of outcomes (cf. Examples 10 and 12 of Alós-Ferrer and Ritzberger, 2008). Hence, this approach, while intuitive, is not suited to model repeated

games in continuous time. In order to be able to “solve” these games, additional assumptions are needed.

A second approach is to view a continuous-time game as the limit of some sequence of discrete-time games and then define continuous-time strategies as limits of sequences of strategies in discrete time. This approach, however, presents difficulties of its own. A particular problem, pointed out by Davidson and Harris (1981) and Fudenberg and Levine (1986), is that sequences of discrete time strategies may not possess a limit (the “chattering problem”). Imagine, for instance, a sequence of discretizations with period length $1/n$ and discrete-time strategies prescribing to cooperate in periods k/n with k odd and defect in other periods.

A third approach is to restrict the sets of strategies in a game, that is, to impose the exogenous constraint that certain strategies cannot be used for e.g. equilibrium analysis. This allows to identify strategy sets which keep the framework tractable (e.g. guaranteeing existence and uniqueness of outcomes), and hence avoids the problems mentioned above. This approach has been pursued in Anderson (1984), Bergin (1992, 2006), Bergin and MacLeod (1993), Perry and Reny (1993), and Perry and Reny (1994), among others (see also Simon and Stinchcombe (1989) for a combination of this approach and discrete-time approximations). Stinchcombe (1992) investigates maximal strategy sets such that a unique outcome can be assigned to every admissible strategy profile, thereby obtaining a setting which is as good as it can be given a potentially problematic extensive form. As mentioned in the introduction, this approach presents conceptual problems because (behavioral) strategies are collections of local decisions, and which decisions are feasible should be solely and completely determined by the extensive form. However, it remains an open question whether the maximal strategy set approach can be reconciled with a pure extensive form approach. This would entail finding a new extensive form such that the unconstrained sets of behavioral strategies are equivalent, in a well-defined sense, to the set of constrained strategies. We will return to this question in Section 1.5.

1.2.3 The Action-Reaction Framework

Our approach to the problem of defining extensive form games in continuous time is different to the ones just mentioned. We will exhibit a specific extensive decision problem capturing repeated decisions in continuous time, such that, for every strategy profile, one and only one associated outcome exists. The basic construction relies on ideas present in the frameworks of Anderson (1984), Stinchcombe (1992), and Bergin (2006). However, the approach is different at a basic level because strategy sets are kept unconstrained; the differences with respect to the “direct approach”-EDP mentioned above are built directly into the construction of the extensive decision problem.

The basic idea of the construction is as follows. At time 0 all players choose a first action that they will have to stick to for some positive amount of time. This amount of time is determined by the choice of “inertia times” during which a player is committed to her current action. After this, whenever a player’s inertia time has run out she can revise her previous action. If she switches to a different one, i.e. “makes a jump”, the players who did not jump can react instantly and choose new actions as well. All players will again have to stick to their new actions for some positive amount of time, i.e. decide on new inertia times. This construction prevents players from jumping again right after an action change and from reacting even though no other player has jumped. The latter is crucial: a direct consequence is that the set of decision points becomes well-ordered, hence eliminating the problems of the direct approach.

We proceed in two steps. First we will describe the set of outcomes/plays of the game. The construction of this set already incorporates the essence of the Action-Reaction Framework. The second step is to appropriately define nodes and choices and show that the resulting structure is indeed an extensive decision problem.

The Outcome Space

Fix a finite set of players I and an arbitrary action space A_i for each player $i \in I$. Assume that A_i is a metric space.

We start by defining the set of plays, i.e. the possible maximal chains of decisions that might actually occur during the game. Ultimately, the history of all decisions taken by a player i builds a function $f_i : \mathbb{R}_+ \mapsto A_i$ as in the direct approach. We will introduce additional constraints to reflect the Action-Reaction Framework.

We require some preliminary notation. First, given a metric space B , call a function $g : \mathbb{R}_+ \mapsto B$ (right-)piecewise constant if for every $t \in \mathbb{R}_+$ there exists $\varepsilon > 0$ such that $g|_{]t, t+\varepsilon[}$ is constant. If g is piecewise constant, $g_+(t) := \lim_{\tau \rightarrow t+} g(\tau)$ exists for all $t \in \mathbb{R}_+$. In this case, define

$$RK(g) = \{t \in \mathbb{R}_+ \mid g_+(t) = g(t)\}$$

to be the set of points where g is right-constant. Second, given any function $g : \mathbb{R}_+ \mapsto B$, let

$$LC(g) = \left\{ t \in]0, +\infty[\mid \exists g_-(t) := \lim_{\tau \rightarrow t-} g(\tau) \wedge g_-(t) = g(t) \right\}$$

denote the set of points where g is left-continuous.⁴

The following definition spells out the first ingredient of our framework.

Definition 3. A *decision path* is a tuple $f = (f_i)_{i \in I}$ such that

(DP.i) for each $i \in I$, $f_i : \mathbb{R}_+ \mapsto A_i$ is piecewise constant,

(DP.ii) for each $i \in I$, $LC(f_i) \cup RK(f_i) = \mathbb{R}_+$,

(DP.iii) for each $t \in \mathbb{R}_+$, if $\exists i \in I$ with $t \in R(f_i)$, then $\exists j \in I$ with $t \in J(f_j)$,

where $J(f_i) := RK(f_i) \setminus LC(f_i)$ and $R(f_i) := LC(f_i) \setminus RK(f_i)$ are the set of *jump points* and *reaction points* of player i , respectively. The set of all decision paths is denoted by F .

Property (DP.i) states that a player's action revision cannot occur arbitrarily close to a previous action revision. A direct consequence (see Lemma

⁴For piecewise constant functions as defined here, a function is right-continuous at t if and only if it is right-constant at t .

A.2 in Section 1.7) is that the set of jump points of any player is well-ordered by the usual order on the real numbers. Property (DP.ii) requires that a player's action revision cannot take the form of an instantaneous change which is then abandoned (i.e. simultaneous failure of left- and right-continuity).⁵ Taken together, (DP.i) and (DP.ii) mean that when a player changes action, be it due to a jump or to a reaction to somebody else's revision, the player is not able to change action again immediately after the change.

Property (DP.iii) is the only condition requiring consistency across players' paths of decisions. Intuitively, jump points are those where a player's decision path has changed discontinuously (a sudden action revision), while t is a reaction point if the player's strategy shifts immediately after t but not at t , in reaction to an observed shift of another player at t : an "instant reaction". (DP.iii) states that a player can change action by instant reaction only if some other player jumped at t .

The second key ingredient of the Action-Reaction Framework are *inertia times*. By (DP.i), after every jump or reaction at t , there exists $\varepsilon > 0$ such that the player is "committed" not to revise action again until at least $t + \varepsilon$ (although a better interpretation is a physical impossibility to revise too often). We will introduce an explicit record of inertia times as part of every play. Formally, let E be the set of all possible functions $\epsilon = (\epsilon_i)_{i \in I}$ with $\epsilon_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\epsilon_i(0) > 0$ for all $i \in I$. The quantity $\epsilon_i(t)$ will play the role of a marker, with the interpretation that $\epsilon_i(t) > 0$ if and only if player i is able to revise her action at t . In that case, $\epsilon_i(t)$ represents the length of time after t for which player i cannot change action again, unless it is as reaction to some other action change.

Define the set of *decision points* of player i as

$$DP(\epsilon_i) := \{t \in \mathbb{R}_+ \mid \epsilon_i(t) > 0\},$$

i.e. the set of times at which player i is able to take a decision. In order to link inertia times with decision paths, we will have to spell out consistency

⁵In particular, (DP.ii) implies that $0 \in RK(f_i)$ for all $i \in I$, i.e. $f_+(0) = f(0)$. That is the players' initial decisions cannot be adjusted arbitrarily close to $t = 0$. Note that this implies $0 \in RK(f_i) \setminus LC(f_i) = J(f_i)$ for all $i \in I$.

conditions. A minimal such condition is that $J(f_i) \cup R(f_i) \subseteq DP(\epsilon_i)$, i.e. whenever a player makes a decision or reacts to another decision at time t , an inertia time $\epsilon_i(t) > 0$ is specified. However, the inclusion will typically be strict, since a player can always decide to keep the previous action, which still requires specifying a (new) inertia time. That is, $t \in DP(\epsilon_i)$ indicates a decision which might not be observable as such (because no action change ensues), while $t \in J(f_i) \cup R(f_i)$ implies an observable action change.

Before introducing the announced consistency conditions, again we require additional notation. Since $\epsilon_i(0) > 0$, for all $i \in I$, $\epsilon \in E$, and $t \in]0, +\infty[$ the intersection $DP(\epsilon_i) \cap]0, t[$ is not empty and hence by the Supremum Axiom we can define

$$\text{Prev}(\epsilon_i, t) := \sup(DP(\epsilon_i) \cap]0, t[),$$

which gives the last time before t that player i has taken a decision. Define $\text{Prev}(\epsilon_i, 0) = 0$ for all $i \in I$ and $\epsilon \in E$. For $i \in I$, $\epsilon \in E$, and $t \in DP(\epsilon_i)$ define

$$\text{Next}(\epsilon_i, t) := t + \epsilon_i(t),$$

which gives the next time after t that player i can initiate an action change if no other player jumps before. Further let

$$PJ(\epsilon_i) := \{t \in \mathbb{R}_+ \mid \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t)) = t\} \cup \{t \in \mathbb{R}_+ \mid t = \text{Prev}(\epsilon_i, t)\}$$

be the set of *potential jumps* for player i , i.e. the set of times where a player is allowed to *initiate* an action change according to the inertia times. Those are of two kinds. The “natural ones” are those where the inertia time since the last time an action change was implemented has “run out”. The second is slightly counterintuitive, and corresponds to points which are the supremum of the set of prior time points where action changes have been initiated, i.e. accumulation points of prior action changes.

Last, for $w = (f, \epsilon) \in F \times E$ and $t \in \mathbb{R}_+$ define (for notational convenience) $IDP(\epsilon, t) := \{i \in I \mid t \in DP(\epsilon_i)\}$, $IJ(f, t) := \{i \in I \mid t \in J(f_i)\}$, and $IPJ(\epsilon, t) := \{i \in I \mid t \in PJ(\epsilon_i)\}$, i.e. the sets of players having decision

points, jumps, and potential jumps at t , respectively.

We are now ready to define the set of plays, which incorporate the connection between decision paths and inertia times.

Definition 4. A *play* is a pair $w = (f, \epsilon) \in F \times E$ such that

(P.i) for each $i \in I$, $J(f_i) \subseteq PJ(\epsilon_i)$;

(P.ii) for each $i \in I$, $J(f_i) \subseteq \bigcap_{j \in I} DP(\epsilon_j)$;

(P.iii) for each $i \in I$, $PJ(\epsilon_i) \subseteq DP(\epsilon_i)$;

(P.iv) for each $i \in I$ and each $t \in DP(\epsilon_i)$ if $\tau \in DP(\epsilon_i) \cap]t, \text{Next}(\epsilon_i, t)[$ then $\bigcup_{j \neq i} J(f_j) \cap]t, \tau] \neq \emptyset$.

The set of all plays is denoted by W .

(P.i) states that a player can jump at t only if t was indeed a potential jump. (P.ii) means that, whenever a player jumps, every player who does not also jump is allowed to react, and all players have to specify inertia times. Note that (P.ii) together with (DP.iii) implies that $J(f_i) \cup R(f_i) \subseteq DP(\epsilon_i)$. (P.iii) requires that every potential jump be a decision point. The interpretation of (P.iv) is as follows. If at time t a player makes a decision with inertia time ϵ , then the only way he can make a decision *before* $t + \epsilon$ is if some other player jumped before $t + \epsilon$.

The Extensive Decision Problem

We first define the decision nodes, and hence the tree.

For every $w = (f, \epsilon) \in W$ and $t \in \mathbb{R}_+$, define the following sets

$$\begin{aligned} x_t(w) &= \{w' = (f', \epsilon') \in W \mid w'(\tau) = w(\tau) \forall \tau \in [0, t[\}, \\ x_t^R(w) &= \{w' = (f', \epsilon') \in x_t(w) \mid f'(t) = f(t) \}, \\ x_t^P(w) &= \{w' = (f', \epsilon') \in x_t^R(w) \mid f'_+(t) = f_+(t) \}. \end{aligned}$$

Nodes of the form $x_t(w)$ are “potential jump nodes” at which a player might make the decision to initiate a change of action. Hence, they will

be part of the tree whenever $t \in \bigcup_{i \in I} PJ(\epsilon_i)$ or, equivalently, whenever $IPJ(\epsilon, t) \neq \emptyset$.

Nodes of the form $x_t^R(w)$ are “reaction nodes” which model the possibility of players to react to a change of action initiated by another player. Hence they are part of the tree whenever $t \in \bigcup_{i \in I} J(f_i)$ but $t \notin J(f_j)$ for some $j \in I$; equivalently, whenever $\emptyset \subsetneq IJ(f, t) \subsetneq I$.

Nodes of the form $x_t^P(w)$ are “peek nodes” where both the actions at t (individual action change initiations) and the immediate reactions to them (the right limits of f), have already been decided, but the times $\epsilon_i(t)$ still have not. Again, they are part of the tree whenever $IPJ(\epsilon, t) \neq \emptyset$.

Note that nodes are independent of the “representant play”. If $w' \in x_t(w)$, then $x_t(w) = x_t(w')$, and analogously for reaction and peek nodes.

Potential jump, reaction, and peek nodes account for all possible decision situations. Note that the root, i.e. the node W containing all plays, is contained in N because $x_0(w) = W$ for all $w \in W$. The root is followed by peek nodes of the form $x_0^P(w)$. The set of nodes is given by

$$\begin{aligned} N = & \{x_t(w) \mid t \geq 0, IPJ(\epsilon, t) \neq \emptyset\} \\ & \bigcup \{x_t^R(w) \mid t > 0, \emptyset \subsetneq IJ(f, t) \subsetneq I\} \\ & \bigcup \{x_t^P(w) \mid t \geq 0, IPJ(\epsilon, t) \neq \emptyset\}. \end{aligned} \quad (1.1)$$

We now specify the choices, and hence the extensive decision problem by reviewing the decisions that have to be taken at each type of node. At potential jump nodes $x_t(w)$, players who are allowed to jump may decide how to continue, i.e. which action to adopt. That is, for every $t \geq 0$, $w = (f, \epsilon) \in W$, $i \in IPJ(\epsilon, t)$, and $a_i \in A_i$, we include the choice

$$\begin{aligned} c_i(x_t(w), a_i) = & \{w' = (f', \epsilon') \in W \mid t \in PJ(\epsilon'_i), f'(\tau) = f(\tau) \forall \tau \in [0, t], \\ & f'_i(t) = a_i\}. \end{aligned}$$

At reaction nodes $x_t^R(w)$, the players who did *not* jump decide on their instant reaction. That is, for every $t > 0$, $w = (f, \epsilon) \in W$ with $IJ(f, t) \neq \emptyset$,

$i \in I \setminus IJ(f, t)$, and $a_i \in A_i$, we include the choice

$$c_i(x_t^R(w), a_i) = \{w' = (f', \epsilon') \in W \mid f'(\tau) = f(\tau) \forall \tau \in [0, t], f'_{i+}(t) = a_i\}.$$

At peek nodes $x_t^P(w)$, all players who either had a potential jump at t or reacted at t decide how long they are going to stick to their action. That is, for every $t \geq 0$, $w = (f, \epsilon) \in W$, $i \in I$ such that $i \in IDP(\epsilon, t)$,⁶ and $\varepsilon_i \in \mathbb{R}_{++}$, we include the choice

$$c_i(x_t^P(w), \varepsilon_i) = \{w' = (f', \epsilon') \in W \mid f'(\tau) = f(\tau) \forall \tau \in [0, t], \\ f'_+(t) = f_+(t), \epsilon'_i(t) = \varepsilon_i\}.$$

Hence, the set of choices of player i is given by

$$C_i = \{c_i(x_t((f, \epsilon)), a_i) \mid t \geq 0, i \in IPJ(\epsilon, t), a_i \in A_i\} \\ \cup \{c_i(x_t^R((f, \epsilon)), a_i) \mid t > 0, i \in I \setminus IJ(f, t), IJ(f, t) \neq \emptyset, a_i \in A_i\} \\ \cup \{c_i(x_t^P((f, \epsilon)), \varepsilon_i) \mid t \geq 0, i \in IDP(\epsilon, t), \varepsilon_i \in \mathbb{R}_{++}\}.$$

Let us now look at information sets. By definition an information set in an EDP is the set of immediate predecessors of a given choice. For a choice $c = c_i(x_t(w), a_i) \in C_i$ with $w = (f, \epsilon)$ we obtain

$$P(c) = \{x_t((f', \epsilon')) \in N \mid f(\tau) = f'(\tau) \forall \tau \in [0, t[, t \in PJ(\epsilon'_i)\}.$$

This means that at a potential jump node $x_t(w)$ a player knows all past actions (i.e. the decision path up to time t) but not the record of inertia times which has led to the particular decision path (with the obvious exception that she knows that the play is such that she is allowed to jump).

For a choice $c = c_i(x_t^R(w), a_i) \in C_i$ with $w = (f, \epsilon)$ (which implies $i \notin IJ(f, t)$) we have

$$P(c) = \{x_t^R((f', \epsilon')) \in N \mid f'(\tau) = f(\tau) \forall \tau \in [0, t]\},$$

⁶This is equivalent to $i \in IPJ(\epsilon, t)$ or $IJ(f, t) \neq \emptyset$ (see Lemma A.5 in the appendix).

i.e. at a reaction node $x_t^R(w)$ the player knows the decision path up to and including time t .

Finally, for a choice $c = c_i(x_t^P(w), \varepsilon_i) \in C_i$ with $w = (f, \epsilon)$ we obtain

$$P(c) = \{x_t^P((f', \epsilon')) \in N \mid f'(\tau) = f(\tau) \forall \tau \in [0, t], f'_+(t) = f_+(t), \\ t \in DP(\epsilon'_i) \}$$

which means that at a peek node $x_t^P(w)$ the player knows the decision path up to and including time t , as well as what all players are “going to do next”, i.e. the right limits at t , and that she took a decision at t (which cannot necessarily be inferred from the decision path).

This completes the specification of the framework. Denote $T := (N, \supseteq)$ and $C := (C_i)_{i \in I}$. We call the pair (T, C) the *Action-Reaction Framework*.

Proposition 1. *The Action-Reaction Framework (T, C) is an extensive decision problem.*

To define an extensive form game on the EDP capturing the Action-Reaction Framework, all what is left is a specification of individual preferences on plays. Plays, however, contain a full specification of inertia times, which are essential to capture the idea that an action initiation cannot occur arbitrarily close to a previous one (as also assumed in Bergin, 1992; Stinchcombe, 1992; Perry and Reny, 1993) but should ultimately be payoff-irrelevant. Hence, one can define a repeated game in continuous time as an EDP as above together with a specification of preferences on plays which does not depend on inertia times, e.g. if utilities on $w = (f, \epsilon)$ only depend on the first argument. As we will clarify below, the information sets described above guarantee that players’ choices only depend on decision paths and not on inertia times.

1.3 A Possibility Result

In this section we aim to show that the framework we have introduced is well-suited to the analysis of repeated games in continuous time. For that,

we need to establish that it is better behaved than general EDPs, since being an EDP does not guarantee that well-specified strategy profiles lead to well-specified outcomes. Fortunately, the conditions guaranteeing outcome existence and uniqueness are already known. We now review them for the general case and then return to our framework.

1.3.1 Strategies and Outcomes in General Extensive Form Games

Given an extensive decision problem, let $X_i := \{x \in X \mid \exists c \in C_i : x \in P(c)\}$ be the *set of moves for player i* , for every $i \in I$.

A *pure strategy* for player $i \in I$ is a function $s_i : X_i \rightarrow C_i$, such that

$$s_i^{-1}(c) = P(c) \text{ for all } c \in s_i(X_i)$$

where $s_i(X_i) \equiv \{s_i(x) \mid x \in X_i\}$.

That is, the function s_i assigns to every move $x \in X_i$ a choice $c \in C_i$ such that (a) choice c is *available* at x , i.e. $s_i(x) = c \Rightarrow x \in P(c)$ or $s_i^{-1}(c) \subseteq P(c)$, and (b) to every move x in an *information set* $P(c)$ the same choice gets assigned, i.e. $x \in P(c) \Rightarrow s_i(x) = c$ or $P(c) \subseteq s_i^{-1}(c)$, for all $c \in C_i$ that are chosen somewhere, viz. $c \in s_i(X_i)$. Let S_i denote the set of all pure strategies for player $i \in I$. A *pure strategy combination* is an element $s = (s_i)_{i \in I} \in S \equiv \times_{i \in I} S_i$.

We want to obtain a framework where every strategy combination induces an outcome/play. Hence, we need to clarify the formal meaning of when a pure strategy combination “induces” a play. Define, for every $s \in S$, the correspondence $R_s : W \rightarrow W$ by

$$R_s(w) = \bigcap \{s_i(x) \mid w \in x \in X, i \in I(x)\}.$$

Say that strategy combination s *induces* the play w if $w \in R_s(w)$, i.e. if it is a fixed point of R_s .

In an arbitrary EDP the correspondence R_s for a given strategy combination $s \in S$ may not have a fixed point at all, or have a whole continuum

thereof. The two basic desiderata on an EDP, expressed in terms of R_s , are as follows.

(A1) For every $s \in S$ there is some $w \in W$ such that $w \in R_s(w)$.

(A2) If for $s \in S$ there is $w \in W$ such that $w \in R_s(w)$, then R_s has no other fixed point and $R_s(w) = \{w\}$.

(A1) says that for every strategy combination $s \in S$ there is an outcome/play $w \in W$ that is induced by s . (A2) requires that the induced outcome is unique. (A1) and (A2) define a function $\phi : S \rightarrow W$ that associates a unique play to each pure strategy combination. (Furthermore, this function is onto by Theorem 4 of Alós-Ferrer and Ritzberger, 2008). These two properties are, therefore, necessary and sufficient to define a normal form (without payoffs).

The main result of Alós-Ferrer and Ritzberger (2008) states that (A1) and (A2) are essentially equivalent to two properties of the tree: “regularity” and “up-discreteness.” Thus, these two properties represent the appropriate restrictions on game trees for a well-founded sequential decision theory.

Definition 5. A game tree (N, \supseteq) is *regular* if $\uparrow x \setminus \{x\}$ has an infimum for every $x \in N$, $x \neq W$. It is *up-discrete* if every (nonempty) chain in N has a maximum.

In the terminology of Alós-Ferrer and Ritzberger (2008), regularity means that there are no *strange* nodes, or, equivalently, that every node other than the root is either *finite* (meaning that it has an immediate predecessor) or *infinite*, meaning that it coincides with the infimum of its strict predecessors. Up-discreteness is equivalent to the chains $\uparrow x$ for $x \in N$ being dually well-ordered (that is, all their subsets have a maximum). This condition is common in order theory and theoretical computer science (see Koppelberg, 1989, chp. 6). It implies that the set of *immediate successors* of a move is nonempty and forms a partition of the move by finite nodes.

Intuitively, up-discreteness should exclude continuous-time examples, since immediate successors can be seen as “the next” decision points. It turns out,

however, that the Action-Reaction Framework fulfills up-discreteness in spite of being a model for decisions in continuous time.

1.3.2 Strategies and Outcomes in the Action-Reaction Framework

In (T, C) the sets of moves and the sets of choices are fixed. As described above this specifies the set of strategies for each player since strategies in an EDP are mappings from the set of moves to the set of choices. Hence in the Action-Reaction Framework there is no freedom in the specification of strategies and in particular players cannot be prevented from using certain strategies. All restrictions on the players' ways to act are already incorporated in the tree and the choice system respectively. Note that due to the structure of the information sets the choices prescribed by strategies only depend on decision paths and not on inertia times.

We denote the set of strategies of player i in the Action-Reaction Framework by S_i . Let further $S := \times_{i \in I} S_i$ denote the set of strategy profiles in (T, C) .

Lemma 1. *The tree of the Action-Reaction Framework is an up-discrete and regular tree.*

By Proposition 1 above and Theorem 4 in Alós-Ferrer and Ritzberger (2008) any decision path in W can be reached by some profile of strategies. Using Proposition 1 above, Lemma 1, and Propositions 6(b) and 9 in Alós-Ferrer and Ritzberger (2008) we obtain that (T, C) is an *Extensive Form* (Alós-Ferrer and Ritzberger, 2008; Alós-Ferrer, Kern, and Ritzberger, 2011). Corollary 5(b) from Alós-Ferrer, Kern, and Ritzberger (2011) then yields the following result.

Theorem 1. *Every strategy profile in the Action-Reaction Framework induces one and only one outcome.*

1.4 An Alternative Approach: Strategy Constraints

In the previous sections, we have established that it is possible to define extensive form games modeling continuous-time problems without the recourse to an artificially constrained strategy set. It is, however, natural to ask whether there is a relation between the Action-Reaction Framework and previous approaches which employed strategy constraints. Indeed, it is possible to embody ideas similar to the ones in the Action-Reaction Framework through strategy constraints. In this section we detail this alternative route and show how these constraints must be imposed to preserve equivalence (in a well-defined sense to be detailed below) with the extensive form approach.

Informally, a *Conditional Response Mapping* is a mapping which specifies, at each time t , an *action* (depending only on the previous history of play) and a *response* which depends on the actions being simultaneously decided by other players. A number of additional conditions must be imposed in order to capture the constraints which are also inherent in the Action-Reaction Framework. Naturally, these additional conditions resemble the restrictions imposed on strategies by e.g. Stinchcombe (1992) and Bergin (2006), among others (see Section 1.6). The reason we refrain from using the term *strategy* is that a priori it is not clear whether the set of Conditional Response Mappings indeed corresponds to the set of strategies in a well-defined extensive form. We shall, however, see that this is the case.

Analogously to the conditions discussed for extensive forms, a coherent framework will be obtained if every profile of mappings induces an outcome contained in the appropriate outcome set and any outcome can be reached by some profile. In order to guarantee these properties, however, it is not sufficient to place restrictions on Conditional Response Mappings only. It is necessary to also constrain the set of possible outcomes, and hence (through the dependence on histories) the domain of these mappings. The appropriate constraints for the set of outcomes are exactly as in the Action-Reaction Framework: outcomes must define decision paths.

Let F denote the set of decision paths as introduced in Definition 3. The formal definition of Conditional Response Mappings is as follows.

Definition 6. A *Conditional Response Mapping* (CRM) for player $i \in I$ is a mapping $\sigma_i : F \times \mathbb{R}_+ \rightarrow A_i^2$, $(f, t) \mapsto (\sigma_i^1(f, t), \sigma_i^2(f, t))$ such that for every $f \in F$ and all $t \in \mathbb{R}_+$

- (CRM.i) if $f(\tau) = f'(\tau)$ for $f' \in F$ and all $\tau \in [0, t[$, then $\sigma_i^1(f, t) = \sigma_i^1(f', t)$;⁷
if $f(\tau) = f'(\tau)$ for $f' \in F$ and all $\tau \in [0, t]$, then $\sigma_i^2(f, t) = \sigma_i^2(f', t)$.
- (CRM.ii) if $t \in \left(\bigcap_{j \in I} LC(f_j) \right) \cup J(f_i)$ then $\sigma_i^2(f, t) = f_i(t)$ and there is $\varepsilon_i(f, t) > 0$ such that $\sigma_i^1(f, \tau) = f_i(t)$ for all $\tau \in]t, t + \varepsilon_i(f, t)[$.
- (CRM.iii) if $t \in LC(f_i) \cap \left(\bigcup_{k \in I} J(f_k) \right)$ then there is $\varepsilon_i(f, t) > 0$ such that $\sigma_i^1(f, \tau) = f_{i+}(t)$ for all $\tau \in]t, t + \varepsilon_i(f, t)[$.

Denote the set of CRMs for player i by Σ_i and let $\Sigma := \times_{i \in I} \Sigma_i$.

For each decision path f and time t , a CRM hence specifies an action, denoted $\sigma_i^1(f, t)$, and an instant response $\sigma_i^2(f, t)$. The first part of condition (CRM.i) specifies that actions depend only on the past history of play, i.e. on the values of f up to (but excluding) t . The second part of this condition stipulates that responses depend only on the values of f up to *and including* t . Equivalently, at time t each player specifies an action and, for any possible profile of actions at t which is part of a decision path, also a conditional response.

Condition (CRM.ii) captures the intuition that, as long as no player has changed action at t (and hence the decision path is left-continuous in all coordinates), then no player can change the current action through a conditional response. That is, “no reaction without a triggering action”. Further, players will be constrained to the current action for some small time interval. Likewise, the same restrictions apply if a given player has changed action at t (“jumped”), which embodies the intuition that two action changes of a given player cannot be arbitrarily close. In particular, all players have to stick to the action picked at time 0 for some positive amount of time.

⁷In particular $\sigma_i^1(f', 0) = \sigma_i^1(f, 0)$ for all $f' \in F$.

Condition (CRM.iii) captures a similar intuition for responses. If a player did not initiate an action change at t , but some other player did, then the original player was allowed to react through the stipulated conditional response (hence no constraint is placed on the second component of the action tuple). The condition requires that the player needs to stick to the action specified as a response for some small time interval, as long as no other player initiates an action change.⁸

The following examples illustrate that the restriction to decision paths is necessary. In other words, CRMs are well-defined mappings only on $F \times \mathbb{R}_+$. The first example shows that a CRM cannot be built by simply gluing together arbitrary chains of decisions.

Example 1. Let $I = \{1\}$ and $A_1 = \{0, 1\}$ and consider the function $h : \mathbb{R}_+ \rightarrow A_1$ defined by

$$h(\tau) = \begin{cases} 0, & \text{if } \tau \in \mathbb{Q}, \\ 1, & \text{if } \tau \in \mathbb{R} \setminus \mathbb{Q}. \end{cases}$$

By (CRM.ii), for any $t \in \mathbb{Q}$, there should exist an $\varepsilon > 0$ such that $\sigma^1(h, \tau) = 0$ for all $\tau \in]t, t + \varepsilon[$. Also by (CRM.ii) for any $t' \in \mathbb{R} \setminus \mathbb{Q} \cap]t, t + \varepsilon[$ there should exist an $\varepsilon' > 0$ such that $\sigma^1(h, \tau) = 1$ for all $\tau \in]t', t' + \varepsilon'[$, which leads to a contradiction.

In this example, the decision maker changes action “too often”, with action changes being arbitrarily close to each other. The next example shows that this problem also arises with more intuitive, “continuous” mappings.

Example 2. Let $I = \{1\}$ and $A_1 = \mathbb{R}_+$. Consider the function $h : \mathbb{R}_+ \rightarrow A_1$ defined by $h(\tau) = \tau$ for all $\tau \in \mathbb{R}_+$. By (CRM.ii) for any $t \in \mathbb{R}_+$, there should exist an $\varepsilon > 0$ such that $\sigma^1(h, \tau) = t$ for all $\tau \in]t, t + \varepsilon[$. Also by (CRM.ii) for any $t' \in]t, t + \varepsilon[$ there should exist an $\varepsilon' > 0$ such that $\sigma^1(h, \tau) = t' > t$ for all $\tau \in]t', t' + \varepsilon'[$, which leads to a contradiction.

The last two conditions in the definition of CRM implicitly incorporate a notion of “inertia” analogous to the inertia times needed for the Action-Reaction Framework. A first step in order to show that the new framework

⁸By (DP.iii), t fulfills either the hypothesis of (CRM.ii) or the hypothesis of (CRM.iii).

is coherent is to make the relationship to inertia times explicit. This corresponds to the following thought experiment. Given a decision path f and a time instant t , imagine the path after t was changed in such a way that nobody changed action after the reactions specified at time t , i.e. the path was fixed at $f_+(t)$. What is the first point in time after t such that a given CRM σ_i would specify a deviation from the new path? If such a first point in time is well defined and equal to $t + \varepsilon$, the quantity ε will fulfill the same role as an inertia time in the Action-Reaction Framework.

Let us formally construct these inertia times for a given CRM $\sigma_i \in \Sigma_i$. For each $f \in F$ and $t \in \mathbb{R}_+$ let f^{t+} be given by $f^{t+}(\tau) = f(\tau)$ for all $\tau \leq t$ and $f^{t+}(\tau) = f_+(t)$ for all $\tau > t$. We call $\varepsilon > 0$ a *deviation point prescribed by σ_i after (f, t)* if $\sigma_i^1(f^{t+}, t+\varepsilon) \neq f_{i+}(t)$. That is, for a deviation point ε , the action prescribed by σ_i at time $t + \varepsilon$ is different from the action/reaction chosen by player i at time t , given that all players stick to their actions/reactions chosen at time t . The following lemma shows that whenever such a deviation point exists, there is a *first* deviation point, which then plays the role of an inertia time.

Lemma 2. *Let $f \in F$, $t \in \mathbb{R}_+$, $i \in I$, $\sigma_i \in \Sigma_i$, and let $E^{\sigma_i}(f, t)$ be the set of all deviation points prescribed by σ_i after (f, t) . If $E^{\sigma_i}(f, t) \neq \emptyset$ then there exists a first deviation point $\varepsilon^{\sigma_i}(f, t) = \min E^{\sigma_i}(f, t)$.*

This property is a consequence of condition (CRM.ii). It should be remarked that the existence of a first deviation point corresponds to the ‘‘Identifiability’’ assumption for admissible strategies imposed by Stinchcombe (1992). The difference is that in Stinchcombe (1992), this property is imposed as one of the conditions restricting the strategy set, while in our framework, it is a property derived from the definition of CRM.

The existence of first deviation points as identified in Lemma 2 is crucial for the framework at hand. It has two important consequences. First, it plays a major role in the proof of outcome existence and uniqueness below. Second, and as already announced, the $\varepsilon^{\sigma_i}(f, t)$ essentially reconstruct inertia times and will allow us to establish the equivalence between the Action-Reaction Framework and the framework based on CRMs.

The next result shows that this framework is coherent, that is, every profile of CRMs induces a unique outcome (after every history) and any outcome in F can be reached by some profile of CRMs.

Definition 7. Let $\sigma \in \Sigma$.

- (i) $f \in F$ is induced by σ if $\sigma_i^1(f, t) = f_i(t)$ and $\sigma_i^2(f, t) = f_{i+}(t)$ for all $i \in I$ and all $t \in \mathbb{R}_+$.
- (ii) Given $\bar{f} \in F$ and $t \in \mathbb{R}_{++}$, $f \in F$ is induced by σ after (\bar{f}, t) if $f(\tau) = \bar{f}(\tau)$ for all $\tau \in [0, t[$, $\sigma_i^1(f, \tau) = f(\tau)$ and $\sigma_i^2(f, \tau) = f_{i+}(\tau)$ for all $\tau \in [t, +\infty[$ and $i \in I$.

Proposition 2. (i) Every $\sigma \in \Sigma$ induces a unique $f \in F$.

(ii) For all $\bar{f} \in F$ and $t \in \mathbb{R}_{++}$ every $\sigma \in \Sigma$ induces a unique $f \in F$ after (\bar{f}, t) .

(iii) Every $f \in F$ is induced by some $\sigma \in \Sigma$.

Properties (i) and (ii) in the last Proposition are comparable to Theorem IV.1 in Stinchcombe (1992), Theorem 3 in Bergin (1992), Lemma A.1 in Perry and Reny (1993), Theorem 2 in Bergin and MacLeod (1993), and Theorem 1 in Bergin (2006). All these results state that, under the constraints of the respective framework, every profile of strategies induces a unique outcome after any history. Property (iii) additionally states that any outcome can be reached by some profile of strategies, a result similar to Theorem IV.2 in Stinchcombe (1992).

The intuition behind the proof of the last result is as follows. Given a profile of CRMs, initial actions are clear. The inertia times identified in Lemma 2 then allow us to identify the (constant) path up to the next deviation point. At that point, the CRMs can be used to establish the new actions/reactions. Applying Lemma 2, the construction can be iterated. Since time is continuous, the exact iterative argument relies on transfinite recursion, which is made possible by the structure of decision paths.

1.5 An Equivalence Result

Proposition 2 shows that the framework based on CRMs is coherent, in the sense that every profile of CRMs induces one and only one outcome. Coherence of the framework, however, is just a necessary prerequisite for extensive form analysis. In this section, we show that CRMs indeed allow for a full-fledged extensive form formulation. In particular, they are shown to be equivalent to the Action-Reaction Framework. Hence, CRMs represent the “translation” into constrained strategy sets of a proper extensive form game modeling continuous-time decisions.

1.5.1 Outcome-Equivalence and Equivalence Classes

We need some additional notation. For $f \in F$ and $t \in \mathbb{R}_{++}$ let $t(f, t) := \sup \bigcup_{i \in I} J(f_i) \cap [0, t[$ and further let $t(f, 0) = 0$. Note that $\bigcup_{i \in I} J(f_i) \cap [0, t[\neq \emptyset$ for $t \in \mathbb{R}_{++}$ as $0 \in J(f_i)$ for all $i \in I$. The time point $t(f, t)$ is essentially the last time strictly before t that some player jumped. In particular (by Lemma C.1 in the appendix), if $t(f, t) < t$ then f is constant on $]t(f, t), t[$. Note, however, that $t' = t(f, t)$ might be an accumulation point of jumps. In this case, necessarily $t(f, t') = t'$. Let $\mathcal{J}(f) := \bigcup_{i \in I} J(f_i) \cup \{t \in \mathbb{R}_+ | t = t(f, t)\}$. Note that $t(f, t) \in \mathcal{J}(f)$ for all $t \in \mathbb{R}_+$.

In order to show the equivalence between the Action-Reaction Framework and the approach based on CRMs, we need to associate a CRM to each strategy $s_i \in S_i$ in the Action-Reaction Framework. The idea is as follows. The construction of a CRM requires to prescribe an action and a reaction for every history-time pair. The structure of the EDP allows for a natural way to define actions $a_i(f, t, s_i)$ and reactions $a_i^R(f, t, s_i)$ that only depend on the history-time pair (f, t) and the strategy s_i . The only difficulty is to determine the set $M(f, s_i) \subseteq \mathbb{R}_+$ of time points t such that the strategy s_i and the past decisions along f imply that i actually has to make a decision at t . Once this is in place, a CRM can be defined by prescribing the action $a_i(f, t, s_i)$ whenever player i has to make a decision, and the left limit of past actions if not (reactions are determined in a similar manner).

First, we identify the natural actions. Given $s_i \in S_i$ and $(f, t) \in F \times \mathbb{R}_+$ let $a_i(f, t, s_i)$ be the action such that $s_i(x_t((f', \epsilon'))) = c_i(x_t((f', \epsilon')), a_i(f, t, s_i))$ for any $(f', \epsilon') \in W$ such that $f'(\tau) = f(\tau)$ for all $\tau \in [0, t[$ and $t \in PJ(\epsilon'_i)$ (provided, of course, some such potential jump node exists). Note that by construction of the EDP, for any $(f'', \epsilon'') \in W$ such that $f''(\tau) = f(\tau)$ for all $\tau \in [0, t[$ and $t \in PJ(\epsilon''_i)$, $x_t(f'', \epsilon'')$ belongs to the same information set as $x_t((f', \epsilon'))$ and hence $s_i(x_t((f', \epsilon'))) = s_i(x_t((f'', \epsilon'')))$. Thus $a_i(f, t, s_i)$ is uniquely determined by the strategy s_i , the time point t , and the decision path f up to t .

Now, analogously to the last paragraph, we determine the natural reactions. Given $s_i \in S_i$ and $(f, t) \in F \times \mathbb{R}_+$ let $a_i^R(f, t, s_i)$ be the action such that $s_i(x_t^R((f', \epsilon'))) = c_i(x_t^R((f', \epsilon')), a_i^R(f, t, s_i))$ for any $(f', \epsilon') \in W$ such that $f'(\tau) = f(\tau)$ for all $\tau \in [0, t]$ and $i \notin IJ(f', t) \neq \emptyset$ (provided, of course, such a reaction node exists). Note that by construction of the EDP, for any $(f'', \epsilon'') \in W$ such that $f''(\tau) = f(\tau)$ for all $\tau \in [0, t]$ and $i \notin IJ(f'', t) \neq \emptyset$, $x_t^R(f'', \epsilon'')$ belongs to the same information set as $x_t^R((f', \epsilon'))$ and hence $s_i(x_t^R((f', \epsilon'))) = s_i(x_t^R((f'', \epsilon'')))$. Thus $a_i^R(f, t, s_i)$ is uniquely determined by the strategy s_i , the time point t , and the decision path f up to and including t .

Now, we proceed to identify the set $M(f, s_i)$ of time points where player i needs to move given f and s_i . The construction of $M(f, s_i)$ requires a definition and a lemma.

Definition 8. Let $i \in I$, $s_i \in S_i$, and $(f, \epsilon) \in W$. For $t_1, t_2 \in \mathbb{R}_+ \cup \{\infty\}$, $t_1 < t_2$, (f, ϵ) agrees with s_i on $[t_1, t_2[$ if for all $\tau \in [t_1, t_2[$,

$$\epsilon_i(\tau) > 0 \Rightarrow s_i(x_\tau^P((f, \epsilon))) = c_i(x_\tau^P((f, \epsilon)), \epsilon_i(\tau)). \quad (1.2)$$

Lemma 3. Let $i \in I$, $s_i \in S_i$, $f \in F$. Then

- (i) for any $t \in \mathcal{J}(f)$ there is $\epsilon \in E$ such that $(f^{t+}, \epsilon) \in W$ and it agrees with s_i on $[t, \infty[$ (in particular, $\epsilon_i(t) > 0$);
- (ii) for any $t \in \mathcal{J}(f)$, if $(f^{t+}, \epsilon), (f^{t+}, \epsilon') \in W$ agree with s_i on $[t, \infty[$ then $\epsilon_i(\tau) = \epsilon'_i(\tau)$ for all $\tau \in [t, \infty[$;

(iii) for any $t \in \mathbb{R}_+$, if $(f^{t(f,t)^+}, \epsilon), (f^{t(f,t)^+}, \epsilon') \in W$ agree with s_i on $[t(f,t), \infty[$ then $PJ(\epsilon_i) \cap [t, \infty[= PJ(\epsilon'_i) \cap [t, \infty[$.

Given any $s_i \in S_i$ and $(f, t) \in F \times \mathbb{R}_+$, we define $\epsilon_i(f, t, s_i) : [t(f, t), \infty[\rightarrow \mathbb{R}_+$ as the unique function given by Lemma 3. That is, if $(f^{t(f,t)^+}, \epsilon) \in W$ agrees with s_i on $[t(f, t), \infty[$ then $\epsilon_i(\tau) = \epsilon_i(f, t, s_i)(\tau)$ for all $\tau \in [t(f, t), \infty[$ and $\epsilon_i(f, t, s_i)(t(f, t)) > 0$. Further, for any $s_i \in S_i$ and $(f, t) \in F \times \mathbb{R}_+$ we define (abusing notation) $PJ(f, t, s_i) = PJ(\epsilon'_i) \cap [t, \infty[$ for any $\epsilon' \in E$ such that $(f^{t(f,t)^+}, \epsilon') \in W$ and agrees with s_i on $[t(f, t), \infty[$. This is well-defined by Lemma 3(iii). Finally, let $M(f, s_i) = \{t \in \mathbb{R}_+ | t \in PJ(f, t, s_i)\}$. The intuition for $M(f, s_i)$ is as follows. Since no player jumps between $t(f, t)$ and t , one can uniquely reconstruct the inertia times chosen by player i between $t(f, t)$ and t according to s_i (Lemma 3(i) and (ii)). This yields a sequence of time points between $t(f, t)$ and t at which player i has to move. If this sequence either includes t or “converges” to it then $t \in M(f, s_i)$.

The considerations above allow us to construct a well-defined CRM given a strategy in (T, C) as follows. Given $i \in I$ and a strategy $s_i \in S_i$ define $\sigma^{s_i} : F \times \mathbb{R}_+ \rightarrow A_i^2$ by

$$\sigma^{s_i, 1}(f, t) := \begin{cases} a_i(f, t, s_i) & \text{if } t \in M(f, s_i), \\ f_{i-}(t) & \text{if } t \notin M(f, s_i) \end{cases}$$

and

$$\sigma^{s_i, 2}(f, t) := \begin{cases} a_i^R(f, t, s_i), & \text{if } t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i), \\ f_i(t), & \text{if } t \in \bigcap_{j \in I} LC(f_j) \cup J(f_i). \end{cases}$$

The intuition behind this construction is as follows. Given a history-time pair (f, t) one first checks whether past decisions prescribe that i should make a decision at t , i.e. whether $t \in M(f, s_i)$. If this is the case, the action chosen at (f, t) is the unique action prescribed by s_i at the corresponding potential jump node. If not, the left-continuous action is chosen.⁹ Reactions are chosen according to the uniquely prescribed reactions at the

⁹If $t \notin M(f, s_i)$, then $t(f, t) < t$ by (P.ii) and the definition of a potential jump. Since f is constant on $]t(f, t), t[$ by Lemma C.1 in the appendix, $f_{i-}(t)$ exists.

corresponding reaction nodes. Note that $t \in M(f, s_i)$ implies that there is $\epsilon \in E$ such that $(f^{t(f,t)^+}, \epsilon) \in W$ and $t \in PJ(\epsilon_i)$ and consequently that $x_t((f^{t(f,t)^+}, \epsilon)) \in X_i$. Hence, in particular $a_i(f, t, s_i)$ is well-defined. Analogously, if $t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ then $i \notin IJ(f, t) \neq \emptyset$. Hence by construction of the game tree (and Lemma C.2 in the appendix), there is $\epsilon \in E$ such that $(f, \epsilon) \in W$ and $x_i^R((f, \epsilon)) \in X_i$ which guarantees that $a_i^R(f, t, s_i)$ is well-defined. As the next proposition shows, the mappings above indeed define CRMs.

Proposition 3. *Let $i \in I$, and $s_i \in S_i$. Then σ^{s_i} is a CRM.*

Given a profile $s \in S$ of strategies in (T, C) , denote by $w^s = (f^s, \varepsilon^s)$ the play induced by s (recall Section 1.3.1) and say that f^s is the decision path induced by s .

The structure of (T, C) allows for a natural way to define an equivalence relation on the set of a player's strategies, which will be used in the sequel.

Definition 9. Let $i \in I$. Two strategies $s_i^1, s_i^2 \in S_i$ are *outcome-equivalent*, $s_i^1 \sim s_i^2$, if they induce the same decision path for any given profile of the other players' strategies, that is $f^{(s_i^1, s_{-i})} = f^{(s_i^2, s_{-i})}$ for all $s_{-i} \in S_{-i}$.

For $s_i \in S_i$ denote the equivalence class of s_i with respect to \sim by $[s_i]$ and let S_i/\sim be the set of equivalence classes.

As the next lemma shows, if two strategies induce the same CRM they are outcome-equivalent.

Lemma 4. *Let $i \in I$ and $s_i^1, s_i^2 \in S_i$. If $\sigma^{s_i^1} = \sigma^{s_i^2}$ then $s_i^1 \sim s_i^2$.*

1.5.2 Equivalence of CRM and Action-Reaction Framework

We now proceed to show that the Action-Reaction Framework and the approach using CRMs are equivalent. Proposition 3 establishes the existence of a well-defined mapping from the set of strategies in (T, C) to the set of CRMs. Next, we construct a mapping from the set of CRMs to the set of (equivalence classes of) strategies in (T, C) and subsequently show that any

profile of strategies in (T, C) induces that same decision path as the associated profile of CRMs and conversely that any profile of CRMs induces the same decision path as the associated strategy profiles in (T, C) .

We first note that every CRM defines a set of associated strategies in (T, C) in a natural way.

Definition 10. Let $\sigma_i \in \Sigma_i$. A strategy $s_i : X_i \rightarrow C_i$ is induced by σ_i if

- (IS.i) $s_i(x_t(w)) = c_i(x_t(w), \sigma_i^1(f, t))$ for all potential jump nodes $x_t(w) \in X_i$ where $w = (f, \epsilon)$,
- (IS.ii) $s_i(x_t^R(w)) = c_i(x_t^R(w), \sigma_i^2(f, t))$ for all reaction nodes $x_t^R(w) \in X_i$ where $w = (f, \epsilon)$,
- (IS.iii) $s_i(x_t^P(w)) = c_i(x_t^P(w), \varepsilon^{\sigma_i}(f, t))$ for all peek nodes $x_t^P(w) \in X_i$ such that $E^{\sigma_i}(f, t) \neq \emptyset$, where $w = (f, \epsilon)$.

Let $S(\sigma_i)$ be the set of all strategies $s_i : X_i \rightarrow C_i$ that are induced by σ_i .

This definition is, for all practical purposes, constructive. At each potential jump node (reaction node) the action prescribed is the action (reaction) chosen by the CRM after the corresponding history-time pair. At peek nodes, the inertia time chosen is the length of the period until the next jump prescribed by the CRM. The only part of the definition which allows for some freedom in the specification of choices corresponds to history-time pairs after which the CRM does not prescribe a jump if no other player jumps. Formally, one then has $s_i(x_t^P((f, \epsilon))) = c_i(x_t^P((f, \epsilon)), \varepsilon)$ for some arbitrary $\varepsilon > 0$ for all peek nodes $x_t^P((f, \epsilon)) \in X_i$ where $E^{\sigma_i}(f, t) = \emptyset$. The next proposition shows that this construction indeed delivers a set of strategies in (T, C) for any given CRM.

Proposition 4. Let $i \in I$, and $\sigma_i \in \Sigma_i$.

- (i) $S(\sigma_i) \neq \emptyset$.
- (ii) If $s_i, s'_i \in S(\sigma_i)$ then $s_i \sim s'_i$.

For a profile $\sigma \in \Sigma$ of CRMs let f^σ denote the outcome induced by σ (recall Definition 7 and Proposition 2).

Theorem 2. *Let $\sigma = (\sigma_i)_{i \in I} \in \Sigma$ be a CRM profile and $s = (s_i)_{i \in I} \in S$ be a strategy profile in (T, C) .*

- (i) $\sigma^{s'_i} = \sigma_i$ for all $i \in I$ and $s'_i \in S(\sigma_i)$.
- (ii) $[s_i] = [s'_i]$ for all $s'_i \in S(\sigma^{s_i})$.
- (iii) $f^{(s'_i)_{i \in I}} = f^\sigma$ for all $(s'_i)_{i \in I} \in \times_{i \in I} S(\sigma_i)$.
- (iv) $f^{(\sigma^{s_i})_{i \in I}} = f^s$.

What the theorem states is the following. By (i) when going from a CRM to a corresponding strategy in the EDP and then from that strategy to the corresponding CRM one obtains the original CRM. Part (ii) says that one obtains an outcome-equivalent strategy when going from a strategy in the EDP to the corresponding CRM and then to a strategy corresponding to that CRM. In (iii) we show that the outcome induced by a profile of CRMs coincides with the decision path of the play induced by any corresponding strategy profile in (T, C) . Part (iv) is the analogous statement for a profile of strategies in (T, C) . These properties show that the approach using the Action-Reaction Framework and the approach using CRMs are equivalent.

1.6 Relation to the Literature

In this section we discuss several frameworks for games in continuous time that have been suggested in the literature, and comment on the relation of those to our approach.

1.6.1 Maximal Strategy Sets

In a remarkable paper, Stinchcombe (1992) proposed a two-step approach in order to obtain a coherent framework for the analysis of continuous-time decision problems. His first step is to reduce the set of possible outcomes, and hence the underlying extensive form. The second step is to restrict the class of “admissible” strategies on that extensive form. This approach is thus

located “in between” the Action-Reaction Framework and the approach using strategy constraints. On the one hand the set of possible outcomes of the game is restricted and a decision tree is used. On the other hand the players’ strategy sets are exogenously restricted.

The construction in Stinchcombe (1992) is as follows. First, the constraints on the set H of possible outcomes guarantee that jumps can only occur on a well-ordered set of time points. An outcome is a list of jump times and actions chosen at these jump times for all players such that the set of jump times is well-ordered by \leq . From this set, the decision nodes and the game tree are defined. Strategies are then mappings from the set of decision nodes to the set of actions. However, players are only allowed to use a strict subset of strategies satisfying two additional assumptions. The first, “identifiability”, requires that the infimum of a set of jump times also has to be a jump time, i.e. at any point in time the next time a strategy prescribes a jump can be identified. The second, “finitely many moves at any point in time”, states that a player is allowed to initiate at most finitely many jumps at any point in time. The main purpose of this condition is to guarantee that profiles of admissible strategies induce outcomes in H .

The results of Stinchcombe (1992) show that every profile of admissible strategies induces a unique outcome after every possible history. Further, every outcome in H can be reached through some profile of admissible strategies. Importantly, the set of admissible strategies is shown to be maximal in the sense that weakening the identifiability condition for any player or the second condition for all players simultaneously would lead to the existence of strategy profiles that induce either no outcome or multiple outcomes.

The identifiability condition is comparable to our conditions (CRM.ii) and (CRM.iii). Both identifiability and our conditions essentially ensure that at any point in time when a player changed his action he has to stick to the new action for some positive amount of time. The effect of either approach is to guarantee that the next point in time when a decision is to be taken is well-defined.

1.6.2 Staying Quiet

Perry and Reny (1993) develop a bargaining model in continuous time. At any point in time players can submit an offer or “stay quiet”, i.e. not make an offer. The strategy sets are then restricted through three conditions. Condition $S1$ requires that once an offer is made by a player, he must stay quiet for an exogenously given strictly positive amount of time. Condition $S2$ specifies that the other player cannot react to the offer for some exogenously given nonnegative amount of time (although he can make an offer himself during that period). The game ends whenever either both players make the same offer at some point in time or at some point a player stays quiet but the other player matches his most recent offer. Perry and Reny (1993) provide an example showing that $S1$ and $S2$ alone do not guarantee outcome existence. This (and outcome uniqueness) is accomplished by condition $S3$, which requires that at any point in time, after making a decision, whether this was making an offer or staying quiet, the player must stay quiet for some strictly positive amount of time. Interestingly, given $S1$, $S3$ is both necessary and sufficient for the existence of an outcome.

Condition $S3$ incorporates an idea akin to inertia times and is related to our condition (CRM.ii), but it is a more stringent constraint in the sense that it does not allow for instant responses. While a player may immediately *learn* about the other player’s offer, by $S3$ he has to stay quiet for some strictly positive amount of time. Hence while the lower bound of possible reaction times is 0, there is a strictly positive delay. A similar but slightly stronger restriction is used in Perry and Reny (1994) where it is required that for every history all points in time where the strategy prescribes something else than staying quiet are isolated points, i.e. for every history and all times t there is an $\varepsilon > 0$ such that the strategy prescribes to stay quiet on $]t - \varepsilon, t + \varepsilon[\setminus \{t\}$.

1.6.3 Conditioning on Counterfactuals

Bergin (1992, 2006) and Bergin and MacLeod (1993) propose an interesting framework for repeated games in continuous time that also relies on reducing the set of allowable strategies. The restrictions imposed in those works result

in a framework guaranteeing outcome existence and uniqueness. In particular, Bergin (2006) presents a general formalization of restricted strategies, and we have drawn from it for the formulation of our Conditional Response Mappings. However, the restrictions imposed in those papers cause problems in a different front, because the framework cannot be captured through an extensive form game. We think it is important to address those here to highlight the kind of problems that can inadvertently be created if a model of continuous time does not rely on an explicit extensive form game.

To illustrate the problems, we focus on Bergin (2006). We first present a brief introduction to the framework in that paper, adapting the original notation to ours. Let I and A_i be as in the Action-Reaction Framework and let $H := \{h = (h_1, \dots, h_{|I|}) \mid h_i : \mathbb{R}_+ \rightarrow A_i \forall i \in I\}$. A “strategy” for a player i in Bergin’s framework is a mapping $b_i : H \times \mathbb{R}_+ \rightarrow A_i$ such that for all $h \in H$ the following conditions hold.

- (B.i) If $h(\tau) = h'(\tau)$ for all $\tau \in [0, t[$ for some $h' \in H$ and some $t \in \mathbb{R}_+$, then $b_i(h, t) = b_i(h', t)$.
- (B.ii) There exists $\varepsilon > 0$ such that $b_i(h, \tau) = b_i(h, 0)$ for all $\tau \in [0, \varepsilon[$.
- (B.iii) If $t \in \bigcap_{j \in I} LC(h_j)$ then there exists $\varepsilon > 0$ such that $b_i(h, \tau) = b_i(h, t)$ for all $\tau \in [t, t + \varepsilon[$.
- (B.iv) If $t \notin LC(h_i)$ then there exists $\varepsilon > 0$ such that $b_i(h, \tau) = b_i(h, t)$ for all $\tau \in [t, t + \varepsilon[$.
- (B.v) If $t \in LC(h_i) \setminus \bigcap_{j \neq i} LC(h_j)$ then there is $\varepsilon_i(h, t) > 0, a_i \in A_i$ such that $b_i(h, \tau) = a_i$ for all $\tau \in]t, t + \varepsilon_i(h, t)[$.

Conditions (B.ii)-(B.v) are similar in spirit to our conditions (CRM.ii) and (CRM.iii). Whenever a player jumps, the other players can react instantly. After that, however, all players have to stick to their new action for some positive amount of time. Bergin (2006) proves that any profile of “strategies” in his framework induces a unique outcome after every history. This approach, however, is problematic for two reasons. First, it can be shown

that the set of outcomes induced by profiles of such “strategies” is equal to our set F of decision paths defined above (see Proposition 5 in Appendix 1.D), i.e. not all elements of H can be reached by such a profile. This is in contrast to the Action-Reaction Framework or the framework of Stinchcombe (1992). Second, the restrictions imposed on the strategies make it impossible to formalize this approach as an extensive form game. Specifically, conditions (B.ii)-(B.v) require the set of choices that are available to a player after a history to not only depend on the history, as it should be in an extensive form, but also on the *chosen strategy*. In a sense, they “depend on a counterfactual future”, because by (B.ii) and (B.iii) after a time t a player is forced to choose what the strategy chose along the (future) outcome path, irrespective of the history after t . Intuitively, the problem is that a “strategy” insists on what “should have been done ” rather than considering the actual path of play.

The following example shows that Bergin’s approach cannot be formalized in the Action-Reaction Framework. This also provides an indication as to why it cannot be formalized as an extensive form game.

Example 3. Let $I = \{1\}$, and $A_1 = \{0, 1\}$. Let h_0 be defined by $h_0(t) = 0$ for all $t \in \mathbb{R}_+$ and

$$h_1(t) = \begin{cases} 0, & \text{if } t < 42, \\ 1, & \text{if } t \geq 42. \end{cases}$$

Let $w_0 := (h_0, \epsilon_0) \in W$ and $w_1 := (h_1, \epsilon_1) \in W$. Given $x \in \uparrow \{w_k\}$ let $c_k(x)$ be the (unique) choice available at x that leads to $\{w_k\}$, $k = 0, 1$. Define a strategy s in (T, C) as follows:

$$s(x) = \begin{cases} c_1(x), & \text{if } x \in \uparrow \{w_1\}, \\ c_0(x), & \text{if } x \in \uparrow \{w_0\} \setminus \uparrow \{w_1\}, \\ c(x), & \text{otherwise,} \end{cases}$$

where $c(x) \in A(x)$ is some arbitrary element of $A(x)$. The corresponding CRM σ satisfies $\sigma^1(h_1, t) = h_1(t)$ for all $t \in \mathbb{R}_+$ and $\sigma^1(h_0, t) = 0$ for all $t > 42$. For any strategy b in Bergin’s framework that induces the same

outcome as σ , i.e. for which $b(h_1, t) = h_1(t)$ for all $t \in \mathbb{R}_+$, (B.i) implies $b(h_0, 42) = 1$. Further, (B.iii) implies that there is an $\varepsilon > 0$ such that $b(h_0, t) = 1$ for all $t \in]42, 42 + \varepsilon[$. Even though $b(h_0, 42) = 1$ means departing from h_0 , (B.iii) implies a condition for actions chosen along h_0 after $t = 42$. Hence the outcome induced by σ and the outcome induced by b after the history-time pair $(h_0, 42 + \varepsilon/2)$ can never be the same. Thus there is no strategy in Bergin's framework that induces the same outcome as σ after every history and could therefore be considered equivalent to σ . In particular there is no strategy in Bergin's framework that could be considered equivalent to the strategy s in the Action-Reaction Framework.

The above example illustrates the problem caused by conditions (B.i)-(B.v). On the one hand a player is required to stick to an action chosen for some positive amount of time. On the other hand this "rule" does not apply to counterfactual histories where the player is forced to immediately switch to the action that was chosen along the actual outcome path. Thus Bergin's framework is an example of a framework where the extensive form does not survive the restrictions imposed on the strategy set.

1.7 Conclusion

Repeated games in continuous time are plagued with problems of outcome nonexistence and nonuniqueness, which amount to various forms of impossibility results and convey the overall message that continuous-time models are not well-founded. In contrast, we provide a *possibility result*. Our approach shows that it is possible to capture continuous-time modeling within the framework of well-defined extensive form games, without any artificial restriction of the associated strategy sets. All the necessary conditions ensuring that every strategy profile induces a unique outcome are incorporated in the game form, which allows for a better understanding of the tradeoffs involved in continuous-time modeling.

Previous work had concentrated on a "second best", placing exogenous restrictions on the players' strategy sets. From a game-theoretic point of

view, however, this is a problematic approach, since it is unclear in which sense a solution concept based on a strategy set restricted for purely technical reasons is related to the original extensive form. What our construction accomplishes is showing that the restrictions for strategy sets considered in the literature (e.g. Stinchcombe, 1992) can be adapted to appropriate conditions formulated from the onset, i.e. incorporated into the game tree and the choice system. The relation to the literature is made clear by showing that the (unrestricted) behavioral strategies from the resulting extensive form are equivalent to those in a restricted class of strategies in a more naïvely specified continuous-time repeated game. In turn, those restricted strategies are closely related to the approaches presented in the literature (Stinchcombe, 1992; Bergin, 2006).

Of course, our results do not mean that naïvely specified continuous time models can be treated as extensive form games, as our initial examples show. Familiarity should not be confused with simplicity, and the continuum is not a simple construction. The accomplishment of this paper is to show that continuous-time modeling is *possible* within the realms of standard game theory. Modeling decisions in continuous time, however, requires a relatively involved framework. The benefits are of two kinds. The first is of practical nature. Once the framework is in place, there is no further question of interpretation of game-theoretic concepts. The game is an extensive form game to which standard ideas apply. The second is more fundamental. In a sense, our construction resolves the tension between technical assumptions imposed for the sake of tractability and conceptual requirements resulting from a well-established theory of strategic interactions. If continuous time is deemed a worthy setting for tractability reasons, it is not necessary to give up the standard decision- and game-theoretic framework in order to use it.

Appendix 1.A: Proofs from Sections 1.2 and 1.3

This appendix contains the proofs of Proposition 1 and Lemma 1, which in turn implies Theorem 1. We start with a few preliminary lemmata which are also used elsewhere. \mathbb{N}_0 will denote the set of natural numbers including 0, i.e. $\mathbb{N}_0 = \{0, 1, 2, \dots\}$.

Lemma A.1. *Let $w = (f, \epsilon) \in W$, $i \in I$, and $t \in \mathbb{R}_+$.*

(i) $\text{Prev}(\epsilon_i, t) \in DP(\epsilon_i)$.

(ii) $\text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t)) \geq t$.

Proof. (i) By definition $\text{Prev}(\epsilon_i, t) = \sup(DP(\epsilon_i) \cap [0, t])$. By contradiction, assume that $\bar{t} := \text{Prev}(\epsilon_i, t) \notin DP(\epsilon_i)$. Then for all $\varepsilon > 0$ there is $\tau \in]\bar{t} - \varepsilon, \bar{t}[\cap DP(\epsilon_i)$. Hence $\sup(DP(\epsilon_i) \cap [0, \bar{t}]) = \bar{t}$ and thus $\bar{t} \in PJ(\epsilon_i)$ by definition of $PJ(\epsilon_i)$. (P.iii) then implies $\bar{t} \in DP(\epsilon_i)$, a contradiction.

(ii) By definition, $\text{Prev}(\epsilon_i, t) \leq t$. If $\text{Prev}(\epsilon_i, t) = t$ then $\epsilon_i(t) > 0$ as $t \in PJ(\epsilon_i) \subseteq DP(\epsilon_i)$ by (P.iii) and hence $\text{Next}(\epsilon_i, t) > t$. If $\text{Prev}(\epsilon_i, t) < t$ assume by contradiction that $\bar{t} := \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t)) < t$. As $\epsilon_i(\text{Prev}(\epsilon_i, t)) > 0$ by (i), $\text{Prev}(\epsilon_i, t) < \bar{t} < t$. By definition of $\text{Prev}(\epsilon_i, t)$, $\epsilon_i(\tau) = 0$ for all $\tau \in]\text{Prev}(\epsilon_i, t), t[$, hence for all $\tau \in]\text{Prev}(\epsilon_i, t), \bar{t}]$, implying $\text{Prev}(\epsilon_i, \bar{t}) = \sup(DP(\epsilon_i) \cap [0, \bar{t}]) = \text{Prev}(\epsilon_i, t)$. Thus $\bar{t} = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, \bar{t}))$ which implies $\bar{t} \in PJ(\epsilon_i) \subseteq DP(\epsilon_i)$ by (P.iii), a contradiction with $\epsilon_i(\bar{t}) = 0$. \square

Lemma A.2. *For every $f \in F$ and all $i \in I$ the set $J(f_i)$ is well-ordered by \leq . In particular the set $\bigcup_{i \in I} J(f_i)$ is well-ordered by \leq .*

Proof. Let $\emptyset \neq U \subseteq J(f_i)$. Since $U \subseteq \mathbb{R}_+$ we have that $t := \inf U$ exists. By (DP.i) there is $\varepsilon > 0$ such that f_i is constant on $]t, t + \varepsilon[$. By contradiction, suppose $t \notin U$, then for every $\varepsilon > 0$ there is $\tau \in]t, t + \varepsilon[$ such that $\tau \in U \subseteq J(f_i)$, contradicting that f_i is constant on $]t, t + \varepsilon[$. \square

Lemma A.3. *Let $w = (f, \epsilon) \in W$.*

(i) *For $i \in I$ the sets $DP(\epsilon_i)$, and $PJ(\epsilon_i)$ are well-ordered by \leq .*

(ii) *The sets $\bigcup_{i \in I} DP(\epsilon_i)$, and $\bigcup_{i \in I} PJ(\epsilon_i)$ are well-ordered by \leq , and hence countable.*

Proof. (i) Fix a player $i \in I$. We will first show that $DP(\epsilon_i)$ is well-ordered. Let $\emptyset \neq U \subseteq DP(\epsilon_i)$. As $U \subseteq \mathbb{R}_+$, $\bar{t} := \inf U$ exists. If $\bar{t} \in U$, we are done. By contradiction, if $\bar{t} \notin U$ then for all $\varepsilon > 0$ there is $t \in]\bar{t}, \bar{t} + \varepsilon[$ such that $t \in U \subseteq DP(\epsilon_i)$. Let

$$t' := \min\left\{ \min_{j \in IDP(\epsilon, \bar{t})} \text{Next}(\epsilon_j, \bar{t}); \min_{j \in I \setminus IDP(\epsilon, \bar{t})} \text{Next}(\epsilon_j, \text{Prev}(\epsilon_j, \bar{t})) \right\}.$$

By Lemma A.1(ii) $\text{Next}(\epsilon_j, \text{Prev}(\epsilon_j, \bar{t})) \geq \bar{t}$ for all $j \in I$. Hence, by (P.iii) $\text{Next}(\epsilon_j, \text{Prev}(\epsilon_j, \bar{t})) > \bar{t}$ for all $j \in I \setminus IDP(\epsilon, \bar{t})$. As $\text{Next}(\epsilon_j, \bar{t}) > \bar{t}$ for all $j \in IDP(\epsilon, \bar{t})$, we obtain $t' > \bar{t}$.

As $f \in F$ by (DP.i) there is $\varepsilon > 0$ such that f is constant on $]\bar{t}, \bar{t} + \varepsilon[$. We claim that $\epsilon_i(\tau) = 0$ for all $\tau \in]\bar{t}, \min\{t', \bar{t} + \varepsilon\}[$. Assume by contradiction that $\epsilon_i(\tau) > 0$ for some $\tau \in]\bar{t}, \min\{t', \bar{t} + \varepsilon\}[$. If $i \in IDP(\epsilon_i, \bar{t})$ then by (P.iv), $\bigcup_{j \neq i} J(f_j) \cap]\bar{t}, \tau] \neq \emptyset$, which contradicts the fact that f is constant on $]\bar{t}, \bar{t} + \varepsilon[$. If $i \notin IDP(\epsilon, \bar{t})$ then as $\text{Prev}(\epsilon_i, \bar{t}) \leq \bar{t} < \tau < t' \leq \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, \bar{t}))$, (P.iv) implies $\bigcup_{j \neq i} J(f_j) \cap]\text{Prev}(\epsilon_i, \bar{t}), \tau] \neq \emptyset$ which again contradicts the fact that f is constant on $]\bar{t}, \bar{t} + \varepsilon[$.

Hence there is $\varepsilon' > 0$ such that $\epsilon_i(\tau) = 0$ for all $\tau \in]t, t + \varepsilon'[$. As $\bar{t} \notin U$, this is a contradiction to the definition of infimum.

(ii) follows from (i) as all sets are finite unions of well-ordered sets. All sets of real numbers which are well-ordered by \leq are countable.¹⁰ \square

¹⁰This is a well-known observation. Let $s(x)$ be the successor of a real number x according to the standard order \leq . The open intervals $(x, s(x))$ for the different elements of the well-ordered set are nonempty and disjoint. Since each such interval contains a different rational number, the well-ordered set must be countable.

Lemma A.4. *Let $w = (f, \epsilon) \in W$, $\bar{t} \in \mathbb{R}_+$, and $\epsilon_i \in \mathbb{R}_{++}$ for all $i \in I$. Further let $a, b \in \times_{i \in I} A_i$ and let $f^{\bar{t}}$ be given by $f^{\bar{t}}(\tau) = f(\tau)$ for all $\tau < \bar{t}$, $f^{\bar{t}}(\bar{t}) = a$, and $f^{\bar{t}}(\tau) = b$ for all $\tau > \bar{t}$. Suppose $f^{\bar{t}} \in F$ and that for all $i \in I$, $\bar{t} \in J(f_i^{\bar{t}})$ only if $i \in IPJ(\epsilon_i, \bar{t})$.*

(i) *There is $\epsilon' \in E$ such that $\epsilon'(t) = \epsilon(t)$ for all $t \in [0, \bar{t}]$, $\epsilon'_i(\bar{t}) = \epsilon_i$ for all $i \in IPJ(\epsilon, \bar{t})$, and $w' = (f^{\bar{t}}, \epsilon') \in W$.*

(ii) *If $\bar{t} \in \mathcal{J}(f^{\bar{t}})$ then there is $\epsilon' \in E$ such that $\epsilon'(t) = \epsilon(t)$ for all $t \in [0, \bar{t}]$, $\epsilon'_i(\bar{t}) = \epsilon_i$ for all $i \in I$, and $w' = (f^{\bar{t}}, \epsilon') \in W$.*

Proof. (i) Let $I(\bar{t}) = \{i \in I \mid \bar{t} \in PJ(\epsilon_i) \cup \bigcup_{j \in I} J(f_j^{\bar{t}-})\}$. For each $i \in I$ define

$$\epsilon'_i(t) = \begin{cases} \epsilon_i(t), & \text{if } t < \bar{t}, \\ \epsilon_i, & \text{if } i \in I(\bar{t}) \text{ and } t = \bar{t} + \epsilon_i n \text{ for some } n \in \mathbb{N}_0, \\ 73, & \text{if } i \notin I(\bar{t}) \text{ and } t = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, \bar{t})) + 73n \text{ for some } n \in \mathbb{N}_0, \\ 0, & \text{otherwise} \end{cases}$$

(Note that the third condition becomes void if $\bar{t} \in \bigcup_{j \in I} J(f_j^{\bar{t}})$).

We will show that $w' = (f^{\bar{t}}, \epsilon') \in W$. To see (P.i) and (P.ii), let $i \in I$ and $t \in J(f_i^{\bar{t}})$. Then $t \leq \bar{t}$ by definition of $f^{\bar{t}}$, and hence $t \in PJ(\epsilon'_i) \cap \bigcap_{j \in I} DP_j(\epsilon'_j)$ by (P.i) and (P.ii) for w and the construction of ϵ . To prove (P.iii), let $i \in I$ and $t \in PJ(\epsilon'_i)$. If $t \leq \bar{t}$ then $t \in DP(\epsilon'_i)$ by (P.iii) for w . If $t > \bar{t}$ then by construction $t \in PJ(\epsilon_i)$ if and only if $t \in DP(\epsilon_i)$.

Finally, we will turn to (P.iv). Let $i \in I$, $t \in DP(\epsilon'_i)$, and $\tau \in DP(\epsilon'_i) \cap [t, \text{Next}(\epsilon'_i, t)[$. If $\text{Next}(\epsilon'_i, \text{Prev}(\epsilon'_i, t)) \leq \bar{t}$, then $\bigcup_{j \neq i} J(f_j^{\bar{t}-}) \cap]t, \tau] \neq \emptyset$ by (P.iv) for w . Hence suppose that $\text{Next}(\epsilon'_i, \text{Prev}(\epsilon'_i, t)) > \bar{t}$. If $t \geq \bar{t}$ then by construction of $f^{\bar{t}}$, $t \in DP(\epsilon'_i)$ implies $DP(\epsilon'_i) \cap]t, \text{Next}(\epsilon'_i, t)[= \emptyset$, a contradiction to our choice of τ . If $\bar{t} > t$ and $\bar{t} > \tau$, applying (P.iv) to w yields $\bigcup_{j \neq i} J(f_j^{\bar{t}-}) \cap]t, \tau] \neq \emptyset$.

If $\tau \geq \bar{t}$ consider three cases. If $\bar{t} \in \bigcup_{j \neq i} J(f_j^{\bar{t}})$, we are done. Therefore suppose $\bar{t} \notin \bigcup_{j \neq i} J(f_j^{\bar{t}})$. Then if $\bar{t} \in PJ(\epsilon_i)$ we get that $\bar{t} \in DP(\epsilon_i)$ and $t < \bar{t} < \text{Next}(\epsilon_i, \bar{t})$. Applying (P.iv) for w then yields $\bigcup_{j \neq i} J(f_j^{\bar{t}-}) \cap]t, \bar{t}] \neq \emptyset$ and hence $\bigcup_{j \neq i} J(f_j^{\bar{t}-}) \cap]t, \tau] \neq \emptyset$. Finally, if $\bar{t} \notin PJ(\epsilon_i)$ then by hypothesis

$\bar{t} \notin J(f_i^{\bar{t}})$. Note that if $t < \text{Prev}(\epsilon_i, \bar{t}) < \bar{t}$, applying (P.iv) for w we obtain $\emptyset \neq \bigcup_{j \neq i} J(f_j^{\bar{t}-}) \cap]t, \text{Prev}(\epsilon_i, \bar{t})] \subseteq]t, \tau]$, as $\text{Prev}(\epsilon_i, \bar{t}) \in DP(\epsilon_i)$ by Lemma A.1(i). Hence w.l.o.g. assume $t = \text{Prev}(\epsilon_i, \bar{t}) < \bar{t}$. By construction of ϵ'_i we get $\tau = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, \bar{t})) + 73n$ for some $n \in \mathbb{N}_0$, which yields $\tau \geq \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, \bar{t})) \geq \text{Next}(\epsilon'_i, t)$, a contradiction to our choice of τ . Hence (P.iv) holds.

(ii) Note that in (i) we actually proved that $\epsilon_i(\bar{t}) = \epsilon_i$ for all $i \in I(\bar{t})$. If $\bar{t} \in \mathcal{J}(f)$ then $I(\bar{t}) = I$ and hence the statement follows. \square

Lemma A.5. *Let $w = (f, \epsilon) \in W$, $t \in \mathbb{R}_+$. Then $i \in IDP(\epsilon, t)$ if and only if $i \in IPJ(\epsilon, t)$ or $IJ(f, t) \neq \emptyset$. In particular, $IDP(\epsilon, t) \neq \emptyset$ implies $IPJ(\epsilon, t) \neq \emptyset$.*

Proof. “If”: If $i \in IPJ(\epsilon, t)$ then $i \in IDP(\epsilon, t)$ by (P.iii). If $IJ(f, t) \neq \emptyset$ then $i \in IDP(\epsilon, t)$ by (P.ii).

“Only if”: Let $i \in IDP(\epsilon, t)$ and suppose that $i \notin IPJ(\epsilon, t)$. Let $\bar{t} = \text{Prev}(\epsilon_i, t)$. Then $\bar{t} < t < \text{Next}(\epsilon_i, \bar{t})$, where the second inequality follows from Lemma A.1(ii). By Lemma A.1(i) $\bar{t} \in DP(\epsilon_i)$. As $t \in DP(\epsilon_i)$, (P.iv) implies that there is $\tilde{t} \in \bigcup_{j \neq i} J(f_j) \cap]\bar{t}, t]$. If $\tilde{t} \in]\bar{t}, t]$ then by (P.ii) $\tilde{t} \in DP(\epsilon_i)$. As $\bar{t} = \text{Prev}(\epsilon_i, t) < \tilde{t} < t$ this would contradict the definition of $\text{Prev}(\epsilon_i, t)$. Hence $t \in \bigcup_{j \neq i} J(f_j)$ and thus $IJ(f, t) \neq \emptyset$. \square

We are now ready to turn to the proof of Proposition 1.

Proof of Proposition 1. For $x \in N$, let $t^x \in \mathbb{R}_+$ and $w^x = (f^x, \epsilon^x) \in W$ be such that $x = x_{t^x}(w^x)$, $x = x_{t^x}^R(w^x)$, or $x = x_{t^x}^P(w^x)$.

We will first show that $T = (N, \supseteq)$ is a game tree (Definition 1).

(TI): Let $x, y \in N$ be such that $x \cap y \neq \emptyset$ and let $w = (f, \epsilon) \in x \cap y$. Then $x_{t^x}^P(w) \subseteq x \subseteq x_{t^x}(w)$ and $x_{t^y}^P(w) \subseteq x \subseteq x_{t^y}(w)$. Without loss of generality, assume $t^x \leq t^y$. Then for any $w^y = (f^y, \epsilon^y) \in y$ we obtain $w^y(\tau) = w(\tau)$ for all $\tau \in [0, t^y[$. If $t^x < t^y$ then $f^y(t^x) = f(t^x)$ and $f_+^y(t^x) = f_+(t^x)$. Hence $w^y \in x_{t^x}^P(w) \subseteq x$, i.e. $y \subseteq x$. If $t^x = t^y$ then $x, y \in \{x_{t^x}(w), x_{t^x}^R(w), x_{t^x}^P(w)\}$ and hence are ordered.

(IR) Let $w, w' \in W$ with $w = (f, \epsilon)$, $w' = (f', \epsilon')$ be such that $w \neq w'$. Then there is $t \in \mathbb{R}_+$ such that $w(t) \neq w'(t)$. Let $t^1 = \min_{i \in I} \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t))$

and $t^2 = \min_{i \in I} \text{Next}(\epsilon'_i, \text{Prev}(\epsilon'_i, t))$. By (P.iv), $\bigcup_{i \in I} DP(\epsilon_i) \cap]\text{Prev}(\epsilon_i, t), t^1[= \emptyset$ and $\bigcup_{i \in I} DP(\epsilon'_i) \cap]\text{Prev}(\epsilon'_i, t), t^2[= \emptyset$. Hence $\text{Prev}(\epsilon_i, t^1) = \text{Prev}(\epsilon_i, t)$ and $\text{Prev}(\epsilon'_i, t^2) = \text{Prev}(\epsilon'_i, t)$ for all $i \in I$ and there are $i, j \in I$ such that $t^1 = \text{Next}(\text{Prev}(\epsilon_i, t^1), t^1)$ and $t^2 = \text{Next}(\text{Prev}(\epsilon'_j, t^2), t^2)$. Thus $t^1 \in PJ(\epsilon_i)$, $t^2 \in PJ(\epsilon'_j)$ and by the construction of the tree, $x_{t^1}(w), x_{t^2}(w') \in N$. By Lemma A.1(ii), $t^1, t^2 \geq t$ and hence by construction of the nodes $w \in x_{t^1}(w)$ and $w \notin x_{t^2}(w')$ and $w' \in x_{t^2}(w')$ and $w' \notin x_{t^1}(w)$.

(BD) Let $h \in 2^N$ be a nonempty chain. Let

$$D := \{t \in \mathbb{R}_+ \mid \exists w \in W : x_t(w) \text{ or } x_t^R(w) \text{ or } x_t^P(w) \in h\}.$$

Note that as in the proof of (TI) above (case $t^x = t^y$) for each $t \in A$ there are at most three nodes $x \in h$ such that $t^x = t$ (a peek node, a reaction node, and a potential jump node).

Suppose first that $\exists \bar{t} = \sup D$ and $\bar{t} \in D$. Let $y \in h$ be the smallest of the nodes in the chain h with $t^y = \bar{t}$. Let $w \in y$. As in the proof of (TI) above (case $t^x < t^y$), it follows that $w \in x$ for all $x \in h$.

Suppose now that either $\exists \bar{t} = \sup D$ and $\bar{t} \notin D$ or $\nexists \sup D$. In the latter case write $\bar{t} = +\infty$ for convenience. For any $0 < K < \bar{t}$ there is $t^K \in]K, \bar{t}[\cap D$ and $w^K = (f^K, \epsilon^K) \in W$ such that $x_{t^K}(w^K)$ or $x_{t^K}^R(w^K)$ or $x_{t^K}^P(w^K) \in h$.

For each $i \in I$ and some $a_i \in A_i$ define $\bar{w}_i = (\bar{f}_i, \bar{\epsilon}_i) \in F \times E$ by

$$\bar{w}_i(\tau) := \begin{cases} w_i^r(\tau), & \text{if } \tau < \bar{t}, \\ (a_i, 0), & \text{if } \tau \geq \bar{t}. \end{cases}$$

Note that (since h is a chain) if $\tau < \bar{t}$ then $w_i^r(\tau) = w_i^K(\tau)$ for any $K \in]0, \bar{t}[$ with $\tau \in [0, t^K[$.

If $\bar{t} < +\infty$ we claim that $\bar{t} \in \bigcup_{j \in I} PJ(\bar{\epsilon}_j)$. Assume by contradiction that $\bar{t} \notin \bigcup_{j \in I} PJ(\bar{\epsilon}_j)$. Then $\text{Prev}(\bar{\epsilon}_j, \bar{t}) < \bar{t}$ and hence $\bar{\epsilon}_j(\tau) = 0$ for all $\tau \in]\text{Prev}(\bar{\epsilon}_j, \bar{t}), \bar{t}[$ and all $j \in I$. Let $K \in]\max_{j \in I} \text{Prev}(\bar{\epsilon}_j, \bar{t}), \bar{t}[$. Then there is $t^K \in]K, \bar{t}[$ and $w^K = (f^K, \epsilon^K) \in W$ such that $x_{t^K}(w^K)$ or $x_{t^K}^R(w^K)$ or $x_{t^K}^P(w^K) \in h$ and hence $IDP(\epsilon^K, t^K) \supseteq IPJ(\epsilon^K, t^K) \neq \emptyset$, where the inclusion follows from (P.iii). By construction $\bar{w}_i(t^K) = w_i^{t^K}(t^K)$ and as $IDP(\epsilon^K, t^K) \neq \emptyset$ this

implies $IDP(\bar{\epsilon}^K, t^K) \neq \emptyset$, which contradicts the fact that $IDP(\bar{\epsilon}, \tau) = \emptyset$ for all $\tau \in]\max_{j \in I} \text{Prev}(\bar{\epsilon}_j, \bar{t}), \bar{t}[$.

For each $i \in IPJ(\bar{\epsilon}, \bar{t})$ choose a'_i such that $a'_i \neq \lim_{t \rightarrow \bar{t}} \bar{f}_{i-}(t)$ if $\lim_{t \rightarrow \bar{t}} \bar{f}_{i-}(t)$ exists and arbitrarily otherwise. For each $i \in I \setminus IPJ(\bar{\epsilon}, \bar{t})$ let $a'_i = \lim_{t \rightarrow \bar{t}} \bar{f}_{i-}(t)$. Note that the limit exists as otherwise \bar{t} would be an accumulation point of jump points of f_i (and hence decision points of $\bar{\epsilon}_i$ by (P.i) and (P.iii)) and hence $i \in IPJ(\bar{\epsilon}, \bar{t})$.

Now define

$$f_i(\tau) := \begin{cases} \bar{f}_i(\tau), & \text{if } \tau < \bar{t} \\ a'_i, & \text{if } \tau \geq \bar{t} \end{cases}$$

and

$$\epsilon_i(\tau) := \begin{cases} \bar{\epsilon}_i(\tau), & \text{if } \tau < \bar{t}, \\ 73, & \text{if } \tau \geq \bar{t} \text{ and } \tau = \bar{t} + 73n \text{ for some } n \in \mathbb{N}_0, \\ 0, & \text{otherwise.} \end{cases}$$

and set $w := (f, \epsilon) \in F \times E$. Note that the construction guarantees the $\bar{t} \in J(f_i)$ if and only if $\bar{t} \in PJ(\bar{\epsilon}_i)$. We will now verify that $w \in W$. To see (P.i), let $\tau \in J(f_i)$. As f is constant on $]\bar{t}, +\infty[$, $\tau \leq \bar{t}$. If $\tau < \bar{t}$ then $\tau \in [0, t^K[$ for some $K > 0$. Then, by (P.i) for w^K , $\tau \in PJ(\epsilon_i^K)$ and hence $\tau \in PJ(\epsilon_i)$. If $\tau = \bar{t}$, by construction of f , $\bar{t} \in J(f_i)$ if and only if $\bar{t} \in PJ(\bar{\epsilon}_i)$ and hence $\bar{t} \in PJ(\epsilon_i)$. This proves (P.i). (P.ii) follows immediately from the construction of ϵ . To see (P.iii) let $\tau \in PJ(\epsilon_i)$. If $\tau \leq \bar{t}$ then $\tau \in DP(\epsilon_i)$ by construction. If $\tau > \bar{t}$ then $\text{Prev}(\epsilon_i, \tau) = \bar{t} + 73n$ for some $n \in \mathbb{N}_0$ implying that $\tau = \bar{t} + 73(n+1)$ and thus $\epsilon_i(\tau) > 0$. Thus (P.iii) is satisfied. To prove (P.iv) let $t \in DP(\epsilon_i)$ and $\tau \in]t, \text{Next}(\epsilon_i, t)[\cap DP(\epsilon_i)$ be such that $\bigcup_{j \in I} J(f_j) \cap]t, \tau] = \emptyset$. If $\tau \leq \bar{t}$ we reach a contradiction with (P.iv) for w^K with $K > \tau$. If $t < \bar{t}$ and $\bar{t} < \tau$ then in particular $\bar{t} \notin \bigcup_{j \in I} J(f_j)$, in contradiction with the construction of f . Now suppose $t \geq \bar{t}$. The construction of ϵ implies that if $\tau \in]t, \text{Next}(\epsilon_i, t)[$ then $\epsilon_i(\tau) = 0$, a contradiction with $\tau \in DP(\epsilon_i)$. Hence (P.iv) holds and we obtain $w \in W$.

Next we will show that (T, C) is an EDP (Definition 2).

(EDP.i) Let $i \in I$ and $c, c' \in C_i$ be such that $P(c) \cap P(c') \neq \emptyset$ and $c \neq c'$. First, let $c = c_i(x_t(w), a_i)$ for some $w = (f, \epsilon) \in W$ and some $a_i \in A_i$. Then

$$P(c) = \{x_t((f', \epsilon')) \in N \mid f(\tau) = f'(\tau) \forall \tau \in [0, t[, t \in PJ(\epsilon'_i)\}.$$

$P(c) \cap P(c') \neq \emptyset$ yields $c' = c_i(x_t(w), a'_i)$ for some $a'_i \in A_i$ and hence $P(c) = P(c')$. As $c \neq c'$, we have $a_i \neq a'_i$ which implies $c \cap c' = \emptyset$. The proofs for the cases $c = c_i(x_t^R(w), a_i)$ and $c = c_i(x_t^P(w), a_i)$ are analogous.

(EDP.ii) Let $x \in X$ and $(c_i)_{i \in I(x)} \in \times_{i \in I(x)} A_i(x)$. If $x = x_{t^x}(w^x)$ let $f(\tau) = f^x(\tau)$ for all $\tau \in [0, t^x[$ and $f_i(\tau) = a_i$ for all $\tau \in [t^x, \infty[$ and all $i \in I$ where the a_i are such that $c_i = c_i(x, a_i)$ if $i \in I(x)$ and $a_i = \lim_{\tau \rightarrow t^x} f_i^x(\tau)$ if $i \notin I(x)$. Applying Lemma A.4(i) there exists $\epsilon \in E$ such that $\epsilon(\tau) = \epsilon^x(\tau)$ for all $\tau \in [0, t^x[$ and $w = (f, \epsilon) \in W$. By construction $w \in x \cap \bigcap_{i \in I(x)} c_i$.

If $x = x_{t^x}^R(w^x)$ let $f(\tau) = f^x(\tau)$ for all $\tau \in [0, t^x]$, and $f_i(\tau) = a_i$ for all $\tau \in]t, \infty[$ and all $i \in I$ where the a_i are such that $c_i = c_i(x, a_i)$ if $i \in I(x)$ and $a_i = f_i^x(t^x)$ if $i \notin I(x)$. By Lemma A.4(i) there is $\epsilon \in E$ such that $\epsilon(\tau) = \epsilon^x(\tau)$ for all $\tau \in [0, t^x[$ and $w = (f, \epsilon) \in W$. By construction $w \in x \cap \bigcap_{i \in I(x)} c_i$.

If $x = x_{t^x}^P(w^x)$ let $f(\tau) = f^x(\tau)$ for all $\tau \in [0, t^x]$, and $f_i(\tau) = f_{i+}(t^x)$ for all $\tau \in]t, \infty[$ and all $i \in I$. By Lemma A.4(i) there is $\epsilon \in E$ such that $\epsilon(\tau) = \epsilon^x(\tau)$ for all $\tau \in [0, t^x[$, $\epsilon_i(t^x) = \epsilon_i$ for all $i \in I(x)$ where ϵ_i is such that $c_i = c_i(x, \epsilon_i)$ if $i \in I(x)$ and $\epsilon_i = 0$ if $i \notin I(x)$ and $w = (f, \epsilon) \in W$. By construction $w \in x \cap \bigcap_{i \in I(x)} c_i$.

(EDP.iii) Let $y, y' \in N$ be such that $y \cap y' = \emptyset$. Let $\bar{t} := \inf\{t \in \mathbb{R}_+ \mid w^y(t) \neq w^{y'}(t)\}$. Then $w^y(\tau) = w^{y'}(\tau)$ for all $\tau \in [0, \bar{t}]$. Note that $\bar{t} \in PJ(\epsilon_i^y)$ if and only if $\bar{t} \in PJ(\epsilon_i^{y'})$ for all $i \in I$ as $w^y(\tau) = w^{y'}(\tau)$ for all $\tau \in [0, \bar{t}]$.

We claim that $\bar{t} \in \bigcup_{i \in I} PJ(\epsilon_i^y)$. By contradiction, suppose that $\bar{t} \notin \bigcup_{i \in I} PJ(\epsilon_i^y)$. If $w^y(\bar{t}) \neq w^{y'}(\bar{t})$ then $\bar{t} \in DP(\epsilon_j^y) \cup DP(\epsilon_j^{y'})$ for some $j \in I$ as either $f^y(\bar{t}) \neq f^{y'}(\bar{t})$ or $\epsilon^y(\bar{t}) \neq \epsilon^{y'}(\bar{t})$. Suppose that $\bar{t} \notin PJ(\epsilon_j^{y'}) \cup PJ(\epsilon_j^y)$ for all $j \in I$. Then by (P.i) and (P.iv), $\bar{t} \notin DP(\epsilon_j^y) \cup DP(\epsilon_j^{y'})$ for all $j \in I$. Hence $\bar{t} \in DP(\epsilon_j^y) \cup DP(\epsilon_j^{y'})$ for some $j \in I$ implies that $\bar{t} \notin PJ(\epsilon_k^{y'}) \cup PJ(\epsilon_k^y)$ for some $k \in I$, a contradiction. If $w^y(\bar{t}) = w^{y'}(\bar{t})$ then $f_+^y(\bar{t}) \neq f_+^{y'}(\bar{t})$ by (DP.i) and (P.iv) as otherwise there is $\varepsilon > 0$ such that $w^y(\tau) = w^{y'}(\tau)$ for

all $\tau \in [0, \bar{t} + \varepsilon[$, contradicting the choice of \bar{t} . But $f_+^y(\bar{t}) \neq f_+^{y'}(\bar{t})$ implies $\bar{t} \in J(f_j^y) \cup J(f_j^{y'})$ for some $j \in I$ and hence $\bar{t} \in \bigcup_{i \in I} PJ(\epsilon_i^y)$, a contradiction.

Thus $\bar{t} \in \bigcup_{i \in I} PJ(\epsilon_i^y) \cap PJ(\epsilon_i^{y'})$ and $x_{\bar{t}}(w^y) = x_{\bar{t}}(w^{y'}) \in N$. If $w^y(\bar{t}) = w^{y'}(\bar{t})$ then $x_{\bar{t}}^R(w^y) = x_{\bar{t}}^R(w^{y'})$ and $f_+^y(\bar{t}) \neq f_+^{y'}(\bar{t})$ by definition of \bar{t} which implies $\emptyset \neq IJ(f^y, \bar{t}) = IJ(f^{y'}, \bar{t}) \subsetneq I$. Let $i \in I \setminus IJ(f^y, \bar{t})$ be such that $f_{i+}^y(\bar{t}) \neq f_{i+}^{y'}(\bar{t})$, and set $c := c_i(x_i^R(w^y), f_{i+}^y(\bar{t}))$ and $c' := c_i(x_i^R(w^y), f_{i+}^{y'}(\bar{t}))$. Then $y \subseteq c$, $y' \subseteq c'$ and $c \cap c' = \emptyset$. If $w^y(\bar{t}) \neq w^{y'}(\bar{t})$ then either $f^y(\bar{t}) \neq f^{y'}(\bar{t})$ or $\epsilon^y(\bar{t}) \neq \epsilon^{y'}(\bar{t})$. If $f^y(\bar{t}) \neq f^{y'}(\bar{t})$ let $i \in I$ be such that $f_i^y(\bar{t}) \neq f_i^{y'}(\bar{t})$ and define $c := c_i(x_{\bar{t}}(w^y), f_i^y(\bar{t}))$ and $c' := (x_{\bar{t}}(w^y), f_i^{y'}(\bar{t}))$. Then $y \subseteq c$, $y' \subseteq c'$ and $c \cap c' = \emptyset$. If $f^y(\bar{t}) = f^{y'}(\bar{t})$ and $\epsilon^y(\bar{t}) \neq \epsilon^{y'}(\bar{t})$ let $i \in I$ be such that $\epsilon_i^y(\bar{t}) \neq \epsilon_i^{y'}(\bar{t})$ and define $c := c_i(x_{\bar{t}}^P(w^y), \epsilon_i^y(\bar{t}))$ and $c' := (x_{\bar{t}}^P(w^{y'}), \epsilon_i^{y'}(\bar{t}))$. Then $y \subseteq c$, $y' \subseteq c'$ and $c \cap c' = \emptyset$.

(EDP.iv) Let $x \supsetneq y \in N$ and $i \in I(x)$. Then $t^x \leq t^y$ which implies $w^x(\tau) = w^y(\tau)$ for all $\tau \in [0, t^x[$. If $x = x_{t^x}(w^x)$ let $c_i = c_i(x, w_i^y(t^x))$. If $x = x_{t^x}^R(w^x)$ then $f^x(t^x) = f^y(t^x)$. Let $c_i = c_i(x, f_+^y i(t^x))$. If $x = x_{t^x}^P(w^x)$ then $f^x(t^x) = f^y(t^x)$ and $f_+^x(t^x) = f_+^y(t^x)$. Let $c_i = c_i(x, \epsilon_i^y(t^x))$. In any case $y \subseteq c_i$. \square

Now we turn to the proof of Lemma 1.

Proof of Lemma 1. As in the proof of Proposition 1, for $x \in N$, let $t^x \in \mathbb{R}_+$ and $w^x = (f^x, \epsilon^x) \in W$ be such that $x = x_{t^x}(w^x)$, $x = x_{t^x}^R(w^x)$, or $x = x_{t^x}^P(w^x)$.

We first show that T is regular. Let $x \in N$. If $x = x_{t^x}^R(w^x)$ then $\emptyset \neq IJ(f^x, t^x) \subseteq IPJ(\epsilon^x, t^x)$ (where the last inclusion follows from (P.i)) and hence $x_{t^x}(w^x) \in N$. Thus $x_{t^x}(w^x) = \min \uparrow x \setminus \{x\}$. If $x = x_{t^x}^P(w^x)$ we distinguish two cases. If $\emptyset \subsetneq IJ(\epsilon, t^x) \subsetneq I$ then $x_{t^x}^R(w^x) \in N$ and $x_{t^x}^R(w^x) = \min \uparrow x \setminus \{x\}$. Otherwise $x_{t^x}^R(w^x) \notin N$ and $IPJ(\epsilon^x, t^x) \neq \emptyset$ as $x \in N$ (recall (1.1)). Then $x_{t^x}(w^x) \in N$ and $x_{t^x}(w^x) = \min \uparrow x \setminus \{x\}$. If $x = x_{t^x}(w^x)$ we again distinguish two cases. If $\text{Prev}(\epsilon_i^x, t^x) < t^x$ for all $i \in I$ let $\bar{t} = \max_{i \in I} \text{Prev}(\epsilon_i^x, t^x)$. By Lemmata A.1(i) and A.5(“only if”) we obtain $IPJ(\epsilon^x, \bar{t}) \neq \emptyset$ and hence $x_{\bar{t}}^P(w^x) \in N$. Then $x_{\bar{t}}^P(w^x) = \min \uparrow x \setminus \{x\}$ as otherwise there would be $x_{\bar{t}}^P(w^x) \supsetneq x_{t'}^P(w^x) \supsetneq x$ implying $\bar{t} < t' < t$. As then $IDP(\epsilon^x, t') \neq \emptyset$ by Lemma A.5(“if”) this would contradict the construction of

\bar{t} . If on the other hand $\text{Prev}(\epsilon_i^x, t^x) = t$ for some $i \in I$, let $y \in \uparrow x \setminus \{x\}$. Since then $t^y < t^x$ and $\text{Prev}(\epsilon_i^x, t^x) = t$ there is $t^y < \bar{t} < t^x$ such that $\bar{t} \in DP(\epsilon_i)$. By Lemma A.5 $IPJ(\epsilon, \bar{t}) \neq \emptyset$ and hence $y \supseteq x_{\bar{t}}(w^x) \in \uparrow x \setminus \{x\}$. As $x \subsetneq y$ for all $y \in \uparrow x \setminus \{x\}$ we obtain $x = \inf \uparrow x \setminus \{x\}$.

It remains to show that T is up-discrete. Let $h \in 2^N$ be a nonempty chain and let $w = (f, \epsilon) \in \bigcap_{x \in h} x$, which exists by (BD). Note that if $x_t(w)$, $x_t^R(w)$, or $x_t^P(w) \in N$ for some $t \in \mathbb{R}_+$ then by construction of T (recall (1.1)), $t \in \bigcup_{i \in I} PJ(\epsilon_i)$. Since $\bigcup_{i \in I} PJ(\epsilon_i)$ is well-ordered by Lemma A.3 we obtain that $\bar{t} := \min\{t \mid x_t(w) \text{ or } x_t^R(w) \text{ or } x_t^P(w) \in h\}$ exists. Hence either $x_{\bar{t}}(w)$, $x_{\bar{t}}^R(w)$ or $x_{\bar{t}}^P(w)$ is a maximum of h . \square

Appendix 1.B: Proofs from Section 1.4

The proofs of results from Sections 1.4 and 1.5 make use of the machinery of ordinal numbers; we refer the reader to Jech (2002, chap. 2).

Let Ord be the class of all ordinal numbers. Given an ordinal $\alpha \in Ord$, a transfinite sequence of (possibly extended) real numbers $(t^\beta)_{\beta < \alpha}$ is a set $\{t^\beta \mid t^\beta \in \mathbb{R} \cup \{\infty\}, \beta < \alpha\}$. A transfinite sequence $(t^\beta)_{\beta < \alpha}$ is increasing if $\gamma < \beta$ implies $t^\gamma \leq t^\beta$ and strictly increasing if $\gamma < \beta$ implies $t^\gamma < t^\beta$. If α is a limit ordinal the limit $\lim_{\beta \rightarrow \alpha} t^\beta$ of the sequence is defined by $\lim_{\beta \rightarrow \alpha} t^\beta = \sup\{t^\beta \mid \beta < \alpha\}$. A sequence $(t^\beta)_{\beta < \alpha}$ is continuous if $t^\gamma = \lim_{\beta \rightarrow \gamma} t^\beta$ for every limit ordinal $\gamma < \alpha$. For the sake of clarity we will write $(t^\beta)_{\beta \leq \alpha}$ for $(t^\beta)_{\beta < \alpha+1}$.

The following definitions and lemmata are used in the proof of Proposition 2 and also elsewhere.

Definition 11. Given $f \in F$, $t \in \mathbb{R}_+$ and $a \in \times_{i \in I} A_i$, define

$$G^-(f, t, a) = \begin{cases} f(\tau), & \text{if } \tau \in [0, t[, \\ a, & \text{if } \tau \geq t \end{cases}$$

and

$$G^+(f, t, a) = \begin{cases} f(\tau), & \text{if } \tau \in [0, t], \\ a, & \text{if } \tau > t \end{cases}$$

Lemma B.1. *Let $f \in F$, $t \in \mathbb{R}_+$, and $a \in \times_{i \in I} A_i$. Then*

(i) $G^-(f, t, a) \in F$.

(ii) *If a is such that for all $i \in I$, $a_i = f_i(t)$ if $t \in \left(\bigcap_{j \in I} LC(f_j) \right) \cup J(f_i)$, then $G^+(f, t, a) \in F$.*

Proof. (i) By construction $G^-(f, t, a)$ is piecewise constant and hence (DP.i) holds. To prove (DP.ii), let $\tau \in \mathbb{R}_+$. If $\tau \notin LC(G_i^-(f, t, a))$, then by construction, $\tau \leq t$ and hence $\tau \in RK(G_i^-(f, t, a))$ by construction. To see (DP.iii), let $\tau \in R(G_i^-(f, t, a))$ for some $i \in I$. Then, as $\tau \notin RK(G_i^-(f, t, a))$, $\tau < t$ and since $f \in F$, there is $j \in I$ such that $\tau \in J(f_j)$ and hence $\tau \in J(G_j^-(f, t, a))$. Thus (DP.iii) holds.

(ii) Let a be as given. By construction $G^+(f, t, a)$ is piecewise constant and hence (DP.i) holds. To prove (DP.ii), let $\tau \in \mathbb{R}_+$. If $\tau \notin LC(G_i^+(f, t, a))$, then by construction $\tau \leq t$. If $\tau < t$, $\tau \in RK(G_i^+(f, t, a))$ since $f \in F$. Suppose $\tau = t$. Then $t \notin LC(G_i^+(f, t, a))$ implies $t \notin LC(f_i)$ and hence by (DP.ii) for f , $t \in J(f_i)$. Then by hypothesis, $a_i = f_i(t)$ and hence $\tau \in RK(G_i^+(f, t, a))$. Hence (DP.ii) holds. To see (DP.iii), let $\tau \in R(G_i^+(f, t, a))$ for some $i \in I$. Then by construction of $G_i^+(f, t, a)$, $\tau \leq t$. If $\tau < t$ then $\tau \in R(f_i)$ and by (DP.iii) for f it follows that $\tau \in J(f_j)$ and hence $\tau \in J(G_j^+(f, t, a))$ for some $j \in I$. Suppose $\tau = t$. That $t \in R(G_i^+(f, t, a))$ implies that $a_i \neq f_i(t)$ which by hypothesis implies $t \notin \left(\bigcap_{j \in I} LC(f_j) \right) \cup J(f_i)$. As $t \notin \bigcap_{j \in I} LC(f_j)$ we get that $t \in J(f_j)$ and hence (since (DP.ii) has already been shown for $G^+(f, t, a)$) $t \in J(G_j^+(f, t, a))$ for some $j \in I$. \square

Lemma B.2. *Let $\alpha \in \text{Ord}$ be a limit ordinal and let $(t^\beta)_{\beta < \alpha} \subseteq \mathbb{R}_+$ be a strictly increasing and continuous transfinite sequence. Then for every $t^0 \leq t < \lim_{\beta \rightarrow \alpha} t^\beta$ there is a unique $\delta \in \text{Ord}$ such that $t \in [t^\delta, t^{\delta+1}[$.*

Proof. Let $t^0 \leq t < \lim_{\beta < \alpha} t^\beta$. Then $\gamma = \min\{\beta < \alpha \mid t^\beta > t\}$ exists as the set $\{\beta \mid \beta < \alpha\}$ is well-ordered. Further, $\gamma > 0$ as $t \geq t^0$ and γ is a successor ordinal as otherwise $\lim_{\beta < \gamma} t^\beta = t^\gamma$ by continuity and hence there would be $\bar{\beta} < \gamma$ with $t < t^{\bar{\beta}} < t^\gamma$. This would contradict that $\gamma = \min\{\beta < \alpha \mid t^\beta > t\}$. Thus $\gamma = \delta + 1$ for some $\delta < \alpha$ and $t \in [t^\delta, t^{\delta+1}[$. \square

Proof of Lemma 2. Note that $f_{i+}^{t+}(t) = f_{i+}(t)$ by definition of f_i^{t+} . If $E^{\sigma_i}(f, t) \neq \emptyset$ then $\bar{\varepsilon} = \inf E^{\sigma_i}(f, t)$ exists. We claim that $\bar{\varepsilon} > 0$. In order to see this we distinguish two cases. First, suppose that $t \in \left(\bigcap_{j \in I} LC(f_j^{t+})\right) \cup J(f_i^{t+})$. In this case, $t \in RK(f_i^{t+})$ by (DP.ii) and (DP.iii) and hence $f_{i+}^{t+}(t) = f_i^{t+}(t)$ which implies $f_i^{t+}(t) = f_{i+}(t)$. Now by (CRM.ii) there is $\varepsilon > 0$ such that $\sigma_i^1(f^{t+}, \tau) = f_i^{t+}(t) = f_{i+}(t)$ for all $\tau \in]t, t + \varepsilon[$. Thus $E^{\sigma_i}(f, t)$ is bounded away from zero, and hence $\bar{\varepsilon} > 0$. Second, suppose that $t \notin \left(\bigcap_{j \in I} LC(f_j^{t+})\right) \cup J(f_i^{t+})$, hence (by (DP.ii)), $t \in LC(f_i^{t+}) \cap \left(\bigcup_{j \in I} J(f_j^{t+})\right)$. In this case, by (CRM.iii) there is $\varepsilon' > 0$ such that $\sigma_i^1(f^{t+}, \tau) = f_{i+}^{t+}(t) = f_{i+}(t)$ for all $\tau \in]t, t + \varepsilon'[$. Again, $E^{\sigma_i}(f, t)$ is bounded away from zero, and hence $\bar{\varepsilon} > 0$. This proves the claim.

Suppose now that $\bar{\varepsilon} \notin E^{\sigma_i}(f, t)$. As f^{t+} is constant on $]t, \infty[$ and $t + \bar{\varepsilon} > t$, $t + \bar{\varepsilon} \in \bigcap_{j \in I} LC(f_j^{t+})$ and hence by (CRM.ii) there is $\varepsilon > 0$ such that $\sigma_i^1(f^{t+}, \tau) = f_i^{t+}(t + \bar{\varepsilon})$ for all $\tau \in]t + \bar{\varepsilon}, t + \bar{\varepsilon} + \varepsilon[$. Since f_i^{t+} is constant on $]t, \infty[$, we have that $f_i^{t+}(t + \bar{\varepsilon}) = f_{i+}^{t+}(t) = f_{i+}(t)$ for all $\tau \in]t + \bar{\varepsilon}, t + \bar{\varepsilon} + \varepsilon[$. This contradicts the construction of $\bar{\varepsilon} = \inf E^{\sigma_i}(f, t)$. \square

Proof of Proposition 2. (i) Let $E^{\sigma_i}(f, t)$ and $\varepsilon^{\sigma_i}(f, t)$ be defined as in Lemma 2. Fix $r \in \mathbb{R}_{++}$ and define $\varepsilon_r^{\sigma_i}(f, t)$ by $\varepsilon_r^{\sigma_i}(f, t) = \varepsilon^{\sigma_i}(f, t)$ if $E^{\sigma_i}(f, t) \neq \emptyset$ and $\varepsilon_r^{\sigma_i}(f, t) = r$ otherwise.

We are going to use transfinite recursion to construct a sequence of functions $(f^\alpha)_{\alpha \in Ord}$ and a sequence of extended real numbers $(t^\alpha)_{\alpha \in Ord} \subseteq \mathbb{R}_+ \cup \{\infty\}$ such that for all $\alpha \in Ord$ the following properties are satisfied.

(TR.i) $f^\alpha \in F$.

(TR.ii) If $\beta < \alpha$ then $f^\alpha(\tau) = f^\beta(\tau)$ for all $\tau \in [0, t^\beta[$.

(TR.iii) $(t^\beta)_{\beta < \alpha}$ is continuous.

(TR.iv) Either $t^\alpha = \infty$ or $(t^\beta)_{\beta < \alpha} \subseteq \mathbb{R}_{++}$ and is strictly increasing.

(TR.v) $\sigma_i^1(f^\alpha, \tau) = f_i^\alpha(\tau)$ and $\sigma_i^2(f^\alpha, \tau) = f_{i+}^\alpha(\tau)$ for all $\tau \in [0, t^\alpha[$ and all $i \in I$.

In order to apply transfinite recursion, we need to complete three steps. First, we will define (f^0, t^0) trivially fulfilling (TR.i)-(TR.v). Second, we will

show that, if (TR.i)-(TR.v) are fulfilled for an ordinal α then $(f^{\alpha+1}, t^{\alpha+1})$ fulfilling (TR.i)-(TR.v) can be defined for the successor ordinal $\alpha+1$. Third, we will show that, for any limit ordinal α , if (f^β, t^β) fulfilling (TR.i)-(TR.v) have been defined for all $\beta < \alpha$, then (f^α, t^α) fulfilling (TR.i)-(TR.v) can be defined. Applying transfinite recursion then yields existence of the full sequences $(f^\alpha)_{\alpha \in Ord}$, $(t^\alpha)_{\alpha \in Ord}$.

Step 1. For all $i \in I$, define f_i^0 by $f_i^0(\tau) = \sigma_i^1(f, 0)$ for all $\tau \in \mathbb{R}_+$ for any $f \in F$. Note that by (CRM.i), $\sigma_i^1(f, 0)$ is independent of f . Set $t^0 := \min_{i \in I} \varepsilon_r^{\sigma_i^1}(f^0, 0)$ which exists and is strictly positive by Lemma 2. For t^0 and f^0 (TR.i)-(TR.iv) are trivially fulfilled. To see that (TR.v) holds, first note that $\sigma_i^1(f^0, \tau) = f_{i+}^0(\tau) = f_i^0(\tau)$ for all $\tau \in [0, t^0[$ and all $i \in I$ by definition of $\varepsilon_r^{\sigma_i^1}(f^0, 0)$. Second, $\sigma_i^2(f^0, \tau) = f_i^0(\tau) = f_{i+}^0(\tau)$ for all $\tau \in \mathbb{R}_+$ and all $i \in I$ by (CRM.ii) since f^0 is a constant function, and so (TR.v) is satisfied.

Step 2. Let $\alpha+1 \in Ord$ be a successor ordinal and suppose that f^α and $t^\alpha \in \mathbb{R}_+ \cup \{\infty\}$ satisfying (TR.i)-(TR.v) have been constructed.

We first construct $f^{\alpha+1}$. For all $i \in I$, define an intermediate function $\bar{f}_i^{\alpha+1} = G_i^-(f^\alpha, t^\alpha, a)$, where for all $i \in I$, $a_i = \sigma_i^1(f^\alpha, t^\alpha)$ and G^- is as in Definition 11. By Lemma B.1(i), $\bar{f}^{\alpha+1} \in F$. Now, for all $i \in I$, define $f_i^{\alpha+1} = G_i^+(\bar{f}^{\alpha+1}, t^\alpha, b)$, where for all $i \in I$, $b_i = \sigma_i^2(\bar{f}^{\alpha+1}, t^\alpha)$ and G^+ is as given in Definition 11. Note that for all $i \in I$, if $t^\alpha \in \bigcap_{j \in I} LC(\bar{f}_j^{\alpha+1}) \cup J(\bar{f}_i^{\alpha+1})$ then $\sigma_i^2(\bar{f}^{\alpha+1}, t^\alpha) = \bar{f}_i^{\alpha+1}(t^\alpha)$ by (CRM.ii). Hence $(\sigma_i^2(\bar{f}^{\alpha+1}, t^\alpha))_{i \in I}$ satisfies the conditions in Lemma B.1(ii) and it follows that $f^{\alpha+1} \in F$, i.e. (TR.i) holds.

To prove that (TR.ii) is satisfied let $\beta < \alpha+1$. By construction, $f^{\alpha+1}(\tau) = f^\alpha(\tau)$ for all $\tau \in [0, t^\alpha[$. If $\beta = \alpha$ this already shows (TR.ii). If $\beta < \alpha$, by (TR.ii) for α , $f^\beta(\tau) = f^\alpha(\tau) = f^{\alpha+1}(\tau)$ for all $\tau \in [0, t^\beta[$, where the last equality holds because $t^\beta \leq t^\alpha$ by (TR.iv) for α , and the conclusion follows.

Now define $t^{\alpha+1} := t^\alpha + \min_{i \in I} \varepsilon_r^{\sigma_i^2}(f^{\alpha+1}, t^\alpha)$. As $\alpha+1$ is a successor ordinal $(t^\beta)_{\beta \leq \alpha+1}$ is continuous if $(t^\beta)_{\beta \leq \alpha}$ is continuous. The latter sequence is continuous by induction hypothesis and hence (TR.iii) holds. To see that (TR.iv) is fulfilled, note that by induction hypothesis either $t^\alpha = \infty$ or $(t^\beta)_{\beta \leq \alpha} \subseteq \mathbb{R}_+$ is strictly increasing. If $t^\alpha = \infty$ then $t^{\alpha+1} = \infty$ by construction. If, on the other hand $(t^\beta)_{\beta \leq \alpha} \subseteq \mathbb{R}_+$ is strictly increasing then by construction $t^\alpha < t^{\alpha+1} < \infty$

because $0 < \varepsilon_r^{\sigma_i}(f^{\alpha+1}, t^\alpha) < \infty$ by Lemma 2. Thus $(t^\beta)_{\beta \leq \alpha+1} \subseteq \mathbb{R}_+$ is strictly increasing.

To prove (TR.v), first note that for all $\tau \in [0, t^\alpha[$ and all $i \in I$, $\sigma_i(f^{\alpha+1}, \tau) = \sigma_i(f^\alpha, \tau) = (f_i^\alpha(\tau), f_{i+}^\alpha(\tau)) = (f_i^{\alpha+1}(\tau), f_{i+}^{\alpha+1}(\tau))$. The first equality follows by construction of $f^{\alpha+1}$ and both parts of (CRM.i), the second from the induction hypothesis, and the third from the construction of $f^{\alpha+1}$. If $t^\alpha = \infty$, this already shows (TR.v). Hence we can now assume that $t^\alpha < \infty$ (and, by construction, $t^{\alpha+1} < \infty$). We now prove the first part of (TR.v). By construction of $f^{\alpha+1}$ and (CRM.i) $\sigma_i^1(f^{\alpha+1}, t^\alpha) = \sigma_i^1(f^\alpha, t^\alpha)$ and since $f_i^{\alpha+1}(t^\alpha) = \sigma_i^1(f^\alpha, t^\alpha)$ by construction of $f^{\alpha+1}$, we obtain $\sigma_i^1(f^{\alpha+1}, t^\alpha) = f_i^{\alpha+1}(t^\alpha)$ for all $i \in I$. Since $t^{\alpha+1} = t^\alpha + \varepsilon_r^{\sigma_i}(f^{\alpha+1}, t^\alpha)$, by definition of $\varepsilon_r^{\sigma_i}(f^{\alpha+1}, t^\alpha)$ it follows that $\sigma_i^1(f^{\alpha+1}, \tau) = f_{i+}^{\alpha+1}(t^\alpha)$ for all $\tau \in]t^\alpha, t^{\alpha+1}[$. Since $f^{\alpha+1}$ is constant on $]t^\alpha, \infty[$, we obtain $f_i^{\alpha+1}(\tau) = f_{i+}^{\alpha+1}(t^\alpha)$ for all $\tau \in]t^\alpha, \infty[$ and all $i \in I$ and hence $\sigma_i^1(f^{\alpha+1}, \tau) = f_i^{\alpha+1}(\tau)$ for all $\tau \in [0, t^{\alpha+1}[$. Now we turn to the second part of (TR.v). By construction of $f^{\alpha+1}$ and (CRM.i) $\sigma_i^2(f^{\alpha+1}, t^\alpha) = \sigma_i^2(\bar{f}^{\alpha+1}, t^\alpha)$ and since $f_{i+}^{\alpha+1}(t^\alpha) = \sigma_i^2(\bar{f}^{\alpha+1}, t^\alpha)$ by construction of $f^{\alpha+1}$, we obtain $\sigma_i^2(f^{\alpha+1}, t^\alpha) = f_{i+}^{\alpha+1}(t^\alpha)$ for all $i \in I$. Since $f^{\alpha+1}$ is constant on $]t^\alpha, \infty[$, (CRM.ii) yields $\sigma_i^2(f^{\alpha+1}, \tau) = f_{i+}^{\alpha+1}(\tau) = f_{i+}^{\alpha+1}(t^\alpha)$ for all $\tau \in]t^\alpha, \infty[$ and hence $\sigma_i^2(f^{\alpha+1}, \tau) = f_{i+}^{\alpha+1}(\tau)$ for all $\tau \in [0, t^{\alpha+1}[$.

Step 3. Let α be a limit ordinal and assume that f^β and $t^\beta \in \mathbb{R}_+ \cup \{\infty\}$ satisfying (TR.i)-(TR.v) have been constructed for all $\beta < \alpha$. Set $t^\alpha := \lim_{\beta \rightarrow \alpha} t^\beta$. We distinguish two cases. Suppose first that $t^\alpha = \infty$. Let $\alpha^* \leq \alpha$ be the first limit ordinal such that $t^{\alpha^*} = \infty$. Then by induction hypothesis, $(t^\beta)_{\beta < \alpha^*} \subseteq \mathbb{R}_+$ is strictly increasing and continuous and hence by Lemma B.2 for every $\tau \in [t^0, \infty[$ there is a unique $\beta < \alpha^*$ such that $\tau \in [t^\beta, t^{\beta+1}[$. Hence every $\tau \in \mathbb{R}_+$ is contained in some interval $[0, t^\beta[$ for some $\beta < \alpha^*$ and by induction hypothesis (TR.ii), f^α defined by $f^\alpha(\tau) = f^\beta(\tau)$ if $\tau \in [0, t^\beta[$ is well-defined. (TR.i)-(TR.iv) hold by induction hypothesis and construction of f^α and t^α . To see (TR.v) let $\tau \in [0, \infty[$. Then there is $\beta < \alpha^*$ such that $\tau \in [0, t^\beta[$. By induction hypothesis $\sigma_i(f^\beta, \tau) = (f_i^\beta(\tau), f_{i+}^\beta(\tau))$. By construction of f^α , $f^\alpha(\tau') = f^\beta(\tau')$ for all $\tau' \in [0, t^\beta[$ and hence by (CRM.i) $\sigma_i(f^\alpha, \tau) = \sigma_i(f^\beta, \tau)$. This yields $\sigma_i(f^\alpha, \tau) = (f_i^\alpha(\tau), f_{i+}^\alpha(\tau))$.

Suppose now that $t^\alpha < \infty$. Then by induction hypothesis $(t^\beta)_{\beta < \alpha} \subseteq \mathbb{R}_+$

is strictly increasing and continuous and by Lemma B.2, for every $\tau \in [t^0, t^\alpha[$ there is a unique $\beta < \alpha$ such that $\tau \in [t^\beta, t^{\beta+1}[$. In particular every $\tau \in [0, t^\alpha[$ is contained in some interval $[0, t^\beta[$ for some $\beta < \alpha$ and by (TR.ii) and for each $i \in I$, the following intermediate function is well-defined.

$$\tilde{f}_i^\alpha(\tau) := \begin{cases} f_i^\beta(\tau), & \text{if } \tau < t^\alpha \text{ and } \tau \in [0, t^\beta[\\ a_i, & \text{if } \tau \geq t^\alpha. \end{cases}$$

(Where $a_i \in A_i$ is arbitrary.) By construction (as in the proof of Lemma B.1(i)) $\tilde{f}^\alpha \in F$. Then for each $i \in I$, define $\bar{f}_i^\alpha = G_i^-(\tilde{f}^\alpha, t^\alpha, a)$, where for all $i \in I$, $a_i = \sigma_i^1(\tilde{f}^\alpha, t^\alpha)$ and G^- is as in Definition 11. Note that by Lemma B.1(i), $\bar{f}^\alpha \in F$. Now, for all $i \in I$, define $f_i^\alpha = G_i^+(\bar{f}^\alpha, t^\alpha, b)$, where for all $i \in I$, $b_i = \sigma_i^2(\bar{f}^\alpha, t^\alpha)$ and G^+ is as given in Definition 11. Note that for all $i \in I$, if $t^\alpha \in \bigcap_{j \in I} LC(\bar{f}^\alpha) \cup J(f_j)$ then $\sigma_i^2(\bar{f}^\alpha, t^\alpha) = \bar{f}_i^\alpha(t^\alpha)$ by (CRM.ii). Hence $(\sigma_i^2(\bar{f}^\alpha, t^\alpha))_{i \in I}$ satisfies the conditions in Lemma B.1(ii) and hence $f^\alpha \in F$, so (TR.i) is satisfied. (TR.ii) and (TR.iii) follow directly by induction hypothesis and the constructions of f^α and t^α . To see (TR.iv) note that by induction hypothesis $(t^\beta)_{\beta < \alpha}$ is strictly increasing and hence, as $t^\alpha = \lim_{\beta \rightarrow \alpha} t^\beta$, $(t^\beta)_{\beta \leq \alpha}$ is strictly increasing. To see (TR.v) let $\tau \in [0, t^\alpha[$. Then there is $\beta < \alpha$ such that $\tau \in [0, t^\beta[$. By induction hypothesis $\sigma_i(f^\beta, \tau) = (f_i^\beta(\tau), f_{i+}^\beta(\tau))$. By construction of f^α , $f^\alpha(\tau') = f^\beta(\tau')$ for all $\tau' \in [0, t^\beta[$ and hence by (CRM.i) $\sigma_i(f^\alpha, \tau) = \sigma_i(f^\beta, \tau)$. This yields $\sigma_i(f^\alpha, \tau) = (f_i^\alpha(\tau), f_{i+}^\alpha(\tau))$ for all $\tau \in [0, t^\alpha[$.

This completes the construction. Transfinite recursion now yields sequences $(f^\alpha)_{\alpha \in Ord}$ and $(t^\alpha)_{\alpha \in Ord}$ satisfying (TR.i)-(TR.v) for all $\alpha \in Ord$. Then there exists a limit ordinal α such that $t^\alpha = \infty$.¹¹ Let α^* be the first limit ordinal such that $t^{\alpha^*} = \infty$. As by (TR.iii) $\lim_{\alpha \rightarrow \alpha^*} t^\alpha = \infty$, Lemma B.2 implies that for all $\tau \in [t^0, \infty[$ there is a unique $\beta < \alpha^*$ such that $\tau \in [t^\beta, t^{\beta+1}[$. Hence for all $\tau \in \mathbb{R}_+$ there is $\beta \in Ord$ such that $\tau \in [0, t^\beta[$. Then, using (TR.ii), f_i given by $f_i(\tau) := f_i^\beta(\tau)$ for $\tau \in [0, t^\beta[$ is well-defined.

From the construction of f and because $f^\alpha \in F$ by (TR.iii) for every

¹¹Otherwise, we would have a strictly increasing mapping from the class of ordinals to \mathbb{R} , which is impossible (e.g. by Lemma III.2 in Stinchcombe, 1992).

$\alpha \in \text{Ord}$, it follows that $f \in F$. To see that $\sigma_i(f, \tau) = (f_i(\tau), f_{i+}(\tau))$ for all $\tau \in \mathbb{R}_+$ and all $i \in I$ let $\tau \in \mathbb{R}_+$. Then there is $\beta < \alpha^*$ such that $\tau \in [0, t^\beta[$. As f^β satisfies (TR.v) it follows that $\sigma_i(f^\beta, \tau) = (f_i^\beta(\tau), f_{i+}^\beta(\tau))$. By construction of f , $f(\tau') = f^\beta(\tau')$ for all $\tau' \in [0, t^\beta[$ and hence by (CRM.i) $\sigma_i(f, \tau) = \sigma_i(f^\beta, \tau)$. This yields $\sigma_i(f, \tau) = (f_i(\tau), f_{i+}(\tau))$ for all $\tau \in [0, \infty[$.

Finally, we will prove that f is unique. Let $f' \in F$ be such that $\sigma_i^1(f', \tau) = f'_i(\tau)$ and $\sigma_i^2(f', \tau) = f'_{i+}(\tau)$ for all $\tau \in \mathbb{R}_+$ and all $i \in I$. Assume $f' \neq f$. Then $\bar{t} := \inf\{\tau \in \mathbb{R}_+ | f'(\tau) \neq f(\tau)\}$ exists. We claim that $\bar{t} > 0$. To see this, note that $f'(0) = f(0)$ by (CRM.i). By (CRM.ii) there is $\varepsilon > 0$ such that $f'(\tau) = f(\tau)$ for all $\tau \in [0, \varepsilon[$. This proves the claim. Because $f'(\tau) = f(\tau)$ for all $\tau \in [0, \bar{t}[$, (CRM.i) implies that $f'_i(\bar{t}) = \sigma_i^1(f', \bar{t}) = \sigma_i^1(f, \bar{t}) = f_i(\bar{t})$ for all $i \in I$. Further, it follows from (CRM.i) that $f'_{i+}(\bar{t}) = \sigma_i^2(f', \bar{t}) = \sigma_i^2(f, \bar{t}) = f_{i+}(\bar{t})$ for all $i \in I$. Hence, as $f', f \in F$, by (DP.i) there is $\varepsilon > 0$ such that $f'(\tau) = f(\tau)$ for all $\tau \in [0, \bar{t} + \varepsilon[$, which contradicts the construction of $\bar{t} = \inf\{\tau \in \mathbb{R}_+ | f'(\tau) \neq f(\tau)\}$.

(ii) Given $\bar{f} \in F$ and $t \in \mathbb{R}_{++}$, define $a_i^0 = \sigma_i(\bar{f}, t)$ for all $i \in I$. Define $f^0 = G^-(\bar{f}, t, a^0)$ and set $t^0 = t + \min_{i \in I} \varepsilon_r^{\sigma_i}(\bar{f}, t)$ where $\varepsilon_r^{\sigma_i}(\bar{f}, t)$ is defined as in the proof of (i). The rest of the proof is analogous to the proof of (i).

(iii) Let $\bar{f} \in F$. For each $i \in I$ fix some arbitrary $a_i \in A_i$. For all $i \in I$, $f \in F$, and $\tau \in \mathbb{R}_+$ define

$$\sigma_i^1(f, t) = \begin{cases} \bar{f}_i(t), & \text{if } f(\tau) = \bar{f}(\tau) \text{ for all } \tau \in [0, t[\\ \lim_{\tau \rightarrow t^-} f_i(\tau), & \text{if } f(\tau) \neq \bar{f}(\tau) \text{ for some } \tau \in [0, t[\text{ and } \exists \lim_{\tau \rightarrow t^-} f(\tau), \\ a_i, & \text{otherwise,} \end{cases}$$

and

$$\sigma_i^2(f, t) = \begin{cases} \bar{f}_{i+}(t), & \text{if } f(\tau) = \bar{f}(\tau) \text{ for all } \tau \in [0, t], \\ f_i(t), & \text{otherwise.} \end{cases}$$

(CRM.i) holds by construction of σ_i .

To see (CRM.ii) let $i \in I$, $f \in F$ and $t \in \bigcap_{j \in I} LC(f_j) \cup J(f_i)$. Note that by (DP.ii) and (DP.iii), $t \in RK(f_i)$. Hence by (DP.i) there is $\varepsilon > 0$ such that f_i is constant on $[t, t + \varepsilon[$ and in particular, $\lim_{r \rightarrow \tau^-} f(r) = f(\tau)$ for all $\tau \in]t, t + \varepsilon[$.

To show the first part of (CRM.ii), we will distinguish two cases. First, if $f(\tau) = \bar{f}(\tau)$ for all $\tau \in [0, t]$ then also $t \in \bigcap_{j \in I} LC(\bar{f}_j) \cup J(\bar{f}_i)$ which implies $t \in RK(\bar{f}_i)$ and hence $\bar{f}_{i+}(t) = \bar{f}_i(t)$. Hence by construction, $\sigma_i^2(f, t) = \bar{f}_{i+}(t) = \bar{f}_i(t) = f_i(t)$. Second, if there is $\tau \leq t$ such that $f(\tau) \neq \bar{f}(\tau)$ then $\sigma_i^2(f, t) = f_i(t)$ by construction. To prove the second part of (CRM.ii), we will again distinguish two cases. First, if there is $\varepsilon' > 0$ such that $f(\tau) = \bar{f}(\tau)$ for all $\tau \in [0, t + \varepsilon'[$ then $\sigma_i^1(f, \tau) = \bar{f}_i(\tau)$ for all $\tau \in]t, t + \varepsilon'[$. Since $t \in RK(f_i)$, $f_i(\tau) = f_i(t)$ for all $\tau \in]t, t + \varepsilon[$ and hence $\sigma_i^1(f, \tau) = \bar{f}_i(\tau) = f_i(\tau) = f_i(t)$ for all $\tau \in]t, t + \min\{\varepsilon, \varepsilon'\}[$. Second, if for every $\tau > t$ there is $\tau' < \tau$ such that $f(\tau') \neq \bar{f}(\tau')$, then by construction $\sigma_i^1(f, \tau) = \lim_{r \rightarrow \tau^-} f_i(r) = f_i(\tau)$ for all $\tau \in]t, t + \varepsilon[$. As $t \in RK(f_i)$, $f_i(\tau) = f_i(t)$ for all $\tau \in]t, t + \varepsilon[$ and hence $\sigma_i^1(f, \tau) = f_i(\tau) = f_i(t)$ for all $\tau \in]t, t + \varepsilon[$.

To establish (CRM.iii) let $i \in I$, $f \in F$, and $t \in LC(f_i) \cap \bigcup_{j \in I} J(f_j)$. Note that by (DP.i) there is $\varepsilon > 0$ such that f is constant on $]t, t + \varepsilon[$ and in particular $\lim_{r \rightarrow \tau^-} f(r) = f(\tau)$ for all $\tau \in]t, t + \varepsilon[$. We distinguish two cases. First, if there is $\varepsilon' > 0$ such that $f(\tau) = \bar{f}(\tau)$ for all $\tau \in [0, t + \varepsilon'[$ then $\sigma_i(f, \tau) = \bar{f}_i(\tau)$ for all $\tau \in]t, t + \varepsilon'[$. Since f is constant on $]t, t + \varepsilon[$, $f_i(\tau) = f_{i+}(t)$ for all $\tau \in]t, t + \varepsilon[$. We thus obtain $\sigma_i(f, \tau) = \bar{f}_i(\tau) = f_i(\tau) = f_{i+}(t)$ for all $\tau \in]t, t + \min\{\varepsilon, \varepsilon'\}[$. Second, if for every $\tau > t$ there is $\tau' < \tau$ such that $f(\tau') \neq \bar{f}(\tau')$, then by construction $\sigma_i^1(f, \tau) = \lim_{r \rightarrow \tau^-} f_i(r) = f_i(\tau)$ for all $\tau \in]t, t + \varepsilon[$. As f is constant on $]t, t + \varepsilon[$, $f_i(\tau) = f_{i+}(t)$ for all $\tau \in]t, t + \varepsilon[$ and hence $\sigma_i^1(f, \tau) = f_i(\tau) = f_{i+}(t)$ for all $\tau \in]t, t + \varepsilon[$.

This shows that σ_i is a CRM for every $i \in I$. Since $\sigma_i^1(\bar{f}, t) = \bar{f}_i(t)$ and $\sigma_i^2(\bar{f}, t) = \bar{f}_{i+}(t)$ for all $t \in \mathbb{R}_+$ and all $i \in I$, σ induces \bar{f} . \square

Appendix 1.C: Proofs from Section 1.5

This appendix contains the proofs of Lemmata 3, and 4, Propositions 3 and 4, and Theorem 2. We start with a few preliminary results.

Lemma C.1. (i) Let $f \in F$ and $t_1, t_2 \in \mathbb{R}_+$ be such that $]t_1, t_2[\cap J(f_i) = \emptyset$ for all $i \in I$. Then f is constant on $]t_1, t_2[$.

(ii) Let $(f, t) \in F \times \mathbb{R}_+$ and $\bar{t} \in [t(f, t), t]$. Then $f^{\bar{t}+}(\tau) = f(\tau)$ for all $\tau \in [0, t]$.

Proof. (i) By (DP.i), f_i is piecewise constant for all $i \in I$. Hence there is $\varepsilon > 0$ and $a \in \times_{i \in I} A_i$ such that $f(\tau) = a$ for all $\tau \in]t_1, t_1 + \varepsilon[$. Suppose by contradiction that there is $t \in [t_1 + \varepsilon, t_2[$ such that $f(t) \neq a$. Then $t^* := \inf\{t \in [t_1 + \varepsilon, t_2[\mid f(t) \neq a\}$ exists. By (DP.ii) $t^* \in LC(f_i) \cup RK(f_i)$ for all $i \in I$. If $t^* \notin RK(f_i)$ for some $i \in I$, then $t^* \in R(f_i) = LC(f_i) \setminus RK(f_i)$ and, by (DP.iii), there is $j \in I$, such that $t^* \in J(f_j)$, which contradicts $]t_1, t_2[\cap J(f_j) = \emptyset$. It follows that $t^* \in RK(f_i)$ for all $i \in I$. Further, $t^* \in LC(f_i)$ for all $i \in I$, as otherwise $t^* \in RK(f_i) \setminus LC(f_i) = J(f_i)$ for some $i \in I$ which again contradicts $]t_1, t_2[\cap J(f_j) = \emptyset$. Hence $t^* \in LC(f_i) \cap RK(f_i)$ for all $i \in I$ and since $f(\tau) = a$ for all $\tau \in]t_1, t^*[$ by construction of t^* , we obtain $f(t^*) = a$. Further, by (DP.i), there is $\varepsilon' > 0$ and $b \in \times_{i \in I} A_i$ such that $f(\tau) = b$ for all $\tau \in]t^*, t^* + \varepsilon'[$. As $t^* \in RK(f_i)$ for all $i \in I$, we obtain $f(t^*) = b$ and hence $a = b$. Thus there is $\varepsilon' > 0$ such that $f(\tau) = a$ for all $]t_1, t^* + \varepsilon'[$, which contradicts the definition of t^* .

(ii) If $t = t(f, t)$ the conclusion follows by construction. Hence suppose that $t > t(f, t)$. By construction $f^{\bar{t}^+}(\tau) = f(\tau)$ for all $\tau \in [0, \bar{t}]$. By (i), f is constant on $]t(f, t), t[$ and hence $f(\tau) = f_+(\bar{t})$ for all $\tau \in]\bar{t}, t[$. Thus $f(\tau) = f^{\bar{t}^+}(\tau)$ for all $\tau \in]\bar{t}, t[$ and the conclusion follows. \square

For $f \in F$ and $t \in \mathbb{R}_+$ define

$$t^+(f, t) := \begin{cases} \min \bigcup_{i \in I} J(f_i) \cap]t, +\infty[, & \text{if } \bigcup_{i \in I} J(f_i) \cap]t, +\infty[\neq \emptyset \\ t, & \text{otherwise.} \end{cases}$$

That is, $t^+(f, t)$ is the next time after t that some player jumps. Note that the minimum used in the construction exists by Lemma A.2.

The following lemma shows that every decision path can be “completed” to a play by appropriately specifying the inertia times.

Lemma C.2. *For every $f \in F$ there is $\epsilon \in E$ such that $(f, \epsilon) \in W$.*

Proof. Fix $f \in F$ and let $t^* = \max \mathcal{J}(f)$ be the *last jump* of f if $\max \mathcal{J}(f)$ exists and $t^* = +\infty$ otherwise. Note that if $t \in \mathcal{J}(f) \setminus \{t^*\}$ then $t^+(f, t) > t$ since then by construction of t^* , $\bigcup_{i \in I} J(f_i) \cap]t, +\infty[\neq \emptyset$. Further, if $t^* < \infty$ then $t^+(f, t^*) = t^*$.

For each $i \in I$ define $\epsilon_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ as follows

$$\epsilon_i(t) := \begin{cases} t^+(f, t) - t, & \text{if } t \in \mathcal{J}(f) \setminus \{t^*\}, \\ 73, & \text{if } t = t^* + 73n \text{ for some } n \in \mathbb{N}_0, \\ 0, & \text{otherwise.} \end{cases}$$

Note that as $0 \in J(f_i)$ for all $i \in I$ it follows that $\epsilon_i(0) > 0$ for all $i \in I$ and hence $\epsilon = (\epsilon_i)_{i \in I} \in E$. By construction, for all $i \in I$,

$$DP(\epsilon_i) = \mathcal{J}(f) \cup \{t \in \mathbb{R}_+ \mid t = t^* + 73n \text{ for some } n \in \mathbb{N}\}. \quad (1.3)$$

It remains to show that $w = (f, \epsilon) \in W$, i.e. that (P.i)-(P.iv) in Definition 4 hold. To prove (P.i), let $t \in J(f_i)$ for some $i \in I$. If $t = t(f, t)$ it follows from (1.3) that $t = \text{Prev}(\epsilon_i, t)$ and hence $t \in PJ(\epsilon_i)$ by definition of the latter. If $t > t(f, t)$, then $t = t^+(t(f, t)) > t(f, t)$ where the inequality holds because $t \leq t^*$. As $t^* \geq t > t(f, t) \in \mathcal{J}(f)$, it follows that $\epsilon_i(t(f, t)) = t^+(t(f, t)) - t(f, t) > 0$, i.e. $t = t(f, t) + \epsilon_i(t(f, t))$. By Lemma C.1, f is constant on $]t(f, t), t[$ and hence $\epsilon_i(\tau) = 0$ for all $\tau \in]t(f, t), t[$. Thus $t(f, t) = \text{Prev}(\epsilon_i, t)$, which implies that $t = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t))$ and $t \in PJ(\epsilon_i)$. To see (P.ii), let again $t \in J(f_i)$ for some $i \in I$. Then, by construction, $\epsilon_j(t) > 0$ for all $j \in I$ and hence by (1.3) $t \in \bigcap_{j \in I} DP(\epsilon_j)$. To establish (P.iii), let $t \in PJ(\epsilon_i)$ for some $i \in I$. If $\text{Prev}(\epsilon_i, t) = t$, which by (1.3) is only possible if $t = t(f, t)$, then $t \in \mathcal{J}(f)$ and hence $\epsilon_i(t) > 0$, i.e. $t \in DP(\epsilon_i)$. Otherwise $t = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t))$. Then, by (1.3), either $\text{Prev}(\epsilon_i, t) \in \mathcal{J}(f) \setminus \{t^*\}$ or $\text{Prev}(\epsilon_i, t) = t^* + 73n$ for some $n \in \mathbb{N}_0$ implying that either $t \in \bigcup_{j \in I} J(f_j)$ or $t = t^* + 73n$ for some $n \in \mathbb{N}_0$. In any case it follows from (1.3) that $\epsilon_i(t) > 0$. To see (P.iv), let $t \in DP(\epsilon_i)$ for some $i \in I$. Then $\bigcup_{j \in I} J(f_j) \cap]t, \text{Next}(\epsilon_i, t)[= \emptyset$ and $\epsilon_i(\tau) = 0$ for all $\tau \in]t, \text{Next}(\epsilon_i, t)[$ by (1.3). Hence (P.iv) holds. \square

Lemma C.3. (i) Let $(f, \epsilon) \in W$, $t \in \mathbb{R}_+$. If $t(f, t) = t$ then $t \in PJ(\epsilon_i)$ for all $i \in I$.

(ii) Let $f \in F$, $i \in I$, $s_i \in S_i$, and $t \in \mathbb{R}_+$. If $t(f, t) = t$ then $t \in PJ(\epsilon_i)$ for any $\epsilon \in E$ such that $(f^{t(f, t)^+}, \epsilon) \in W$ and $(f^{t(f, t)^+}, \epsilon)$ agrees with s_i on $[t(f, t), \infty[$.

Proof. (i) If $t = 0$ then $\text{Prev}(\epsilon_i, 0) = 0$ and hence $0 \in PJ(\epsilon_i)$ by definition. If $t > 0$ then, since $t(f, t) = \sup \left(\bigcup_{j \in I} J(f_j) \cap [0, t[\right) = t$ and $\bigcup_{j \in I} J(f_j) \cap [0, t[\subseteq DP(\epsilon_i) \cap [0, t[$ by (P.ii), $\text{Prev}(\epsilon_i, t) = \sup(DP(\epsilon_i) \cap [0, t[) = t$ and hence $t \in PJ(\epsilon_i)$.

(ii) This follows from (i) since $(f^{t(f,t)^+}, \epsilon) \in W$ by hypothesis. \square

Proof of Lemma 3. (i) First note that $f^{t^+} = G^+(f, t, f_+(t))$ (recall Definition 11 in Section 1.7). Since $f \in F$, it follows that $f_{j^+}(t) = f_j(t)$ if $t(f, t) \in \bigcap_{k \in I} LC(f_k) \cup J(f_j)$ for all $j \in I$, and hence $f^{t^+} \in F$ by Lemma B.1(ii).

We are going to use transfinite recursion to construct a sequence of functions $(\epsilon^\alpha)_{\alpha \in \text{Ord}}$ and a sequence of extended real numbers $(t^\alpha)_{\alpha \in \text{Ord}} \subseteq \mathbb{R}_+ \cup \{\infty\}$ such that for all $\alpha \in \text{Ord}$ the following properties are satisfied.

(TR.i) $(f^{t^+}, \epsilon^\alpha) \in W$.

(TR.ii) If $\beta < \alpha$ then $\epsilon^\alpha(\tau) = \epsilon^\beta(\tau)$ for all $\tau \in [0, t^\beta[$.

(TR.iii) $(t^\beta)_{\beta \leq \alpha}$ is continuous.

(TR.iv) Either $t^\alpha = \infty$ or $(t^\beta)_{\beta \leq \alpha} \subseteq]t, \infty[\cap \left(\bigcap_{j \in I} PJ(\epsilon_j^\alpha) \right)$ and is strictly increasing.

(TR.v) If $\epsilon_i^\alpha(\tau) > 0$ then $s_i(x_\tau^P((f^{t^+}, \epsilon^\alpha))) = c_i(x_\tau^P((f^{t^+}, \epsilon^\alpha)), \epsilon^\alpha(\tau))$ for all $\tau \in [t, t^\alpha[$

In order to apply transfinite recursion, we need to complete three steps. First, we will define (ϵ^0, t^0) trivially fulfilling (TR.i)-(TR.v). Second, we will show that, if (TR.i)-(TR.v) are fulfilled for an ordinal α then $(\epsilon^{\alpha+1}, t^{\alpha+1})$ fulfilling (TR.i)-(TR.v) can be defined for the successor ordinal $\alpha + 1$. Third, we will show that, for any limit ordinal α , if $(\epsilon^\beta, t^\beta)$ fulfilling (TR.i)-(TR.v) have been defined for all $\beta < \alpha$, then $(\epsilon^\alpha, t^\alpha)$ fulfilling (TR.i)-(TR.v) can be defined. Applying transfinite recursion then yields existence of the full sequences $(\epsilon^\alpha)_{\alpha \in \text{Ord}}$, $(t^\alpha)_{\alpha \in \text{Ord}}$.

Step 1. Let $\bar{\epsilon} \in E$ be such that $(f^{t^+}, \bar{\epsilon}) \in W$ which exists by Lemma C.2. Note that $x_t^P((f^{t^+}, \bar{\epsilon})) \in N$ and player i has to make a choice at this node. This is because either $t \in J(f_j)$ for some $j \in I$, in which case $i \in$

$IDP(\bar{\epsilon}, t)$ by (P.ii) or $t = t(f, t)$ in which case $t = \text{Prev}(\bar{\epsilon}, t)$ by (P.ii) and definition of $\text{Prev}(\cdot, \cdot)$ and hence $i \in IPJ(\bar{\epsilon}, t)$. In both cases, by Lemma A.5, $IPJ(\bar{\epsilon}, t) \neq \emptyset$ and $x^P((f^{t+}, \bar{\epsilon})) \in X_i$. Set $t^0 := t + \varepsilon^0$, where $\varepsilon^0 > 0$ is such that $s_i(x_i^P((f^{t+}, \bar{\epsilon}))) = c_i(x_i^P((f^{t+}, \bar{\epsilon})), \varepsilon^0)$. Since $(f^{t+}, \bar{\epsilon}) \in W$, by (P.i) $t \in J(f_j^{t+})$ only if $t \in PJ(\bar{\epsilon}_j)$. Further, since $t \in \mathcal{J}(f)$, it follows that $t \in \mathcal{J}(f^{t+})$ by construction of f^{t+} . Then, by Lemma A.4(ii), there is $\epsilon^0 \in E$ such that $\epsilon^0(\tau) = \bar{\epsilon}(\tau)$ for all $\tau \in [0, t]$, $\epsilon_j^0(t) = \varepsilon^0$ for all $j \in I$, and $(f^{t+}, \epsilon^0) \in W$. Thus (TR.i) holds. (TR.ii) and (TR.iii) are trivially fulfilled. (TR.iv) holds by construction, since $(f^{t+}, \epsilon^0) \in W$ and hence by (P.iv) $\epsilon_j^0(\tau) = 0$ for all $\tau \in]t, t^0[$ and all $j \in I$. (TR.v) holds since by construction $\epsilon_i^0(t) = \varepsilon^0$ and $\epsilon_i^0(\tau) = 0$ for all $\tau \in]t, t^0[$.

Step 2. Let $\alpha + 1 \in \text{Ord}$ be a successor ordinal and suppose that ϵ^α and $t^\alpha \in \mathbb{R}_+ \cup \{\infty\}$ satisfying (TR.i)-(TR.v) have been constructed. We distinguish two cases. Suppose first that $t^\alpha = \infty$. In this case set $t^{\alpha+1} = \infty$ and define $\epsilon^{\alpha+1} = \epsilon^\alpha$. Then (TR.i)-(TR.v) are satisfied by induction hypothesis. Suppose now that $t^\alpha < \infty$. Since $t^\alpha \in \bigcap_{j \in I} PJ(\epsilon_j^\alpha)$ by induction hypothesis, it follows from (P.iii) that $x_{t^\alpha}^P((f^{t+}, \epsilon^\alpha)) \in X_i$. Set $t^{\alpha+1} = t^\alpha + \varepsilon^\alpha$, where $\varepsilon^\alpha > 0$ is such that $s_i(x_{t^\alpha}^P((f^{t+}, \epsilon^\alpha))) = c_i(x_{t^\alpha}^P((f^{t+}, \epsilon^\alpha)), \varepsilon^\alpha)$. Applying Lemma A.4(i) to f^{t+} , t^α , and ε^α yields that there is $\epsilon^{\alpha+1} \in E$ such that $\epsilon^{\alpha+1}(\tau) = \epsilon^\alpha(\tau)$ for all $\tau \in [0, t^\alpha]$, $\epsilon_j^{\alpha+1}(t^\alpha) = \varepsilon^\alpha$ for all $j \in IPJ(\epsilon^\alpha, t^\alpha) = I$ and $(f^{t+}, \epsilon^{\alpha+1}) \in W$. Hence (TR.i) holds. To see (TR.ii) let $\beta < \alpha + 1$. By construction, $\epsilon^{\alpha+1}(\tau) = \epsilon^\alpha(\tau)$ for all $\tau \in [0, t^\alpha]$. If $\beta = \alpha$ this already shows (TR.ii). If $\beta < \alpha$, by (TR.ii) for α , $\epsilon^\beta(\tau) = \epsilon^\alpha(\tau) = \epsilon^{\alpha+1}(\tau)$ for all $\tau \in [0, t^\beta]$, where the last equality holds because $t^\beta \leq t^\alpha$ by (TR.iv) for α , and the conclusion follows. To see (TR.iii) note that $(t^\beta)_{\beta \leq \alpha}$ is continuous by induction hypothesis and hence $(t^\beta)_{\beta \leq \alpha+1}$ is continuous as $\alpha + 1$ is a successor ordinal. To see (TR.iv), note that since $t^\alpha < \infty$ by induction hypothesis $(t^\beta)_{\beta \leq \alpha} \subseteq]t, \infty[\cap \left(\bigcap_{j \in I} PJ(\epsilon_j^\alpha) \right)$ is strictly increasing. Then by construction $t^\alpha < t^{\alpha+1} < \infty$ because $0 < \varepsilon^\alpha < \infty$ by construction of C_i . Further, that $t^\alpha \in \bigcap_{j \in I} PJ(\epsilon_j^\alpha)$ implies $t^\alpha \in \bigcap_{j \in I} PJ(\epsilon_j^{\alpha+1})$ by construction. Since $(f^{t+}, \epsilon^{\alpha+1}) \in W$, it follows by (P.iv) that $\epsilon^{\alpha+1}(\tau) = 0$ for all $\tau \in]t^\alpha, t^{\alpha+1}[$, and hence $t^{\alpha+1} \in]t, \infty[\cap \left(\bigcap_{j \in I} PJ(\epsilon_j^{\alpha+1}) \right)$. Thus $(t^\beta)_{\beta \leq \alpha+1} \subseteq]t, \infty[\cap \left(\bigcap_{j \in I} PJ(\epsilon_j^\alpha) \right)$ is strictly increasing. (TR.v) is satisfied by construction and induction hypothesis.

Step 3. Let α be a limit ordinal and assume that ϵ^β and $t^\beta \subseteq \mathbb{R}_+ \cup \{\infty\}$ satisfying (TR.i)-(TR.v) have been constructed for all $\beta < \alpha$. Set $t^\alpha = \lim_{\beta \rightarrow \alpha} t^\beta$. Suppose first that $t^\alpha = \infty$. Let $\alpha^* \leq \alpha$ be the first limit ordinal such that $t^{\alpha^*} = \infty$. Then by induction hypothesis, $(t^\beta)_{\beta < \alpha^*} \subseteq]t, \infty[$ is strictly increasing and continuous and hence by Lemma B.2 for every $\tau \in [t^0, \infty[$ there is a unique $\beta < \alpha^*$ such that $\tau \in [t^\beta, t^{\beta+1}[$. Hence every $\tau \in \mathbb{R}_+$ is contained in some interval $[0, t^\beta[$ for some $\beta < \alpha^*$ and by induction hypothesis (TR.ii), ϵ^α defined by $\epsilon^\alpha(\tau) = \epsilon^\beta(\tau)$ if $\tau \in [0, t^\beta[$ is well-defined. (TR.i)-(TR.v) hold by induction hypothesis and construction of ϵ^α and t^α .

Suppose now that $t^\alpha < \infty$. Then by induction hypothesis $(t^\beta)_{\beta < \alpha} \subseteq]t, \infty[$ is strictly increasing and continuous and by Lemma B.2, for every $\tau \in [t^0, t^\alpha[$ there is a unique $\beta < \alpha$ such that $\tau \in [t^\beta, t^{\beta+1}[$. In particular every $\tau \in [0, t^\alpha[$ is contained in some interval $[0, t^\beta[$ for some $\beta < \alpha$ and by (TR.ii) and for each $i \in I$, the following function is well-defined.

$$\epsilon_j^\alpha(\tau) := \begin{cases} \epsilon_j^\beta(\tau), & \text{if } \tau < t^\alpha \text{ and } \tau \in [0, t^\beta[\\ 73, & \text{if } \tau = t^\alpha + 73n \text{ for some } n \in \mathbb{N}_0 \\ 0, & \text{otherwise.} \end{cases}$$

To see (TR.i), let $\tau \in J(f_j^{t^+})$ for some $j \in I$. By definition of f^{t^+} and construction of t^α , $\tau \leq t < t^\alpha$. There is $\beta < \alpha$ such $\tau \in [0, t^\beta[$ and we obtain $\tau \in PJ(\epsilon_j^\beta) \cap \bigcap_{k \in I} DP(\epsilon_k^\beta)$ by (TR.i) for $(f^{t^+}, \epsilon^\beta)$. Hence $\tau \in PJ(\epsilon_j^\alpha) \cap \bigcap_{k \in I} DP(\epsilon_k^\alpha)$ by construction. This establishes (P.i) and (P.ii). To see (P.iii) let $\tau \in PJ(\epsilon_j^\alpha)$ for some $j \in I$. We distinguish two cases. First, if $\tau < t^\alpha$ then $\tau \in DP(\epsilon_j^\beta)$ for some $\beta < \alpha$ by induction hypothesis. Hence by construction $\tau \in DP(\epsilon_j^\alpha)$. Second, if $\tau \geq t^\alpha$ then $\tau \in PJ(\epsilon_j^\alpha)$ implies that τ is of the form $t^\alpha + 73n$ for some $n \in \mathbb{N}_0$ and hence $\tau \in DP(\epsilon_j^\alpha)$ by construction of ϵ^α . To prove (P.iv), let $\tau \in DP(\epsilon_j^\alpha)$. We again distinguish two cases. First, if $\tau < t^\alpha$ then $\tau \in [0, t^\beta[$ for some $\beta < \alpha$. By (P.iv) for $(f^{t^+}, \epsilon^\beta)$, $\bigcup_{k \neq j} J(f_k^{t^+}) \cap]\tau, \tau'] \neq \emptyset$ for any $\tau' \in DP(\epsilon^\beta) \cap]\tau, \text{Next}(\epsilon_j^\beta, \tau)[$. Hence by construction of ϵ^α , $\bigcup_{k \neq j} J(f_k^{t^+}) \cap]\tau, \tau'] \neq \emptyset$ for any $\tau' \in DP(\epsilon^\alpha) \cap]\tau, \text{Next}(\epsilon_j^\alpha, \tau)[$. Second, if $\tau \geq t^\alpha$ note that $\bigcup_{j \in I} J(f_j^{t^+}) \cap]\tau, \infty[= \emptyset$ by definition of f^{t^+} . Since $DP(\epsilon_j^\alpha) \cap]\tau, \text{Next}(\epsilon_j^\alpha, \tau)[= \emptyset$ by construction of ϵ^α , the conclusion follows.

(TR.ii) and (TR.iii) follow directly by induction hypothesis and the constructions of ϵ^α and t^α . To see (TR.iv) note that by induction hypothesis $(t^\beta)_{\beta < \alpha}$ is strictly increasing and hence, as $t^\alpha = \lim_{\beta \rightarrow \alpha} t^\beta$, $(t^\beta)_{\beta \leq \alpha}$ is strictly increasing. Further, by induction hypothesis $t^\beta \in \bigcap_{j \in I} PJ(\epsilon_j^\alpha)$ for all $\beta < \alpha$ and hence by construction $t^\beta \in \bigcap_{j \in I} PJ(\epsilon_j^\alpha)$ for all $\beta < \alpha$. It follows from (P.iii) that for all $j \in I$, $t^\alpha = \text{Prev}(\epsilon_j^\alpha, t^\alpha)$ and thus that $t^\alpha \in \bigcap_{j \in I} PJ(\epsilon_j^\alpha)$. (TR.v) holds by construction and induction hypothesis.

This completes the construction of the sequences. Transfinite recursion now yields sequences $(\epsilon^\alpha)_{\alpha \in Ord}$ and $(t^\alpha)_{\alpha \in Ord}$ satisfying (TR.i)-(TR.v) for all $\alpha \in Ord$. Then there exists a limit ordinal α such that $t^\alpha = \infty$.¹² Let α^* be the first limit ordinal such that $t^{\alpha^*} = \infty$. As by (TR.iii) $\lim_{\beta \rightarrow \alpha^*} t^\beta = \infty$, Lemma B.2 implies that for all $\tau \in [t^0, \infty[$ there is a unique $\beta < \alpha^*$ such that $\tau \in [t^\beta, t^{\beta+1}[$. Hence for all $\tau \in \mathbb{R}_+$ there is $\beta \in Ord$ such that $\tau \in [0, t^\beta[$. Then, using (TR.ii), ϵ_j defined by $\epsilon_j(\tau) := \epsilon_j^\beta(\tau)$ if $\tau \in [0, t^\beta[$ is well-defined.

From the construction of ϵ and because $(f^{t^+}, \epsilon^\beta) \in W$ by (TR.iii) for every $\beta \in Ord$, it follows that $(f^{t^+}, \epsilon) \in W$. To see that (f^{t^+}, ϵ) agrees with s_i on $[t, \infty[$ let $\tau \in [t, \infty[$ be such that $\epsilon_i(\tau) > 0$. Then there is $\beta < \alpha^*$ such that $\tau \in [0, t^\beta[$ and $\epsilon_i(\tau) = \epsilon_i^\beta(\tau)$ for all $\tau' \in [0, t^\beta[$. Further, by (TR.v) $s_i(x_\tau^P((f^{t^+}, \epsilon^\beta))) = c_i(x_\tau^P((f^{t^+}, \epsilon^\beta)), \epsilon_i^\beta(\tau))$. Because $x_\tau^P((f^{t^+}, \epsilon^\beta)) = x_\tau^P((f^{t^+}, \epsilon))$, it follows that $s_i(x_\tau^P((f^{t^+}, \epsilon))) = s_i(x_\tau^P((f^{t^+}, \epsilon^\beta)))$ and as $\epsilon_i(\tau) = \epsilon_i^\beta(\tau)$ the conclusion follows.

It remains to show that $\epsilon_i(t) > 0$. If $t \in \bigcup_{j \in I} J(f_j)$ then $t \in \bigcup_{j \in I} J(f_j^{t^+})$ and hence $\epsilon_i(t) > 0$ by (P.ii). If $t = t(f, t) = t(f^{t^+}, t)$ then $\epsilon_i(t) > 0$ by (P.i) and (P.iii).

(ii) Let $\epsilon, \epsilon' \in E$ be such that $w = (f^{t^+}, \epsilon)$ and $w' = (f^{t^+}, \epsilon')$ agree with s_i on $[t, \infty[$ and assume $\epsilon_i(\tau) \neq \epsilon'_i(\tau)$ for some $\tau \in [t, \infty[$. Since $t \in \mathcal{J}(f)$ and $w, w' \in W$ it follows from (P.i) and (P.iii) that $t \in DP(\epsilon_i) \cap DP(\epsilon'_i)$. Further, $s_i(x_t^P(w)) = s_i(x_t^P(w'))$ since w and w' have the same decision path and $t \in DP(\epsilon_i) \cap DP(\epsilon'_i)$. Hence, as w and w' agree with s_i on $[t, \infty[$, it follows that $\epsilon_i(t) = \epsilon'_i(t)$. Thus $\tau > t$ and $K := \{\tau' > t \mid \epsilon_i(\tau') \neq \epsilon'_i(\tau')\} \subseteq DP(\epsilon_i) \cup DP(\epsilon'_i)$ is non-empty and well-ordered which by Lemma A.3(i) implies that $\bar{t} := \min K$ exists. By Lemma A.5, because $\bar{t} \in DP(\epsilon_i) \cup DP(\epsilon'_i)$

¹²Recall footnote 11.

and f^{t+} is constant on $]t, +\infty[$, we obtain $\bar{t} \in PJ(\epsilon_i) \cup PJ(\epsilon'_i)$. Note that $\bar{t} \in PJ(\epsilon_i)$ if and only if $\bar{t} \in PJ(\epsilon'_i)$ as $\epsilon_i(\tau) = \epsilon'_i(\tau)$ for all $\tau \in [t, \bar{t}]$. Hence $\bar{t} \in PJ(\epsilon_i) \cap PJ(\epsilon'_i)$. Then, by (P.iii), $\bar{t} \in DP(\epsilon_i) \cap DP(\epsilon'_i)$ which yields $\epsilon_i(\bar{t}) = \epsilon'_i(\bar{t})$. This contradicts the construction of \bar{t} .

(iii) Let $\epsilon, \epsilon' \in E$ be such that $(f^{t(f,t)^+}, \epsilon)$ and $(f^{t(f,t)^+}, \epsilon')$ agree with s_i on $]t(f,t), \infty[$. We start by claiming that $t' \in PJ(\epsilon_i)$ if and only if $t' \in PJ(\epsilon'_i)$ for all $t' \in]t(f,t), \infty[$. To prove this, let $t' > t(f,t)$ be such that $t' \in PJ(\epsilon_i)$. By (i), $t(f,t) \in DP(\epsilon_i) \cap DP(\epsilon'_i)$ and hence $\text{Prev}(\epsilon_i, t') \geq t(f,t)$ and $\text{Prev}(\epsilon'_i, t') \geq t(f,t)$. By (ii), $\epsilon_i(\tau) = \epsilon'_i(\tau)$ for all $\tau \in [t(f,t), \infty[$ and hence $\text{Prev}(\epsilon_i, t') = \text{Prev}(\epsilon'_i, t')$. Since $t' \in PJ(\epsilon_i)$, either $t = \text{Prev}(\epsilon_i, t') = \text{Prev}(\epsilon'_i, t')$ or $t = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, t')) = \text{Next}(\epsilon'_i, \text{Prev}(\epsilon'_i, t'))$ where the second equality holds because $\epsilon_i(\tau) = \epsilon'_i(\tau)$ for all $\tau \in [t(f,t), \infty[$ and $\text{Prev}(\epsilon_i, t') \geq t(f,t)$. In either case $t' \in PJ(\epsilon'_i)$ which proves the claim. If $t(f,t) < t$ this already proves $PJ(\epsilon_i) \cap [t, \infty[= PJ(\epsilon'_i) \cap [t, \infty[$. If $t(f,t) = t$ then by Lemma C.3 $t \in PJ(\epsilon_i) \cap PJ(\epsilon'_i)$ and it follows that $PJ(\epsilon_i) \cap [t, \infty[= PJ(\epsilon'_i) \cap [t, \infty[$. \square

Lemma C.4. *Let $f \in F$ and $t \in \mathbb{R}_+$. If $t(f,t) = t$ then $t \in M(f, s_i)$ for all $i \in I$ and all $s_i \in S_i$.*

Proof. By Lemma 3(i) there is $\epsilon \in E$ such that $(f^{t(f,t)^+}, \epsilon) \in W$ and it agrees with s_i on $]t(f,t), \infty[$. Hence, by Lemma C.3(ii), $t \in M(f, s_i)$. \square

Proof of Proposition 3. To see the first part of (CRM.i) consider first the case $t = 0$. Note that by definition $t(\hat{f}, 0) = 0$ for all $\hat{f} \in F$ and hence by Lemma C.4, $0 \in M(\hat{f}, s_i)$. Thus, for any $f, \hat{f} \in F$, by construction of $\sigma^{s_i, 1}$, we obtain $\sigma^{s_i, 1}(f, 0) = a_i(f, 0, s_i) = a_i(\hat{f}, 0, s_i) = \sigma^{s_i, 1}(\hat{f}, 0)$, where the second equality follows from the fact that $x_0((f', \epsilon)) = W$ for any $(f', \epsilon) \in W$. Now let $t \in \mathbb{R}_{++}$ and $f, \hat{f} \in F$ be such that $f(\tau) = \hat{f}(\tau)$ for all $\tau \in [0, t[$. Then $t(f,t) = t(\hat{f}, t)$. We claim that $t \in M(f, s_i)$ if and only if $t \in M(\hat{f}, s_i)$. If $t > t(f,t) = t(\hat{f}, t)$ then $f^{t(f,t)^+} = \hat{f}^{t(\hat{f}, t)^+}$ and hence $t \in M(f, s_i)$ if and only if $t \in M(\hat{f}, s_i)$ by Lemma 3(iii). If $t = t(f,t) = t(\hat{f}, t)$ then $t \in M(f, s_i) \cap M(\hat{f}, s_i)$ by Lemma C.4. This proves the claim. In view of this claim, we distinguish two cases. First, if $t \notin M(f, s_i) \cup M(\hat{f}, s_i)$ then by construction $\sigma_i^{s_i, 1}(f, t) = f_{i-}(t) = \hat{f}_{i-}(t) = \sigma_i^{s_i, 1}(\hat{f}, t)$. Second, if $t \in M(f, s_i) \cap M(\hat{f}, s_i)$ then $\sigma_i^{s_i, 1}(f, t) = a_i(f, t, s_i)$ and $\sigma_i^{s_i, 1}(\hat{f}, t) = a_i(\hat{f}, t, s_i)$. Further, there are

$\epsilon, \hat{\epsilon} \in E$ such that $(f^{t(f,t)^+}, \epsilon), (\hat{f}^{t(f,t)^+}, \hat{\epsilon}) \in W$ and $t \in PJ(\epsilon_i) \cap PJ(\hat{\epsilon}_i)$. Note that by Lemma C.1(ii) $f^{t(f,t)^+}(\tau) = f(\tau) = \hat{f}(\tau) = \hat{f}^{t(f,t)^+}(\tau)$ for all $\tau \in [0, t]$. By definition $s_i(x_t((f^{t(f,t)^+}, \epsilon))) = c_i(x_t((f^{t(f,t)^+}, \epsilon)), a_i(f, t, s_i))$ and $s_i(x_t((\hat{f}^{t(f,t)^+}, \hat{\epsilon}))) = c_i(x_t((\hat{f}^{t(f,t)^+}, \hat{\epsilon})), a_i(\hat{f}, t, s_i))$. It follows from the construction of the EDP that $x_t((f^{t(f,t)^+}, \epsilon))$ and $x_t((\hat{f}^{t(f,t)^+}, \hat{\epsilon}))$ are in the same information set and hence that $a_i(f, t, s_i) = a_i(\hat{f}, t, s_i)$ implying that $\sigma_i^{s_i,1}(f, t) = \sigma_i^{s_i,1}(\hat{f}, t)$.

To see the second part of (CRM.i), let $f, \hat{f} \in F$, and $t \in \mathbb{R}_+$ be such that $f(\tau) = \hat{f}(\tau)$ for all $\tau \in [0, t]$. Then $t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ if and only if $t \in \bigcup_{j \in I} J(\hat{f}_j) \cap LC(\hat{f}_i)$. We distinguish two cases. First, if $t \notin \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ (in which case $t \notin \bigcup_{j \in I} J(\hat{f}_j) \cap LC(\hat{f}_i)$) by construction $\sigma^{s_i,2}(f, t) = f_i(t) = \hat{f}_i(t) = \sigma^{s_i,2}(\hat{f}, t)$. Second, if $t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ (in which case $t \in \bigcup_{j \in I} J(\hat{f}_j) \cap LC(\hat{f}_i)$) we obtain $\sigma_i^{s_i,2}(f, t) = a_i^R(f, t, s_i)$ and $\sigma_i^{s_i,2}(\hat{f}, t) = a_i^R(\hat{f}, t, s_i)$. By Lemma C.2 there are $\epsilon, \hat{\epsilon} \in E$ such that $(f, \epsilon), (\hat{f}, \hat{\epsilon}) \in W$. By construction of the game tree $x_t^R((f, \epsilon)), x_t^R((\hat{f}, \hat{\epsilon})) \in X_i$. Further, by definition $s_i(x_t^R((f, \epsilon))) = c_i(x_t^R((f, \epsilon)), a_i^R(f, t, s_i))$ and $s_i(x_t^R((\hat{f}, \hat{\epsilon}))) = c_i(x_t^R((\hat{f}, \hat{\epsilon})), a_i^R(\hat{f}, t, s_i))$. Since $f(\tau) = \hat{f}(\tau)$ for all $\tau \in [0, t]$, $x_t^R((f, \epsilon))$ and $x_t^R((\hat{f}, \hat{\epsilon}))$ are in the same information set. Hence $a_i^R(f, t, s_i) = a_i^R(\hat{f}, t, s_i)$ and it follows that $\sigma^{s_i,2}(f, t) = \sigma^{s_i,2}(\hat{f}, t)$.

To show (CRM.ii) let $f \in F$ and $t \in \bigcap_{j \in I} LC(f_j) \cup J(f_i)$. The first part holds because by construction $\sigma^{s_i,2}(f, t) = f_i(t)$. To see the second part, note that by (DP.i) there is $\varepsilon > 0$ such that f is constant on $]t, t + \varepsilon[$. We distinguish two cases. First, if $t \in \bigcap_{j \in I} LC(f_j)$ then $t(f, \tau) = t(f, t)$ for all $\tau \in]t, t + \varepsilon[$. We claim that $PJ(f, \tau, s_i) = PJ(f, t, s_i) \cap [\tau, \infty[$ for all $\tau \in]t, t + \varepsilon[$. To see this let $\tau \in]t, t + \varepsilon[$. Since $t(f, \tau) = t(f, t)$ it follows that $f^{t(f,\tau)^+} = f^{t(f,t)^+}$. Hence, if $\epsilon \in E$ is such that $(f^{t(f,\tau)^+}, \epsilon) \in W$ and $(f^{t(f,\tau)^+}, \epsilon)$ agrees with s_i on $[t(f, t), \infty[$ then also $(f^{t(f,t)^+}, \epsilon)$ agrees with s_i on $[t(f, t), \infty[$. By definition it follows that $PJ(f, \tau, s_i) = PJ(f, t, s_i) \cap [\tau, \infty[$, proving the claim. As $f^{t(f,t)^+}$ is constant on $]t(f, t), \infty[$, if $\epsilon \in E$ is such that $(f^{t(f,t)^+}, \epsilon) \in W$ and $(f^{t(f,t)^+}, \epsilon)$ agrees with s_i on $[t(f, t), \infty[$ then by (P.iv) there is $\varepsilon' > 0$ such that $PJ(\epsilon_i) \cap]t, t + \varepsilon'[= \emptyset$. Hence, by Lemma 3(ii), $PJ(\epsilon_i) \cap]t, t + \varepsilon'[= \emptyset$ for all $\epsilon \in E$ is such that $(f^{t^+}, \epsilon) \in W$ and (f^{t^+}, ϵ) agrees with s_i on $[t, \infty[$ implying that $PJ(f, t, s_i) \cap]t, t + \varepsilon'[= \emptyset$. We obtain that for

all $\tau \in]t, t + \min\{\varepsilon, \varepsilon'\}[$, $PJ(f, \tau, s_i) \cap]t, t + \min\{\varepsilon, \varepsilon'\}[= \emptyset$ and hence $\tau \notin M(f, s_i)$. Thus $\sigma^{s_i^1}(f, \tau) = f_{i-}(\tau) = f_i(t)$ for all $\tau \in]t, t + \min\{\varepsilon, \varepsilon'\}[$, where the first equality holds by construction and the second because $t \in RK(f_i)$ and f is constant on $]t, t + \varepsilon[$. The latter in turn follows from $t \in \bigcap_{j \in I} LC(f_j)$ and (DP.iii). Second, if $t \in J(f_i)$ then $t(f, \tau) = t$ for all $\tau \in]t, t + \varepsilon[$. As f^{t+} is constant on $]t, \infty[$, if $\epsilon \in E$ is such that $(f^{t+}, \epsilon) \in W$ and (f^{t+}, ϵ) agrees with s_i on $[t, \infty[$ then by (P.iv) there is $\varepsilon' > 0$ such that $PJ(\epsilon_i) \cap]t, t + \varepsilon'[= \emptyset$. Hence, by Lemma 3(ii), $PJ(\epsilon_i) \cap]t, t + \varepsilon'[= \emptyset$ for all $\epsilon \in E$ is such that $(f^{t+}, \epsilon) \in W$ and (f^{t+}, ϵ) agrees with s_i on $[t, \infty[$. As $f^{t(f, \tau)+} = f^{t+}$ for all $\tau \in]t, t + \varepsilon[$, this implies that $PJ(f, \tau, s_i) \cap]t, t + \varepsilon'[= \emptyset$ for all $\tau \in]t, t + \varepsilon[$ and hence $\tau \notin M(f, s_i)$. We thus obtain $\sigma^{s_i^1}(f, \tau) = f_{i-}(\tau) = f_i(t)$ for all $\tau \in]t, t + \min\{\varepsilon, \varepsilon'\}[$, where the second equality holds because $t \in J(f_i)$ and f is constant on $]t, t + \varepsilon[$.

To prove (CRM.iii), let $f \in F$ and $t \in LC(f_i) \cap \bigcup_{j \in I} J(f_j)$. By (DP.i) there is $\varepsilon > 0$ such that f is constant on $]t, t + \varepsilon[$ which implies that $t(f, \tau) = t$ for all $\tau \in]t, t + \varepsilon[$. Since f^{t+} is constant on $]t, \infty[$ if $\epsilon \in E$ is such that $(f^{t+}, \epsilon) \in W$ and (f^{t+}, ϵ) agrees with s_i on $[t, \infty[$ then by (P.iv) there is $\varepsilon' > 0$ such that $PJ(\epsilon_i) \cap]t, t + \varepsilon'[= \emptyset$. As $f^{t(f, \tau)+} = f^{t+}$ for all $\tau \in]t, t + \varepsilon[$, this implies that $PJ(f, \tau, s_i) \cap]t, t + \varepsilon'[= \emptyset$ for all $\tau \in]t, t + \varepsilon[$. Hence $\sigma_i^{s_i^1}(f, \tau) = f_{i-}(\tau) = f_{i+}(t)$ for all $\tau \in]t, t + \min\{\varepsilon, \varepsilon'\}[$, where the second equality holds because f is constant on $]t, t + \varepsilon[$. \square

The remaining proofs in this section are simplified if one relies on the following auxiliary concept and its characterization in Lemma C.5 below.

Definition 12. For player i $s_i^1, s_i^2 \in S_i$ are *CRM-equivalent* if $\sigma^{s_i^1} = \sigma^{s_i^2}$.

Lemma C.5. Let $i \in I$. Then $s_i^1, s_i^2 \in S_i$ are *CRM-equivalent* if and only if

(O.i) For all potential jump nodes $x = x_t((f, \epsilon)) \in X_i$, if $t \in M(f, s_i^1) \cap M(f, s_i^2)$ then $s_i^1(x) = s_i^2(x)$.

(O.ii) For all potential jump nodes $x = x_t((f, \epsilon)) \in X_i$, if $t \in M(f, s_i^k) \setminus M(f, s_i^l)$ for $k \neq l$ then $s_i^k(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon)), f_{i-}(t))$.

(O.iii) For all reaction nodes $x = x_t^R((f, \epsilon)) \in X_i$, $s_i^1(x) = s_i^2(x)$.

Proof of Lemma C.5. “If”: Let s_i^1 and s_i^2 satisfy (O.i)-(O.iii) and let $(f, t) \in F \times \mathbb{R}_+$. To prove that $\sigma^{s_i^1,1}(f, t) = \sigma^{s_i^2,1}(f, t)$ we distinguish three cases. First, if $t \notin M(f, s_i^1) \cup M(f, s_i^2)$ then $\sigma^{s_i^1,1}(f, t) = f_{i-}(t) = \sigma^{s_i^2,1}(f, t)$ by construction. Second, if $t \in M(f, s_i^1) \cap M(f, s_i^2)$, then by (O.i) we obtain $s_i^1(x_t((f, \epsilon))) = s_i^2(x_t((f, \epsilon)))$ for all potential jump nodes $x_t(f, \epsilon) \in X_i$. In particular $a_i(f, t, s_i^1) = a_i(f, t, s_i^2)$ and hence by construction $\sigma^{s_i^1,1}(f, t) = \sigma^{s_i^2,1}(f, t)$. Third, if $t \in M(f, s_i^1) \setminus M(f, s_i^2)$ (and analogously if superindices are exchanged) then by (O.ii) we obtain $s_i^1(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon)), f_{i-}(t))$ for all potential jump nodes $x_t((f, \epsilon)) \in X_i$. Then by construction $\sigma^{s_i^1,1}(f, t) = f_{i-}(t)$ and $\sigma^{s_i^2,1}(f, t) = f_{i-}(t)$ and hence $\sigma^{s_i^1,1}(f, t) = \sigma^{s_i^2,1}(f, t)$.

To see that $\sigma^{s_i^1,2}(f, t) = \sigma^{s_i^2,2}(f, t)$ we distinguish two cases. First, if $t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ then by construction $\sigma^{s_i^1,2}(f, t) = a_i^R(f, t, s_i^1)$ and $\sigma^{s_i^2,2}(f, t) = a_i^R(f, t, s_i^2)$. By (O.iii) it follows that $a_i^R(f, t, s_i^1) = a_i^R(f, t, s_i^2)$ and hence $\sigma^{s_i^1,2}(f, t) = \sigma^{s_i^2,2}(f, t)$. Second, if $t \in \bigcap_{j \in I} LC(f_j) \cup J(f_i)$ then by construction $\sigma^{s_i^1,2}(f, t) = f_i(t) = \sigma^{s_i^2,2}(f, t)$.

“Only if”: Let $s_i^1, s_i^2 \in S_i$ be such that $\sigma^{s_i^1} = \sigma^{s_i^2}$. To prove (O.i) let $x = x_t((f, \epsilon)) \in X_i$ be such that $t \in M(f, s_i^1) \cap M(f, s_i^2)$. We have $s_i^1(x) = c_i(x, a_i(f, t, s_i^1))$ and $s_i^2(x) = c_i(x, a_i(f, t, s_i^2))$. Further, since $t \in M(f, s_i^1) \cap M(f, s_i^2)$, by construction $\sigma^{s_i^1,1}(f, t) = a_i(f, t, s_i^1)$ and $\sigma^{s_i^2,1}(f, t) = a_i(f, t, s_i^2)$. Since $\sigma^{s_i^1}(f, t) = \sigma^{s_i^2}(f, t)$ we obtain $a_i(f, t, s_i^1) = a_i(f, t, s_i^2)$ and hence $s_i^1(x) = s_i^2(x)$. To see (O.ii) let $x = x_t((f, \epsilon)) \in X_i$ be such that $t \in M(f, s_i^1) \setminus M(f, s_i^2)$. Since $t \notin M(f, s_i^2)$, we obtain $\sigma^{s_i^1,1}(f, t) = \sigma^{s_i^2,1}(f, t) = f_{i-}(t)$. Since $t \in M(f, s_i^1)$, $\sigma^{s_i^1,1}(f, t) = a_i(f, t, s_i^1)$ and hence $s_i^1(x) = c_i(x, a_i(f, t, s_i^1)) = c_i(x, f_{i-}(t))$. The case where $t \in M(f, s_i^2) \setminus M(f, s_i^1)$ works analogously. Finally, to see (O.iii) let $x = x_t^R((f, \epsilon)) \in X_i$ be a reaction node. Then $s_i^1(x) = c_i(x, a_i^R(f, t, s_i^1))$ and $s_i^2(x) = c_i(x, a_i^R(f, t, s_i^2))$. Since x is a reaction node for player i , $t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ and hence $\sigma^{s_i^1,2}(f, t) = a_i^R(f, t, s_i^1)$ and $\sigma^{s_i^2,2}(f, t) = a_i^R(f, t, s_i^2)$. Hence, since $\sigma^{s_i^1,2}(f, t) = \sigma^{s_i^2,2}(f, t)$, we obtain $s_i^1(x) = s_i^2(x)$. \square

Proof of Lemma 4. We will rely on Lemma C.5 and prove that if s_i^1 and s_i^2 satisfy (O.i)-(O.iii) then $s_i^1 \sim s_i^2$. Fix $s_{-i} \in S_{-i}$. Let $w^k = (f^k, \epsilon^k)$ be the play induced by (s_i^k, s_{-i}) for $k = 1, 2$. Then, by construction w^k

agrees with s_i^k on $[t(f, t), \infty[$ for every $t \in \mathbb{R}_+$, i.e. $\epsilon_i^k(\tau) > 0$ implies that $s_i(x_\tau^P(w^k)) = c_i(x_\tau^P(w^k), \epsilon_i(\tau))$ for all $\tau \in [t(f, t), \infty[$.

Assume by contradiction that $f^1 \neq f^2$. Since $t(f^1, 0) = 0$ by definition, Lemma C.4 implies $0 \in M(f^1, s_i^1) \cap M(f^1, s_i^2)$. Since $W = x_0(w^1)$, (O.i) yields $s_i^1(W) = s_i^2(W)$. Since for all $j \neq i$ the action prescribed at W is the same in both strategy profiles, we obtain $f^1(0) = f^2(0)$. Then, by (DP.ii), $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, \varepsilon[$ for some $\varepsilon > 0$. Since $f^1 \neq f^2$, $\bar{t} := \inf\{\tau > 0 \mid f^1(\tau) \neq f^2(\tau)\}$ exists and $\bar{t} > 0$.

We claim that $\bar{t} \in \bigcup_{j \in I} J(f_j^1) \cup \bigcup_{j \in I} J(f_j^2)$. Suppose $\bar{t} \notin \bigcup_{j \in I} J(f_j^1) \cup \bigcup_{j \in I} J(f_j^2)$. Then $\bar{t} \in \bigcap_{j \in I} LC(f_j^1) \cap \bigcap_{j \in I} LC(f_j^2)$ which implies $f^1(\bar{t}) = f^2(\bar{t})$. By (DP.i) and (DP.iii) there is $\varepsilon' > 0$ such that $f^1(\tau) = f^1(\bar{t}) = f^2(\bar{t}) = f^2(\tau)$ for all $\tau \in [\bar{t}, \bar{t} + \varepsilon[$, which contradicts the definition of \bar{t} . This proves the claim. By (P.i) it follows that $\bar{t} \in \bigcup_{j \in I} PJ(\epsilon_j^1) \cup \bigcup_{j \in I} PJ(\epsilon_j^2)$.

Claim A. $\epsilon_j^1(\tau) = \epsilon_j^2(\tau)$ for all $\tau \in [0, \bar{t}[$ and all $j \in I \setminus \{i\}$.

To see this suppose that there is $j \in I \setminus \{i\}$ such that $\epsilon_j^1(\tau) \neq \epsilon_j^2(\tau)$ for some $\tau \in [0, \bar{t}[$. Then $\{\tau < \bar{t} \mid \epsilon_j^1(\tau) \neq \epsilon_j^2(\tau)\} \subseteq DP(\epsilon_j^1) \cup DP(\epsilon_j^2)$ is nonempty and well-ordered by Lemma A.3(i) and hence $t^* := \min\{\tau < \bar{t} \mid \epsilon_j^1(\tau) \neq \epsilon_j^2(\tau)\}$ exists. Note that since $f^1(0) = f^2(0)$ and $IPJ((\epsilon^1, 0) = IPJ((\epsilon^2, 0) = I$, it follows from the construction of the EDP that $x_0^P(w^1)$ and $x_0^P(w^2)$ are in the same information set and hence that $s_j(x_0^P(w^1)) = s_j(x_0^P(w^2))$ which implies $\epsilon_j^1(0) = \epsilon_j^2(0)$ and hence $t^* > 0$. Further, note that $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, t^*]$, since $t^* < \bar{t}$ and that $t^* \in PJ(\epsilon_j^1)$ if and only if $t^* \in PJ(\epsilon_j^2)$ since $\epsilon_j^1(\tau) = \epsilon_j^2(\tau)$ for all $\tau \in [0, t^*[$. Suppose that $t^* \in PJ(\epsilon_j^1) \cap PJ(\epsilon_j^2)$. Then by construction of the EDP $x_{t^*}^P(w^1)$ and $x_{t^*}^P(w^2)$ are in the same information set of player j . Hence $s_j(x_{t^*}^P(w^1)) = s_j(x_{t^*}^P(w^2))$ and it follows that $\epsilon_j^1(t^*) = \epsilon_j^2(t^*)$ which contradicts the construction of t^* . Hence $t^* \notin PJ(\epsilon_j^1) \cup PJ(\epsilon_j^2)$. If $\epsilon_j^1(t^*) = 0$ then $t^* \notin \bigcup_{k \in I} J(f_k^1)$ by (P.ii) and since $t^* < \bar{t}$, this implies $t^* \notin \bigcup_{k \in I} J(f_k^2)$. Further, note that $t(f^2, t^*) < t^*$ by Lemma C.3(i) because $t^* \notin PJ(\epsilon_j^2)$. Hence by Lemma C.1(i) f^2 is constant on $]t(f^2, t^*), t^*[$. Since $t^* \notin \bigcup_{k \in I} J(f_k^2)$ it follows that f^2 is constant on $]t(f^2, t^*), t^*[$. Since $t^* \notin PJ(\epsilon_j^2)$, $t(f^2, t^*) \leq \text{Prev}(\epsilon_j^2, t^*) < t^* < \text{Next}(\epsilon_j^2, \text{Prev}(\epsilon_j^2, t^*))$ and hence (P.iv) yields $\epsilon_j^2(t^*) = 0$, which contradicts the definition of t^* . We thus obtain $\epsilon_j^1(t^*) > 0$ and analogously $\epsilon_j^2(t^*) > 0$. Hence $t^* \in DP(\epsilon_j^1) \cap DP(\epsilon_j^2)$ and

in particular $x_{t^*}^P(w^1), x_{t^*}^P(w^2) \in X_j$. Further, $t^* < \bar{t}$ yields $f_+^1(t^*) = f_+^2(t^*)$ and from the construction of the EDP it follows that $x_{t^*}^P(w^1)$ and $x_{t^*}^P(w^2)$ are in the same information set which implies $s_j(x_{t^*}^P(w^1)) = s_j(x_{t^*}^P(w^2))$. Thus $\epsilon_j^1(t^*) = \epsilon_j^2(t^*)$, a contradiction to the construction of t^* . This proves the claim.

For each $j \in I \setminus \{i\}$, since $\epsilon_j^1(\tau) = \epsilon_j^2(\tau)$ for all $\tau \in [0, \bar{t}]$, we obtain that $\bar{t} \in PJ(\epsilon_j^1)$ if and only if $\bar{t} \in PJ(\epsilon_j^2)$.

Claim B. For $k = 1, 2$, $\bar{t} \in PJ(\epsilon_i^k)$ if and only if $\bar{t} \in M(f^k, s_i^k)$.

If $t(f^k, \bar{t}) = \bar{t}$, then by Lemma C.3(i) and Lemma C.4, $\bar{t} \in PJ(\epsilon_i^k) \cap M(f^k, s_i^k)$. Hence, suppose $t(f^k, \bar{t}) < \bar{t}$. Let $\epsilon \in E$ be such that $((f^k)^{t(f^k, \bar{t})+}, \epsilon)$ agrees with s_i^k on $[t(f^k, \bar{t}), \infty[$ and suppose there is $\tau \in [t(f^k, \bar{t}), \bar{t}]$ such that $\epsilon_i(\tau) \neq \epsilon_i^k(\tau)$. Then, let $t' := \min\{\tau \in [t(f^k, \bar{t}), \bar{t}] \mid \epsilon_i(\tau) \neq \epsilon_i^k(\tau)\}$, which exists since $\{\tau \in [t(f^k, \bar{t}), \bar{t}] \mid \epsilon_i(\tau) \neq \epsilon_i^k(\tau)\} \subseteq DP(\epsilon_i) \cup DP(\epsilon_i^k)$ is well-ordered by Lemma A.3(i). Note that $t' < t(f^k, \bar{t})$ because $\epsilon_i(t(f^k, \bar{t})) > 0$ by Lemma 3(i) and then by definition of agreeing (Definition 8, equation (1.2)) $\epsilon_i(t(f^k, \bar{t})) = \epsilon_i^k(t(f^k, \bar{t}))$. Since $\epsilon_i(\tau) = \epsilon_i^k(\tau)$ for all $\tau \in [t(f^k, \bar{t}), t']$ it follows that $t' \in PJ(\epsilon_i)$ if and only if $t' \in PJ(\epsilon_i^k)$. Since $t' \in DP(\epsilon_i) \cup DP(\epsilon_i^k)$ and both $(f^k)^{t(f^k, \bar{t})+}$ and f^k are constant on $]t(f^k, \bar{t}), \bar{t}[$ by Lemma C.1(i), Lemma A.5 yields $t' \in PJ(\epsilon_i) \cup PJ(\epsilon_i^k)$ and hence $t' \in PJ(\epsilon_i) \cap PJ(\epsilon_i^k)$. Thus $x_{t'}^P(((f^k)^{t(f^k, \bar{t})+}, \epsilon))$ and $x_{t'}^P(w^k)$ are in the same information set of player i , thus $s_i^k(x_{t'}^P(((f^k)^{t(f^k, \bar{t})+}, \epsilon))) = s_i^k(x_{t'}^P(w^k))$ and $\epsilon_i(t') = \epsilon_i^k(t')$, a contradiction with our choice of t' . Thus $\epsilon_i(\tau) = \epsilon_i^k(\tau)$ for all $\tau \in [t(f^k, \bar{t}), \bar{t}]$ implying that $\bar{t} \in PJ(\epsilon_i^k)$ if and only if $\bar{t} \in PJ(\epsilon_i)$. This proves the claim.

Claim C. $f^1(\bar{t}) = f^2(\bar{t})$.

We first prove $f_j^1(\bar{t}) = f_j^2(\bar{t})$ for all $j \in I \setminus \{i\}$. By Claim A, either $\bar{t} \in PJ(\epsilon_j^1) \cap PJ(\epsilon_j^2)$ or $\bar{t} \notin PJ(\epsilon_j^1) \cup PJ(\epsilon_j^2)$. In the first case $s_j(x_{\bar{t}}(w^1)) = s_j(x_{\bar{t}}(w^2))$ and hence $f_j^1(\bar{t}) = f_j^2(\bar{t})$. In the second case, by (P.i) $f_j^1(\bar{t}) = f_{j-}^1(\bar{t}) = f_{j-}^2(\bar{t}) = f_j^2(\bar{t})$.

It remains to show that $f_i^1(\bar{t}) = f_i^2(\bar{t})$. We distinguish three cases. First, if $\bar{t} \notin PJ(\epsilon_i^1) \cup PJ(\epsilon_i^2)$, then by (P.i) $f_i^1(\bar{t}) = f_{i-}^1(\bar{t})$ and $f_i^2(\bar{t}) = f_{i-}^2(\bar{t})$. Since $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, \bar{t}]$, $f_{i-}^1(\bar{t}) = f_{i-}^2(\bar{t})$ and we obtain $f_i^1(\bar{t}) = f_i^2(\bar{t})$.

Second, suppose that $\bar{t} \in PJ(\epsilon_i^1) \cap PJ(\epsilon_i^2)$. Note that since $f^1(\tau) = f^2(\tau)$

for all $\tau \in [0, \bar{t}[$, by construction of the EDP $s_i^1(x_{\bar{t}}(w^1)) = s_i^1(x_{\bar{t}}(w^2))$. By Claim B, we obtain $\bar{t} \in M(f^1, s_i^1) \cap M(f^2, s_i^2)$. We now prove that $\bar{t} \in M(f^1, s_i^2)$. Since $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, \bar{t}[$, it follows that $t(f^1, \bar{t}) = t(f^2, \bar{t})$. We distinguish two cases. If $t(f^1, \bar{t}) = t(f^2, \bar{t}) < \bar{t}$ then $(f^1)^{t(f^1, \bar{t})} = (f^2)^{t(f^1, \bar{t})}$ and by Lemma 3(iii), $PJ(f^2, \bar{t}, s_i^2) = PJ(f^1, \bar{t}, s_i^2)$ implying that $\bar{t} \in M(f^1, s_i^2)$. If on the other hand $t(f^1, \bar{t}) = t(f^2, \bar{t}) = \bar{t}$ then by Lemma C.3(i) $\bar{t} \in PJ(\epsilon_i)$ for any $\epsilon \in E$ such that $(f^1, \epsilon) \in W$. In particular $\bar{t} \in PJ(\epsilon_i)$ for any $\epsilon \in E$ such that $(f^1, \epsilon) \in W$ and (f^1, ϵ) agrees with s_i^2 on $[t(f^1, \bar{t}), \infty[$ implying that $\bar{t} \in PJ(f^1, \bar{t}, s_i^2)$. This proves that $\bar{t} \in M(f^1, s_i^2)$. Thus $\bar{t} \in M(f^1, s_i^1) \cap M(f^1, s_i^2)$ and (O.i) yields $s_i^1(x_{\bar{t}}(w^2)) = s_i^2(x_{\bar{t}}(w^2))$ which implies $s_i^1(x_{\bar{t}}(w^1)) = s_i^2(x_{\bar{t}}(w^2))$; hence $f_i^1(\bar{t}) = f_i^2(\bar{t})$.

Third, if $\bar{t} \in PJ(\epsilon_i^1) \setminus PJ(\epsilon_i^2)$ (and analogously if $\bar{t} \in PJ(\epsilon_i^2) \setminus PJ(\epsilon_i^1)$), we obtain $\bar{t} \in M(f^1, s_i^1) \setminus M(f^2, s_i^2)$ by Claim B. We claim that $\bar{t} \notin M(f^1, s_i^2)$. Since $t(f^1, \bar{t}) = t(f^2, \bar{t}) = \bar{t}$ would imply $\bar{t} \in M(f^2, s_i^2)$ by Lemma C.4, we obtain $t(f^1, \bar{t}) = t(f^2, \bar{t}) < \bar{t}$. By Lemma 3(iii) $PJ(f^1, \bar{t}, s_i^2) = PJ(f^2, \bar{t}, s_i^2)$ and since $\bar{t} \notin PJ(f^2, \bar{t}, s_i^2)$ this implies $\bar{t} \notin PJ(f^1, \bar{t}, s_i^2)$ which proves the claim. (O.ii) now yields $s_i^1(x_{\bar{t}}(w^1)) = c_i(x_{\bar{t}}(w^1), f_{i-}^1(\bar{t}))$ implying that $f_i^1(\bar{t}) = f_{i-}^1(\bar{t})$. Further, since $\bar{t} \notin PJ(\epsilon_i^2)$, it follows by (P.i) that $f_i^2(\bar{t}) = f_{i-}^2(\bar{t})$. Since $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, \bar{t}[$, $f_{i-}^1(\bar{t}) = f_{i-}^2(\bar{t})$ holds and we obtain $f_i^1(\bar{t}) = f_i^2(\bar{t})$. This proves the claim.

Hence $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, \bar{t}]$ and in particular $IJ(f^1, \bar{t}) = IJ(f^2, \bar{t})$. Since $\bar{t} \in \bigcup_{j \in I} J(f_j^1) \cup \bigcup_{j \in I} J(f_j^2)$ we obtain $IJ(f^1, \bar{t}) = IJ(f^2, \bar{t}) \neq \emptyset$ and by (DP.ii) it follows that $f_{j+}^1(\bar{t}) = f_j^1(\bar{t}) = f_j^2(\bar{t}) = f_{j+}^2(\bar{t})$ for all $j \in IJ(f^1, \bar{t}) = IJ(f^2, \bar{t})$. Further, by construction of the EDP, $s_j(x_{\bar{t}}^R(w^1)) = s_j(x_{\bar{t}}^R(w^2))$ and hence $f_{j+}^1(\bar{t}) = f_{j+}^2(\bar{t})$ for all $j \in I \setminus IJ(f^1, \bar{t})$, $j \neq i$. We now distinguish two cases. If $i \in IJ(f^1, \bar{t})$ then by (DP.ii) $f_{i+}^1(\bar{t}) = f_{i+}^2(\bar{t})$. If on the other hand, $i \notin IJ(f^1, \bar{t})$ then $x_{\bar{t}}^R(w^1), x_{\bar{t}}^R(w^2) \in X_i$ which by (O.iii) implies that $s_i^1(x_{\bar{t}}^R(w^1)) = s_i^2(x_{\bar{t}}^R(w^1))$. Since $f^1(\tau) = f^2(\tau)$ for all $\tau \in [0, \bar{t}]$, by construction of the EDP $s_i^2(x_{\bar{t}}^R(w^1)) = s_i^2(x_{\bar{t}}^R(w^2))$. This yields $s_i^1(x_{\bar{t}}^R(w^1)) = s_i^2(x_{\bar{t}}^R(w^2))$ and hence $f_{i+}^1(\bar{t}) = f_{i+}^2(\bar{t})$. In both cases, we obtain $f_{i+}^1(\bar{t}) = f_{i+}^2(\bar{t})$, which by (DP.i) contradicts our choice of \bar{t} . \square

Proof of Proposition 4. (i) Fix $\varepsilon > 0$ and Let $s_i : X_i \rightarrow C_i$ be a mapping fulfilling (IS.i)-(IS.iii) and $s_i(x_t^P((f, \epsilon))) = c_i(x_t^P((f, \epsilon)), \varepsilon)$ for all peek nodes $x_t^P((f, \epsilon)) \in X_i$ such that $E^{\sigma_i}(f, t) = \emptyset$. We will show that $s_i^{-1}(c) = P(c)$ for all $c \in s_i(X_i)$. Let $(f, \epsilon) \in W$. Note that $c_i(x_t((f, \epsilon)), \sigma_i^1(f, t)) = c_i(x_{t'}((f', \epsilon')), \sigma_i^1(f', t')) \in C_i$ if and only if $t = t'$, $f(\tau) = f'(\tau)$ for all $\tau \in [0, t[$, and $t \in PJ(\epsilon_i)$ and $t' \in PJ(\epsilon'_i)$. Hence if $c = c_i(x_t((f, \epsilon)), \sigma_i^1(f, t)) \in C_i$ then by (IS.i)

$$s_i^{-1}(c) = \{x_t((f', \epsilon')) \in N \mid f(\tau) = f'(\tau) \forall \tau \in [0, t[, t \in PJ(\epsilon'_i)\} = P(c).$$

Analogously $s_i^{-1}(c) = P(c)$ follows if $c = c_i(x_t^R(w), \sigma_i^2(f, t)) \in C_i$, $c = c_i(x_t^P(w), \varepsilon^{\sigma_i}(f, t)) \in C_i$, or $c = c_i(x_t^P(w), \varepsilon) \in C_i$.

(ii) Let $s_i, s'_i \in S(\sigma_i)$. We will show that s_i and s'_i are CRM-equivalent by showing that (O.i)-(O.iii) are satisfied (Lemma C.5). Lemma 4 then yields $s_i \sim s'_i$. By (IS.i) $s_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon)), \sigma_i^1(f, t)) = s'_i(x_t((f, \epsilon)))$ for all potential jump nodes $x_t((f, \epsilon)) \in X_i$ and hence in particular (O.i) holds.

To see (O.ii), let $x_t((f, \epsilon)) \in X_i$ be such that $t \in M(f, s_i) \setminus M(f, s'_i)$. As $t \notin M(f, s'_i)$, $t(f, t) < t$ (by Lemma C.4) and there is $\tau' \in [t(f, t), t[$ such that $\epsilon_i(f, t, s_i)(\tau') \neq \epsilon_i(f, t, s'_i)(\tau')$. We claim that $E^{\sigma_i}(f^{t(f, t)^+}, \bar{\tau}) = \emptyset$ for some $\bar{\tau} \in [t(f, t), t[$. Suppose $E^{\sigma_i}(f^{t(f, t)^+}, \tau) \neq \emptyset$ for all $\tau \in [t(f, t), t[$. By definition of agreeing (Definition 8, Equation 1.2) and Lemma 3(ii) for all $\tau \in [t(f, t), t[$ if $\epsilon_i(f, t, s_i)(\tau) > 0$ then $s_i(x_\tau^P((f^{t(f, t)^+}, \epsilon))) = c_i(x_\tau^P((f^{t(f, t)^+}, \epsilon)), \epsilon_i(f, t, s_i))$ for any $\epsilon \in E$ such that $(f^{t(f, t)^+}, \epsilon) \in W$ agrees with s_i on $[t(f, t), \infty[$. By (IS.iii), since $E^{\sigma_i}(f^{t(f, t)^+}, \tau) \neq \emptyset$, this implies $\epsilon_i(f, t, s_i)(\tau) = \varepsilon^{\sigma_i}(f^{t(f, t)^+}, \tau)$ for all $\tau \in [t(f, t), t[$ such that $\epsilon_i(f, t, s_i)(\tau) > 0$. An analogous argument yields $\epsilon_i(f, t, s'_i)(\tau) = \varepsilon^{\sigma_i}(f^{t(f, t)^+}, \tau)$ for all $\tau \in [t(f, t), t[$ such that $\epsilon_i(f, t, s'_i)(\tau) > 0$. Hence $\epsilon_i(f, t, s_i)(\tau) = \epsilon_i(f, t, s'_i)(\tau)$ for all $\tau \in [t(f, t), t[$, a contradiction. This proves the claim that $E^{\sigma_i}(f^{t(f, t)^+}, \bar{\tau}) = \emptyset$ for some $\bar{\tau} \in [t(f, t), t[$. Since $f^{t(f, t)^+}$ is constant on $]t(f, t), \infty[$ and $\bar{\tau} \in [t(f, t), t[$, $E^{\sigma_i}(f^{t(f, t)^+}, \bar{\tau}) = \emptyset$ implies that $\sigma_i^1(f^{t(f, t)^+}, \bar{\tau}) = f_{i+}^{t(f, t)^+}(\bar{\tau})$ for all $\tau \in]\bar{\tau}, \infty[$. Since by Lemma C.1(ii) $f^{t(f, t)^+}(\tau) = f(\tau)$ for all $\tau \in [0, t[$ we hence obtain $\sigma_i^1(f, t) = \sigma_i^1(f^{t(f, t)^+}, t) = f_{i+}^{t(f, t)^+}(\bar{\tau})$ where the first equality follows by (CRM.i). Further $f_{i+}^{t(f, t)^+}(\bar{\tau}) = f_{i+}^{t(f, t)^+}(t)$ and by Lemma C.1(ii) $f_{i-}^{t(f, t)^+}(t) = f_{i-}(t)$. Hence $\sigma_i^1(f, t) = f_{i-}(t)$.

Finally (O.iii) holds since by (IS.ii), $s_i(x_t^R((f, \epsilon))) = c_i(x_t^R((f, \epsilon)), \sigma_i^2(f, t)) = s'_i(x_t^R((f, \epsilon)))$ for all reaction nodes $x_t^R((f, \epsilon)) \in X_i$. \square

Lemma C.6. *Let $(f, t) \in F \times \mathbb{R}_+$, $i \in I$, $\sigma_i \in \Sigma_i$, and $s_i \in S(\sigma_i)$. If $t \notin M(f, s_i)$ then $\sigma_i^1(f, t) = f_{i-}(t)$.*

Proof. By Lemma 3(i) there is $\epsilon \in E$ such that $(f^{t(f,t)^+}, \epsilon) \in W$ agrees with s_i on $[t(f, t), \infty[$. By Lemma 3(iii), $t \notin M(f, s_i)$ yields $t \notin PJ(\epsilon_i)$ and hence $\bar{t} := \text{Prev}(\epsilon_i, t) < t < \text{Next}(\epsilon_i, \bar{t})$. By Lemma A.1(i), $\bar{t} \in DP(\epsilon_i)$. Since $\epsilon_i(t(f, t)) > 0$ by Lemma 3(i), $t(f, t) \leq \bar{t} < t$ and hence $f(\tau) = f^{\bar{t}^+}(\tau)$ for all $\tau \in [0, t[$ by Lemma C.1(ii). We distinguish two cases. First, if $E^{\sigma_i}(f, \bar{t}) = \emptyset$, then $\sigma_i^1(f, t) = \sigma_i^1(f^{\bar{t}^+}, t) = f_{i+}(\bar{t}) = f_{i-}(t)$. The first equality holds by (CRM.i), the second because $E^{\sigma_i}(f, \bar{t}) = \emptyset$ and the third because f is constant on $]t(f, t), t[$ by Lemma C.1(i). Second, if $E^{\sigma_i}(f, \bar{t}) \neq \emptyset$, then, since $\bar{t} < t$ and $f(\tau) = f^{t(f,t)^+}(\tau)$ for all $\tau \in [0, t[$ by Lemma C.1(ii), we obtain $E^{\sigma_i}(f, \bar{t}) = E^{\sigma_i}(f^{t(f,t)^+}, \bar{t})$ and (IS.iii) yields $s_i(x_{\bar{t}}^P(f^{t(f,t)^+}, \epsilon)) = c_i(x_{\bar{t}}((f^{t(f,t)^+}, \epsilon)), \varepsilon^{\sigma_i}(f^{t(f,t)^+}, \bar{t}))$. Since $(f^{t(f,t)^+}, \epsilon)$ agrees with s_i on $[t(f, t), \infty[$ and $\bar{t} \in DP(\epsilon_i)$, it follows that $\epsilon_i(\bar{t}) = \varepsilon^{\sigma_i}(f^{t(f,t)^+}, \bar{t})$ and $\text{Next}(\epsilon_i, \bar{t}) = \bar{t} + \varepsilon^{\sigma_i}(f^{t(f,t)^+}, \bar{t})$ implying that $\bar{t} < t < \bar{t} + \varepsilon^{\sigma_i}(f^{t(f,t)^+}, \bar{t})$. Since $E^{\sigma_i}(f, \bar{t}) = E^{\sigma_i}(f^{t(f,t)^+}, \bar{t}) \neq \emptyset$ we obtain $\varepsilon^{\sigma_i}(f, \bar{t}) = \varepsilon^{\sigma_i}(f^{t(f,t)^+}, \bar{t})$ and hence $\sigma_i^1(f, t) = \sigma_i^1(f^{\bar{t}^+}, t) = f_{i+}(\bar{t}) = f_{i-}(t)$. The first equality follows from (CRM.i), the second from the definition of $\varepsilon^{\sigma_i}(f, t)$ and $t < \bar{t} + \varepsilon^{\sigma_i}(f, \bar{t})$, and the third holds because f is constant on $]t(f, t), t[$ by Lemma C.1(i). \square

Proof of Theorem 2. (i) Let $(f, t) \in F \times \mathbb{R}_+$, $\sigma_i \in \Sigma_i$, and $s_i \in S(\sigma_i)$. To show $\sigma^{s_i,1}(f, t) = \sigma_i^1(f, t)$ we distinguish two cases. First, if $t \notin M(f, s_i)$ then $\sigma_i^1(f, t) = f_{i-}(t)$ by Lemma C.6. Since $\sigma^{s_i,1}(f, t) = f_{i-}(t)$ by construction, we obtain $\sigma^{s_i,1}(f, t) = \sigma_i(f, t)$. Second, if $t \in M(f, s_i)$ then by construction $\sigma^{s_i,1}(f, t) = a_i(f, t, s_i)$. As by (IS.i) $s_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon)), \sigma_i^1(f, t))$ for all potential jump nodes $x_t((f, \epsilon)) \in X_i$, it follows that $a_i(f, t, s_i) = \sigma_i^1(f, t)$ since by Lemma 3(i) there is $\epsilon \in E$ such that $(f^{t(f,t)^+}, \epsilon) \in W$ agrees with s_i on $[t(f, t), \infty[$. Hence $\sigma^{s_i,1}(f, t) = \sigma_i^1(f, t)$.

To prove $\sigma_i^{s_i,2}(f, t) = \sigma_i^2(f, t)$, we again distinguish two cases. First, if $t \in \bigcap_{j \in I} LC(f_j) \cup J(f_i)$ then $\sigma_i^2(f, t) = f_i(t) = \sigma^{s_i,2}(f, t)$ where the first equality holds by (CRM.ii) and the second by construction. Second, if $t \notin$

$\bigcap_{j \in I} LC(f_j) \cup J(f_i)$ then by (IS.ii) $s_i(x_t^R((f, \epsilon))) = c_i(x_t^R((f, \epsilon))), \sigma_i^2(f, t)$ for all reaction nodes $x_t^R((f, \epsilon)) \in X_i$. Hence $\sigma^{s_i, 2}(f, t) = \sigma_i^2(f, t)$ by construction of $\sigma^{s_i, 2}$.

(ii) Let $s_i \in S_i$ and $s'_i \in S(\sigma^{s_i})$. We will show that s_i, s'_i satisfy (O.i)-(O.iii) and hence are CRM-equivalent by Lemma C.5. By Lemma 4 $s_i \sim s'_i$. To see (O.i), let $x_t((f, \epsilon)) \in X_i$ with $t \in M(f, s_i) \cap M(f, s'_i)$. By construction, $\sigma^{s_i, 1}(f, t) = a_i(f, t, s_i)$. By (IS.i), $s'_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon))), \sigma^{s_i, 1}(f, t)$. Hence $s'_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon))), a_i(f, t, s_i) = s_i(x_t((f, \epsilon)))$.

To prove (O.ii), first let $x_t((f, \epsilon)) \in X_i$ with $t \in M(f, s'_i) \setminus M(f, s_i)$. By construction, $\sigma^{s_i, 1}(f, t) = f_{i-}(t)$. Since $s'_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon))), \sigma^{s_i, 1}(f, t)$ by (IS.i), we have $s'_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon))), f_{i-}(t)$. Second, let $x_t((f, \epsilon)) \in X_i$ with $t \in M(f, s_i) \setminus M(f, s'_i)$. By Lemma C.6, $\sigma^{s_i, 1}(f, t) = f_{i-}(t)$. Since $t \in M(f, s_i)$, $\sigma^{s_i, 1}(f, t) = a_i(f, t, s_i)$ by construction of σ^{s_i} and it follows by construction of $a_i(f, t, s_i)$ that $s_i(x_t((f, \epsilon))) = c_i(x_t((f, \epsilon))), f_{i-}(t)$.

Finally, we prove (O.iii). Let $x_t^R((f, \epsilon)) \in X_i$ be a reaction node. Then $t \in \bigcup_{j \in I} J(f_j) \cap LC(f_i)$ by construction of the game tree, implying that $\sigma^{s_i, 2}(f, t) = a_i^R(f, t, s_i)$ by construction of σ^{s_i} . By (IS.ii), $s'_i(x_t^R((f, \epsilon))) = c_i(x_t^R((f, \epsilon))), \sigma^{s_i, 2}(f, t)$ which yields $a_i^R(f, t, s'_i) = a_i^R(f, t, s_i)$. We thus obtain $s'_i(x_t^R((f, \epsilon))) = s_i(x_t^R((f, \epsilon)))$.

(iii) Let $\sigma = (\sigma_i)_{i \in I} \in \Sigma$ and $(s_i)_{i \in I} \in \times_{i \in I} S(\sigma_i)$. Further denote by $w = (f, \epsilon) = w^{(s_i)_{i \in I}}$ the play induced by $(s_i)_{i \in I}$. Suppose that $f \neq f^\sigma$. Then $\bar{t} := \inf\{t | f(t) \neq f^\sigma(t)\}$ exists. Since $W = x_0(f, \epsilon')$ for any $\epsilon' \in E$ such that $(f, \epsilon') \in W$, by (IS.i) $s_i(W) = c_i(W, \sigma_i^1(f, 0))$ for all $i \in I$ and we obtain $f_i(0) = \sigma_i^1(f, 0) = \sigma_i^1(f^\sigma, 0) = f_i^\sigma(0)$ for all $i \in I$ where the first equality follows from the fact that f is the outcome induced by $(s_i)_{i \in I}$, the second equality follows from (CRM.i) and the third from the definition of f^σ . By (DP.ii), $f_+(0) = f(0)$ and $f_+^\sigma(0) = f^\sigma(0)$. Hence (by (DP.i)) there is $\varepsilon > 0$ such that $f(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \varepsilon[$ which implies $\bar{t} > 0$. We claim that $\bar{t} \in \bigcup_{i \in I} J(f_i) \cup \bigcup_{i \in I} J(f_i^\sigma)$. Otherwise $\bar{t} \in LC(f_i) \cap LC(f_i^\sigma)$ for all $i \in I$ and hence by (DP.iii) $\bar{t} \in RK(f_i) \cap RK(f_i^\sigma)$ for all $i \in I$. This implies (by (DP.i)) that there is $\varepsilon > 0$ such that $f(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \bar{t} + \varepsilon[$ contradicting the definition of \bar{t} . This proves the claim.

Claim: $\bar{t} \in \bigcup_{i \in I} PJ(\epsilon_i)$.

If $\bar{t} \in \mathcal{J}(f)$ this immediately follows from (P.i) and Lemma C.3(i). Hence, suppose $\bar{t} \notin \mathcal{J}(f)$. Then $\bar{t} \in \bigcup_{i \in I} J(f_i^\sigma)$. Since $f(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \bar{t}[$ it follows that $t(f^\sigma, \bar{t}) = t(f, \bar{t}) < \bar{t}$. By Lemma C.1(ii), $(f^\sigma)^{t(f, \bar{t})+}(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \bar{t}[$ implying that $\sigma_i^1((f^\sigma)^{t(f, \bar{t})+}, \bar{t}) = \sigma_i^1(f^\sigma, \bar{t})$ by (CRM.i). Now let $i \in I$ be such that $\bar{t} \in J(f_i^\sigma)$. Then $\sigma_i^1(f^\sigma, \bar{t}) = f_i^\sigma(\bar{t}) \neq f_{i-}^\sigma(\bar{t}) = f_{i+}^\sigma(t(f^\sigma, \bar{t}))$ where the inequality follows because $f_{i-}^\sigma(\bar{t})$ exists since by Lemma C.1(i) f^σ is constant on $]t(f^\sigma, \bar{t}), \bar{t}[$ and the last equality holds for the same reason. This implies that $E^{\sigma_i}(f^\sigma, t(f, \bar{t})) \neq \emptyset$ and that $\varepsilon^{\sigma_i}(f^\sigma, t(f, \bar{t})) = \min E^{\sigma_i}(f^\sigma, t(f, \bar{t}))$ exists and is strictly positive (by Lemma 2). Further, that $(f^\sigma)^{t(f, \bar{t})+}(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \bar{t}[$ implies that $\sigma_i^1((f^\sigma)^{t(f, \bar{t})+}, \tau) = \sigma_i^1(f^\sigma, \tau) = f_i^\sigma(\tau) = f_{i+}(\tau)$ for all $\tau \in]t(f, \bar{t}), \bar{t}[$ where the first equality follows from (CRM.i), the second follows because f^σ is the outcome induced by σ , and the third follows because f is constant on $]t(f, \bar{t}), \bar{t}[$. This yields $\varepsilon^{\sigma_i}(f^\sigma, t(f, \bar{t})) = \bar{t} - t(f^\sigma, \bar{t})$. Since $s_i \in S(\sigma_i)$ and $E^{\sigma_i}(f, t(f, \bar{t})) \neq \emptyset$ we obtain $s_i(x_{t(f, \bar{t})}^P(w)) = c_i(x_{t(f, \bar{t})}^P(w), \varepsilon^{\sigma_i}(f, t(f, \bar{t})))$ by (IS.iii) and hence that $\epsilon_i(t(f, \bar{t})) = \varepsilon^{\sigma_i}(f, t(f, \bar{t}))$. This yields $\epsilon_i(t(f, \bar{t})) = \bar{t} - t(f, \bar{t}) > 0$. By (P.iv) $\epsilon_i(\tau) = 0$ for all $\tau \in]t(f, \bar{t}), \bar{t}[$ because f is constant on $]t(f, \bar{t}), \bar{t}[$. Thus $t(f, \bar{t}) = \text{Prev}(\epsilon_i, \bar{t})$ implying that $\bar{t} = \text{Next}(\epsilon_i, \text{Prev}(\epsilon_i, \bar{t}))$. This proves the claim.

Claim: $f(\bar{t}) = f^\sigma(\bar{t})$.

To prove this, we distinguish two cases. First, if $IPJ(\epsilon, \bar{t}) = I$, then $s_i(x_{\bar{t}}(w)) = c_i(x_{\bar{t}}(w), \sigma_i^1(f, \bar{t}))$ holds by (IS.i) implying that $f_i(\bar{t}) = f_i^\sigma(\bar{t})$ for all $i \in I$. Second, if $IPJ(\epsilon, \bar{t}) \subsetneq I$, then by Lemma C.3(i), $t(f, \bar{t}) < \bar{t}$. For all $i \in IPJ(\epsilon, \bar{t})$, $s_i(x_{\bar{t}}(w)) = c_i(x_{\bar{t}}(w), \sigma_i^1(f, \bar{t}))$ holds by (IS.i) and we obtain that $f_i(\bar{t}) = f_i^\sigma(\bar{t})$ for all $i \in IPJ(\epsilon, \bar{t})$. By (P.i) $f_i(\bar{t}) = f_{i-}(\bar{t})$ for all $i \notin IPJ(\epsilon, \bar{t})$. Further since $\bar{t} \notin PJ(\epsilon_i)$ and w agrees with s_i on $[t(f, \bar{t}), \infty[$ we obtain (by Lemma 3(iii)) $\bar{t} \notin M(f, s_i)$ for all $i \notin IPJ(\epsilon, \bar{t})$. By Lemma C.6, $\sigma_i^1(f, \bar{t}) = f_{i-}(\bar{t})$ and hence $f_i(\bar{t}) = f_i^\sigma(\bar{t})$ for all $i \notin IPJ(\epsilon, \bar{t})$. This proves the claim.

As $\bar{t} \in \bigcup_{i \in I} (J(f_i) \cup J(f_i^\sigma))$, by the last claim $\bar{t} \in \bigcup_{i \in I} (J(f_i) \cap J(f_i^\sigma))$ and $IJ(f, \bar{t}) = IJ(f^\sigma, \bar{t})$. For all $i \in IJ(f, \bar{t})$ we obtain $f_{i+}(\bar{t}) = f_i(\bar{t}) = f_i^\sigma(\bar{t}) = f_{i+}^\sigma(\bar{t})$, where the first and third equalities hold by (DP.ii). For all $i \notin$

$IJ(f, \bar{t})$, $x_{\bar{t}}^R(w) \in X_i$ and by (IS.ii) $s_i(x_{\bar{t}}^R(w)) = c_i(x_{\bar{t}}^R(w), \sigma_i^2(f, \bar{t}))$. Hence $f_{i+}(\bar{t}) = \sigma_i^2(f, \bar{t}) = \sigma_i^2(f^\sigma, \bar{t}) = f_{i+}^\sigma(\bar{t})$, where the first and third equalities hold since f and f^σ are the outcomes induced by s and σ , respectively, and the second follows from (CRM.i). Thus $f(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \bar{t}]$ and $f_+(\bar{t}) = f_+^\sigma(\bar{t})$. By (DP.i) there is $\varepsilon > 0$ such that $f(\tau) = f^\sigma(\tau)$ for all $\tau \in [0, \bar{t} + \varepsilon]$, a contradiction with the choice of \bar{t} . Hence $f = f^\sigma$.

(iv) Let $s \in S$ and denote $f := f^{(\sigma^{s_i})_{i \in I}}$. For each $i \in I$, let $\bar{s}_i \in S(\sigma^{s_i})$. By (ii), $\bar{s}_i \sim s_i$ for all $i \in I$ implying that $f^{\bar{s}} = f^{(s_1, \bar{s}_{-1})}$. Applying (ii) again, yields $f^{(s_1, \bar{s}_2, (\bar{s}_i)_{i \neq 1, 2})} = f^{(s_1, s_2, (\bar{s}_i)_{i \neq 1, 2})}$. Iteratively proceeding this way, we obtain $f^{(\bar{s}_i)_{i \in I}} = f^s$. By (iii), $f^{(\bar{s}_i)_{i \in I}} = f$ implying that $f = f^s$. \square

Appendix 1.D: Proposition 5

Proposition 5. *Let $B = \{(b_i)_{i \in I} | b_i : H \times \mathbb{R}_+ \rightarrow A_i \text{ satisfies (B.i)-(B.v)} \forall i \in I\}$. Then $\{h \in H | \exists b \in B \text{ s.t. } b_i(h, t) = h_i(t) \text{ for all } t \in \mathbb{R}_+, i \in I\} = F$.*

Proof of Proposition 5. Let

$$H' := \{h \in H | \exists b \in B \text{ s.t. } b_i(h, t) = h_i(t) \text{ for all } t \in \mathbb{R}_+, i \in I\}.$$

“ \subseteq ” Let $h = (h_i)_{i \in I} \in H'$ and $b = (b_i)_{i \in I} \in B$ be such that $b_i(h, t) = h_i(t)$ for all $t \in \mathbb{R}_+$ and all $i \in I$. (DP.i) holds because by (B.ii)-(B.v) for every $t \in \mathbb{R}_+$ and every $i \in I$ there is $\varepsilon > 0$ such that $h_i|_{(t, t+\varepsilon)}$ is constant. To see (DP.ii), let $t \in \mathbb{R}_+$ and $i \in I$ be such that $t \notin LC(h_i)$. If $t = 0$ then $t \in RK(h_i)$ by (B.ii). If $t > 0$, by (B.iv) there is $\varepsilon > 0$ such that $h_i(\tau) = b_i(h, \tau) = b_i(h, t) = h_i(t)$ for all $\tau \in [t, t+\varepsilon[$, i.e. $t \in RK(h_i)$. To prove (DP.iii) we will show the contrapositive. Let $t \in \mathbb{R}_+$ be such that $t \notin \bigcup_{i \in I} J_i(h_i)$. Then by (B.iii) there is $\varepsilon > 0$ such that $h_i(\tau) = b_i(h, \tau) = b_i(h, t) = h_i(t)$ for all $\tau \in [t, t+\varepsilon[$ and all $i \in I$, i.e. $t \in \bigcap_{i \in I} RK(h_i)$ and hence $t \notin \bigcup_{i \in I} R(h_i)$.

“ \supseteq ” Fix $f \in F$. For $h \in H$, define

$$t_f(h) := \begin{cases} \inf\{t \in \mathbb{R}_+ | h(t) \neq f(t)\} & \text{if } h \neq f, \\ \infty, & \text{if } h = f. \end{cases}$$

For each $i \in I$ define

$$b_i(h, t) = \begin{cases} f_i(t), & \text{if } h(\tau) = f(\tau) \text{ for all } \tau \in [0, t[\\ f_i(t_f(h)), & \text{otherwise.} \end{cases}$$

Fix $i \in I$. Property (B.i) is satisfied by construction. To see (B.ii), note that since $0 \in RK(f_i)$ there is $\varepsilon > 0$ such that $f_i(\tau) = f_i(0)$ for all $\tau \in [0, \varepsilon[$. If $t_f(h) = 0$ then $b_i(h, \tau) = f_i(0) = b_i(h, 0)$ for all $\tau \in [0, \varepsilon[$. If $t_f(h) > 0$ then $b_i(h, \tau) = f_i(\tau) = f_i(0) = b_i(h, 0)$ for all $\tau \in [0, \min\{\varepsilon, t_f(h)\}[$.

To prove (B.iii), let $h \in H$ and $t \in \mathbb{R}_+$ be such that $t \in \bigcap_{j \in I} LC(h_j)$. We distinguish two cases. First, if $t_f(h) \leq t$, then by construction $b_i(h, \tau) = f_i(t_f(h)) = b_i(h, t)$ for all $\tau \in [t, \infty[$. Second, if $t_f(h) > t$ then by construction there is $\varepsilon > 0$ such that $b_i(h, \tau) = f_i(\tau)$ for all $\tau \in [t, t + \varepsilon[$. Further $t \in \bigcap_{j \in I} LC(f_j)$ since $f(\tau) = h(\tau)$ for all $\tau \in [0, t_f(h)[$. By (DP.ii) and (DP.iii) it follows that $t \in RK(f_i)$ and hence there is $\varepsilon' > 0$ such that $b_i(h, \tau) = f_i(\tau) = f_i(t) = b_i(h, t)$ for all $\tau \in [t, t + \varepsilon'[$.

To show (B.iv), let $h \in H$ and $t \in \mathbb{R}_+$ be such that $t \notin LC(h_i)$. If $t_f(h) \leq t$ then $b_i(h, \tau) = f_i(t_f(h)) = b_i(h, t)$ for all $\tau \in [t, \infty[$. If $t < t_f(h)$ then there is $\varepsilon > 0$ such that $b_i(h, \tau) = f_i(\tau)$ for all $\tau \in [t, t + \varepsilon[$. Since $t \notin LC(f_i)$, by (DP.ii) $t \in RK(f_i)$. Hence there is $\varepsilon' > 0$ such that $b_i(h, \tau) = f_i(\tau) = f_i(t) = b_i(h, t)$ for all $\tau \in [t, t + \varepsilon'[$.

Finally, we will prove (B.v). Let $h \in H$ and $t \in \mathbb{R}_+$ be such that $t \in LC(h_i) \setminus \bigcap_{j \neq i} LC(h_j)$. Once more, we distinguish two cases. First, if $t_f(h) \leq t$ then $b_i(h, \tau) = f_i(t_f(h)) = b_i(h, t)$ for all $\tau \in [t, \infty[$. Second, if $t < t_f(h)$ then there is $\varepsilon > 0$ such that $b_i(h, \tau) = f_i(\tau)$ for all $\tau \in [t, t + \varepsilon[$. Since by (DP.i) f is piecewise constant, the conclusion follows.

Thus $b_i \in B_i$ for every $i \in I$ and by construction $b_i(f, t) = f_i(t)$ for all $t \in \mathbb{R}_+$ and all $i \in I$. Thus $f \in H'$. \square

References Chapter 1

- ALÓS-FERRER, C., AND J. KERN (2013): “On the Definition of Continuous-Time Repeated Games with Imperfect Monitoring,” mimeo.
- ALÓS-FERRER, C., J. KERN, AND K. RITZBERGER (2011): “Comment on ‘Trees and Extensive Forms’,” *Journal of Economic Theory*, 146(5), 2165–2168.
- ALÓS-FERRER, C., AND K. RITZBERGER (2005): “Trees and Decisions,” *Economic Theory*, 25(4), 763–798.
- (2008): “Trees and Extensive Forms,” *Journal of Economic Theory*, 43(1), 216–250.
- (2013a): “Large Extensive Form Games,” *Economic Theory*, 52(1), 75–102.
- (2013b): “Characterizing Existence of Equilibrium for Large Extensive Form Games,” mimeo.
- ANDERSON, R. M. (1984): “Quick-Response Equilibrium,” Working Papers in Economic Theory and Econometrics, #IP-323, C.R.M, University of California at Berkeley.
- BERGIN, J. (1992): “A model of strategic behaviour in repeated games,” *Journal of Mathematical Economics*, 21(2), 113–153.
- (2006): “The folk theorem revisited,” *Economic Theory*, 27(2), 321–332.
- BERGIN, J., AND W. B. MACLEOD (1993): “Continuous Time Repeated Games,” *International Economic Review*, 34(1), 21–37.
- DAVIDSON, R., AND R. HARRIS (1981): “Nonconvexities in Continuous-Time Investment,” *Review of Economic Studies*, 48(2), 235–253.

- FRIEDMAN, A. (1994): “Differential Games,” in *Handbook of Game Theory*, ed. by R. J. Aumann, and S. Hart, vol. 2, pp. 781–799. Elsevier, Amsterdam, Princeton, NJ.
- FUDENBERG, D., AND D. LEVINE (1986): “Limit Games and Limit Equilibria,” *Journal of Economic Theory*, 38(2), 261–279.
- JECH, T. (2002): *Set Theory: The Third Millennium Edition, revised and expanded*. Springer Verlag, Berlin Heidelberg New York.
- KELLER, G., AND S. RADY (1999): “Optimal Experimentation in a Changing Environment,” *Review of Economic Studies*, 66(3), 475–507.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73(1), 39–68.
- KUHN, H. W. (1953): “Extensive Games and the Problem of Information,” in *Contributions to the Theory of Games*, ed. by H. Kuhn, and A. Tucker, vol. II. Princeton University Press, Princeton, NJ.
- PERRY, M., AND P. J. RENY (1993): “A Non-cooperative Bargaining Model with Strategically Timed Offers,” *Journal of Economic Theory*, 59(1), 50–77.
- (1994): “A Noncooperative View of Coalition Formation and The Core,” *Econometrica*, 62(4), 795–817.
- PICCIONE, M., AND A. RUBINSTEIN (1997): “On the Interpretation of Decision Problems with Imperfect Recall,” *Games and Economic Behavior*, 20(1), 3–24.
- SANNIKOV, Y. (2007): “Games with Imperfectly Observable Actions in Continuous Time,” *Econometrica*, 75(5), 1243–1523.
- SIMON, L., AND M. B. STINCHCOMBE (1989): “Extensive Form Games in Continuous Time, Part I: Pure Strategies,” *Econometrica*, 57(5), 1171–1214.

CHAPTER 1

REPEATED GAMES IN CONTINUOUS TIME AS EXTENSIVE FORM GAMES

STINCHCOMBE, M. B. (1992): “Maximal Strategy Sets for Continuous-Time Game Theory,” *Journal of Economic Theory*, 56(2), 235–265.

VON NEUMANN, J., AND O. MORGENSTERN (1944): *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

CHAPTER 2

COMMENT ON “TREES AND EXTENSIVE FORMS”

2.1 Introduction

This paper corrects the formulation of a property in Alós-Ferrer and Ritzberger (2008) (henceforth referred to as AR) which determines when an Extensive Decision Problem¹ (EDP) is called an Extensive Form (EF). We present a corrected formulation of the property and show which and how results in AR are affected by the reformulation. We further present a counterexample which shows that some of the original results do not hold under the restated version of the property.

The rest of the paper is organized as follows. Section 2.2 introduces the necessary notation. Section 2.3 presents the correctly stated version of the property and illustrates in detail which and how results in AR change under the new formulation. Finally, Section 2.4 contains a counterexample for the results that do not hold under the new version of the property.

2.2 Preliminaries

We will rely on the notation and concepts introduced in Sections 1.2.1 and 1.3.1. Some additional notation is required, however. For a game tree $T = (N, \supseteq)$, a *history* is a nonempty chain h in N that is not maximal in T and for which $\uparrow x \subseteq h$ for all $x \in h$. For as history h in T a *continuation* is the complement of h in a play that contains h . A game tree is *weakly up-discrete* if all maximal chains in $\downarrow x \setminus \{x\}$ have maxima, for all nodes $x \in N$ for which $\downarrow x \setminus \{x\} \neq \emptyset$. A game tree is *coherent* if every history without

¹See Definition 2 in Section 1.2.1.

minimum has at least one continuation with a maximum.² Given a move $x \in X$ and a play $w \in x$ the *perfect information choice* $\gamma(x, w) \subseteq W$ is the set of plays $\gamma(x, w) = \bigcup \{z \mid w \in z \in \downarrow x \setminus \{x\}\}$. A game tree (N, \supseteq) has *available choices* if $\gamma(x, w) \subsetneq x$ for all $w \in x$ and all $x \in X$. Let $\Gamma(T) = \{\gamma(x, w) \mid w \in x \in X\}$ be the set of perfect information choices and $S(N)$ be the set of strange nodes in N . If a game tree T has available choices then $\Pi(T) = (T, C_1)$ where $C_1 = \Gamma(T) \cup S(N)$ is a single-player EDP (AR, Theorem 1).³ A game tree is *selective*, if for all $w, w' \in W$, $w \neq w'$ implies that there is $x \in X$ such that $w, w' \in x$ and $\gamma(x, w) \neq \gamma(x, w')$.

For an EDP, given a strategy profile $s = (s_i)_{i \in I} \in S$ and a history h , a node $x \in N$ is *discarded* at h , if $x \subsetneq W(h) = \bigcap_{y \in h} y$ and there are $z \in \uparrow x \setminus \{x\}$, $i \in I(z)$, and $c \in A_i(z)$ such that $z \subseteq W(h)$ and $x \subseteq c \neq s_i(z)$. $D^h(s)$ denotes the set of all nodes discarded at h and $U^h(s) = \{x \mid x \subseteq W(h)\} \setminus D^h(s)$ is the set of *undiscarded* nodes at h . The strategy profile $(s_i)_{i \in I} \in S$ *induces an outcome after history h* if there is $w \in W(h)$ such that $w \in R_s^h(w)$, where $R_s^h(w) = \bigcap \{s_i(x) \mid w \in x \subseteq W(h), x \in X, i \in I(x)\}$. If every strategy profile induces an outcome after every history then the EDP is *everywhere playable*.

2.3 Corrected Formulation and Changes

In AR, an EF is defined as an EDP which satisfies a stronger version of property (EDP.iii), namely

(EDP.iii') for all $y, y' \in N$, if $y \cap y' = \emptyset$ then there are $i \in I$ and $c, c' \in C_i$ such that $y \subseteq c$, $y' \subseteq c'$, $c \cap c' = \emptyset$, and $P(c) \cap P(c') \neq \emptyset$.

This property is misstated in AR. The correct formulation is as follows:

(EDP.iii') for all $y, y' \in N$, if $y \cap y' = \emptyset$ then there are $x \in X$, $i \in I(x)$ and $c, c' \in C_i$ such that $x \in P(c) \cap P(c')$, $y \subseteq x \cap c$, $y' \subseteq x \cap c'$, and $c \cap c' = \emptyset$.

²It can be shown (AR, Corollary 3) that for regular game trees weak up-discreteness and coherence is equivalent to up-discreteness.

³Theorem 1 in AR actually states that a game tree T having available choices is *equivalent* to $\Pi(T)$ being a single-player EDP.

In what follows the numbering of results and definitions corresponds to that in AR. With the corrected formulation of (EDP.iii') provided here, Proposition 7, Proposition 9, Theorem 5, and Corollary 4 are true as stated. Propositions 5 and 8(b) do not hold with the new version of (EDP.iii'). Proposition 10 remains true as stated, but requires a different proof. Finally, Theorem 6 and Corollary 5 remain true with a (slight) change of the hypotheses. We now explain the necessary changes in more detail.

Proposition 7, Proposition 9, Theorem 5, and Corollary 4

All those results are true as stated, with the corrected formulation of (EDP.iii') stated here. The proofs (with minor, straightforward adaptations) remain as in AR.

Proposition 7. *An EDP (T, C) satisfies (EDP.iii') if and only if T is selective and*

(EDP.ii') $x \cap [\bigcap_{i \in I(x)} c_i] = \gamma(x, w)$ for some $w \in x \cap [\bigcap_{i \in I(x)} c_i]$ for all $(c_i)_{i \in I(x)} \in \times_{i \in I(x)} A_i(x)$ and for all $x \in X$.

This is not true for the version of (EDP.iii') incorrectly stated in the paper, but holds under the new formulation.

Proposition 9. *An EDP (T, C) with a weakly up-discrete tree $T = (N, \supseteq)$ is an EF if and only if T is selective.*

Theorem 5. *Consider an EF and fix a pure strategy combination $s \in S$. If $w \in R_s(w)$ then (a) $R_s(w) = \{w\}$, and (b) if $w' \in R_s(w')$ then $w' = w$.*

Corollary 4. *The tree of an EF is selective and, hence, regular.*

Propositions 5 and 8(b)

These results are not true for the corrected version of (EDP.iii') (see Example 14 in Section 2.4). There is a common mistake in the proofs of Propositions 5 and 8(b) that is as follows. The construction of a strategy s selecting both w and w' fails when, for a node x with $w \in x$ but $w' \notin x$, $s_i(x)$ is required to pick up the choice leading to w . For a different node x' with $w' \in x'$ but

$w \notin x'$, $s_i(x')$ will be required to pick up the choice leading to w' . However, it might be the case that x and x' belong to the same information set of player i , in which case an incompatibility arises.

This problem cannot appear under perfect information. Therefore, the statement remains true for the game $\Pi(T)$. In this case, it follows from Proposition 7 that the corrected version of (EDP.iii') reduces to selectiveness, because (EDP.ii') is always fulfilled for $\Pi(T)$. The resulting property coincides with the original formulation of Proposition 8(a) in AR. The following result replaces the original versions of both Proposition 5 and 8.

Proposition 5. *Consider a game tree T with available choices. If T is not selective, then the perfect information EDP $\Pi(T)$ fails outcome uniqueness.*

Consider the class of weakly up-discrete trees. This includes the class on which every EDP is everywhere playable. By Theorem 5, Proposition 9, and the new version of Proposition 5 above, a weakly up-discrete tree T is selective if and only if every EDP (T, C) satisfies outcome uniqueness.

Proposition 10

The statement of Proposition 10 is the following.

Proposition 10. *Fix a history h for a game tree $T = (N, \supseteq)$. If for an arbitrary EF (T, C) every strategy combination induces outcomes after h , then for the problem $\Pi(T)$ every strategy induces outcomes after h .*

Proposition 10 is correct as stated also with the version of (EDP.iii') given here. But its proof contains the same mistake pointed out for Propositions 5 and 8(b). The correct proof is as follows..

Proof of Proposition 10. Suppose for some history h there is a strategy s' for $\Pi(T)$ that does not induce an outcome after h . Let h' be a maximal chain in $U^h(s')$ (the undiscarded nodes after h) and $W(h') = \bigcap_{x \in h'} x$. Fix a play w as follows. If h' has a minimum z (which then cannot be a terminal node by hypothesis), let $w \in s'(z)$. Otherwise, let $w \in W(h')$. If $U^h(s') = \emptyset$ (and hence there is no such chain h'), fix an arbitrary $w \in W(h)$.

Consider now an arbitrary EF (T, C) . By Proposition 7 T is selective. We construct a strategy profile in (T, C) as follows. For all $y \in N$ such that $w \in y \subseteq W(h)$ and $i \in I(y)$, choose $s_i(y)$ such that $y \cap \left[\bigcap_{i \in I(y)} s_i(y) \right] = s'(y)$. This is possible by (EDP.ii') and (EDP.iv). If $y \in h'$, then of course $s'(y) = \gamma(y, w)$. For any other node, specify the strategy profile s arbitrarily. Notice that we determine s only along a play, which is possible by (EDP.iv). We claim that $w \notin R_s^h(w)$. For, if it were, by construction of s , we would obtain $w \in R_{s'}^h(w)$ for $\Pi(T)$, a contradiction.

Let $w' \in W(h)$ with $w' \neq w$. By selectiveness there exists $x \in X$ such that $w, w' \in x$ and $\gamma(x, w) \cap \gamma(x, w') = \emptyset$ (by Proposition 1(a) in AR). Notice that, necessarily, $x \subseteq W(h)$. There are two possibilities. If $x \notin h'$, $x \in D^h(s')$ for $\Pi(T)$ which implies by construction that $x \in D^h(s)$ for (T, C) . Hence $w' \notin R_s^h(w')$. If $x \in h'$, then $s'(x) = \gamma(x, w) \neq \gamma(x, w')$. Since $s'(x) = x \cap \left[\bigcap_{i \in J(x)} s_i(x) \right]$, it follows that $w' \notin R_s^h(w')$. Since $w' \in W(h)$ was arbitrary, we conclude that s does not induce an outcome after h in (T, C) . \square

Theorem 6 and Corollary 5

The “only if” direction of Theorem 6 relied on (the original version of) Proposition 5 and needs to be (slightly) reformulated by changing the hypothesis that (T, C) is an EDP to the hypothesis that it is an EF:

Theorem 6. *An EF (T, C) satisfies (A1) and (A2) if and only if the (rooted) game tree $T = (N, \supseteq)$ is regular, weakly up-discrete, and coherent.⁴*

The proof of the “if” implication remains essentially the same as in AR. The only caveat is that Proposition 9 refers to the corrected version of (EDP.iii') provided here and hence Theorem 5, which requires this version, can be used. To see the “only if” direction, note that Proposition 7 implies that T is selective. Hence it is also regular by Proposition 6(a) in AR. By Proposition 10 the game $\Pi(T)$ fulfills (A1). By Corollary 2 in AR, every EDP defined on T satisfies (A1). Theorem 3 in AR then implies that T is up-discrete, hence

⁴See Section 1.3.1 for the statements of (A1) and (A2).

weakly up-discrete and coherent by Corollary 3 in AR. This argument does not make use of Proposition 5.

The statement of Corollary 5 remains true when EDP is replaced by EF in its formulation:

Corollary 5. (a) *If an EF satisfies (A1) and (A2), then so does every EF with the same tree.*

(b) *An EF satisfies (A1) and (A2) if and only if its tree is regular and up-discrete. Furthermore, the EDP is then everywhere playable.*

2.4 Example

Consider the direct approach to modeling repeated games in continuous time presented in Section 1.2.2 where W is the set of functions $f : \mathbb{R}_+ \rightarrow A$, and A is some fixed set of actions containing at least two elements. Defining $N = \{x_t(f) \mid t \in \mathbb{R}_+, f \in W\}$, where $x_t(f) = \{g \in W \mid g(\tau) = f(\tau) \forall \tau \in [0, t]\}$ for $f \in W$ and $t \in \mathbb{R}_+$ it can be shown that $T = (N, \supseteq)$ is a game tree (Alós-Ferrer and Ritzberger, 2005). One can define an EDP on this tree using choices $c_t(f, a) = \{g \in x_t(f) \mid g(t) = a\}$ for every $t \in \mathbb{R}_+$, $f \in W$, and $a \in A$ (Alós-Ferrer and Ritzberger, 2005, Example 16). The resulting EDP is referred to as the *differential game*.

There is a mistake in Example 7 of AR concerning the computation of the perfect information choices for the tree of the differential game. There it is falsely stated that for a play $g \in x_t(f) \in N$, the perfect information choice $\gamma(x_t(f), g) = c_t(f, a)$. The correct expression, however, is as follows:

$$\gamma(x_t(f), g) = \{h \in W \mid \exists \tau > t \text{ such that } h|_{[0, \tau[} = g|_{[0, \tau[}\}$$

The differential game is hence different from $\Pi(T)$ for the tree T . In Example 13 in AR, this gives rise to another mistake. There it is stated that the tree of the differential game is not selective. In fact, the tree of the differential game is selective. To see this, let $f, g \in W$ with $f \neq g$. If $f(0) \neq g(0)$, then $\gamma(W, f) \neq \gamma(W, g)$. If $f(0) = g(0)$, let $t^* = \sup\{\tau > 0 \mid f|_{[0, \tau[} = g|_{[0, \tau[}\}$. Then $x_{t^*}(f) = x_{t^*}(g)$ and $\gamma(x_{t^*}(f), f) \neq \gamma(x_{t^*}(f), g)$.

The point of Example 13 in AR was to provide a counterexample showing that a regular game tree is not necessarily selective, i.e. that the converse of Proposition 6(a) in AR does not hold. Yet, this is already accomplished by Example 4 in AR.

The differential game actually fails (EDP.ii') and hence (by Proposition 7) also the corrected version of (EDP.iii'). On the other hand, the game $\Pi(T)$ based on the same tree is, in fact, an EF. The comment after the statement on Proposition 7 needs to be adjusted accordingly (AR, p. 240-241).

We now provide a common counterexample to the statements of Propositions 5 and 8(b) under the corrected formulation of (EDP.iii') given here.

Example 14. Let T be the tree of the differential game as above. Consider an EDP (T, C) based on this tree as follows. There is a continuum of players, $I = \mathbb{R}_+$. Each player chooses an action $a \in A$. Player t is the only player who plays at time t . All nodes at period t belong to the same information set, i.e. no player ever learns any previous decision. That is, the choices of player t are of the form $c_t(a) = \{f \in W \mid f(t) = a\}$. The set of nodes where player t is active is the “slice” $X_t = \{x_t(f) \mid f \in W\}$. Further, each such slice is the only information set of the corresponding player, where all choices of the form $c_t(a)$ are available, $P(c_t(a)) = X_t$ for all $a \in A$.

This game is just the “cascading information sets” version of the normal-form game where each player in $I = \mathbb{R}_+$ chooses an action $a \in A$. The strategy of player t is simply an action $a \in A$ and the outcomes (plays) of the game are simply functions $f : \mathbb{R}_+ \rightarrow A$. Hence, (A1) and (A2) follow immediately.

Recalling the expression of $\gamma(x_t(f), g)$ given above, it is immediate that this EDP fails (EDP.ii'), hence (by Proposition 7) also (EDP.iii'). Since outcome uniqueness (A2) is satisfied, this shows that Propositions 5 and 8(b) as stated in AR do not hold with the corrected version of (EDP.iii').

References Chapter 2

ALÓS-FERRER, C., AND K. RITZBERGER (2005): “Trees and Decisions,”
Economic Theory, 25(4), 763–798.

——— (2008): “Trees and Extensive Forms,” *Journal of Economic Theory*,
43(1), 216–250.

CHAPTER 3
CIRCULANT GAMES

3.1 Introduction

Games with cyclical structures are ubiquitous in game theory. Simple examples like Matching Pennies and Rock-Paper-Scissors are routinely used to illustrate the concepts of mixed strategies and mixed strategy Nash equilibria in any introductory class to game theory. Beyond their pedagogical value, these simple examples have a wide range of application in game theory. Evolutionary game theory is one prominent example and, e.g., the mating strategies of the common side-blotched lizard have been shown to follow a rock-paper-scissors pattern (Sinervo and Lively, 1996). A cyclical game structure can be captured by circulant payoff matrices, in which each row vector is rotated by one element relative to the preceding row vector (Hofbauer, Schuster, Sigmund, and Wolff, 1980; Diekmann and van Gils, 2009). Games with circulant payoff matrices have been studied extensively in the literature on evolutionary game theory (Hofbauer and Sigmund, 1998) and population dynamics (Hofbauer, Schuster, Sigmund, and Wolff, 1980; Diekmann and van Gils, 2009). Circulant payoff matrices also underly certain classes of coordination games, starting with matching games, that have been studied in the literature on symmetries and focal points (Casajus, 2000; Janssen, 2001).¹

The class of games we study here is important for at least two fields of applications. First, the analysis of the convergence properties of various evolutionary dynamics for cyclical game structures has often focused on uniformly mixed strategies. Games in which this strategy profile is the unique equilibrium constitute important examples of convergence failure (see, e.g.,

¹The simplest example of a matching game is Heads and Tails. If both players match the strategy of the other player each player gets a payoff of 1, otherwise each player receives a payoff of zero.

Sandholm, 2010, Chapter 9.2.1, pp. 327-330). Still, many games with a cyclical structure have more than one equilibrium and the non-convergence to one particular equilibrium may not be conclusive for the convergence properties of the whole system. Second, matching games and more general coordination games constitute an archetypal framework to analyze features external to the games' formal structure. The cyclical game structure provides a framework where strategies cannot be differentiated according to differences in payoffs. Yet, matching games are just one particular representation of such symmetric frameworks and many different, equally appropriate cyclical game structures may exist (see, e.g., Alós-Ferrer and Kuzmics, 2013). A rigorous characterization of the set of Nash equilibria of cyclical game structures in general is still missing.

The aim of this paper is to bridge these gaps and provide a more general analysis of games with a cyclical structure. More precisely, we investigate a class of finite two-player normal-form $n \times n$ games we coin *circulant games*, in which the players' payoff matrices are circulant. We also require that the first row of each matrix is ordered. This approach allows us to integrate classical examples from Game Theory into one single class of games. Well-known games such as the ones mentioned above, as well as subclasses of common-interest and coordination games (including matching games) belong to the class of circulant games.

Our results shed new light on the common features shared by these games. Our main results identify the *exact* number of (pure or mixed) Nash equilibria in circulant games. We also obtain necessary and sufficient conditions for the existence of pure strategy Nash equilibria and, in case of non-existence, for the uniqueness of the uniformly mixed Nash equilibrium (a profile which we show to be a Nash equilibrium for all circulant games). As a consequence of our main results we obtain that the maximal number of Nash equilibria in these games is exactly $2^n - 1$. The number of pure strategy Nash equilibria is either 0, 1, 2, or n . Further, we are also able to characterize the structure of the set of mixed Nash equilibria. The best response correspondences induce an equivalence relation on each player's set of pure strategies. In any Nash equilibrium all strategies within one equivalence class are either played

with strictly positive or with zero probability. We show how to derive the equivalence classes, allowing for a characterization of the support of all Nash equilibrium strategies.

Our results also contribute to the literature on the number of Nash equilibria in finite two-player normal-form $n \times n$ games. Provided that such a game is non-degenerate the number of Nash equilibria is finite and odd (see, e.g., Shapley, 1974). Quint and Shubik (1997) show that for any odd integer number y between 1 and $2^n - 1$, there exists a game with exactly y Nash equilibria. However, as shown in von Stengel (1997), $2^n - 1$ is not an upper bound on the number of Nash equilibria in such games. New upper bounds on the number of distinct Nash equilibria are established in Keiding (1998) and von Stengel (1999). For the class of coordination games $2^n - 1$ is the (tight) upper bound on the number of equilibria (Quint and Shubik, 2002). Our results show that this is also true for the class of circulant games.

Recently, several other articles have analyzed subclasses of games with a special focus on different notions of cyclicity. Duersch, Oechssler, and Schipper (2012) consider symmetric two-player zero-sum normal-form games and define generalized rock-paper-scissors matrices (*gRPS*) in terms of best response cycles. In their setting, a game has a pure strategy Nash equilibrium if and only if it is not a *gRPS*. Bahel (2012) and Bahel and Haller (2013) examine zero-sum games that are based on cyclic preference relations on the set of actions and characterize the set of Nash equilibria. In the former paper, actions are distinguishable, i.e., one specific action is the beginning of the cyclic relation, and there exists a unique Nash equilibrium. In the latter, actions are anonymous, i.e., each action can be seen as the beginning of the cycle without affecting the relation, and depending on the number of actions the Nash equilibrium is unique or there exists an infinite number of Nash equilibria.

The remainder of this paper is structured as follows. Section 3.2 introduces the class of circulant games. Section 3.3 states the main results and presents a recipe to characterize the support of all Nash equilibrium strategies for a given circulant game. Section 3.4 presents generalizations of circulant games and Section 3.5 concludes. All proofs are relegated to the appendix.

3.2 Circulant Games

Let $\Gamma = ((S_1, S_2), (\pi_1, \pi_2))$ be a finite two-player normal-form game where $S_i = \{0, 1, \dots, n_i - 1\}$ denotes player i 's set of pure strategies and $\pi_i : S_1 \times S_2 \rightarrow \mathbb{R}$ denotes player i 's payoff function for $i = 1, 2$.² We will write player i 's payoff function as the $n_1 \times n_2$ matrix $A_i = (a_{kl}^i)_{k \in S_1, l \in S_2}$ given by $a_{kl}^i = \pi_i(k, l)$. Thus in both matrices each row corresponds to a pure strategy of player 1 and each column to a pure strategy of player 2. Following the notation in e.g. Alós-Ferrer and Kuzmics (2013), we will also write $\pi_i(s|s')$ for player i 's payoff if he chooses a strategy s and player $-i$ chooses strategy s' . The set of mixed strategies for player i is denoted by Σ_i . For $\sigma_i \in \Sigma_i$, $\sigma_i(s)$ denotes the probability that σ_i places on the pure strategy $s \in S_i$. The set of all pure strategies played with strictly positive probability is denoted by $\text{supp}(\sigma_i)$. Payoff functions are extended to the sets of mixed strategies through expected payoffs. Given a mixed strategy σ_{-i} of player $-i$, a *best response* for player i against σ_{-i} is a strategy σ_i such that $\pi_i(\sigma_i|\sigma_{-i}) \geq \pi_i(\sigma'_i|\sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$. The set of best responses for player i against a strategy σ_{-i} of the other player is denoted by $BR_i(\sigma_{-i})$. A finite two-player normal-form game is *non-degenerate* (Quint and Shubik, 1997) if for any mixed strategy σ_i of player i with $|\text{supp}(\sigma_i)| = m$, player $-i$ has at most m pure strategy best responses against σ_i . In what follows Γ_n denotes a finite two-player normal-form game in which $S_1 = S_2 = S^n = \{0, \dots, n - 1\}$.

The following two results are well-known and will be used throughout the paper.

Proposition 1 (Best Response Condition, Nash, 1951). *Let Γ be a finite two-player normal-form game. Then $\sigma_i \in \Sigma_i$ is a best response to $\sigma_{-i} \in \Sigma_{-i}$ if and only if for all $s_i \in S_i$*

$$\sigma_i(s_i) > 0 \Rightarrow \pi_i(s_i|\sigma_{-i}) = \max_{s \in S_i} \pi_i(s|\sigma_{-i}).$$

²We choose to label players' strategies from 0 to $n_i - 1$ as this will later simplify notation significantly.

Proposition 2 (Shapley, 1974; Quint and Shubik, 1997). *Let Γ be a finite non-degenerate two-player normal-form game with strategy set $S_1 = S_2 = S$. Then*

- (i) Γ has a finite and odd number of Nash equilibria.
- (ii) if $T_1, T_2 \subseteq S$ then Γ has at most one Nash equilibrium (σ_1, σ_2) such that $\text{supp}(\sigma_1) = T_1$ and $\text{supp}(\sigma_2) = T_2$.

Circulant games will be defined through circulant matrices (see Davis, 1979) which we introduce now.

Definition 1. A matrix $A \in \mathbb{R}^{n \times n}$ is *circulant* if it has the form

$$A = \begin{pmatrix} a_0 & a_1 & a_2 & \cdots & a_{n-1} \\ a_{n-1} & a_0 & a_1 & \cdots & a_{n-2} \\ a_{n-2} & a_{n-1} & a_0 & \cdots & a_{n-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_1 & a_2 & a_3 & \cdots & a_0 \end{pmatrix}$$

and *anti-circulant* if

$$A = \begin{pmatrix} a_0 & \cdots & a_{n-3} & a_{n-2} & a_{n-1} \\ a_1 & \cdots & a_{n-2} & a_{n-1} & a_0 \\ a_2 & \cdots & a_{n-1} & a_0 & a_1 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ a_{n-1} & \cdots & a_{n-4} & a_{n-3} & a_{n-2} \end{pmatrix}.$$

Circulant and anti-circulant matrices are hence fully specified by the first row vector. Each remaining row vector is rotated by one element relative to the preceding row vector. We are now ready to define a circulant game.

Definition 2. A two-player normal-form game Γ_n is a *circulant game* if

- (i) each player's payoff matrix is either circulant or anti-circulant,
- (ii) $a_0^1 > a_1^1 \geq \cdots \geq a_{n-1}^1$, and

- (iii) either $a_{n-k}^2 > a_{n-k+1}^2 \geq \cdots \geq a_{n-1}^2 \geq a_0^2 \geq a_1^2 \geq \cdots \geq a_{n-k-1}^2$ or
 $a_{n-k}^2 > a_{n-k-1}^2 \geq \cdots \geq a_1^2 \geq a_0^2 \geq a_{n-1}^2 \geq \cdots \geq a_{n-k+1}^2$ for some
 $1 \leq k \leq n$.

The parameter k is called the *shift* of Γ_n .

The shift describes the position of player 2's largest payoff in the first row of his payoff matrix. As we will see later, knowing the shift and the number of pure strategies suffices to determine the exact number and structure of Nash equilibria in circulant games.

Note that if A_i is circulant then $a_{ij} = a_{j-i}$ and if A_i is anti-circulant then $a_{ij} = a_{i+j}$ where the indices are to be read modulo n , e.g. $-1 = n - 1$, $n + 1 = 1$, etc. In a circulant game, if player 1's payoff matrix is circulant then $\pi_1(s|s') = a_{s'-s}^1$ and if player 1's payoff matrix is anti-circulant then $\pi_1(s|s') = a_{s+s'}^1$. Similarly if player 2's payoff matrix is circulant then $\pi_2(s|s') = a_{s-s'}^2$ and if player 2's payoff matrix is anti-circulant then $\pi_2(s|s') = a_{s+s'}^2$. Throughout the paper the sum and difference of two strategies (and the multiplication of a strategy with an integer) in a circulant game is to be read modulo n .

In a circulant game the entries in the first row of player 1's payoff matrix (weakly) decrease when moving from left to right with a_0^1 being the unique maximum payoff. The entries in the first row of player 2's payoff matrix (weakly) decrease either when moving from the largest payoff to the right, or when moving from the largest payoff to the left. The shift k is determined by the position of the unique maximum payoff in the first row of player 2's payoff matrix. A shift of $k = n$ corresponds to a_0^2 being player 2's largest payoff. A shift of $k = 0$ is of course possible but for notational convenience is formally represented by a shift of $k = n$.

Since in a circulant game the sum of the payoffs in each row and each column is constant, if one player plays the completely uniformly mixed strategy, then all of the other player's pure strategies yield the same payoff. An immediate consequence of this is the following.

Lemma 1. *Let Γ_n be a circulant game. Then $\sigma^* = (\sigma_1^*, \sigma_2^*)$ where $\sigma_i^*(s) = 1/n$ for all $s \in S^n$, $i = 1, 2$, is a Nash equilibrium of Γ_n .*

We can classify circulant games according to whether the players' payoff matrices "rotate" in the same or in opposite directions.

Definition 3. A circulant game is *iso-circulant* if the players' payoff matrices are either both circulant or both anti-circulant matrices. It is *counter-circulant* if one player's payoff matrix is circulant and the other player's payoff matrix is anti-circulant.

For $n = 2$ every iso-circulant game is also counter-circulant and vice versa, as any circulant 2×2 matrix is also anti-circulant. For $n \geq 3$, however, the class of iso-circulant games is disjoint from the class of counter-circulant games. Iso-circulant games with shift $k = n$ capture the class of (weakly ordered) circulant coordination games.

Example 15 (Matching Pennies).

The game given by

$$A_1 = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$$

is the well-known Matching Pennies game. Both players' payoff matrices are circulant (and anti-circulant) and for player 2, $a_{n-1}^2 = a_1^2 = 1$ is the largest payoff. Hence, it is an iso-circulant (and also a counter-circulant) game with shift $k = 1$. $[(1/2, 1/2), (1/2, 1/2)]$ is a Nash equilibrium of this game. As we will show later it is the unique one.

Example 16 (Rock-Paper-Scissors).

The game given by

$$A_1 = \begin{pmatrix} 3 & 2 & 1 \\ 2 & 1 & 3 \\ 1 & 3 & 2 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \\ 3 & 1 & 2 \end{pmatrix}.$$

is Rock-Paper-Scissors. Strategies are labeled such that for player 1, strategy 0 is 'Rock', strategy 1 is 'Scissors', and strategy 2 is 'Paper' and for player 2, strategy 0 is 'Scissors', strategy 1 is 'Rock', and strategy 2 is 'Paper'. Both players' payoff matrices are anti-circulant and for player 2, $a_{n-1}^2 =$

$a_2^2 = 3$ is the largest payoff. This is an iso-circulant game with shift $k = 1$. $[(1/3, 1/3, 1/3), (1/3, 1/3, 1/3)]$ is a Nash equilibrium of this game. As we will see later it is the unique one.

Example 17 (4×4 Coordination Game).

The game given by

$$A_1 = \begin{pmatrix} 5 & 4 & 3 & 2 \\ 2 & 5 & 4 & 3 \\ 3 & 2 & 5 & 4 \\ 4 & 3 & 2 & 5 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 5 & 4 & 3 & 2 \\ 2 & 5 & 4 & 3 \\ 3 & 2 & 5 & 4 \\ 4 & 3 & 2 & 5 \end{pmatrix}$$

is an example of an iso-circulant game with shift $k = 4$ as both players' payoff matrices are circulant and for player 2, $a_{n-4}^2 = a_0^2 = 5$ is the largest payoff. The uniform probability distribution over all pure strategies, $[(1/4, 1/4, 1/4, 1/4), (1/4, 1/4, 1/4, 1/4)]$, constitutes a Nash equilibrium. It is, however, not the only one. As we will see later, our results immediately imply that this game has 15 Nash equilibria.

The following two games are examples of counter-circulant games. In both games player 1's payoff matrix is anti-circulant and player 2's payoff matrix is circulant.

Example 18.

$$A_1 = \begin{pmatrix} 4 & 3 & 2 & 1 \\ 3 & 2 & 1 & 4 \\ 2 & 1 & 4 & 3 \\ 1 & 4 & 3 & 2 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 1 & 4 & 3 & 2 \\ 2 & 1 & 4 & 3 \\ 3 & 2 & 1 & 4 \\ 4 & 3 & 2 & 1 \end{pmatrix}$$

This is a counter-circulant game with shift $k = 3$ as for player 2, $a_{n-3}^2 = a_1^2 = 4$ is the largest payoff. The uniform probability distribution over all pure strategies $[(1/4, 1/4, 1/4, 1/4), (1/4, 1/4, 1/4, 1/4)]$ is a Nash equilibrium of this game. As we will see later this game has 3 Nash equilibria.

Example 19.

$$A_1 = \begin{pmatrix} 5 & 4 & 3 & 2 & 1 \\ 4 & 3 & 2 & 1 & 5 \\ 3 & 2 & 1 & 5 & 4 \\ 2 & 1 & 5 & 4 & 3 \\ 1 & 5 & 4 & 3 & 2 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 3 & 2 & 1 & 5 & 4 \\ 4 & 3 & 2 & 1 & 5 \\ 5 & 4 & 3 & 2 & 1 \\ 1 & 5 & 4 & 3 & 2 \\ 2 & 1 & 5 & 4 & 3 \end{pmatrix}$$

This is a counter-circulant game with shift $k = 2$ as for player 2, $a_{n-2}^2 = a_3^2 = 5$ is the largest payoff. The uniform probability distribution over all pure strategies $[(1/5, 1/5, 1/5, 1/5, 1/5), (1/5, 1/5, 1/5, 1/5, 1/5)]$ is a Nash equilibrium of this game. As we will see later this game has 7 Nash equilibria.

3.3 Main Results

In this section we present the main results on the number and the structure of Nash equilibria in circulant games. We start by presenting some preliminary lemmata. All proofs are relegated to the appendix.

3.3.1 Preliminaries

Lemma 2. *Let Γ_n be a circulant game with shift k in which player 1's payoff matrix is anti-circulant and let $d = \gcd(k, n)$.*

- (i) *If Γ_n is iso-circulant, then in any Nash equilibrium (σ_1, σ_2) , for all $s \in S^n$, $\sigma_i(s) = 0$ if and only if $\sigma_i(s + km) = 0$ for all $m = 0, \dots, \frac{n}{d} - 1$, $i = 1, 2$.*
- (ii) *If Γ_n is counter-circulant, then in any Nash equilibrium (σ_1, σ_2) , for all $s \in S^n$, $\sigma_1(s) = 0$ if and only if $\sigma_1(-s + k) = 0$ and $\sigma_2(s) = 0$ if and only if $\sigma_2(-s - k) = 0$*

Given an iso-circulant game Γ_n , we can define an equivalence relation \sim on the set S^n by $s \sim s'$ if and only if $s = s' + mk$ for some $0 \leq m \leq \frac{n}{d} - 1$, where $d = \gcd(n, k)$. Denote the equivalence class of $s \in S^n$ by $I(s)$.

Note that, $s' + m_1k \neq s' + m_2k$ for all $0 \leq m_1 < m_2 \leq \frac{n}{d} - 1$. Hence $I(s) = \{s + mk | 0 \leq m \leq \frac{n}{d} - 1\}$ contains n/d elements and there are d different equivalence classes. Let $I(S^n) = \{I(s) | s \in S^n\}$ be the set of equivalence classes. Suppose player 1's payoff matrix is anti-circulant. By Lemma 2(i) two strategies are equivalent if and only if in any Nash equilibrium either both are simultaneously played with positive probability or both are simultaneously played with zero probability.

For a counter-circulant game let $C_1(s) = \{s, -s+k\}$ and $C_2(s) = \{s, -s-k\}$ for all $s \in S^n$. Note that any class $C_1(s)$ contains at least one and at most two elements. It contains one element if $-s+k \equiv s \pmod{n}$ and two elements if $-s+k \not\equiv s \pmod{n}$. The former occurs if and only if either $2s = k$ or $2s = n+k$. Thus there is a singleton class if and only if either $\frac{k}{2} \in S^n$ or $\frac{(n+k)}{2} \in S^n$, i.e. if either k or $(n+k)$ is an even number. In particular there can be at most two singleton classes. Similarly, any class $C_2(s)$ contains one element if $-s-k \equiv s \pmod{n}$ and two elements if $-s-k \not\equiv s \pmod{n}$. The former occurs if and only if either $2s = n-k$ or $2s = 2n-k$. Thus there is a singleton class if and only if either $n-k$ or $2n-k$ is an even number, which holds if and only if either k or $(n+k)$ is an even number, i.e. if and only if $\frac{k}{2} \in S^n$ or $\frac{(n+k)}{2} \in S^n$. We define $C_i(S^n) := \{C_i(s) | s \in S^n\}$, $i = 1, 2$. Suppose player 1's payoff matrix is anti-circulant. Then, by Lemma 2(ii), $s' \in C_i(s)$ if and only if in any Nash equilibrium either both s and s' are simultaneously played with positive probability or both are simultaneously played with zero probability. It can be shown (Lemma B.3 in the appendix) that the sets $C_i(S^n)$, $i = 1, 2$, form a partition of S^n .

The following lemma covers the connection between the support of a strategy of player i and the best response of player $-i$ against that strategy.

Lemma 3. *Let Γ_n be a circulant game in which player 1's payoff matrix is anti-circulant.*

- (i) *If Γ_n is iso-circulant then if $\sigma_i \in \Sigma_i$ and $I(s) \in I(S^n)$ are such that $\text{supp}(\sigma_i) \cap I(s) = \emptyset$ then $BR_{-i}(\sigma_i) \cap I(-s) = \emptyset$.*
- (ii) *If Γ_n is counter-circulant then if $\text{supp}(\sigma_{-i}) \cap C_{-i}(s) = \emptyset$ for $C_{-i}(s) \in C_{-i}(S^n)$ then $BR_i(\sigma_{-i}) \cap C_i(-s) = \emptyset$.*

3.3.2 The Number of Nash Equilibria

Theorem 1. *Let Γ_n be an iso-circulant game with shift k and let $d = \gcd(k, n)$ denote the greatest common divisor of k and n . Then Γ_n has $2^d - 1$ Nash equilibria.*

Since by definition $k \leq n$, necessarily $\gcd(k, n) \leq n$. It follows that an iso-circulant game can have at most $2^n - 1$ Nash equilibria. Further, an iso-circulant game has a unique Nash equilibrium if and only if $\gcd(k, n) = 1$. Together with Lemma 1, this implies that if $\gcd(k, n) = 1$ then the unique Nash equilibrium is the one where both players choose the uniformly mixed strategy. Some immediate consequences of these results are the following.

Matching Pennies (Example 15) is an iso-circulant game with shift $k = 1$. Hence, $[(1/2, 1/2), (1/2, 1/2)]$ is the unique Nash equilibrium. Rock-Paper-Scissors (Example 16) is an iso-circulant game with shift $k = 1$. Hence, the unique Nash equilibrium is $[(1/3, 1/3, 1/3), (1/3, 1/3, 1/3)]$.

Proposition 3. *Let Γ_n be an iso-circulant game with shift k . Γ_n has n pure strategy Nash equilibria if and only if $k = n$. Further, Γ_n has no pure strategy Nash equilibrium if and only if $k \neq n$.*

By the last proposition an iso-circulant game Γ_n has either 0 or n pure strategy Nash equilibria. The 4×4 coordination game in Example 17 is an iso-circulant game with shift $k = 4$. As $\gcd(4, 4) = 4$, by Theorem 1, this game has $2^4 - 1 = 15$ Nash equilibria. By Proposition 3 four of these are in pure strategies.

Theorem 2. *Let Γ_n be a counter-circulant game with shift k .*

- (i) *If n is odd, then Γ_n has exactly $2^{\frac{n+1}{2}} - 1$ Nash equilibria.*
- (ii) *If both n and k are even, then Γ_n has exactly $2^{\frac{n}{2}+1} - 1$ Nash equilibria.*
- (iii) *If n is even and k is odd, then Γ_n has exactly $2^{\frac{n}{2}} - 1$ Nash equilibria.*

It follows that a counter-circulant game can have at most $2^{\frac{n}{2}+1} - 1$ Nash equilibria. Further, a counter-circulant game has a unique Nash equilibrium

if and only if $n = 2$ and $k = 1$. Example 18 is a counter-circulant game with shift $k = 3$. As n is even and k is odd, by Theorem 2(iii) the game has $2^2 - 1 = 3$ Nash equilibria. Example 19 is a counter-circulant game with shift $k = 2$. As n is odd, by Theorem 2(i) the game has $2^3 - 1 = 7$ Nash equilibria.

Proposition 4. *Let Γ_n be a counter-circulant game with shift k .*

- (i) Γ_n has exactly one pure strategy Nash equilibrium if and only if n is odd.
- (ii) Γ_n has exactly two pure strategy Nash equilibria if and only if both n and k are even.
- (iii) Γ_n has no pure strategy Nash equilibrium if and only if n is even and k is odd.

In Example 18 n is even and k is odd, hence by Proposition 4(iii) none of its three Nash equilibria are in pure strategies. In Example 19 n is odd, hence by Proposition 4(i) one of its seven Nash equilibria is in pure strategies.

It follows from (i) and (ii) in Proposition 4 that the class of counter-circulant games with even shift is a class of games for which a *pure* strategy Nash equilibrium always exists.

3.3.3 The Structure of Nash Equilibria

The next lemma shows that only specific subsets of S^n can arise as the support of a Nash equilibrium strategy of player 1.

Lemma 4. *Let Γ_n be a circulant game in which player 1's payoff matrix is anti-circulant.*

- (i) *If Γ_n is iso-circulant then for any union $U = \bigcup_{j=1}^m I(s^j)$ of elements of $I(S^n)$ there is a unique Nash equilibrium (σ_1, σ_2) such that $\text{supp}(\sigma_1) = U$. Further, for any Nash Equilibrium (σ_1, σ_2) there is a union $U = \bigcup_{j=1}^m I(s^j)$ of elements of $I(S^n)$ such that $\text{supp}(\sigma_1) = U$.*

(ii) If Γ_n is counter-circulant then for any union $U = \bigcup_{j=1}^m C_1(s^j)$ of elements of $C_1(S^n)$ there is a unique Nash equilibrium (σ_1, σ_2) such that $\text{supp}(\sigma_1) = U$. Further, for any Nash Equilibrium (σ_1, σ_2) there is a union $U = \bigcup_{j=1}^m C_1(s^j)$ of elements of $C_1(S^n)$ such that $\text{supp}(\sigma_1) = U$.

By Lemma 4, there exists a straightforward way to characterize the support of all Nash equilibrium strategies for a given circulant game. Moreover, once we know what to look for the weights of the strategies in the support can be easily derived.

Consider first the case of an iso-circulant game with n and k , and let $d = \text{gcd}(n, k)$. We can transform the game so that player 1's payoff matrix is anti-circulant (see Lemma A.1(i) in the appendix). Recall that by Lemma 2(i) the circulant structure of the payoff matrices allows us to define an equivalence relation on the set of pure strategies S^n for each player. For a pure strategy $s \in S^n$, the corresponding equivalence class $I(s) = \{s + mk \mid 0 \leq m \leq \frac{n}{d} - 1\}$ contains n/d elements and there are d different equivalence classes. In any Nash equilibrium all strategies within one equivalence class are either played with strictly positive or with zero probability. It follows from Lemma 4(i) that in any Nash equilibrium the support of either player's strategy is the union of classes in $I(S^n) = \{I(s) \mid s \in S^n\}$ and further that for any such union of classes in $I(S^n)$ there is a unique Nash equilibrium in which player 1's strategy has this union as its support. Further, if the mixed strategy profile (σ_1, σ_2) is a Nash equilibrium with $\text{supp}(\sigma_1) = \bigcup_{j=1}^m I(s^j)$ for some strategies $s^1, \dots, s^m \in S^n$ then by Lemma 3(i) it follows that $\text{supp}(\sigma_2) = \bigcup_{j=1}^m I(-s^j)$. The actual probabilities for each pure strategy of course depend on the actual payoffs, however, the structure of the supports is the same for all iso-circulant games with the same shift and the same number of pure strategies.

Let us revisit the 4×4 Coordination game from Example 17. We can transform this game so that both payoff matrices are anti-circulant (see Table 3.1 in appendix 3.C and Lemma A.1(i) in Appendix 3.A). In this game $n = k = d = 4$ and hence there are four (singleton) classes: $I(0) = \{0\}$, $I(1) = \{1\}$, $I(2) = \{2\}$, and $I(3) = \{3\}$. Each class is part of a (pure strategy) Nash equilibrium in which $\text{supp}(\sigma_1) = I(s)$ and $\text{supp}(\sigma_2) = I(-s)$, and there are four such combinations. For instance, in one Nash equilibrium

player 1 plays the strategy $s = 1$, i.e. chooses support $I(1)$ and player 2 plays $s = 3$, chooses support $I(-1) = I(3)$. Analogously, the three remaining pure strategy Nash equilibria are given by the profiles $(0, 0)$, $(2, 2)$, and $(3, 1)$. Further, each union of two classes is part of a (mixed strategy) Nash equilibrium in which $\text{supp}(\sigma_1) = I(s^1) \cup I(s^2)$ and $\text{supp}(\sigma_2) = I(-s^1) \cup I(-s^2)$. There are six such combinations, e.g., in one Nash equilibrium player 1 puts positive probability only on $I(0)$ and $I(1)$ and player 2 puts positive probability on $I(-0) \cup I(-1) = I(0) \cup I(3)$. The probabilities are easily derived from the corresponding indifference conditions and the Nash equilibrium strategy profile is $[(1/4, 3/4, 0, 0), (3/4, 0, 0, 1/4)]$. Similarly, there are four Nash equilibria in which the support of player 1's (and player 2's) strategy is the union of three classes, e.g., $[(1/4, 1/4, 1/2, 0), (1/2, 0, 1/4, 1/4)]$. Finally, there is one Nash equilibrium where player 1's (and player 2's) strategy put positive probability on all four equivalence classes, i.e. plays a completely mixed strategy: $[(1/4, 1/4, 1/4, 1/4), (1/4, 1/4, 1/4, 1/4)]$.³

Consider now the case of a counter-circulant game with given n and k . We can transform this game so that player 1's payoff matrix is anti-circulant (see Lemma A.1(ii) in the appendix). Recall that by Lemma 2(ii) we can define an equivalence relation on set of pure strategies for each player. For all $s \in S$ let $C_1(s) = \{s, -s + k\}$ denote the corresponding equivalence class of player 1 and $C_2(s) = \{s, -s - k\}$ the one of player 2. Note that any class $C_1(s)$, $C_2(s)$ contains at least one and at most two elements. It follows from Lemma 4(ii) that in any Nash equilibrium the support of player 1's strategy is a union of classes in $C_1(S^n) = \{C_1(s) | s \in S^n\}$ and that for any union of classes in $C_1(S^n)$ there is a Nash equilibrium in which the support of player 1's strategy has this union as its support. Further, if (σ_1, σ_2) is a Nash equilibrium with $\text{supp}(\sigma_1) = \bigcup_{j=1}^m C_1(s^j)$ for some strategies $s^1, \dots, s^m \in S^n$ then by Lemma 3(ii) it follows that $\text{supp}(\sigma_2) = \bigcup_{j=1}^m C_2(-s^j)$.

Let us revisit the game in Example 18. Here, $n = 4$ and $k = 3$. There are two classes for player 1: $C_1(0) = C_1(3) = \{0, 3\}$ and $C_1(1) = C_1(2) = \{1, 2\}$. Correspondingly there are two classes for player 2: $C_2(0) = C_2(1) = \{0, 1\}$ and $C_2(2) = C_2(3) = \{2, 3\}$. There are two Nash equilibria in which the

³Table 3.1 in the appendix contains the remaining Nash equilibrium profiles.

support of player 1's (and player 2's) strategy consists of a single class, e.g. $[(1/4, 0, 0, 3/4), (1/4, 3/4, 0, 0)]$. Further there is one equilibrium in which both players play the completely mixed strategy $[(1/4, 1/4, 1/4, 1/4)]$.⁴

3.4 Generalizations

By our definition there are games that are not circulant games, but can be transformed into one by a simple relabeling of strategies. We chose to exclude those games from our definition for ease of exposition. However, the results presented above also apply for these games.

It is not necessary to insist on each row containing the same entries. All our proofs go through if payoffs are transformed in a way that preserves the order of entries in each row and in each column of the payoff matrices.

Example 20. In the 3×3 game with payoff matrices

$$A_1 = \begin{pmatrix} 3.1 & 1.9 & 0.8 \\ 1.5 & 0.9 & 3.4 \\ 0.5 & 3.2 & 2.1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 0.7 & 2.2 & 3.5 \\ 1.8 & 2.6 & 0.1 \\ 3.0 & 0.5 & 2.8 \end{pmatrix}.$$

the order of payoffs in each row and in each column is the same as in Rock-Paper-Scissors (Example 16). The proof of Theorem 1 can easily be generalized to this case to show that this game has a unique Nash equilibrium. As the sum of payoffs in each row is not constant, however, the unique Nash equilibrium is not the strategy profile in which both players play the uniformly mixed strategies.

In this sense, our results on the number and the structure of Nash equilibria only depend on the order of payoffs in the rows and columns of the payoff matrices.

Our results further generalize to coordination games in which players obtain a strictly positive payoff if and only if they use the same strategy

⁴Table 3.2 in the appendix shows the Nash equilibria and the equivalence classes for the two counter-circulant games we introduced in Example 18 and 19.

and a payoff of 0 otherwise i.e., so-called games of pure coordination. The resulting payoff matrices are of the form

$$A_1 = \begin{pmatrix} a_0 & 0 & 0 & \cdots & 0 \\ 0 & a_1 & 0 & \cdots & 0 \\ 0 & 0 & a_2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & a_{n-1} \end{pmatrix}, \quad A_2 = \begin{pmatrix} a_0 & 0 & 0 & \cdots & 0 \\ 0 & a_1 & 0 & \cdots & 0 \\ 0 & 0 & a_2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & a_{n-1} \end{pmatrix}$$

Proving that such games have $2^n - 1$ Nash equilibria works analogously to the proof of Theorem 1.

3.5 Conclusion

In this paper we introduce and investigate a class of two-player normal-form games we coin circulant games. Such games have a straightforward representation in form of circulant matrices. Each player's payoff matrix is fully characterized by a single row vector, which is rotated to obtain the rest of the matrix. All circulant games have a Nash equilibrium where players randomize between all pure strategies with equal probability (uniformly mixed Nash equilibrium), but might have many other pure and mixed Nash equilibria.

The circulant structure underlying the payoff matrices has interesting implications. First, the best response correspondences induce a partition on each players' set of pure strategies into equivalence classes. In any Nash Equilibrium all strategies within one class are either played with strictly positive or with zero probability. Second, there exists a simple one-to-one correspondence between the players' respective equivalence classes. If some player puts zero probability on one class, the other has one corresponding equivalence class he plays with zero probability. Finally, a single parameter k fully determines the strategy classes and the relation between the players' classes. The parameter itself only depends on the position of the largest payoff in the first row of a player's payoff matrix. For a given circulant game, knowing k and the number of pure strategies n suffices to calculate

the exact number of Nash equilibria and to describe the support of all Nash equilibrium strategies. As an immediate consequence of our main results we establish $2^n - 1$ as the tight upper bound on the number of Nash equilibria in these games.

The class of circulant games contains a large variety of games with cyclical payoff structures including well-known games such as Matching Pennies, Rock-Paper-Scissors or subclasses of coordination and common interest games. We shed new light on the features these games have in common focusing on the circulant structure of their payoff matrices. For example Matching Pennies is the two-strategy variant of Rock-Paper-Scissors. Beyond their zero-sum property the two games belong to the same sub-class circulant games. Both are characterized by $k = 1$ and the only Nash equilibrium is the uniformly mixed one. The common denominator that connects these games is the balanced payoff structure induced by the circulant matrices with a shift of $k = 1$. Moreover, this reinterpretation is robust in the sense that only relative payoffs matter. We can write down many variants of Rock-Paper-Scissors, including asymmetric evaluations of wins or losses and variants that cannot be transformed into zero-sum games. Yet, the balanced structure is preserved and the best players can do is to randomize between all pure strategies with equal probability.

Appendix 3.A: Transformation of Games

Lemma A.1. (i) *Let Γ_n be an iso-circulant game in which both players' payoff matrices are circulant. There is a permutation of row vectors that fixes the first row in both matrices and transforms both players' payoff matrices into anti-circulant matrices.*

(ii) *Let Γ_n be a counter-circulant game in which player 1's payoff matrix is circulant. There is a permutation of row vectors that fixes that first row in both matrices and transforms player 1's payoff matrix into an anti-circulant matrix and player 2's matrix into a circulant matrix.*

Proof. (i) A matrix A is anti-circulant if and only if $A = PC$, where C is a circulant matrix and

$$P = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 & 1 \\ 0 & 0 & \cdots & 0 & 1 & 0 \\ 0 & 0 & \cdots & 1 & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots & \vdots \\ 0 & 1 & \cdots & 0 & 0 & 0 \end{pmatrix}$$

(Davis, 1979, p. 162, Corollary). The matrix P switches rows i and $n + 1 - i$ and fixes the first row. Using this result, we obtain that PA_1 and PA_2 are anti-circulant matrices since both A_1 and A_2 are circulant matrices.

(ii) Using the matrix P defined as in (i), we obtain that PA_1 is anti-circulant (Davis, 1979, p. 162, Corollary). As A_2 is anti-circulant, $A_2 = PC$ for some circulant matrix C (Davis, 1979, p. 162, Corollary). Hence $PA_2 = P(PC)$ and since $P = P^{-1}$ (Davis, 1979, p.28, equ. (2.4.22)), we obtain that PA_2 is a circulant matrix.

□

Appendix 3.B: Proofs of Main Results

We remind the reader that the sum and the difference of strategies in a circulant game as well as multiplications of integers with strategies are read modulo n . Central to the proofs of our main results is Proposition 5 below.

Proposition 5 identifies sufficient conditions under which the number of Nash equilibria of a finite two player normal-form game can be calculated by merely identifying one parameter of the game. Under the hypotheses of Proposition 5, each Nash equilibrium strategy of a player corresponds to one specific combination of elements of a partition of that player's strategy set. Moreover, for each possible combination of elements of the partition there exists exactly one corresponding Nash equilibrium strategy. The parameter necessary to determine the number of Nash equilibria is the cardinality of the partition.

The proof of Theorem 1 (Theorem 2) first establishes that iso-circulant (counter-circulant) games satisfy the hypotheses of Proposition 5. Determining the cardinality of the partitions is then merely a counting exercise.

Proposition 5. *For the two-player normal-form game Γ_n let $\bar{S}_1 = \{[s]_1 | s \in S^n\}$ and $\bar{S}_2 = \{[s]_2 | s \in S^n\}$ be partitions of S^n such that $|\bar{S}_1| = |\bar{S}_2|$. If Γ_n , \bar{S}_1 , and \bar{S}_2 satisfy*

- (a) *for all Nash equilibria (σ_1, σ_2) , and all $s, s' \in S^n$, if $s' \in [s]_i$ then $\sigma_i(s) = 0$ if and only if $\sigma_i(s') = 0$,*
- (b) *for all $\sigma_i \in \Sigma_i$, $i = 1, 2$, $\text{supp}(\sigma_i) \cap [s]_i = \emptyset$ for $[s]_i \in \bar{S}_i$ implies $BR_{-i}(\sigma_i) \cap [-s]_{-i} = \emptyset$,*
- (c) *for all $s \in S^n$, Γ_n has a Nash equilibrium (σ_1, σ_2) with $\text{supp}(\sigma_1) = [s]_1$ and $\text{supp}(\sigma_2) = [-s]_2$,*

then

- (i) *for any $M \subseteq \bar{S}_1$ Γ_n has a unique Nash equilibrium (σ_1, σ_2) with $\text{supp}(\sigma_1) = \bigcup_{[s]_1 \in M} [s]_1$;*
- (ii) *Γ_n has exactly $2^{|\bar{S}_1|} - 1$ Nash equilibria.*

Proof. (i) Given $\emptyset \neq M \subseteq \bar{S}_1$ let $-M := \{[-s]_2 | [s]_1 \in M\} \subseteq$ and let Γ_n^M be the reduced game where player 1's set of strategies is $\bigcup_{[s]_1 \in M} [s]_1$ and player 2's set of strategies is $\bigcup_{[s]_1 \in M} [-s]_2$ (and the payoff functions are restricted accordingly).

Claim A: Let $M' \subseteq M \subseteq \bar{S}_1$ be a nonempty subset of \bar{S}_1 and let $(\sigma_1^{M'}, \sigma_2^{M'})$ be a completely mixed Nash equilibrium of $\Gamma_n^{M'}$. Then (σ_1^M, σ_2^M) defined by $\sigma_1^M(s) = \sigma_1^{M'}(s)$ if $[s]_1 \in M'$ and $\sigma_1^M(s) = 0$ otherwise, and $\sigma_2^M(s) = \sigma_2^{M'}(s)$ if $[s]_2 \in -M'$ and $\sigma_2^M(s) = 0$ otherwise is a Nash equilibrium in Γ_n^M .

Since $(\sigma_1^{M'}, \sigma_2^{M'})$ is a completely mixed Nash equilibrium of $\Gamma_n^{M'}$, all strategies in $\bigcup_{[s]_1 \in M'} [-s]_2$ yield the same payoff for player 2 against $\sigma_1^{M'}$. By hypothesis (b), since $\text{supp}(\sigma_1^{M'}) = \bigcup_{[s]_1 \in M'} [s]_1$, no strategy outside $\bigcup_{[s]_1 \in M'} [-s]_2$ can be a best response for player 2 against $\sigma_1^{M'}$. Analogously all strategies in $\bigcup_{[s]_1 \in M'} [s]_1$ yield the same payoff for player 1 against $\sigma_2^{M'}$, and since $\text{supp}(\sigma_2^{M'}) = -\bigcup_{[s]_1 \in M'} [-s]_2$, no strategy outside $\bigcup_{[s]_1 \in M'} [s]_1$ is a best response for player 1 against $\sigma_2^{M'}$. Hence, by Proposition 1, (σ_1^M, σ_2^M) is a Nash equilibrium in Γ_n^M . This proves the claim.

Claim B: For any $\emptyset \neq M \subseteq \bar{S}_1$, the reduced game Γ_n^M has exactly one completely mixed Nash equilibrium.

Let $\emptyset \neq M \subseteq \bar{S}_1$ be such that $|M| = m$. We will prove the claim by induction over m . Note first, that by hypothesis (b), in any Nash equilibrium (σ_1, σ_2) of Γ_n^M , $\text{supp}(\sigma_1)$ is a union of elements of M .

For $m = 1$, this follows by hypothesis (c). For $m > 1$, by induction hypothesis we obtain that for all $\emptyset \neq M' \subsetneq M$ the reduced game $\Gamma_n^{M'}$ has a unique completely mixed Nash equilibrium. By Claim A, for every $\emptyset \neq M' \subsetneq M$ there is a Nash equilibrium $(\sigma_1^{M'}, \sigma_2^{M'})$ in Γ_n^M with $\text{supp}(\sigma_1^{M'}) = \bigcup_{[s]_1 \in M'} [s]_1$. As by Proposition 2(ii) for any $\emptyset \neq M' \subsetneq M$ there can be at most one Nash equilibrium (σ_1, σ_2) in Γ_n^M with $\text{supp}(\sigma_1) = M'$ we obtain that there is exactly one such Nash equilibrium. This implies that Γ_n^M has at least $2^m - 2$ Nash equilibria.

Suppose there is no completely mixed Nash equilibrium in Γ_n^M . Then Γ_n^M has exactly $2^m - 2$ Nash equilibria. From hypotheses (a) and (b) it follows that

Γ_n is non-degenerate and hence that Γ_n^M is non-degenerate. By Proposition 2(i) Γ_n^M must have an odd number of Nash equilibria, which contradicts the fact that $2^m - 2$ is even. Hence there is at least one completely mixed Nash equilibrium and again because Γ_n^M is non-degenerate by Proposition 2(ii) there is exactly one. This proves the claim.

By Claim B, for $\emptyset \neq M \subseteq \overline{S}_1$, Γ_n^M has exactly one completely mixed Nash equilibrium (σ_1^M, σ_2^M) . By Claim A, this induces a Nash equilibrium (σ_1, σ_2) in Γ_n with $\text{supp}(\sigma_1) = \bigcup_{[s]_1 \in M} [s]_1$. Any Nash equilibrium $(\sigma'_1, \sigma'_2) \neq (\sigma_1, \sigma_2)$ with $\text{supp}(\sigma'_1) = \bigcup_{[s]_1 \in M} [s]_1$ would induce a completely mixed Nash equilibrium in Γ_n^M different from (σ_1^M, σ_2^M) , a contradiction. Hence Γ_n has exactly one Nash equilibrium (σ_1, σ_2) with $\text{supp}(\sigma_1) = \bigcup_{[s]_1 \in M} [s]_1$.

(ii) From (i) it follows that for any $\emptyset \neq M \subseteq \overline{S}_1$ there is a unique Nash equilibrium (σ_1, σ_2) in Γ_n such that $\text{supp}(\sigma_1) = \bigcup_{[s]_1 \in M} [s]_1$. Further, by hypothesis (a), for any Nash equilibrium (σ_1, σ_2) of Γ_n there is $\emptyset \neq M \subseteq \overline{S}_1$ such that $\text{supp}(\sigma_1) = \bigcup_{[s]_1 \in M} [s]_1$. As \overline{S}_1 has $2^{|\overline{S}_1|} - 1$ nonempty subsets, Γ_n has exactly $2^{|\overline{S}_1|} - 1$ Nash equilibria. □

The following lemma is required in the proofs of Lemmata 2 and 3.

Lemma B.1. *Let Γ_n be a circulant game with shift k in which player 1's payoff matrix is anti-circulant.*

- (i) *For all $\sigma_2 \in \Sigma_2$ and all $s \in S^n$ if $\sigma_2(s) = 0$ then $-s \notin BR_1(\sigma_2)$.*
- (ii) *If Γ_n is iso-circulant, then for all $\sigma_1 \in \Sigma_1$ and all $s \in S^n$ if $\sigma_1(s) = 0$ then $(-s - k) \notin BR_2(\sigma_1)$.*
- (iii) *If Γ_n is counter-circulant, then for all $\sigma_1 \in \Sigma_1$ and all $s \in S^n$ if $\sigma_1(s) = 0$ then $(s - k) \notin BR_2(\sigma_1)$.*

Proof. (i) Let $\sigma_2 \in \Sigma_2$ be such that $\sigma_2(s) = 0$ for some $s \in S^n$. Since player 1's payoff matrix is anti-circulant $\pi_1(s|s') = a_{s+s'}^1$. We will show that there exists a strategy for player 1 that yields a strictly higher payoff against σ_2 than strategy $-s$. Let $l := \min\{s < l' \leq s+n-1 | \sigma_2(l') > 0\}$. Since $n > 1$ the set $\{s < l' \leq s+n-1 | \sigma_2(l') > 0\}$ is non-empty and l exists. By construction

of l , $\sigma_2(s) = \cdots = \sigma_2(l-1) = 0$. We claim that $\pi_1(-s|\sigma_2) < \pi_1(-l|\sigma_2)$. To see this, note that

$$\pi_1(-s|\sigma_2) = \sum_{t=l}^{s+n-1} \sigma_2(t) a_{t-s}^1$$

and

$$\pi_1(-l|\sigma_2) = \sum_{t=l}^{s+n-1} \sigma_2(t) a_{t-l}^1.$$

Comparing these payoffs for $t = l$ we obtain that $a_{t-l}^1 = a_0^1 > a_{t-s}^1 = a_{l-s}^1$, where the strict inequality holds by part (ii) of Definition 2. Further, for $l < t \leq s+n-1$ we have $0 \leq t-l < t-s \leq n-1$ and hence that $a_{t-l}^1 \geq a_{t-s}^1$ again by part (ii) of Definition 2. Since by construction of l , $\sigma_2(l) > 0$ we obtain $\pi_1(-s|\sigma_2) < \pi_1(-l|\sigma_2)$ which proves the claim. Hence $-s \notin BR_1(\sigma_2)$.

(ii) Let $\sigma_1 \in \Sigma_1$ and $s \in S^n$ be such that $\sigma_1(s) = 0$. Since player 2's payoff matrix is anti-circulant, $\pi_2(s|s') = a_{s'+s}^2$ for $s, s' \in S$. Since Γ_n is a circulant game, by part (iii) of Definition 2 either $a_{n-k}^2 > a_{n-k+1}^2 \geq \cdots \geq a_{n-1}^2 \geq a_0^2 \geq a_1^2 \geq \cdots \geq a_{n-k-1}^2$ or $a_{n-k}^2 > a_{n-k-1}^2 \geq \cdots \geq a_1^2 \geq a_0^2 \geq a_{n-1}^2 \geq \cdots \geq a_{n-k+1}^2$. We will only prove the result for the former case as the proof for the latter works analogously .

Let $l := \min\{s < l' \leq s+n-1 | \sigma_1(l') > 0\}$ which exists since $\{s < l' \leq s+n-1 | \sigma_1(l') > 0\} \neq \emptyset$. Then $\sigma_1(s) = \cdots = \sigma_1(l-1) = 0$. We claim that $\pi_2(-s-k|\sigma_1) < \pi_2(-l-k|\sigma_1)$. To see this, note that

$$\pi_2(-s-k|\sigma_1) = \sum_{t=l}^{s+n-1} \sigma_1(t) a_{t-s-k}^2$$

and

$$\pi_2(-l-k|\sigma_1) = \sum_{t=l}^{s+n-1} \sigma_1(t) a_{t-l-k}^2.$$

For $t = l$ we have $a_{t-l-k}^2 = a_{n-k}^2 > a_{t-s-k}^2 = a_{l-s-k}^2$, where the strict inequality holds by part (iii) of Definition 2. Further, for $l < t \leq s+n-1$ we

have $a_{t-l-k}^2 \geq a_{t-s-k}^2$ by part (iii) of Definition 2 since $t-l-k < t-s-k$, $-k \leq t-l-k < n-k-1$, and $-k < t-s-k \leq n-k-1$. Since by construction of l , $\sigma_1(l) > 0$ we obtain that $\pi_2(-s-k|\sigma_1) < \pi_2(-l-k|\sigma_1)$ which proves the claim. Hence $(-s-k) \notin BR_2(\sigma_1)$.

(iii) Let $\sigma_1 \in \Sigma_1$ and $s \in S^n$ be such that $\sigma_1(s) = 0$. Since player 2's payoff matrix is circulant, $\pi_2(s|s') = a_{s-s'}^2$ for $s, s' \in S$. Since Γ_n is a circulant game, by definition either $a_{n-k}^2 > a_{n-k+1}^2 \geq \dots \geq a_{n-1}^2 \geq a_0^2 \geq a_1^2 \geq \dots \geq a_{n-k-1}^2$ or $a_{n-k}^2 > a_{n-k-1}^2 \geq \dots \geq a_1^2 \geq a_0^2 \geq a_{n-1}^2 \geq \dots \geq a_{n-k+1}^2$. We will only prove the result for the former case as the proof for the latter works analogously. Let $l := \min\{s < l' \leq s+n-1 | \sigma_1(l') > 0\}$ which exists since $\{s < l' \leq s+n-1 | \sigma_1(l') > 0\} \neq \emptyset$. Then $\sigma_1(s) = \dots = \sigma_1(l-1) = 0$. We claim that $\pi_2(s-k|\sigma_1) < \pi_2(-l-k|\sigma_1)$. To see this, note that

$$\pi_2(s-k|\sigma_2) = \sum_{t=l}^{s+n-1} \sigma_1(t) a_{s-k-t}^2$$

and

$$\pi_2(l-k|\sigma_2) = \sum_{t=l}^{s+n-1} \sigma_1(t) a_{l-k-t}^2.$$

For $t = l$ we have $a_{l-k-t}^2 = a_{n-k}^2 > a_{s-k-t}^2 = a_{s-k-l}^2$. Further, for $l < t \leq s+n-1$ we have $a_{l-k-t}^2 \geq a_{s-k-t}^2$ by part (iii) of the definition of circulant game since $l-k-t > s-k-t$, $-k \geq l-k-t > -n-k+1$, and $-k > s-k-t \geq -n-k+1$. Since by construction of l , $\sigma_1(l) > 0$ we obtain that $\pi_2(s-k|\sigma_1) < \pi_2(l-k|\sigma_1)$ which proves the claim. Hence $(s-k) \notin BR_2(\sigma_1)$. \square

Lemma B.1 allows us to rule out certain strategies as best responses for player i if player $-i$ plays some strategy with zero probability in the case that player 1's payoff matrix is anti-circulant. By (i) if player 2 plays a strategy s with probability 0 then for player 1 strategy $-s$ cannot be a best response. Similarly, (ii) and (iii) state that if in an iso-circulant (counter-circulant) game player 1 places probability 0 on strategy s then $-s-k$ ($s-k$) cannot be a best response for player 2.

We are now ready to prove Lemmata 2 and 3. It follows from Lemma 2(i) and Lemma 3(i) that iso-circulant games satisfy hypotheses (a) and (b) in Proposition 5. Analogously, Lemma 2(ii) and Lemma 3(ii) establish that counter-circulant games fulfill (a) and (b) in Proposition 5.

Proof of Lemma 2. (i) The “if” part is trivial. To see the “only if” part let (σ_1, σ_2) be a Nash equilibrium of Γ_n and let $s \in S^n$ be such that $\sigma_1(s) = 0$. By Lemma B.1(ii), $\sigma_2(-s-k) = 0$ and consequently by Lemma B.1(i) $\sigma_1(s+k) = 0$. Iterating this argument yields $\sigma_1(s+mk) = 0$ for all $m = 0, \dots, \frac{n}{d} - 1$. If $\sigma_2(s) = 0$ the argument works analogously.

(ii) By Lemma B.1(i) and (iii) for any Nash equilibrium (σ_1, σ_2) and any $s \in S^n$ we obtain

$$\sigma_1(s) = 0 \Rightarrow \sigma_2(s-k) = 0 \Rightarrow \sigma_1(-s+k) = 0$$

and

$$\sigma_1(-s+k) = 0 \Rightarrow \sigma_2(-s) = 0 \Rightarrow \sigma_1(s) = 0.$$

Analogously, for player 2, we obtain

$$\sigma_2(s) = 0 \Rightarrow \sigma_1(-s) = 0 \Rightarrow \sigma_2(-s-k) = 0$$

and

$$\sigma_2(-s-k) = 0 \Rightarrow \sigma_1(s+k) = 0 \Rightarrow \sigma_2(s) = 0.$$

□

Proof of Lemma 3. (i) First, let $s \in S^n$ be such that $\text{supp}(\sigma_1) \cap I(s) = \emptyset$. By Lemma B.1(ii), $-s - (m+1)k \notin BR_2(\sigma_1)$ for all $0 \leq m \leq n/d - 1$. As $\{-s - (m+1)k | 0 \leq m \leq n/d - 1\} = I(-s)$ we obtain $BR_2(\sigma_1) \cap I(-s) = \emptyset$.

Next, let $s \in S^n$ be such that $\text{supp}(\sigma_2) \cap I(s) = \emptyset$. By Lemma B.1(i), $-s - mk \notin BR_1(\sigma_2)$ for all $0 \leq m \leq n/d - 1$. As $\{-s - mk | 0 \leq m \leq n/d - 1\} = I(-s)$ we obtain $BR_1(\sigma_2) \cap I(-s) = \emptyset$.

(ii) If $\text{supp}(\sigma_{-i}) \cap C_{-i}(s) = \emptyset$ for $C_{-i}(s) \in C_{-i}(S^n)$, then, since $C_{-i}(s) = \{s, -s + (-1)^{i-1}k\}$, by Lemma B.1(i) and (iii), $-s, s + (-1)^{i-1}k \notin BR_i(\sigma_{-i})$. Hence $BR_i(\sigma_{-i}) \cap C_i(-s) = \emptyset$. □

The following Lemma B.2 establishes that iso-circulant games fulfill hypothesis (c) in Proposition 5 and is used in the proofs of Theorem 1 and Proposition 3.

Lemma B.2. *Let Γ_n be an iso-circulant game in which both players' payoff matrices are anti-circulant. For every $s \in S^n$, there is a Nash equilibrium (σ_1, σ_2) such that $\text{supp}(\sigma_1) = I(s)$ and $\text{supp}(\sigma_2) = I(-s)$.*

Proof. Given $\bar{s} \in S^n$, define $\sigma_1(s) = d/n$ for all $s \in I(\bar{s})$ and $\sigma_2(s) = d/n$ for all $s \in I(-\bar{s})$. By construction $\text{supp}(\sigma_1) = I(\bar{s})$ and $\text{supp}(\sigma_2) = I(-\bar{s})$. By Lemma 3(i), no strategy outside $I(\bar{s})$ can be a best response for player 1 against σ_2 and no strategy outside $I(-\bar{s})$ can be a best response for player 2 against σ_1 . Further, $\pi_1(s|\sigma_2) = \sum_{m=0}^{n/d-1} \frac{d}{n} a_{s+\bar{s}+mk} = \pi_1(s'|\sigma_2)$ for all $s, s' \in I(\bar{s})$ and analogously $\pi_2(s|\sigma_1) = \pi_2(s'|\sigma_1)$ for all $s, s' \in I(-\bar{s})$. Proposition 1 yields that (σ_1, σ_2) is a Nash equilibrium of Γ_n . \square

We are now ready to prove Theorem 1 and Proposition 3.

Proof of Theorem 1. If Γ_n is an iso-circulant game in which both players' payoff matrices are anti-circulant then by Lemma 2(i), Lemma 3(i) and Lemma B.2, Γ_n and $\bar{S}_1 = \bar{S}_2 = I(S^n)$ as defined in section 3.3.1 then satisfy the hypotheses of Proposition 5. As $|I(S^n)| = d$, it follows that Γ_n has $2^d - 1$ Nash equilibria. If Γ_n is an iso-circulant game in which both players' payoff matrices are circulant, there is a permutation of row vectors that transforms both players' payoff matrices into anti-circulant matrices while fixing the first row in both matrices (Lemma A.1(i)). This permutation, which is essentially a relabeling of the players' strategies, does not affect the number of equilibria. Hence, the proof is complete. \square

Proof of Proposition 3. Note first that if both players' payoff matrices are circulant then by Lemma A.1(i) the game can be transformed into a different version of the same game in which both players' payoff matrices are anti-circulant by a permutation of row vectors. Since such a permutation does not affect the number of pure strategy Nash equilibria, we assume wlog that both players' payoff matrices are anti-circulant.

To see the “if” part suppose $k = n$. Then by construction, each class $I(s)$ is a singleton set and there are n disjoint classes. Hence by Lemma B.2, Γ_n has at least n pure strategy Nash equilibria. By Lemma 2(i), in any pure strategy Nash equilibrium (σ_1, σ_2) , $\text{supp}(\sigma_1) = I(s)$ for some $s \in S$ and hence Γ_n has exactly n pure strategy Nash equilibria.

To prove the “only if” part let Γ_n have n pure strategy Nash equilibria and let (s_1, s_2) be one of them. By Lemma 2(i), $I(s_1)$ must be a singleton set. By construction, $I(s_1)$ is a singleton set if and only if $k = n$.

This proves the first part of the theorem.

To see the second part, note that by construction of the classes $I(s)$ is a singleton set if and only if $k = n$ for any $s \in S$. Further by Lemma 2(i) and Lemma B.2, Γ_n has a pure strategy Nash equilibrium if and only if there is a singleton equivalence class $I(s)$. Hence, Γ_n has no pure strategy Nash equilibrium if and only if $k \neq n$. \square

Before we can turn to the proofs of Theorem 2 and Proposition 4 we require a couple more preliminary lemmata. One hypothesis in Proposition 5 requires the sets \bar{S}_1 and \bar{S}_2 to be partitions of the strategy set. While this is true by construction for $I(S^n)$ in the case of iso-circulant games, the following Lemma B.3 shows that the $C_1(S^n)$ and $C_2(S^n)$ form a partition of S^n .

Lemma B.3. *Let Γ_n be a counter-circulant game. For $i = 1, 2$ the set $C_i(S^n)$ is a partition of S^n .*

Proof. We will prove the result for $i = 1$ as the proof for $i = 2$ works analogously. Since $s \in C_1(s)$ for all $s \in S^n$, it follows that $\bigcup_{s \in S^n} C_1(s) = S^n$. If there is $\bar{s} \in C_1(s) \cap C_1(s')$ for some $s, s' \in S^n$, then since $\bar{s} \in C_1(s)$ either $\bar{s} = s$ or $\bar{s} = -s + k$. If $\bar{s} = s$ then $C_i(s) = C_i(\bar{s})$. If $\bar{s} = -s + k$ then $-\bar{s} + k = s - k + k = s$. In any case it follows that $C_1(\bar{s}) = C_1(s)$. Using the same argument one obtains $C_1(\bar{s}) = C_1(s')$ and hence that $C_1(s) = C_1(s')$. \square

The following Lemma B.4 establishes that counter-circulant games fulfill property (c) in Proposition 5.

Lemma B.4. *Let Γ_n be a counter-circulant game in which player 1's payoff matrix is anti-circulant and let $\sigma = (\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$.*

- (i) $C_i(s)$ is a singleton set if and only if $C_{-i}(-s)$ is a singleton set.
- (ii) For every $s \in S^n$, there is a Nash equilibrium (σ_1, σ_2) such that $\text{supp}(\sigma_1) = C_1(s)$ and $\text{supp}(\sigma_2) = C_2(-s)$.

Proof. (i) Suppose that $C_i(s)$ is a singleton. By construction, $s \equiv -s + (-1)^{i-1}k \pmod n$ which is equivalent to $-s \equiv s + (-1)^i k \pmod n$. This holds if and only if $C_{-i}(-s)$ is a singleton.

(ii) Note that this follows from (i) and Lemma 3(ii) if $C_1(s)$ is a singleton set. Hence, suppose that $C_1(s) = \{s, -s + k\}$ contains two elements. Then, by (i), $C_2(-s) = \{-s, s - k\}$ contains two elements and neither $2s = k$ nor $2s = n + k$. Choose $\sigma_1(s)$ as the solution to $xa_{-2s}^2 + (1-x)a_{-k}^2 = xa_{-k}^2 + (1-x)a_{2s-2k}^2$, i.e.

$$\sigma_1^s(s) = \frac{a_{2s-2k}^2 - a_{n-k}^2}{a_{2s-2k}^2 - a_{n-k}^2 + a_{n-2s}^2 - a_{n-k}^2}.$$

By definition a_{n-k}^2 is player 2's largest payoff implying that $a_{2s-2k}^2 - a_{n-k}^2 < 0$ since $2s \neq n + k$ and that $a_{n-2s}^2 - a_{n-k}^2 < 0$ since $2s \neq k$. Hence $\sigma_1(s) \in]0, 1[$.

Choose $\sigma_2^s(-s)$ as the solution to $xa_0^1 + (1-x)a_{2s-k}^1 = xa_{-2s+k}^1 + (1-x)a_0^1$, i.e.

$$\sigma_2^s(-s) = \frac{a_0^1 - a_{2s-k}^1}{a_0^1 - a_{2s-k}^1 + a_0^1 - a_{-2s+k}^1}.$$

By definition a_0^1 is player 1's largest payoff. Hence as $2s \neq k$ $a_0^1 - a_{2s-k}^1 > 0$ and $a_0^1 - a_{-2s+k}^1 > 0$ implying that $\sigma_2(-s) \in]0, 1[$. By Lemma 3(ii) and Proposition 1, (σ_1, σ_2) is a Nash equilibrium. \square

The set $C_1(S^n)$ is a partition of the strategy set for player 1 while $C_2(S^n)$ is a partition of the strategy set for player 2. By Lemma 3(ii) a class $C_1(s)$ of player 1 "corresponds" to a class $C_2(-s)$ of player 2 in the sense that if player 1 puts probability 0 on all strategies in $C_1(s)$ then none of the strategies in $C_2(-s)$ are a best response for player 2 and vice versa. Part (i) of Lemma B.4 states that two corresponding classes contain the same number of elements.

By (ii) for every class $C_1(s)$ there is always a Nash Equilibrium such that player 1's strategy has this class as its support while player 2's strategy has support $C_2(-s)$. The equilibrium constructed to prove (ii) is such that player 1 chooses his strategy (with support $C_1(s)$) such that player 2 is indifferent between all strategies in $C_2(-s)$ (and vice versa). As Γ_n is a non-degenerate game, by Proposition 2(ii) this is the unique equilibrium (σ_1, σ_2) such that $\text{supp}(\sigma_1) = C_1(s)$ and $\text{supp}(\sigma_2) = C_2(-s)$.

We are now ready to prove Theorem 2 and Proposition 4.

Proof of Theorem 2. If Γ_n is a counter-circulant game in which player 1's payoff matrix is anti-circulant and player 2's payoff matrix is circulant then by Lemma B.3, $C_1(S^n)$ and $C_2(S^n)$ as defined in section 3.3.1 are partitions of S^n . Further, by Lemma B.4(i), $|C_1(S^n)| = |C_2(S^n)|$ and by Lemmata 2(ii), 3(ii), and B.4(ii), Γ_n , $\bar{S}_1 = C_1(S^n)$, and $\bar{S}_2 = C_2(S^n)$ satisfy properties (a)-(c) in Proposition 5 and hence Γ_n has $2^{|C_1(S^n)|} - 1$ Nash equilibria.

To prove (i)-(iii) it hence suffices to determine $|C_1(S^n)|$. Note that any class $C_1(s)$ contains either one or two elements. It contains one element if and only if $-s+k \equiv s$ which occurs if and only if either $2s = k$ or $2s = n+k$. Further, there are at most two singleton classes.

(i) If n is odd, then either $n-k$ is odd (if k is even) or $2n-k$ is odd (if k is odd). Hence there is one singleton class in $C_1(S^n)$ and since all other elements of $C_1(S^n)$ contain two elements, $|C_1(S^n)| = (n-1)/2+1 = (n+1)/2$.

(ii) If both n and k are even, then both k and $n+k$ are even and $k/2, (n+k)/2 \in S^n$. Hence there are two singleton classes in $C_1(S^n)$ and since all other elements of $C_1(S^n)$ contain two elements, $|C_1(S^n)| = (n-2)/2+2 = (n+2)/2$.

(iii) If n is even and k is odd, then $n+k$ is odd and hence neither $k/2 \in S^n$ nor $(n+k)/2 \in S^n$. Hence there is no singleton class and hence all elements of $C_1(S^n)$ contain 2 elements, implying that $|C_1(S^n)| = n/2 = n/2$.

If Γ_n is a counter-circulant game in which player 1's payoff matrix is circulant and player 2's payoff matrix is anti-circulant, there is a permutation of row vectors that transforms player 1's payoff matrix into an anti-circulant matrix. Applying the same permutation of row vectors to player 2's payoff matrix yields a different version of the same game in which strategies are

differently labeled and player 1's payoff matrix is anti-circulant and player 2's payoff matrix is circulant (Lemma A.1(ii)). This permutation does not affect the number of Nash equilibria and hence the proof of Theorem 2 is complete. \square

Proof of Proposition 4. Note first that if player 1's payoff matrix is circulant then by Lemma A.1(i) the game can be transformed into a different version of the same game in which player 1's payoff matrix is anti-circulant by a permutation of row vectors. Since such a permutation does not affect the number of pure strategy Nash equilibria, we assume wlog that player 1's payoff matrix is anti-circulant.

(i) By Lemmata 2(ii) and B.4(ii), Γ_n has one pure strategy Nash equilibrium if and only if one of the classes $C_1(s)$ is a singleton set, which by construction happens if and only if n is odd.

(ii) By Lemmata 2(ii) and B.4(ii), Γ_n has two pure strategy Nash equilibria if and only if two of the classes $C_1(s)$ are singleton sets, which by construction happens if and only if both n and k are even.

(iii) By Lemmata 2(ii) and B.4(ii), Γ_n has no pure strategy Nash equilibrium if and only if none of the classes $C_1(s)$ is a singleton set, which by construction happens if and only if n is even and k is odd. \square

Finally, we prove Lemma 4.

Proof of Lemma 4. (i) To see the first part, let $M = \bigcup_{j=1}^m I(s^j)$ be a union of elements of $I(S^n)$. By Lemma 2(i) and Lemma B.2, Γ_n and $\overline{S}_1 = \overline{S}_2 = I(S^n)$ as defined in section 3.3.1 then satisfy the hypotheses of Proposition 5. Hence, there is a unique Nash equilibrium (σ_1, σ_2) with $\text{supp}(\sigma_1) = M$.

To prove the second part, let (σ_1, σ_2) be a Nash equilibrium. By Lemma 2(i), $\text{supp}(\sigma_1)$ is a union of elements in $I(S^n)$.

(ii) To see the first part, let $M = \bigcup_{j=1}^m C_1(s^j)$ be a union of elements of $C_1(S^n)$. By Lemma B.3, $C_1(S^n)$ and $C_2(S^n)$ as defined in section 3.3.1 are partitions of S^n . Further, by Lemma B.4(i), $|C_1(S^n)| = |C_2(S^n)|$ and by Lemma 2(ii), Lemma 3(ii), and B.4(ii), Γ_n , $\overline{S}_1 = C_1(S^n)$, and $\overline{S}_2 = C_2(S^n)$

satisfy properties (a)-(c) in Proposition 5. It follows that there is a unique Nash equilibrium (σ_1, σ_2) with $\text{supp}(\sigma_1) = M$.

To prove the second part, let (σ_1, σ_2) be a Nash equilibrium. By Lemma 2(ii), $\text{supp}(\sigma_1)$ is a union of elements in $C_1(S^n)$. \square

Appendix 3.C: Tables

Table 3.1: Examples of iso-circulant games.

	Matching Pennies	Rock-Paper-Scissors	4×4 Coordination
Matrix Player 1	$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$	$\begin{pmatrix} 3 & 2 & 1 \\ 2 & 1 & 3 \\ 1 & 3 & 2 \end{pmatrix}$	$\begin{pmatrix} 5 & 4 & 3 & 2 \\ 4 & 3 & 2 & 5 \\ 3 & 2 & 5 & 4 \\ 2 & 5 & 4 & 3 \end{pmatrix}$
Matrix Player 2	$\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$	$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \\ 3 & 1 & 2 \end{pmatrix}$	$\begin{pmatrix} 5 & 4 & 3 & 2 \\ 4 & 3 & 2 & 5 \\ 3 & 2 & 5 & 4 \\ 2 & 5 & 4 & 3 \end{pmatrix}$
Shift k	1	1	4
$\gcd(n, k)$	1	1	4
Number of Equilibria	1	1	15
Equivalence Classes	$I(0)=\{0\}$ $I(1)=\{1\}$	$I(0)=\{0\}$ $I(1)=\{1\}$ $I(2)=\{2\}$	$I(0)=\{0\}$ $I(1)=\{1\}$ $I(2)=\{2\}$ $I(3)=\{3\}$
Nash Equilibria			
Pure			$s_1 = 0, s_2 = 0$ $s_1 = 1, s_2 = 3$ $s_1 = 2, s_2 = 2$ $s_1 = 3, s_2 = 1$
Support 2 Classes	$\sigma_1 = (1/2, 1/2)$ $\sigma_2 = (1/2, 1/2)$		$\sigma_1 = (1/4, 3/4, 0, 0)$ $\sigma_2 = (3/4, 0, 0, 1/4)$ $\sigma_1 = (1/2, 0, 1/2, 0)$ $\sigma_2 = (1/2, 0, 1/2, 0)$ $\sigma_1 = (3/4, 0, 0, 1/4)$ $\sigma_2 = (1/4, 3/4, 0, 0)$ $\sigma_1 = (0, 1/4, 3/4, 0)$ $\sigma_2 = (0, 0, 1/4, 3/4)$ $\sigma_1 = (0, 1/2, 0, 1/2)$ $\sigma_2 = (0, 1/2, 0, 1/2)$ $\sigma_1 = (0, 0, 1/4, 3/4)$ $\sigma_2 = (0, 1/4, 3/4, 0)$
Support 3 Classes		$\sigma_1 = (1/3, 1/3, 1/3)$ $\sigma_2 = (1/2, 1/3, 1/3)$	$\sigma_1 = (1/4, 1/4, 1/2, 0)$ $\sigma_2 = (1/2, 0, 1/4, 1/4)$ $\sigma_1 = (1/4, 1/2, 0, 1/4)$ $\sigma_2 = (1/4, 1/2, 0, 1/4)$ $\sigma_1 = (1/2, 0, 1/4, 1/4)$ $\sigma_2 = (1/4, 1/4, 1/2, 0)$ $\sigma_1 = (0, 1/4, 1/4, 1/2)$ $\sigma_2 = (0, 1/4, 1/4, 1/2)$
Support 4 Classes			$\sigma_1 = (1/4, 1/4, 1/4, 1/4)$ $\sigma_2 = (1/4, 1/4, 1/4, 1/4)$

Table 3.2: Examples of counter-circulant games.

	Example 4	Example 5
Matrix Player 1	$\begin{pmatrix} 4 & 3 & 2 & 1 \\ 3 & 2 & 1 & 4 \\ 2 & 1 & 4 & 3 \\ 1 & 4 & 3 & 2 \end{pmatrix}$	$\begin{pmatrix} 5 & 4 & 3 & 2 & 1 \\ 4 & 3 & 2 & 1 & 5 \\ 3 & 2 & 1 & 5 & 4 \\ 2 & 1 & 5 & 4 & 3 \\ 1 & 5 & 4 & 3 & 2 \\ 3 & 2 & 1 & 5 & 4 \\ 4 & 3 & 2 & 1 & 5 \\ 5 & 4 & 3 & 2 & 1 \\ 1 & 5 & 4 & 3 & 2 \\ 2 & 1 & 5 & 4 & 3 \end{pmatrix}$
Matrix Player 2	$\begin{pmatrix} 1 & 4 & 3 & 2 \\ 2 & 1 & 4 & 3 \\ 3 & 2 & 1 & 4 \\ 4 & 3 & 2 & 1 \end{pmatrix}$	
Shift k	3	2
n	4	5
Number of Equilibria	3	7
Equivalence Classes		
Player 1		
Singleton		$C_1(1) = \{1\}$
2 Elements	$C_1(0) = C_1(3) = \{0, 3\}$ $C_1(1) = C_1(2) = \{1, 2\}$	$C_1(0) = C_1(2) = \{0, 2\}$ $C_1(3) = C_1(4) = \{3, 4\}$
Player 2		
Singleton		$C_2(4) = \{4\}$
2 Elements	$C_2(0) = C_2(1) = \{0, 1\}$ $C_2(2) = C_2(3) = \{2, 3\}$	$C_2(0) = C_2(3) = \{0, 3\}$ $C_2(1) = C_2(2) = \{0, 3\}$
Nash Equilibria		
Pure		$s_1 = 1, s_2 = 4$
Support 1 Class mixed	$\sigma_1 = (1/4, 0, 0, 3/4)$ $\sigma_2 = (1/4, 3/4, 0, 0)$ $\sigma_1 = (0, 3/4, 1/4, 0)$ $\sigma_2 = (0, 0, 1/4, 3/4)$	$\sigma_1 = (3/5, 0, 2/5, 0, 0)$ $\sigma_2 = (3/5, 0, 0, 2/5, 0)$ $\sigma_1 = (0, 0, 0, 4/5, 1/5)$ $\sigma_2 = (0, 1/5, 4/5, 0, 0)$
Support 2 Classes	$\sigma_1 = (1/4, 1/4, 1/4, 1/4)$ $\sigma_2 = (1/4, 1/4, 1/4, 1/4)$	$\sigma_1 = (3/5, 1/5, 1/5, 0, 0)$ $\sigma_2 = (3/5, 0, 0, 1/5, 1/5)$ $\sigma_1 = (1/5, 0, 2/5, 1/5, 1/5)$ $\sigma_2 = (1/5, 1/5, 1/5, 2/5, 0)$ $\sigma_1 = (0, 2/5, 0, 2/5, 1/5)$ $\sigma_2 = (0, 1/5, 2/5, 0, 2/5)$
Support 3 Classes		$\sigma_1 = (1/5, 1/5, 1/5, 1/5, 1/5)$ $\sigma_2 = (1/5, 1/5, 1/5, 1/5, 1/5)$

References Chapter 3

- ALÓS-FERRER, C., AND C. KUZMICS (2013): “Hidden symmetries and focal points,” *Journal of Economic Theory*, 148(1), 226–258.
- BAHEL, E. (2012): “Rock-paper-scissors and cycle-based games,” *Economics Letters*, 115(3), 401–403.
- BAHEL, E., AND H. HALLER (2013): “Cycles with undistinguished actions and extended Rock-Paper-Scissors games,” *Economics Letters*, 120(3), 588–591.
- CASAJUS, A. (2000): “Focal Points in Framed Strategic Forms,” *Games and Economic Behavior*, 32(3), 263–291.
- DAVIS, P. J. (1979): *Circulant Matrices*. Wiley, New York.
- DIEKMANN, O., AND S. A. VAN GILS (2009): “On the Cyclic Replicator Equation and the Dynamics of Semelparous Populations,” *SIAM Journal on Applied Dynamical Systems*, 8(3), 1160–1189.
- DUERSCH, P., J. OECHSSLER, AND B. C. SCHIPPER (2012): “Pure strategy equilibria in symmetric two-player zero-sum games,” *International Journal of Game Theory*, 41(3), 553–564.
- HOFBAUER, J., P. SCHUSTER, K. SIGMUND, AND R. WOLFF (1980): “Dynamical Systems Under Constant Organization II: Homogeneous Growth Functions of Degree $p=2$,” *SIAM Journal on Applied Mathematics*, 38(2), 282–304.
- HOFBAUER, J., AND K. SIGMUND (1998): *Evolutionary Games and Population Dynamics*. Cambridge University Press.
- JANSSEN, M. C. (2001): “Rationalizing Focal Points,” *Theory and Decision*, 50(3), 119–148.
- KEIDING, H. (1998): “On the structure of the set of Nash equilibria of weakly nondegenerate bimatrix games,” *Annals of Operations Research*, 84(0), 231–238.

NASH, J. F. (1951): “Non-cooperative Games,” *Annals of Mathematics*, 54(2), 286–295.

QUINT, T., AND M. SHUBIK (1997): “A theorem on the number of Nash equilibria in a bimatrix game,” *International Journal of Game Theory*, 26(3), 353–359.

QUINT, T., AND M. SHUBIK (2002): “A bound on the number of Nash equilibria in a coordination game,” *Economics Letters*, 77(3), 323–327.

SANDHOLM, W. (2010): *Population Games and Evolutionary Dynamics*, Economic Learning and Social Evolution. MIT Press.

SHAPLEY, L. (1974): “A note on the Lemke-Howson algorithm,” *Mathematical Programming Study*, 1, 175–189.

SINERVO, B., AND C. M. LIVELY (1996): “The rock-paper-scissors game and the evolution of alternative male strategies,” *Nature*, 380(6571), 240–243.

VON STENGEL, B. (1997): “New lower bounds for the number of equilibria in bimatrix games,” *Technical Report 264, Dept. of Computer Science, ETH Zürich*.

——— (1999): “New maximal numbers of equilibria in bimatrix games,” *Discrete & Computational Geometry*, 21(4), 557–568.

CHAPTER 4

PREFERENCE REVERSALS: TIME AND AGAIN

4.1 Introduction

The concept of preference is of fundamental importance for decision theory and economic analysis. Yet, preferences are not a primitive but a derived object which structures choices as long as they exhibit some basic consistency, e.g. in the form of the weak axiom of revealed preference. If choices are consistent, a number of elementary predictions can be derived, which form the basis for decision theory, microeconomics, consumer research, and judgment and decision making. One such prediction is that choices should agree with valuations: if a decision maker chooses one option over another, he should value the former more than the latter.

This common-sense prediction is at odds with observed decisions under risk. The preference reversal phenomenon, first documented in psychology by Slovic and Lichtenstein (1968) and Lindman (1971), describes a situation in which participants are asked to state monetary valuations for a series of lotteries (usually through minimum selling prices), and separately choose from pairs of those lotteries. The pairs consist of a *P-bet*, which has a high probability of paying a moderate amount of money, and a *\$-bet*, which has a low probability of paying a high amount of money. A *preference reversal* occurs if either the *P-bet* is chosen from a pair in which the *\$-bet* is priced higher or the *\$-bet* is chosen from a pair in which the *P-bet* is priced higher. The preference reversal phenomenon is characterized by a high rate of reversals of the first type (between 40 and 80 percent in most experiments), which are called *predicted reversals*. Reversals of the second type, termed *unpredicted*, are less frequent (between 5 and 30 percent). The asymmetry between both types of reversals is especially problematic, for, if reversals were due to e.g. participants' errors, one should expect similar numbers of both types. In

other words, while one could explain away unpredicted reversals as noisy observations, predicted reversals remain a serious challenge to basic economic analysis.

It is no surprise that preference reversals have received a great deal of attention in the last half century. After the first replication in economics by Grether and Plott (1979), a large number of experimental and theoretical studies has shown that the phenomenon is extremely stable. It has been replicated in various ways using hypothetical and real payments, different payment schemes, and different elicitation methods for lottery prices (for a survey, see e.g. Seidl, 2002). Preference reversals of this particular form have been documented beyond lottery choice, e.g. in the field of health utility measurements (Stalmeier, Wakker, and Bezembinder, 1997; Bleichrodt and Pinto Prades, 1994; Oliver, 2013). They have also been shown to be relevant for decision making under ambiguity (Maafi, 2011; Trautmann, Vieider, and Wakker, 2011; Ball, Bardsley, and Ormerod, 2012). Furthermore, other forms of inconsistencies between different preference elicitation methods have been established in the literature, including reversals between pricing and rating (Schkade and Johnson, 1989) as well as discrepancies between certainty and probability equivalents (Hershey and Schoemaker, 1985; Johnson and Schkade, 1989; Delquié, 1993). In addition to their conceptual importance for decision analysis, these phenomena are of great relevance for applied economics, since they cast doubts on the validity of e.g. consumer valuations, and, accordingly, on demand estimations and policy decisions based on those valuations.

The present research provides new evidence on the determinants of preference reversals. We propose a simple, process-based model which predicts the observed pattern of reversals. Specifically, we disentangle the causes behind the *existence* of reversals and their *asymmetry*, i.e. the predominance of predicted preference reversals. The key determinant behind the existence of reversals of both types is the presence of noise in the evaluation phase, or, in other words, imprecise preferences (Schmidt and Hey, 2004; Butler and Loomes, 2007). The asymmetry of reversals, on the other hand, is caused by an overpricing phenomenon due to anchoring of evaluations on the largest

monetary outcomes of a lottery (Tversky, Sattath, and Slovic, 1988; Tversky, Slovic, and Kahneman, 1990). This phenomenon is itself a consequence of the cardinal/monetary framing of the evaluation phase.

Received evidence on preference reversals could potentially be explained by a number of alternative, “as if” models. Our model, however, delivers additional, testable predictions on *decision times*. In particular, choices associated to reversals of either type are predicted to be slower than corresponding non-reversals. Measuring decision times hence allows us to put our model to a more stringent test than if we relied on choice data only, and we consequently do so in two experiments.¹

Our first experiment confirmed the predictions of the model, both for choices and decision times. We established the basic effects using different payment methods to incentivize pricing tasks. Specifically, we employed the BDM procedure (Becker, DeGroot, and Marschak, 1964) and an ordinal payment scheme (Goldstein and Einhorn, 1987; Tversky, Slovic, and Kahneman, 1990; Cubitt, Munro, and Starmer, 2004). The aim of our second experiment was to disentangle the two causes of preference reversals. To do so, we set out to eliminate the overpricing phenomenon by moving away from cardinal elicitation tasks. Instead, we employed two different ranking methods (plus a control BDM replication), one with a price framing, and one where we carefully removed all references to prices. In terms of our model, eliminating overpricing in the lottery evaluation phase should reduce the occurrence of predicted reversals should be reduced. However, the basic predictions for decision times remain unaffected as they arise from the assumption of noise in the evaluation phase only. As hypothesized, predicted reversals were greatly reduced, but decisions times associated with reversals remained significantly

¹The measurement of decision times or response times is a standard tool in psychology (see, e.g., Bargh and Chartrand, 2000). To our knowledge, the first studies employing them in economics were those of Wilcox (1993, 1994), who related them to decision costs in the context of risky choice. Decision times were also used by Moffatt (2005) relying on risky-choice data from Hey (2001). More recently, Piovesan and Wengström (2009) measured response times in a dictator game. Rubinstein (2007) advocated the measurement of decision times in large-scale, web-based experiments to better understand the process of reasoning behind economic decisions. Achtziger and Alós-Ferrer (2013) measured response times within a Bayesian-updating paradigm in order to study intuitive decision making in economic contexts.

longer than those associated with non-reversals.

Our research also delivers additional theoretical and methodological insights. A first, interesting prediction was unexpected before the development of the model. On the basis of our assumptions, we are able to prove that decisions where the riskier $\$$ -bet is chosen without giving rise to a reversal should be slower than those non-reversals where the P -bet is chosen. This nontrivial prediction arises as a consequence of the conjunction of imprecise preferences and the overpricing phenomenon, and hence was predicted for (and observed in) the first experiment but not for the second. A further, striking observation was that choices in the treatment with unframed ranking-based evaluations were much faster than those in other treatments, in spite of the fact that choice phases were identical across treatments. This fact has a simple process-based explanation within our model. Last, our design specifically allowed comparing the number of preference reversals occurring when prices are elicited before the choice phase to the number of preference reversals occurring when prices are elicited after the choice phase. This comparison was motivated by evidence from psychology (see Section 4.2.3 below) indicating that choices might sharpen and even modify previously imprecise preferences. In agreement with this literature, we show that ordering effects, although small, are present in the measurement of reversals.

The remainder of the paper is organized as follows. Section 4.2 spells out our model and derives its predictions and corresponding experimental hypotheses. Sections 4.3 and 4.4 describe the first and second experiments and their results, respectively. Section 4.5 concludes.

4.2 A Simple Model of Preference Reversals and Decision Times

In this section we present our formal model, which is meant to be as simple as possible. We first state and discuss the underlying assumptions, and then derive a number of predictions concerning preference reversals and the associated decision times.

The building blocks of our model are grounded on received evidence from the literature on preference reversals. First, evidence by Schmidt and Hey (2004) suggested that part of the preference reversal phenomenon might be due to pricing errors, while choice errors play a minor role. Butler and Loomes (2007) found that subjects in preference reversal experiments exhibit imprecise monetary valuations of lotteries.² Our model incorporates these observations by assuming a noisy evaluation phase, in comparison to a relatively noise-free choice phase. Second, we rely on the *compatibility hypothesis* proposed by Tversky, Sattath, and Slovic (1988) and further investigated by Tversky, Slovic, and Kahneman (1990), according to which attributes that naturally map onto the evaluation scale are given predominant weight in the evaluation phase. Since the evaluation scale usually refers to prices, the monetary outcomes of the lotteries might anchor valuations, giving rise to an overpricing of the \$-bet, where a large monetary outcome is salient.³

4.2.1 Model and Rationale

We consider a choice between a P -bet and a \$-bet and the pricing decisions for both bets. Let u_P and $u_\$$ denote the “true” utilities of the P -bet and the \$-bet, respectively. Denote by CE_P and $CE_\$$ the elicited certainty equivalents of the P -bet and the \$-bet, respectively.

Relying on evidence by Schmidt and Hey (2004) and Butler and Loomes (2007), we assume that the price elicitation phase is noisier than the choice phase. This is formalized in two parts. The first assumption states that the pricing of lotteries is a noisy process.

²See Blavatsky (2009) for a formal model focused on those findings.

³Tversky, Slovic, and Kahneman (1990) used a design with additional choices between the bets and cash amounts and showed that at least part of the predicted reversals arise because of an overpricing of \$-bets. Tversky, Sattath, and Slovic (1988) also proposed the *prominence hypothesis*, which assumes a bias in the choice stage rather than in the evaluation stage (see also Fischer, Carmon, Ariely, and Zauberman, 1999). Cubitt, Munro, and Starmer (2004) investigated a number of alternative hypotheses including prominence and compatibility and dismissed each of them in isolation, concluding that a combination of hypotheses would be a more reasonable explanation of their findings.

Assumption 1. $CE_P = u_P + \zeta_P$ and $CE_\$ = u_\$ + \zeta_\$$, where ζ_P and $\zeta_\$$ are independent error terms with everywhere positive density functions.⁴

In contrast, the choice phase should be comparatively noise-free. For simplicity, the second assumption postulates that choices follow the underlying utilities. Write $c(P, \$) = P$ if the P -bet was chosen in the choice task and $c(P, \$) = \$$ if the $\$$ -bet was chosen.

Assumption 2. $c(P, \$) = P$ whenever $u_P > u_\$$ and $c(P, \$) = \$$ whenever $u_P < u_\$$.

The main element of our model relies on the compatibility hypothesis (Tversky, Sattath, and Slovic, 1988; Tversky, Slovic, and Kahneman, 1990). It implies that, when pricing lotteries, it is likely that subjects focus their attention on the salient monetary outcomes. Since the $\$$ -bet yields a large outcome with moderate probability and the P -bet pays a moderate outcome with high probability, subjects will tend to state a higher price for the $\$$ -bet. This overpricing phenomenon can be captured by simply assuming a strictly positive mean for the error term associated with the valuation of the $\$$ -bet.

Assumption 3. *There is a tendency to overprice the $\$$ -bet, i.e. $E[\zeta_\$] = K > 0$ but $E[\zeta_P] = 0$. Further, the densities of ζ_P and $\zeta_\$$ are symmetric around the means and unimodal.*⁵

It is a well-established fact that decision times reflect preferences in the sense that hard choices, where the decision maker is close to being indifferent, results in longer decision times than easy choices, where one option is clearly better (Wilcox, 1993; Shultz, Léveillé, and Lepper, 1999; Moffatt, 2005; Chabris, Laibson, Morris, Schuldt, and Taubinsky, 2009; Sharot, DeMartino, and Dolan, 2009; Alós-Ferrer, Granić, Shi, and Wagner, 2012). To model this effect in a simple way, we postulate that the choice time DT_C only

⁴The second part of the assumption is for technical convenience. The analysis goes through, with more cumbersome proofs, if the error terms have bounded support.

⁵A density function is unimodal with mode m if it is nondecreasing for all $x < m$ and nonincreasing for all $x > m$. For example, normally distributed error terms fulfill our assumptions.

depends on the utility difference $|u_P - u_S|$. To avoid unnecessarily complicating the model, we make the simplifying assumption that decisions are of two kinds, easy and hard. Easy decisions correspond to utility pairs (u_P, u_S) such that $|u_P - u_S| \geq \delta$ for some $\delta > 0$, while utility pairs (u_P, u_S) with $|u_P - u_S| < \delta$ lead to hard decisions. Denote by $T_E = E[DT_C | |u_P - u_S| \geq \delta]$ and $T_H = E[DT_C | |u_P - u_S| < \delta]$ the expected choice times for easy and hard decisions, respectively. The next assumption captures the idea that choice decisions in which a subject is close to indifference between two items are harder than “obvious” choices.

Assumption 4. *Hard choices take longer than easy choices, i.e. $T_H > T_E$.*

Our assumptions are meant to reflect the basic principles involved in preference reversal experiments without unnecessarily complicating the exposition and the analysis. Of course, one could postulate more involved formulations, as e.g. a continuously monotonic relation between choice times and closeness to indifference. The next section shows that the simple versions postulated above are enough to provide testable hypotheses.

4.2.2 Predictions

In preference reversal experiments, results refer to a relatively large number of evaluation and choice decisions. Systematic biases are avoided, e.g. by offering choices between lotteries of similar expected values, or counterbalancing the difference in expected values across pairs. Hence, to obtain experimental hypotheses, it is reasonable to treat the utilities u_S and u_P as random variables. Specifically, we assume that the utilities of the lotteries in an experiment are drawn from i.i.d. continuous random variables with some fixed distribution. Since, in our model, both choices and decision times are assumed to depend on utility differences only, the analysis relies on the distribution of $u_P - u_S$. We assume that this distribution has an everywhere positive density h .⁶

⁶Since u_P and u_S are i.i.d, $u_P - u_S$ and $u_S - u_P$ have the same distribution. If the distribution of u_P and u_S has density v then $h = (v * v^-)$, where $v^-(s) = v(-s)$ for all s and the symbol $*$ denotes the convolution operator.

Our model makes four predictions which can be experimentally tested. The first one concerns a well-established observation in the literature, namely that predicted reversals are more frequent than unpredicted ones.

Proposition 1. *Under Assumptions 1, 2, and 3,*

- (i) *there are more predicted than unpredicted preference reversals, i.e. $Pr(CE_{\$} > CE_P, c(P, \$) = P) > Pr(CE_P > CE_{\$}, c(P, \$) = \$)$;*
- (ii) *and the reversal rate is higher for predicted preference reversals than for unpredicted preference reversals, i.e. $Pr(CE_{\$} > CE_P | c(P, \$) = P) > Pr(CE_P > CE_{\$} | c(P, \$) = \$)$.*

The intuition for this result is straightforward. Both kinds of reversals result from noise in the evaluation phase shifting the evaluations of the lotteries in opposite directions. A reversal occurs when, due to noisy realizations, the evaluation ranking is reversed with respect to the one derived from utilities. The overpricing phenomenon helps produce predicted reversals: initially, the \$-bet is ranked lower than the P -bet ($u_{\$} < u_P$), but overpricing tends to shift the valuation of the \$-bet higher than that of the P -bet. Overpricing, however, makes unpredicted reversals harder: the \$-bet is initially ranked higher and overpricing tends to increase its evaluation with respect to the P -bet even more.

We can reformulate the predictions arising from the last proposition straight away as experimental hypotheses.

H1a. The average number of predicted preference reversals per subject is larger than the average number of unpredicted reversals.

H1b. The average rate of predicted reversals (i.e. percentage of reversals over all P -choices) per subject is larger than the average rate of unpredicted reversals (i.e. percentage of reversals over all \$-choices).

These predictions fit received evidence in the literature on preference reversals, and are hence a first validation of the model. We will, of course, also test them with our own data. The added value of the model, however, is given by the following, novel predictions, which concern decision times

in the choice task. The first refers to decision times in “conflict situations” as compared to those in “non-conflict situations”, i.e. for choices leading to preference reversals vs. choices not leading to preference reversals.

Proposition 2. *Let DT_C denote the decision time in the choice phase. Under Assumptions 1, 2, and 4,*

- (i) *the decision time for a P -bet leading to a preference reversal is longer than the decision time for a P -bet that does not lead to a preference reversal, i.e. $E[DT_C | CE_{\$} > CE_P, c(P, \$) = P] > E[DT_C | CE_P > CE_{\$}, c(P, \$) = P]$;*
- (ii) *and the decision time for a $\$$ -bet leading to a preference reversal is longer than the decision time for a $\$$ -bet that does not lead to a preference reversal, i.e. $E[DT_C | CE_P > CE_{\$}, c(P, \$) = \$] > E[DT_C | CE_{\$} > CE_P, c(P, \$) = \$]$.*

The intuition for this result is again simple. Since the origin of reversals lies in the noise arising in the evaluation process, it is clear that reversals are more likely when utilities were close, and hence errors in the evaluation phase are more likely to reverse the order of the lotteries. Decisions where utilities are close are comparatively harder and hence take longer. In other words, reversals are more likely to involve hard choices than non-reversals, which leads to longer decision times.

This proposition translates into the following experimental hypotheses.

H2a. The average decision time for predicted preference reversals is longer than the average decision time for comparable non-reversals (i.e. non-reversals where the P -bet is chosen).

H2b. The average decision time for unpredicted preference reversals is longer than the average decision time for comparable non-reversals (i.e. non-reversals where the $\$$ -bet is chosen).

The next prediction is orthogonal to preference reversals. At the same time, it represents an *a priori* unexpected feature of the model and is hence especially valuable for its validation. It concerns decision times when the

\$-bet was chosen given that it was priced higher compared to decision times when the P -bet was chosen given that it was priced higher.

Proposition 3. *Under Assumptions 1, 2, 3, and 4, the decision time for a \$-bet that does not lead to a preference reversal is longer than the decision time for a P -bet that does not lead to a preference reversal, i.e. $E[DT_C | CE_{\$} > CE_P, c(P, \$) = \$] > E[DT_C | CE_P > CE_{\$}, c(P, \$) = P]$.*

This result seems less intuitive. On the one hand, under overpricing, it is more likely that \$-bets will be priced higher than P -bets than the other way around. Hence, the probability that a hard (slow) \$-bet-choice will result in a non-reversal is larger than the probability that a hard P -bet-choice will result in a non-reversal. On the other hand, an easy (fast) \$-bet-choice is also more likely to result in a non-reversal than an easy P -bet-choice. The reason for this is that in the first case $u_{\$} > u_p$ and overpricing pushes the evaluations further apart, while in the second case $u_{\$} < u_p$ and overpricing pushes the evaluations together. Proposition 3 holds because the relative likelihood for a hard choice to lead to a non-reversal compared to the likelihood for an easy choice to lead to a non-reversal is larger for \$-bets than for P -bets.⁷

This result leads to our next experimental prediction.

H3. The average decision time for non-reversals where the \$-bet is chosen is longer than the average decision time for non-reversals where the P -bet is chosen.

4.2.3 Order Effects and Preference Reappraisal

Preference-reversal experiments include a pricing/evaluation phase and a choice phase. Up to date, the literature has remained silent on order effects, i.e. on whether there is any difference between experiments where the choice phase precedes the evaluation phase, and experiments where the order of the tasks is the opposite. Preference reversals have been established in experiments using either of the two possible orderings.

⁷In particular, the arguments in the proof of this result hold only for non-reversals. No analogous version of Proposition 3 for reversals can be established.

We argue, however, that order effects need to be taken into account. The reason is that, as discussed above, imprecise preferences have been identified as one of the factors driving preference reversals. If preferences are imprecise, a large literature in psychology indicates that they might become more precise, or be generally altered, by the mere act of making choices. In the classical Free-Choice Paradigm (Brehm, 1956), subjects first face a rating (ranking) task, then a choice task, and finally another rating (ranking) task identical to the first one. The chosen options are usually evaluated more positively in the second rating (ranking) task while the options that were not chosen tend to be evaluated more negatively.⁸ According to Cognitive Dissonance Theory (Festinger, 1957), this happens because in the reevaluation phase subjects attempt to reduce the tension between the negative aspects of a chosen option and the positive aspects of an option that was not chosen. Self-Perception Theory (Bem, 1967), on the other hand, attributes this phenomenon to the fact that subjects learn their preferences better by making choices and hence ratings (rankings) in the second phase more closely resemble the “true preferences”. This raises the question of whether preference reversals are affected by the order of valuation and choices. More precisely, if preference reappraisal occurs during the choice phase, there should be fewer preference reversals if the valuation task follows the choice task. This yields an additional hypothesis.

H4. Preference reversals are reduced if the valuation task follows the choice task, compared to the opposite ordering.

More generally, and in view of the discussion above, our expectation was that effects would in general be more clear when considering post-choice evaluations than when relying on pre-choice ones. For example, if one relies on self-perception theory, classifying choices as reversals or non-reversals on the basis of pre-choice evaluations will result in some false classifications, effectively adding more noise to all measurements. However, since preference reversals have been established in the literature using both possible task

⁸Although this classical task has recently been shown to be affected by statistical biases, improved versions of the task have meanwhile re-established the basic effect. See e.g. Alós-Ferrer, Granić, Shi, and Wagner (2012).

orderings, we expected order effects to be subtle. The difference should be more clear for unpredicted reversals, because, if those are purely due to noise, any reduction of noise in the evaluation task should eliminate at least part of them.

Finally, it should be noted that there is evidence from fMRI (Jarcho, Berkman, and Lieberman, 2011) as well as response time studies (Alós-Ferrer, Granić, Shi, and Wagner, 2012) indicating that preference reappraisal and process conflict resolution in the Free-Choice Paradigm occur during the choice phase. Incorporating this additional observation into our model would not affect our predictions, as discussed in the next subsection.

4.2.4 Process Conflict and Decision Times

In psychological terms, the compatibility hypothesis suggests that several decision processes might be at work when choosing from a pair of lotteries. Overpricing might result from a process focusing on monetary outcomes only, which competes with a more global decision process that evaluates lotteries by taking both outcomes and winning probabilities into account.

Taking the process view into account is important because this view delivers standard predictions regarding process data (specifically, decision times). According to dual-process models from psychology (Schneider and Shiffrin, 1977; Strack and Deutsch, 2004; Rottenstreich, Sood, and Brenner, 2007; Evans, 2008; Weber and Johnson, 2009; Alós-Ferrer and Strack, 2013) decision processes can be either automatic and fast, corresponding to quick heuristics, or controlled and slow. In our interpretation, overpricing should result from an automatic (impulsive) process, while a more global valuation should be associated with a more cognitive, controlled process. The quick “look-at-monetary-outcomes” process is more prevalent in the pricing task and causes an overpricing of the \$-bets. We hypothesize that this process is also active in the choice task, but there it is often inhibited, which leads to a choice causing a preference reversal.

A basic prediction of dual-process models is that inhibiting automatic processes costs time and cognitive resources. More generally, conflict de-

tection and resolution is time consuming, that is, decision times are longer when several processes conflict than when processes are aligned. If reversals result from an automatic process which affects the pricing of lotteries but is inhibited in the choice phase, preference reversals should be associated with longer decision times in the choice phase.

This observation can be incorporated into our model by postulating that decision times in the choice task, DT , are noisy and consist of two components, choice time DT_C and conflict resolution time DT_R , i.e. $DT = DT_C + DT_R$. The next assumption reflects the considerations above and concerns conflict resolution time only.

Assumption 5. *Conflict resolution is time-consuming, i.e. conflict resolution time is longer for a reversal than for a comparable non-reversal:*

$$E[DT_R | CE_P > CE_{\$}, c(P, \$) = \$] > E[DT_R | CE_{\$} > CE_P, c(P, \$) = \$] \text{ and} \\
 E[DT_R | CE_{\$} > CE_P, c(P, \$) = P] > E[DT_R | CE_P > CE_{\$}, c(P, \$) = P].$$

How does this assumption affect our predictions? Hypotheses H1a and H1b do not concern decision times and are hence unaffected. Hypothesis H3 is equally unaffected since this prediction does not concern preference reversals. The additional assumption affects the interpretation of Hypotheses H2a and H2b. Since total decision time is now viewed as the sum of choice time and conflict resolution time, Proposition 2, which states that choice times are longer for reversals than for non-reversals, does not directly translate into experimental hypotheses anymore. However, by Assumption 5, also conflict resolution time is on average larger for reversals. Hence, both effects are aligned and Hypotheses H2a and H2b still obtain.

4.3 Experiment 1: Preference Reversals and Decision Times

The objective of our first experiment was to test the predictions of the model with regard to both choices and decision times. This would allow us to conclude that the combination of imprecise preferences in the evaluation phase

and an overpricing phenomenon arising from the compatibility hypothesis is able to explain received evidence on preference reversals while at the same time fitting novel evidence on process data.

4.3.1 Experimental Design and Procedures

We followed a between-subject design comprised of three independent, consecutive single-decision making parts. The first and third phases were evaluation tasks, while the second, intermediate phase contained the choice task. This way, we can consider two kinds of preference reversals. On the one hand, we have “Price-Choice Reversals” which occur comparing the evaluations in the first phase and the choices in the second phase. On the other hand, we have “Choice-Price Reversals” which occur comparing the evaluations in the third phase and the choices in the second phase. Each of our Hypotheses H1 to H3 can be tested either on Price-Choice or Choice-Price reversals (or non-reversals), and we will report the results for both possibilities, keeping in mind that we expect clearer results for the Choice-Price ordering. Comparing both orderings allows us to test Hypothesis H4.

The stimuli were 40 different lotteries, which are presented in Table 4.5 in Appendix 4.B. Each of the pairs in the choice task contained one P -bet and one $\$$ -bet from this set of lotteries, with the former being defined by a high probability of winning a moderate amount of money and the latter being defined by a low probability of winning a high amount of money.⁹

We employed a pricing method for the evaluation of lotteries in phases 1 and 3. In these two pricing tasks participants were asked to state their minimum selling price for each of the 40 lotteries which were presented sequentially in fully randomized order (“State the lowest price for which you

⁹Of the 20 lottery pairs, pairs 3 to 8 were such that the expected value of the P -bet was higher than the expected value of the $\$$ -bet (with a difference between €1.00 to €3.40). Pairs 9 to 14, which most closely resemble the ones commonly used in the literature, had roughly equal expected values. In pairs 15 to 20, the $\$$ -bet had a higher expected value than the P -bet (difference between €1.60 to €4.80). Finally, lottery pairs 1 and 2 were such that one bet dominated the other strictly and were only included as a basic rationality check. Only 2 out of 141 subjects chose one of the two strictly dominated lotteries in phase 2. These two lottery pairs are therefore excluded from the analysis.

are just willing to sell the presented lottery.”). Subjects were only allowed to state prices between €2 (the lower amount to win) and the higher amount to win. An example screen display for the pricing tasks is shown in Figure 4.6(a) in Appendix 4.C. The colors in the pie charts (green and blue) were counterbalanced across subjects. In phase two, the choice task, subjects faced the 20 lottery pairs sequentially and had to choose the lottery they would prefer to play out. See Figure 4.6(c) in Appendix 4.C for an example screen of the choice task. The order of the pairs and the onscreen position of the P -bet (i.e. left or right) was randomized. For each choice, we recorded the individuals’ decision times as the time elapsed between the presentation of the lottery pair and clicking the button (“This lottery”) underneath one of the two lotteries.

After the three tasks, participants filled in a questionnaire containing various questions on their statistical knowledge, sociodemographic background, and personality attitudes.

All three tasks were incentivized. Payoffs were determined independently for each task after completion of the ex-post questionnaire to prevent spillover effects between tasks (e.g. through wealth effects). The two treatments in this experiment, *BDM* and *OrdPM*, differed only in the payment scheme used in the pricing tasks (phases 1 and 3). The former used a BDM payment scheme (Becker, DeGroot, and Marschak, 1964), and the latter a variant of the Ordinal Payment Method (Goldstein and Einhorn, 1987; Tversky, Slovic, and Kahneman, 1990; Cubitt, Munro, and Starmer, 2004). We included these two treatments to ensure that our results were robust with respect to the elicitation method.

The two schemes determined the payment in an evaluation task as follows. In the BDM treatment, after one of the 40 lotteries was picked at random the computer drew a price from a uniform distribution over the interval $]2, A[$, where A denotes the higher of the two amounts to win. If this price was higher than or equal to the price stated by the subject, the subject received this amount. If it was lower, the subject played the lottery and the payment was the realized outcome of that lottery. This was done separately for each pricing task. In the OrdPM treatment, two lotteries were chosen at random.

The more highly priced lottery of the two was then played out and the realized outcome was the payoff for this phase (in case of a tie, the computer chose one at random). As in the BDM treatment, this procedure was conducted separately for the two pricing phases. Note that under the ordinal payment scheme, the absolute prices do not play a role, but only the induced ordering matters.

The payment method for the choice task in phase 2 was identical in both treatments. One of the 20 lottery pairs was picked at random, then the lottery the subject had chosen from this pair was played out and the realized outcome of that lottery was the payment for this round. The total payment a participant received in the experiment was the sum of realized payoffs in the three decision tasks.

4.3.2 Procedures

Before the experiment started, participants were briefly informed that the session consisted of three decision tasks, that payment for each task was partly determined by their decisions and partly by luck, that the tasks were paid independently of one another and that lotteries from each phase were not played out before the end of the experiment. In addition, four control questions had to be answered, using pencil and paper, before the start of the experiment to ensure that participants understood the concept of a lottery and its pie chart representation. Detailed instructions about each individual decision-making task (phase 1 to 3) and how payments would be determined in each phase were handed to the participants prior to the start of each phase.

The experiment was programmed in z-tree (Fischbacher, 2007). Participants were university students with majors other than psychology and economics. Each student participated in only one session. We conducted 7 sessions with a total of 141 participants (91 female). Of those, 67 were allocated to the BDM treatment and 74 to the OrdPM treatment. A session lasted about 2 hours with average earnings of €24.76 in the BDM treatment and of €23.03 in the OrdPM treatment.

Table 4.1: Preference reversal rates, Experiment 1.

Treatment	Predicted Reversals		Unpredicted Reversals	
	Price-Choice	Choice-Price	Price-Choice	Choice-Price
BDM	48.75	47.20	18.80	8.99
OrdPM	40.96	35.47	17.69	11.14

Note: Predicted (resp. unpredicted) reversal rates computed as percentage of reversals over all P -bet-choices (resp. \$-bet-choices).

4.3.3 Results of Experiment 1

As a first illustration, Table 4.1 contains the average reversal rates in the BDM and the OrdPM treatments in Experiment 1. The rate of predicted (unpredicted) reversals is computed as the number of predicted (unpredicted) reversals divided by the number of P -bet choices (\$-bet choices). Figure 4.1 further depicts the average number of preference reversals per participant for both treatments. It is already apparent from Table 4.1 and Figure 4.1 that predicted reversals are more frequent than unpredicted reversals, that measuring reversals with respect to post-choice attitudes reduces their quantity, possibly by reducing noise, and that there might be some minor differences between treatments. We now proceed to test for these observations and our experimental hypotheses.

Predicted vs. unpredicted reversals (H1). We conducted two-sided Wilcoxon Signed-Rank (hereafter WSR) tests to assess whether participants generated more predicted than unpredicted reversals. Tests were highly significant both for the BDM (Price-Choice, $N = 67$, $z = 6.060$, $p < 0.0001$; Choice-Price, $N = 67$, $z = 6.439$, $p < 0.0001$) and the OrdPM treatments (Price-Choice, $N = 74$, $z = 6.177$, $p < 0.0001$; Choice-Price, $N = 74$, $z = 5.770$, $p < 0.0001$). This confirms our Hypothesis H1a. To confirm Hypothesis H1b, we computed the predicted and unpredicted preference reversal rates for each subject individually as the percentage of P -bet/\$-bet

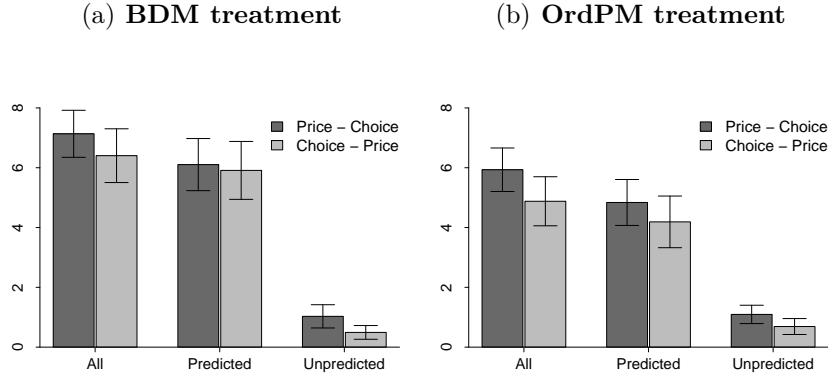


Figure 4.1: Average number of reversals per subject, Experiment 1.

Note: Reversals for the Price-Choice (dark bars) and Choice-Price (light bars) task orderings. Error bars depict the 95 percent confidence interval.

choices (respectively) resulting in reversals. The rates for predicted reversals were significantly higher than the rates of unpredicted reversals for both treatments and both possible task orderings (BDM Price-Choice, $N = 60$, $z = 4.170$, $p < 0.0001$; BDM Choice-Price, $N = 60$, $z = 5.140$, $p < 0.0001$; OrdPM Price-Choice, $N = 69$, $z = 4.585$, $p < 0.0001$; OrdPM Choice-Price, $N = 69$, $z = 3.595$, $p < 0.0005$).¹⁰

Order effects (H4). In both treatments, there were significantly fewer unpredicted reversals when prices are elicited after choices (Choice-Price) than when they are elicited before choices (Price-Choice) according to two-sided WSR tests (BDM, $N = 67$, $z = -3.487$, $p < 0.0005$; OrdPM, $N = 74$, $z = -2.858$, $p = 0.004$). There were no significant differences in the number of predicted reversals, although there seems to be a trend towards fewer predicted Choice-Price reversals in the OrdPM treatment (BDM, $N = 67$, $z = -0.169$, $p = 0.865$; OrdPM, $N = 74$, $z = -1.526$, $p = 0.127$). Since unpredicted reversals are essentially due to noise, this is consistent with the interpretation that measuring reversals through post-choice evaluations reduces noise.

¹⁰The tests for reversal rates include of course only the participants for which both rates can be computed. For instance, if a participant never chose a \$-bet, no rate of unpredicted reversals can be computed.

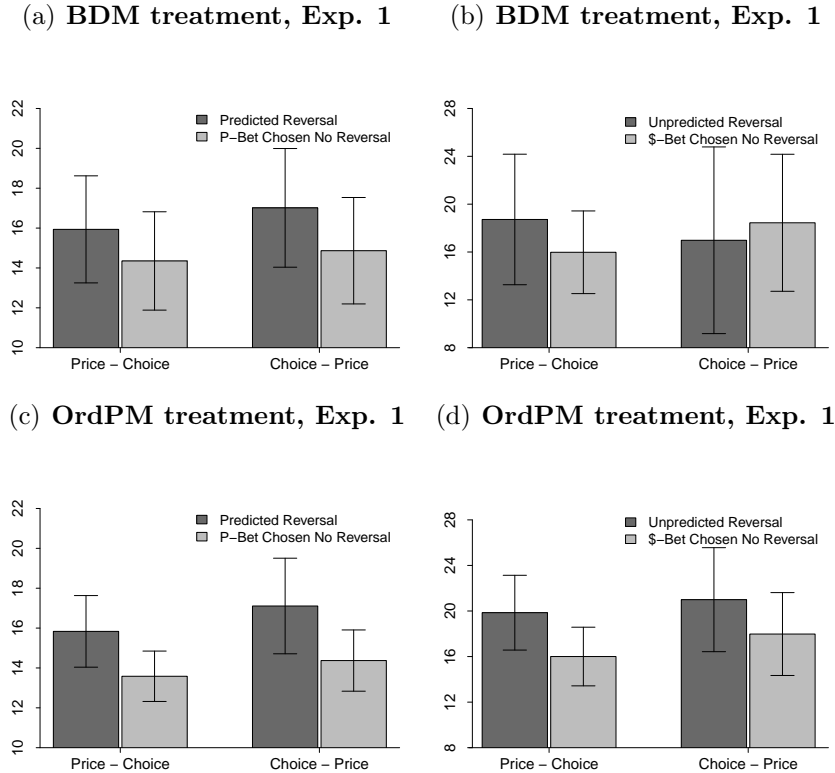


Figure 4.2: Average decision time per individual in the choice task, Experiment 1.

Note: Predicted reversals are compared to non-reversals where the P -bet was chosen, unpredicted reversals to non-reversals where the $\$$ -bet was chosen. Error bars depict the 95 percent confidence interval.

Treatment effects (reversals). We compared the individual numbers of reversals across treatments using Mann-Whitney-U (MWU) tests. We found significantly fewer predicted reversals in the OrdPM treatment than in the BDM treatment, for both task orderings (Price-Choice, $z = -2.101$, $p = 0.036$; Choice-Price, $z = -2.688$, $p = 0.007$). There were, however, no significant differences for unpredicted reversals (Price-Choice, $z = -0.735$, $p = 0.462$; Choice-Price, $z = 1.067$, $p = 0.286$).

Decision times and reversals (H2). Figure 4.2 displays the decision times for reversals and comparable non-reversals for both treatments and both task orderings. Each type of reversal is compared with the correct counterfactual,

i.e. predicted reversals are compared with non-reversals where the P -bet was chosen, and unpredicted reversals with non-reversals where the $\$$ -bet was chosen.

Two-sided WSR tests confirmed that predicted reversals involved significantly longer decision times than comparable non-reversals, both for Price-Choice and for Choice-Price, both for the BDM (Price-Choice $N = 61$, $z = 2.758$, $p = 0.006$; Choice-Price, $N = 54$, $z = 3.625$, $p < 0.0005$) and the OrdPM treatments (Price-Choice, $N = 66$, $z = 2.894$, $p = 0.004$; Choice-Price, $N = 57$, $z = 2.987$, $p = 0.003$).¹¹ Unpredicted reversals were also associated with significantly longer decision times in the OrdPM treatment (Price-Choice, $N = 39$, $z = 2.854$, $p = 0.004$; Choice-Price, $N = 22$, $z = 1.883$, $p = 0.060$), but there were no significant differences for unpredicted reversals in the BDM treatment (Price-Choice, $N = 31$, $z = 0.950$, $p = 0.342$; Choice-Price, $N = 17$, $z = -0.947$, $p = 0.344$).

Decision times and non-reversals (H3). Non-reversals were clearly slower when the $\$$ -bet was chosen than when the P -bet was chosen. The difference was highly significant independently of whether choices were declared non-reversals according to pre-choice or post-choice evaluations, for both the BDM treatment (Price-Choice, $N = 56$, $z = 3.242$, $p = 0.001$; Choice-Price, $N = 51$, $z = 2.995$, $p = 0.003$) and the OrdPM treatment (Price-Choice, $N = 64$, $z = -3.681$, $p < 0.0005$; Choice-Price, $N = 59$, $z = -3.204$, $p = 0.001$). The differences are illustrated in Figure 4.3.

4.3.4 Regression Analysis for Experiment 1

We also conducted a random effects panel regression analysis (with standard errors clustered at the subject level) to further investigate the relation between preference reversals and decision times, and to further test our hypotheses while controlling for a number of natural variables, e.g. individual and lottery-pair covariates. Since decision times are always positive, we used the log of decision times ($\log DT$) as the dependent variable. The main re-

¹¹Every test on decision times was conducted for the population of subjects for which the involved average decision times could be computed. For instance, if a subject did not display any unpredicted reversal, no decision time can be computed for this category.

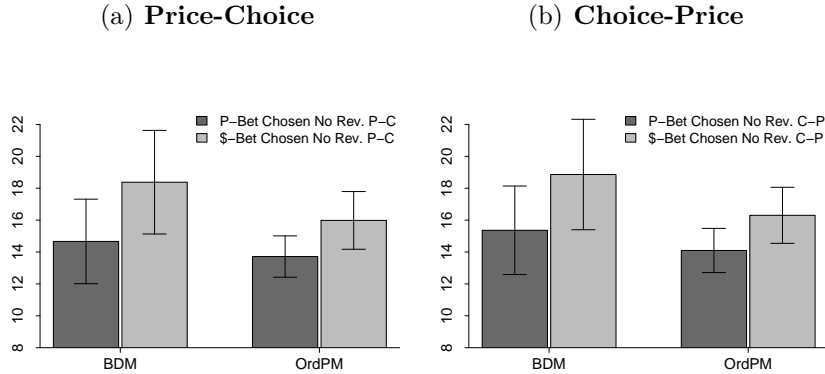


Figure 4.3: Average non-reversal decision time per individual in the choice task, Experiment 1.

Note: Choices classified as non-reversals according to the indicated task ordering, Price-Choice (left) and Choice-Price (right). Error bars depict the 95 percent confidence interval.

sults of these regressions are displayed in Table 4.2. For each treatment, we report a regression including a dummy variable for Price-Choice reversals and an analogous one with a dummy variable for Choice-Price reversals. We also ran a number of additional regressions and found the main effects to be robust (in magnitude and significance) to the inclusion or exclusion of additional control variables.

The regressions include dummies for choices which were part of reversals, for \$-bet-choices, and the interaction thereof. Hence we can make any comparison among reversals and non-reversals where the \$-bet or the P -bet was chosen, either directly through specific regression coefficients or via appropriate postestimation tests, which are also reported in the table.

Predicted reversals vs. non-reversals. Hypothesis H2a states that decision times for predicted preference reversals should be longer on average than decision times for comparable non-reversals, i.e. non-reversals where the P -bet was chosen. Since a \$-choice dummy is included, the comparison between predicted reversals and non-reversals where the P -bet was chosen corresponds to the reversal dummy in the regression, which is highly significant and positive for both regressions for the OrdPM treatment, and for the Choice-Price

CHAPTER 4
PREFERENCE REVERSALS: TIME AND AGAIN

Table 4.2: Random effects panel regressions for decision times, Experiment 1.

Treatment	BDM	BDM	OrdPM	OrdPM
Order	P-C	C-P	P-C	C-P
ReversalPC	0.018 (0.029)		0.078*** (0.025)	
ReversalCP		0.083*** (0.031)		0.109*** (0.027)
\$-Choice	0.127*** (0.054)	0.155*** (0.031)	0.126*** (0.027)	0.151*** (0.026)
\$-Choice ×ReversalPC	-0.097 (0.073)		0.107** (0.052)	
\$-Choice ×ReversalCP		-0.189** (0.095)		0.058 (0.064)
DiffEV	-0.023* (0.012)	-0.022* (0.012)	-0.021** (0.009)	-0.023** (0.009)
Ratio	0.041*** (0.011)	0.036*** (0.011)	0.038*** (0.008)	0.038*** (0.008)
StatedDiff-1	-0.010 (0.007)	-0.009 (0.007)	-0.008** (0.003)	-0.009*** (0.003)
StatedDiff-3	-0.010 (0.006)	-0.011* (0.006)	-0.011*** (0.003)	-0.011*** (0.003)
Round	-0.008*** (0.003)	-0.008*** (0.003)	-0.005*** (0.002)	-0.005*** (0.002)
Female	-0.286** (0.113)	-0.290** (0.113)	-0.156** (0.072)	-0.151** (0.071)
Position	0.016 (0.023)	0.015 (0.023)	-0.012 (0.019)	-0.009 (0.019)
Color	-0.021 (0.111)	-0.019 (0.112)	0.088 (0.070)	0.084 (0.068)
Constant	2.666*** (0.106)	2.649*** (0.105)	2.563*** (0.073)	2.556*** (0.072)
Nr. Obs.	1340	1340	1480	1480
Nr. Groups	67	67	74	74
R2-Overall	0.103	0.101	0.118	0.119
Wald test	0.000	0.000	0.000	0.000
Postestimation tests				
Reversal	-0.079	-0.106	0.185***	0.167***
+($\$$ -Choice×Reversal)	(0.072)	(0.094)	(0.045)	(0.056)

Note: All regressions are random-effects panel estimations, with log decision time as dependent variable. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

regression for the BDM treatment. This indicates that predicted reversals took longer than comparable non-reversals, confirming Hypothesis H2a.

Unpredicted reversals vs. non-reversals. Hypothesis H2b states that unpredicted reversals should take longer than non-reversals where the \$-bet was chosen. The difference between both types of choices corresponds to $\beta_{Reversal} + \beta_{\$-Choice \times Reversal}$, which is highly significant and positive in both regressions for the OrdPM, confirming Hypothesis 2b. However, the postestimation tests are not significant for the BDM treatment.

Comparison of non-reversals. According to Hypothesis H3, non-reversals where the \$-bet was chosen should take longer than non-reversals where the P -bet was chosen. Since reversals dummies are included, this comparison corresponds to the \$-choice dummy, which is highly significant and positive for all four regressions. Hence, conditional on the absence of a preference reversal, \$-bet-choices took longer, confirming Hypothesis H3.

Controls: Lotteries. We included a number of covariates in order to control for differences in the lottery pairs. The ratio of the two higher amounts to win in the \$-bet and the P -bet (Ratio) had a significant positive effect in both treatments. The absolute value of the difference in expected values of the P -bet and the \$-bet (DiffEV) had a weakly significant negative effect in both treatments. We further included the absolute difference in the prices stated for the lotteries in phases one and three (StatedDiff-1, StatedDiff-3) as a rough measure of how similar (or different) the participant viewed the lotteries within a pair. Both were highly significant in the OrdPM treatment, but essentially not significant in the BDM treatment.

Other controls. Decision time measurements in repeated tasks usually capture a learning effect as participants gain familiarity with the interface. We controlled for this effect by including the round in which the choice was made as a regressor (Round). This was significantly positive in both treatments. A dummy variable controlling for gender (Female) was also significant in both treatments. Finally, we controlled for onscreen position (Position) of the P -bet and the \$-bet and for the colors used in the pie-chart (Color) to verify that these factors did not influence the results. As expected, these variables never had significant effects.

4.3.5 Discussion of Experiment 1

The analysis of the data confirms our predictions as derived from the model in Section 4.2. First, predicted reversals are clearly more frequent than un-predicted ones, in agreement with previous experiments. Second, preference reversals appear to involve longer decision times. This effect is clear for an ordinal-based elicitation of prices; in the BDM treatment, the effect is also present albeit less pronounced. Third, in both treatments we found that \$-bet-choices which are part of non-reversals take significantly longer than P -bet-choices part of non-reversals.

In view of the evidence, we conclude that the data is compatible with the idea that preference reversals arise from the combination of two factors. First, as pointed out by Schmidt and Hey (2004) and Butler and Loomes (2007), monetary valuations of lotteries are typically imprecise, and hence preference elicitation through pricing tasks is much noisier than actual choices. Second, as summarized by the compatibility hypothesis (Tversky, Sattath, and Slovic, 1988; Tversky, Slovic, and Kahneman, 1990), the use of pricing tasks causes an overpricing phenomenon which anchors up the evaluation of bets where a relatively high monetary outcome is salient. These observations produce testable hypotheses for both choice data and decision times once we incorporate the observation that easier choices (where the alternatives are farther away from indifference) take longer (e.g. Wilcox, 1993; Shultz, Léveillé, and Lepper, 1999; Moffatt, 2005).

Regarding ordering effects, we observe small but systematic differences suggesting that a Price-Choice ordering, where the evaluation task precedes actual choices, might be noisier than the opposite order, hence producing both more reversals and slightly less clear effects. This is compatible with self-perception theory (Bem, 1967), which holds that actual choices serve as “self-signals” which help reduce noise in future evaluations of alternatives.

Last, we observe small but definite treatment effects, pointing out that price evaluations conducted through the BDM “price-list” scheme might be noisier than those conducted according to a more intuitive, ordinal-like scheme. This is reflected by the fact that preference reversals (and especially unpre-

dicted ones, which are presumably due to noise) are more frequent in the BDM case. It is also compatible with the general observation that effects are often more clearly observed in the OrdPM treatment than in the BDM one.

4.4 Experiment 2: Eliminating Reversals

The objective of our second experiment was twofold. First, we wanted to show that the overpricing phenomenon can be next to eliminated by using ordinal, ranking-based evaluation tasks. Second, this manipulation would allow us to disentangle the two building blocks of our model. The absence of the overpricing phenomenon should result in a reduction of predicted preference reversals, while the assumption of imprecise preferences still delivers predictions on decision times.

4.4.1 Motivation and Hypotheses

In our first experiment we found that the method used to elicit participants' minimum selling prices affects the rate of preference reversals. According to the compatibility hypothesis, predicted reversals appear because participants focus more on monetary outcomes when their preferences are elicited through prices. Notably, preference reversals were also present in the OrdPM treatment, where the use of prices in the evaluation task was simple framing, with no direct monetary consequences. This raises the natural hypothesis that the overpricing phenomenon predicted by the compatibility hypothesis arises due to a price-based, cardinal framing (i.e., a "rating task") in the evaluation phases. As a consequence, moving away from a cardinal evaluation task towards a more natural, ordinal-based one (a "ranking task") should greatly reduce preference reversals.

Specifically, suppose that, by employing a ranking-based evaluation task, we were able to shut down the decision process responsible for the overpricing phenomenon. In terms of the model in Section 4.2, this would imply $K = 0$ in Assumption 3. It is easy to revisit our theoretical predictions and derive new experimental hypotheses for such a situation. First, Propo-

sition 1 crucially depends on Assumption 3, and hence we would not expect Hypotheses H1a/H1b to hold in this setting. Although from the point of view of the model we would expect no differences in reversal rates, this rests upon the implicit assumption that there is no other (second-order) latent process causing unpredicted reversals. Even if this was the case, a conservative hypothesis derived from our theoretical analysis is that the number and frequency of predicted preference reversals should be greatly reduced in comparison to treatments with price-framed evaluations.

H5. There will be fewer predicted preference reversals if ordinal, ranking-based evaluation tasks are used than if rating-based tasks are used.

The first decision-times predictions spelled out in Proposition 2, however, do *not* depend on Assumption 3. Hence, independently of whether evaluation tasks are based on ratings or rankings, we would expect Hypotheses H2a/b to hold.

H6a/b. Even if ordinal, ranking-based evaluation tasks are used, choices associated with predicted preference reversals take longer than P -bet-choices associated with non-reversals, and choices associated with unpredicted preference reversals take longer than $\$$ -bet-choices associated with non-reversals.

Proposition 3 depends on Assumption 3. If $K = 0$, we would *a priori* expect no differences in the decision times associated with non-reversals where the P -bet or the $\$$ -bet was chosen.

H7. If ordinal, ranking-based evaluation tasks are used, the average decision time for non-reversals where the P -bet is chosen is not different from the average decision time for non-reversals where the $\$$ -bet is chosen.

4.4.2 Design of Experiment 2

The basic setup of our second experiment was almost identical to Experiment 1, with the exception that we used different evaluation tasks. We used two different ranking-based tasks and one BDM task. The former were meant to shut down the overpricing decision processes; the latter was intended as a control treatment. In each of the three treatments, presentation of lotteries

was such that participants faced a total of three blocks consisting each of six lotteries, i.e. a total of 18 pairs.¹² In the *Rank-Unframed treatment*, we used a purely ranking-based task. Participants were asked to assign ranks (from most preferred to least preferred) to the lotteries according to how much they would like to play each lottery, separately for each block. Most importantly, we did not make any reference to prices (see Figure 4.6(b) in Appendix 4.C for an example screen display of the two ranking treatments). In this sense, the task was unframed. The *Rank-Framed treatment* was identically programmed. The only difference was in the experimental instructions. Participants were asked to rank the lotteries (from 1 to 6) according to their minimum selling price, separately for each block. However, they were *not* asked to type in or otherwise state the prices, but merely to think about them and use them for the ranking. Finally, in the *BDM2 treatment*, participants had to complete a pricing task that was identical to the one in the BDM treatment in Experiment 1, with the only exception that (for comparability with the other treatments) lotteries were presented one after another in three blocks of six lotteries each. Again, colors and onscreen positions of the lotteries were completely randomized in all treatments.

As in Experiment 1, all three tasks were incentivized and payoffs for each task were determined independently. Payoffs for the evaluation task of the BDM2 and the choice tasks of all three treatments were determined in the same way as in Experiment 1. Payoffs for the evaluation phases for Treatments Rank-Unframed and Rank-Framed were determined as follows. First, the computer picked one of the six blocks at random. From the six lotteries contained in that block, the computer again randomly picked two. The one that had been ranked higher by the participant was then played out and the participant received the outcome of that lottery as payment for that round. In all three treatments, payments were determined and presented to participants only after all three tasks had been completed. Since in both ranking treatments there was no actual “pricing” task, we will refer to the two possible task orderings for these treatments as “Rank-Choice” and “Choice-Rank”.

¹²We only used 18 of the lottery pairs that had been used in Experiment 1 (pairs 3-20 in Table 4.5), excluding pairs 1 and 2 which contained stochastically dominated lotteries.

4.4.3 Procedures

We followed the same procedures as in Experiment 1. We conducted 12 sessions with a total of 215 participants (102 female). Of those, 73 were allocated to the Rank-Unframed treatment, 73 to the Rank-Framed treatment, and 69 to the BDM2 treatment. Sessions in the Rank-Unframed treatment lasted roughly an hour with average earnings of €23.36. Sessions in the Rank-Framed treatment lasted one hour and 20 minutes with average earnings of €24.07, while sessions in the BDM2 treatment lasted about 2 hours with average earnings of €28.44.

4.4.4 Results of Experiment 2

Table 4.3 shows the average reversal rates for all three treatments for both Price/Rank-Choice and Choice-Price/Rank reversals. As before, the percentage of predicted (unpredicted) reversals is computed as the number of predicted (unpredicted) reversals divided by the number of P -bet choices ($\$$ -bet choices). Figure 4.4 shows the average number of reversals per subject in the three treatments. The basic trends are already apparent. Predicted reversals were enormously reduced in both ranking treatments, and especially in the rank-unframed one, to the extent of dropping below the levels of unpredicted reversals. Further, as in Experiment 1 we observe that measuring reversals with respect to post-choice attitudes reduces their quantity.

BDM replication. The first observation is that, as expected, there is no qualitative difference between the results of Treatment BDM2 and Treatment BDM of Experiment 1. For instance, in Treatment BDM2 the number of predicted reversals is significantly higher than the number of unpredicted ones (WSR tests; Price-Choice, $N = 69$, $z = 6.658$, $p < 0.0001$; Choice-Price, $N = 69$, $z = 6.680$, $p < 0.0001$). Likewise, the rates of predicted reversals (relative to the number of P -bet-choices) are significantly higher than the rates of unpredicted reversals (relative to the number of $\$$ -bet-choices) (Price-Choice, $N = 68$, $z = 4.495$, $p < 0.0001$; Choice-Price, $z = 4.585$, $p < 0.0001$).

Reduction of predicted reversals (H5). Kruskal-Wallis tests confirmed that the number of predicted reversals was significantly different across treatments

Table 4.3: Preference reversal rates, Experiment 2.

Treatment	Predicted Reversals		Unpredicted Reversals	
	Price-Choice	Choice-Price	Price-Choice	Choice-Price
BDM	46.87	44.66	16.05	11.32
Rank-Framed	17.67	13.57	34.78	32.95
Rank-Unfr.	12.64	8.39	49.32	45.95

Note: Predicted (resp. unpredicted) reversal rates computed as percentage of reversals over all P -bet-choices (resp. $\$$ -bet-choices).

(Price/Rank-Choice, $\chi^2 = 71.304$, $df = 2$, $p < 0.0001$; Choice-Price/Rank, $\chi^2 = 81.095$, $df = 2$, $p < 0.0001$). To confirm that the differences were between the ranking treatments and the control BDM2 treatment, we conducted two-sided MWU tests with Holm-Bonferroni correction to account for multiple comparisons (p -values below are the adjusted values). Both ranking treatments generated significantly fewer predicted reversals than the BDM2 treatment (Rank-Framed Price/Rank-Choice, $z = -6.769$, $p < 0.0001$; Rank-Framed Choice-Price/Rank, $z = -7.040$, $p < 0.0001$; Rank-Unframed Price/Rank-Choice, $z = -7.745$, $p < 0.0001$; Rank-Unframed Choice-Price/Rank, $z = -8.210$, $p < 0.0001$). The difference in the number of predicted reversals across both ranking treatments was not significant for the Rank-Choice ordering ($z = -0.824$, $p = 0.410$), but for the Choice-Rank ordering there were significantly fewer predicted reversals in the Rank-Unframed treatment than in the Rank-Framed treatment ($z = -2.248$, $p = 0.025$). This last result agrees with the idea that the Rank-Unframed treatment goes one step further in the elimination of the overpricing process than a ranking-based but still price-framed approach.

Order effects. As in Experiment 1, there were significantly fewer unpredicted reversals in the BDM2 treatment when prices were elicited after choices than when they were elicited before choices ($N = 69$, $z = -1.884$, $p = 0.059$), but no significant differences for predicted reversals ($N = 67$,

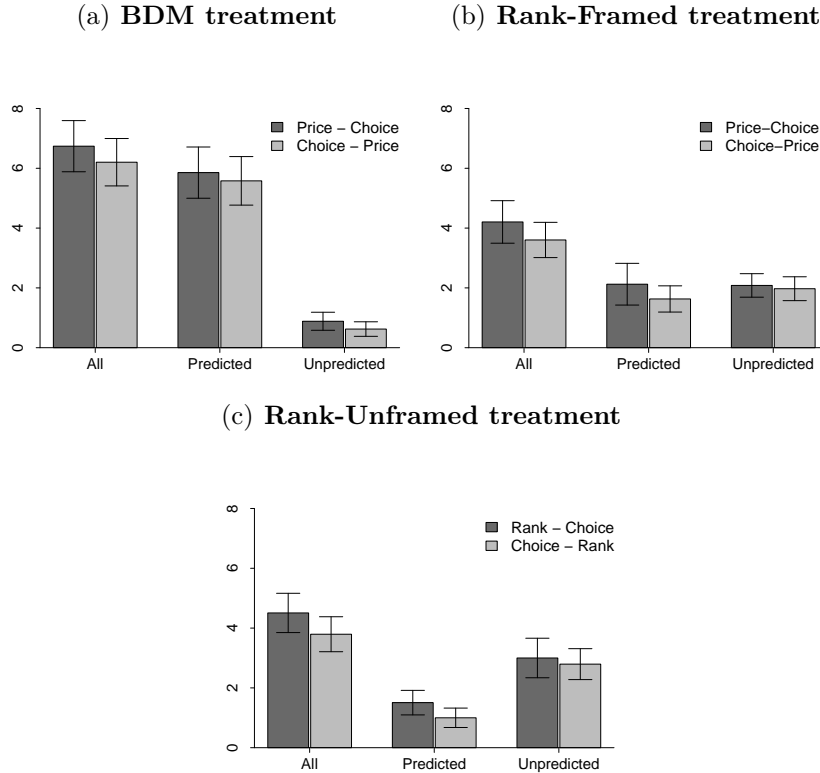


Figure 4.4: Average number of reversals per subject, Experiment 2.

Note: Reversals for the Price/Rank-Choice (dark bars) and Choice-Price/Rank (light bars) task orderings. Error bars depict the 95 percent confidence interval.

$z = -0.470$, $p = 0.638$). We found no differences in the Rank-Framed treatment (unpredicted reversals, $N = 73$, $z = -0.532$, $p = 0.595$; predicted reversals, $N = 73$, $z = -1.154$, $p = 0.248$). For the Rank-Unframed treatment, we only found differences for predicted reversals (unpredicted reversals, $N = 73$, $z = -0.513$, $p = 0.608$; predicted reversals, $N = 73$, $z = -2.245$, $p = 0.025$).

Decision Times and Reversals (H2/H6). Figure 4.5 displays the decision times for reversals and comparable non-reversals for all three treatments and both task orderings, comparing each type of reversal with the appropriate non-reversals. Two-sided WSR tests confirmed that predicted reversals involved longer decision times than comparable non-reversals, both for

Price/Rank-Choice and for Choice-Price/Rank, for all treatments. For both ranking treatments, the differences were highly significant (Rank-Framed Rank-Choice, $N = 42$, $z = 3.551$, $p < 0.0005$; Rank-Framed Choice-Rank, $N = 45$, $z = 2.743$, $p = 0.006$; Rank-Unframed Rank-Choice, $N = 43$, $z = 2.614$, $p = 0.009$); Rank-Unframed Choice-Rank, $N = 34$, $z = 3.163$, $p = 0.002$). This confirms that the decision times effect predicted by our model, which is independent of the overpricing assumption, is still present under ordinal (ranking) evaluation tasks. In the case of the BDM2 treatment, the test missed significance for the Price-Choice ordering ($N = 64$, $z = 1.595$, $p = 0.111$), but the difference was significant for Choice-Price ($N = 58$, $z = 3.004$, $p = 0.003$).

For both ranking treatments unpredicted reversals were again significantly slower than comparable non-reversals independently of task ordering (Rank-Framed Rank-Choice, $N = 49$, $z = 2.875$, $p = 0.004$; Rank-Framed Choice-Rank, $N = 45$, $z = 3.014$, $p = 0.003$; Rank-Unframed Rank-Choice, $N = 49$, $z = 1.930$, $p = 0.054$); Rank-Unframed Choice-Rank, $N = 47$, $z = 3.656$, $p < 0.0005$). In the BDM2 treatment the decision time differences were not significant for the Price-Choice ordering ($N = 30$, $z = 1.131$, $p = 0.258$), but unpredicted reversals were significantly slower for the Choice-Price ordering ($N = 23$, $z = 1.992$, $p = 0.046$).

Decision times and non-reversals (H3/H7). Treatment BDM2 successfully replicated the finding that non-reversals are slower when the \$-bet is chosen than when the P -bet is chosen, as predicted in Hypothesis H3 (WSR tests; Price-Choice, $N = 60$, $z = 1.984$, $p = 0.047$; Choice-Price, $N = 58$, $z = 2.609$, $p = 0.009$). However, for ranking treatments we expected no differences (Hypothesis H7). There is still a significant difference for the Rank-Choice ordering (Rank-Framed, $N = 57$, $z = 1.835$, $p = 0.066$; Rank-Unframed, $N = 54$, $z = 1.825$, $p = 0.068$), but there is clearly no significant difference for the (presumably more appropriate) classification according to the Choice-Rank ordering (Rank-Framed, $N = 55$, $z = 0.733$, $p = 0.463$; Rank-Unframed, $N = 48$, $z = 0.385$, $p = 0.701$). Figure 4.6 illustrates these results.

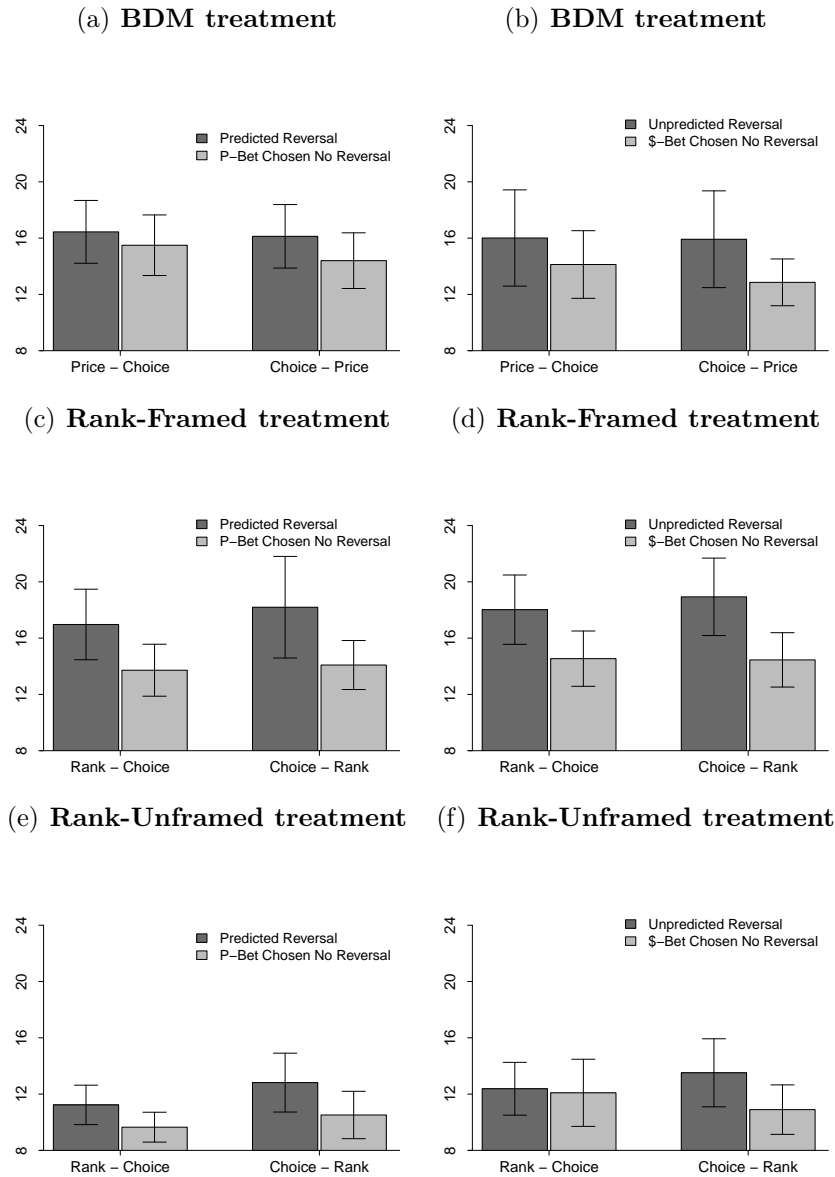


Figure 4.5: Average decision time per individual in the choice task, Experiment 2.

Note: Predicted reversals are compared to non-reversals where the P -bet was chosen, unpredicted reversals to non-reversals where the $\$$ -bet was chosen. Error bars depict the 95 percent confidence interval.

Decision times in the Rank-Unframed Treatment. As can be seen in Figures 4.5 and 4.6, all decisions in the Rank-Unframed treatment were signifi-

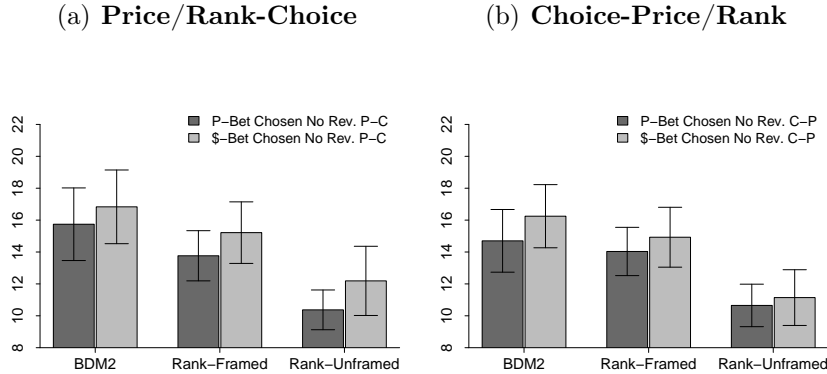


Figure 4.6: Average non-reversal decision time per individual in the choice task, Experiment 2.

Note: Choices classified as non-reversals according to the indicated task ordering. Error bars depict the 95 percent confidence interval.

cantly quicker than in the other two treatments. The difference is substantial: the median decision time over all choices was 13.41 s in BDM2, 12.52 s in Rank-Framed, and only 9.61 s in Rank-Unframed. This difference is remarkable, because the choice phases in which the decision times were measured were completely identical across treatments; the differences across treatments concerned only the evaluation phases. We will discuss this observation in detail below.

A Kruskal-Wallis test confirmed that the decision times were significantly different across treatments ($\chi^2 = 35.545$, $df = 2$, $p < 0.0001$). Two-sided MWU tests with Holm-Bonferroni correction to account for multiple comparisons showed that decisions were faster in the Rank-Unframed treatment than in both of the other treatments (BDM2, $z = -5.722$, $p < 0.0001$; Rank-Framed, $z = -4.225$, $p < 0.0001$).¹³

¹³The difference between decision times in Treatments Rank-Framed and BDM2 missed significance, $z = -1.596$, $p = 0.111$.

4.4.5 Regression Analysis for Experiment 2

As for Experiment 1, we conducted a random effects panel regression analysis on the log of decision times from Experiment 2. The objective was to confirm and clarify our results while controlling for natural individual and lottery-pair characteristics; specifically, we included the same controls as in Experiment 1.¹⁴ Table 4.4 contains the main results of all treatments. For each treatment, in the first regression reversals are classified as such according to the Price/Rank-Choice task ordering, while in the second one the Choice-Price/Rank is used. We present a single regression for each treatment and task ordering, but the results are robust with respect to the control variables.

Predicted reversals vs. non-reversals. The reversal dummies were highly significant in all treatments and task orderings, except for the “noisiest” Price-Choice in Treatment BDM2. This indicates that, as in Experiment 1, predicted reversals took longer than comparable non-reversals, confirming Hypothesis H2a.

Unpredicted reversals vs. non-reversals. Hypothesis H2b states that unpredicted reversals should take longer than non-reversals where the \$-bet was chosen. The difference corresponds to $\beta_{Reversal} + \beta_{\$-Choice \times Reversal}$, which is indeed highly significant and positive in all four regressions for the ranking treatments. The postestimation tests are not significant for the BDM2 treatment.

Comparison of non-reversals. The \$-choice dummy is significant and positive for Treatment BDM2. That is, as in Experiment 1, non-reversals where the \$-bet was chosen took longer than non-reversals where the P -bet was chosen in this treatment (Hypothesis H3). As stated in Hypothesis H7, we expected this effect to disappear for the purely ordinal, unframed treatment Rank-Unframed. Indeed, the dummy is not significant in any of the regressions for this treatment. The prediction is less clear for the “intermediate” treatment Rank-Framed, where the evaluation task was also

¹⁴For the two ranking treatments, StatedDiff-1 and StatedDiff-3 refer to the difference in stated ranks between the two lotteries within a pair in phases 1 and 3, respectively.

CHAPTER 4
PREFERENCE REVERSALS: TIME AND AGAIN

Table 4.4: Random effects panel regressions for decision times, Experiment 2.

Treatment Order	BDM2		RankFramed		RankUnframed	
	P-C	C-P	R-C	C-R	R-C	C-R
ReversalPC	0.049 (0.031)		0.135*** (0.042)		0.108*** (0.044)	
ReversalCP		0.100*** (0.032)		0.155*** (0.043)		0.199*** (0.051)
\$-Choice	0.072** (0.033)	0.098*** (0.033)	0.073** (0.033)	0.062* (0.033)	0.030 (0.036)	-0.006 (0.035)
\$-Choice ×ReversalPC	0.026 (0.069)		0.014 (0.062)		0.014 (0.061)	
\$-Choice ×ReversalCP		0.011 (0.077)		0.039 (0.064)		0.008 (0.067)
DiffEV	-0.022** (0.031)	-0.020* (0.011)	-0.026** (0.010)	-0.025** (0.010)	-0.013 (0.010)	-0.008 (0.010)
Ratio	0.034*** (0.011)	0.033*** (0.011)	0.038*** (0.010)	0.038*** (0.010)	0.044*** (0.009)	0.043*** (0.009)
StatedDiff-1	-0.005 (0.005)	-0.004 (0.005)	-0.011 (0.010)	-0.014 (0.010)	-0.024** (0.010)	-0.026*** (0.010)
StatedDiff-3	-0.003 (0.005)	-0.005 (0.005)	-0.046*** (0.010)	-0.041*** (0.010)	-0.045*** (0.011)	-0.042*** (0.011)
Round	-0.007*** (0.002)	-0.007*** (0.002)	-0.008*** (0.002)	-0.007*** (0.002)	-0.003 (0.002)	-0.003 (0.002)
Female	-0.190** (0.089)	-0.181** (0.087)	-0.238*** (0.077)	-0.237*** (0.077)	-0.217** (0.087)	-0.212** (0.085)
Position	0.004 (0.023)	0.008 (0.023)	0.008 (0.022)	0.004 (0.022)	0.030 (0.021)	0.022 (0.021)
Color	0.108 (0.087)	0.112 (0.086)	0.088 (0.077)	0.085 (0.077)	-0.034 (0.085)	-0.040 (0.083)
Constant	2.624*** (0.095)	2.595*** (0.094)	2.649*** (0.087)	2.645 (0.087)	2.391*** (0.090)	2.388*** (0.088)
Nr. Obs.	1242	1244	1314	1314	1314	1314
Nr. Groups	69	69	73	73	73	73
R2-Overall	0.066	0.069	0.128	0.140	0.104	0.122
Wald test	0.000	0.000	0.000	0.000	0.000	0.000
Postestimation tests						
Reversal +(\$-Ch×Rev)	0.075 (0.060)	0.111 (0.068)	0.149*** (0.043)	0.194*** (0.044)	0.121*** (0.031)	0.207*** (0.041)

Note: All regressions are random-effects panel estimations, with log decision time as dependent variable. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

ordinal but there was an indirect framing in terms of prices.¹⁵ For this treatment, the \$-choice was significantly positive, but e.g. only at the 10% level for the Choice-Rank ordering.

Controls: Lotteries. As in Experiment 1, the ratio of the two higher amounts to win in the \$-bet and the P -bet (Ratio) had a significant positive effect throughout. Likewise, the absolute difference in expected values of the P -bet and the \$-bet (DiffEV) had a weakly significant negative effect, but not in the Rank-Unframed treatment. The absolute difference in the prices/ranks stated for the lotteries in phases one and three (StatedDiff-1, StatedDiff-3) was significant for Rank-Unframed but not for BDM2 (and only the second measure was significant for Rank-Framed).

Other controls. As in Experiment 1, we controlled for learning and familiarity effects by including the round in which the choice was made as a regressor. Also as in Experiment 1, female participants were significantly quicker in all treatments and task orderings. The onscreen position (Position) of the P -bet and the \$-bet and the colors used in the pie-chart (Color) had, as expected, no effect.

4.4.6 Discussion of Experiment 2

The analysis of the data confirms our predictions, strengthening our interpretation that preference reversals arise from the combination of noisy evaluations and an overpricing phenomenon. The almost-complete disappearance of predicted reversals in the ranking treatments (especially when ordinally framed) confirms that the overpricing phenomenon appears due to the cardinal, rating-based frame used in standard evaluation tasks as those employed in Experiment 1. The fact that reversals are still associated with longer decision times (a prediction our model derives from noisy evaluations) even though the overpricing process has been impaired is further evidence that both noisy evaluations and the overpricing phenomenon need to be taken into account as different ingredients in order to model preference reversals.

¹⁵We consider the framing “indirect” because, contrary to the tasks in Experiment 1 or Treatment BDM2, participants did not actually write down prices.

An important observation is that decision times in Treatment Rank-Unframed were significantly lower than those in other treatments. To understand this effect, recall our dual-process interpretation as sketched in Section 4.2.4. In this treatment, we removed all references to prices, and it is easy to argue that the decision process which usually causes overpricing was simply not activated at all. Hence, in the choice phase there was no process conflict, and no additional time was spent in conflict resolution. It is especially interesting to observe that in the Treatment Rank-Framed, where the evaluation task was also ordinal but the frame made a reference to prices, decision times were closer to those of the rating treatments, even though predicted reversals were also greatly reduced. Again, the interpretation is simple. The price frame generally activated the process behind overpricing, but the fact that the task was ultimately a purely ordinal one made it less likely that this process actually shaped the decision in the evaluation tasks, hence reducing reversals. However, since the process had been activated, it needed to be inhibited in the choice phase, causing longer decision times.

Our results are consistent with evidence from Bateman, Day, Loomes, and Sugden (2007). These authors also observed a reduction in predicted preference reversal rates in an experiment where lotteries were ranked within sets which also contained sure amounts. Their ranking task is not directly comparable to ours because \$-bets and P -bets were ranked separately, i.e. within different sets, and the ranks of P -bets relative to \$-bets were inferred indirectly. Oliver (2013) used a similar method for the measurement of preferences in the health domain (life expectancy).

Our main object of study have been *predicted* preference reversals, since they are empirically more relevant and the compatibility hypothesis points to an overpricing phenomenon as a reason for the predominance of these reversals, while the origin of unpredicted ones might be just noisy evaluations. Nonetheless, it is interesting to observe that the number and rate of *unpredicted* reversals increased in the ranking treatments with respect to the control (BDM2) treatment. We hypothesize that, when the cues on which the overpricing process acts are removed, attention is diverted to probabilities instead. Following the compatibility hypothesis, this would result in an

over-evaluation of P-bets, for which a high probability is salient. However, this process is weaker than the one causing overpricing of \$-bets with pricing frames, simply because monetary rewards are a more immediately accessible concept than probabilities. Thus, in a standard preference-reversal study, this second, probability-based process is overshadowed by the overpricing of \$-bets. Our evidence in this respect is consistent with Cubitt, Munro, and Starmer (2004), where the rate of unpredicted reversals increased when subjects were asked for “probabilistic valuations” instead of prices, trying to induce a probability anchor and shift the predictions of the compatibility hypothesis to unpredicted, rather than predicted reversals. However, the rates of predicted reversals remained relatively high, suggesting that such valuation tasks, being still cardinal, do not completely remove the salience of monetary outcomes.¹⁶ Casey (1991, 1994) observed a higher rate of unpredicted reversals compared to predicted ones using very high payoffs and maximum buying prices (rather than minimum selling prices). Again, however, predicted reversal rates remained comparatively high. Casey (1994) argues that high stakes might induce buyers to anchor on the smallest monetary outcome of a lottery, adjusting the valuation upwards on the basis of probabilities, and hence resulting in an overpricing of P-bets. In our terms, the setting of Casey (1991, 1994) might correspond to a combination of elements enhancing the second process mentioned above. If such a second process is assumed, the increase of unpredicted reversals in our ranking treatments, in Cubitt, Munro, and Starmer (2004), and in Casey (1991, 1994) can be easily explained within our model.

Last, we observe order effects similar to those already seen in Experiment 1, again supporting our view that post-choice elicitation tasks carry less noise than pre-choice analogues, possibly due to “preference sharpening” or reappraisal in the sense of self-perception theory.

¹⁶Participants were asked for the probability p making them indifferent between a given lottery and receiving a fixed, high monetary outcome X with probability p . Hence monetary outcomes remained an important part of the frame.

4.5 General Discussion and Conclusion

We propose a simple, parsimonious model which predicts both preference reversals and a clear pattern of decision times in choices among lotteries. We conducted two experiments which confirm the predictions derived from the model. The consideration of decision times allows us to put our model to a more stringent test than if we had relied exclusively on choice data. At the same time, the insights provided by the analysis of decision times allow us to deepen our understanding of the actual decision processes behind preference reversals. Our model, which is based on insights from the previous literature, postulates that reversals arise due to the interaction of noise in the evaluation phases and a psychological process (or set thereof) which causes an overpricing phenomenon of lotteries with a salient monetary outcome. In our second experiment, we have been able to effectively shut down that process, resulting in the practical elimination of predicted preference reversals and a notable reduction of decision times.

Our experimental design also allowed us to evaluate different experimental possibilities with regard to the amount of noise they induce. By using two evaluation phases, one pre-choice and one post-choice, we are able to conclude that post-choice evaluation tasks are in general more appropriate for preference elicitation, in accordance with evidence on preference reappraisal from psychology. By using different evaluation tasks across treatments, we conclude that tasks based on the BDM procedure might add additional, unwanted noise and other tasks, as e.g. the Ordinal Payment Method, might be more accurate. Finally, if one is interested in preferences rather than certainty equivalents, our second experiment shows that the most accurate evaluation method (in the sense of inducing fewer reversals) is to rely on purely ordinal, ranking-based tasks.

Our research investigated (theoretically and experimentally) the mechanisms and processes behind the preference reversal phenomenon. Previous research (see e.g. Cubitt, Munro, and Starmer, 2004) has pointed out that a combination of psychological mechanisms might be the simplest explanation of the phenomenon. Given the fundamental importance of preference (and

consumer demand) elicitation methods for both decision theory and applied economics, and the amount of attention dedicated to the preference reversal phenomenon in the last half century, we believe that fleshing out these mechanisms is an important step. At the same time, we show that a simple parsimonious model can account for received evidence and provide new, testable hypotheses. By using process data (decision times), we are able to show that our model is more than an *as if* construction and, in spite of its simplicity, is able to capture the essential features of the actual mechanisms behind the phenomenon.

Appendix 4.A: Proofs

Throughout the appendix, let $\Delta\zeta = \zeta_P - \zeta_\$ + K$. Under Assumption 3, ζ_P and $\zeta_\$$ are i.i.d. and unimodal, implying that $\Delta\zeta$ is symmetrically distributed around 0 and unimodal (cf. Purkayastha, 1998, Theorem 2.1).

Proof of Proposition 1. (i) Since $K > 0$ by Assumption 3, $Pr(\Delta\zeta < -K - s) < Pr(\Delta\zeta < K - s)$ for all $s \in [0, \infty[$ and the conclusion follows from the following computations.

$$\begin{aligned} & Pr(CE_\$ > CE_P, C(P, \$) = P) \\ &= \int_0^\infty Pr(CE_\$ > CE_P | u_P - u_\$ = s) h(s) ds = \int_0^\infty Pr(\Delta\zeta < K - s) h(s) ds \end{aligned}$$

$$\begin{aligned} Pr(CE_P > CE_\$, C(P, \$) = \$) &= \int_0^\infty Pr(CE_P > CE_\$ | u_\$ - u_P = s) h(s) ds \\ &= \int_0^\infty Pr(\Delta\zeta > K + s) h(s) ds = \int_0^\infty Pr(\Delta\zeta < -K - s) h(s) ds. \end{aligned}$$

(ii) Note that $Pr(CE_\$ > CE_P | c(P, \$) = P) = \frac{Pr(CE_\$ > CE_P, c(P, \$) = P)}{Pr(u_P > u_\$)}$, and $Pr(CE_P > CE_\$ | c(P, \$) = \$) = \frac{Pr(CE_P > CE_\$, c(P, \$) = \$)}{Pr(u_\$ > u_P)}$. Since $Pr(u_P > u_\$) = Pr(u_\$ > u_P)$, the conclusion follows from (i). \square

The next lemma is used in the proof of Proposition 2.

Lemma A.1. *Under Assumption 1, the following hold.*

$$(i) \ Pr(CE_\$ > CE_P | 0 < u_P - u_\$ < \delta) > Pr(CE_\$ > CE_P | u_P - u_\$ > \delta).$$

$$(ii) \ Pr(CE_P > CE_\$ | 0 < u_\$ - u_P < \delta) > Pr(CE_P > CE_\$ | u_\$ - u_P > \delta).$$

Proof. We prove part (i). The proof of part (ii) is analogous. We have

$Pr(CE_{\$} > CE_P | u_P - u_{\$} = s) = Pr(\Delta\zeta < K - s)$ and

$$\begin{aligned} & Pr(CE_{\$} > CE_P | 0 < u_P - u_{\$} < \delta) \\ &= \frac{1}{Pr(0 < u_P - u_{\$} < \delta)} \int_0^{\delta} Pr(CE_{\$} > CE_P | u_P - u_{\$} = s) h(s) ds \\ &> \frac{1}{Pr(0 < u_P - u_{\$} < \delta)} \int_0^{\delta} Pr(\Delta\zeta < K - \delta) h(s) ds = Pr(\Delta\zeta < K - \delta). \end{aligned}$$

Similarly $Pr(CE_{\$} > CE_P | u_P - u_{\$} = s) = Pr(\Delta\zeta < K - s)$ and

$$\begin{aligned} & Pr(CE_{\$} > CE_P | u_P - u_{\$} > \delta) \\ &= \frac{1}{Pr(u_P - u_{\$} > \delta)} \int_{\delta}^{\infty} Pr(CE_{\$} > CE_P | u_P - u_{\$} = s) h(s) ds \\ &< \frac{1}{Pr(u_P - u_{\$} > \delta)} \int_{\delta}^{\infty} Pr(\Delta\zeta < K - \delta) h(s) ds = Pr(\Delta\zeta < K - \delta) \end{aligned}$$

and the conclusion follows. \square

Proof of Proposition 2. (i) To shorten notation let $\Delta_0 = Pr(CE_{\$} > CE_P | 0 < u_P - u_{\$} < \delta)$, $\Delta_1 = Pr(CE_{\$} > CE_P | u_P - u_{\$} > \delta)$, $P^{\delta} = Pr(0 < u_P - u_{\$} < \delta | 0 < u_P - u_{\$})$, and $P = Pr(CE_{\$} > CE_P | u_P > u_{\$})$.

With these definitions, $P = \Delta_0 P^{\delta} + \Delta_1 (1 - P^{\delta})$. We obtain $E[DT_C | CE_{\$} > CE_P, c(P, \$) = P] = \frac{1}{P} [\Delta_0 P^{\delta} T_H + \Delta_1 (1 - P^{\delta}) T_E]$, and $E[DT_C | CE_P > CE_{\$}, c(P, \$) = P] = \frac{1}{1-P} [(1 - \Delta_0) P^{\delta} T_H + (1 - \Delta_1) (1 - P^{\delta}) T_E]$. A simple calculation shows that

$$\begin{aligned} E[DT_C | CE_{\$} > CE_P, c(P, \$) = P] &> E[DT_C | CE_P > CE_{\$}, c(P, \$) = P] \\ &\Leftrightarrow P^{\delta} T_H [\Delta_0 - P] > (1 - P^{\delta}) T_E [P - \Delta_1] \end{aligned}$$

As $P = \Delta_0 P^{\delta} + \Delta_1 (1 - P^{\delta})$, we obtain $\Delta_0 - P = (1 - P^{\delta})(\Delta_0 - \Delta_1)$ and $P - \Delta_1 = P^{\delta}(\Delta_0 - \Delta_1)$. Hence $E[DT_C | CE_{\$} > CE_P, c(P, \$) = P] > E[DT_C | CE_P > CE_{\$}, c(P, \$) = P]$ holds if and only if $T_H(\Delta_0 - \Delta_1) > T_E(\Delta_0 - \Delta_1)$. By Lemma A.1(i), $\Delta_0 > \Delta_1$ and hence the inequality holds if and only if $T_H > T_E$, which is true by Assumption 4.

(ii) is analogous to (i), using part (ii) of Lemma A.1 instead of (i). \square

The next lemma is used in the proof of Proposition 3.

Lemma A.2. $Pr(0 < u_P - u_\$ < \delta | 0 < u_P - u_\$) = Pr(0 < u_\$ - u_P < \delta | 0 < u_\$ - u_P)$.

Proof. First note that since u_P and $u_\$$ are i.i.d, $u_P - u_\$$ and $u_\$ - u_P$ are identically distributed and $Pr(u_\$ - u_P > 0) = Pr(u_\$ - u_P < 0) = 1/2$. Then $Pr(0 < u_\$ - u_P < \delta | 0 < u_\$ - u_P) = \frac{Pr(0 < u_\$ - u_P < \delta)}{Pr(u_\$ > u_P)} = \frac{Pr(0 < u_P - u_\$ < \delta)}{Pr(u_P > u_\$)} = Pr(0 < u_P - u_\$ < \delta | 0 < u_P - u_\$)$. \square

Proof of Proposition 3. To shorten notation let $\Delta_0 = Pr(CE_\$ > CE_P | 0 < u_\$ - u_P < \delta)$, $\Delta_1 = Pr(CE_\$ > CE_P | u_\$ - u_P > \delta)$, $\Delta_2 = Pr(CE_P > CE_\$ | 0 < u_P - u_\$ < \delta)$, $\Delta_3 = Pr(CE_P > CE_\$ | u_P - u_\$ > \delta)$, $P_1 = Pr(CE_\$ > CE_P | u_\$ > u_P)$, $P_2 = Pr(CE_P > CE_\$ | u_P > u_\$)$. Let also P^δ be the probability given in Lemma A.2.

With these definitions, we have that $P_1 = \Delta_0 P^\delta + \Delta_1 (1 - P^\delta)$ and $P_2 = \Delta_2 P^\delta + \Delta_3 (1 - P^\delta)$.

We obtain $E[DT_C | CE_\$ > CE_P, c(P, \$) = \$] = \frac{1}{P_1} [\Delta_0 P^\delta T_H + \Delta_1 (1 - P^\delta) T_E]$ and $E[DT_C | CE_P > CE_\$, c(P, \$) = P] = \frac{1}{P_2} [\Delta_2 P^\delta T_H + \Delta_3 (1 - P^\delta) T_E]$. This yields.

$$\begin{aligned} E[DT_C | CE_\$ > CE_P, c(P, \$) = \$] &> E[DT_C | CE_P > CE_\$, c(P, \$) = P] \\ \Leftrightarrow P^\delta (1 - P^\delta) T_H [\Delta_0 \Delta_3 - \Delta_1 \Delta_2] &> (1 - P^\delta) P^\delta T_E [\Delta_0 \Delta_3 - \Delta_1 \Delta_2] \end{aligned}$$

Since $T_H > T_E$ by Assumption 4, the claim holds if $\Delta_1 \Delta_2 < \Delta_0 \Delta_3$. The rest of the proof is devoted to establish this fact. For this, we rely on ideas taken from Wijsman (1985).

First, note that

$$\begin{aligned}\Delta_0 &= \frac{1}{Pr(0 < u_S - u_P < \delta)} \int_0^\delta Pr(\Delta\zeta < K + s)h(s)ds, \\ \Delta_1 &= \frac{1}{Pr(u_S - u_P > \delta)} \int_\delta^\infty Pr(\Delta\zeta < K + s)h(s)ds, \\ \Delta_2 &= \frac{1}{Pr(0 < u_P - u_S < \delta)} \int_0^\delta Pr(\Delta\zeta < -K + s)h(s)ds, \text{ and} \\ \Delta_3 &= \frac{1}{Pr(u_P - u_S > \delta)} \int_\delta^\infty Pr(\Delta\zeta < -K + s)h(s)ds.\end{aligned}$$

Now let $f_1(s) := Pr(\Delta\zeta < K + s)$, $f_2(s) = Pr(\Delta\zeta < -K + s)$,

$$g_1(s) = \begin{cases} h(s) & \text{if } s \in]\delta, \infty[, \\ 0 & \text{otherwise,} \end{cases} \text{ and } g_2(s) = \begin{cases} h(s) & \text{if } s \in [0, \delta], \\ 0 & \text{otherwise.} \end{cases}$$

As u_S and u_P are i.i.d $Pr(0 < u_S - u_P < \delta) = Pr(0 < u_P - u_S < \delta)$ and $Pr(u_S - u_P > \delta) = Pr(u_P - u_S > \delta)$ and hence showing that $\Delta_1\Delta_2 < \Delta_0\Delta_3$ boils down to showing that

$$\int_0^\infty f_1(s)g_1(s)ds \int_0^\infty f_2(s)g_2(s)ds < \int_0^\infty f_2(s)g_1(s)ds \int_0^\infty f_1(s)g_2(s)ds.$$

To see that this is true note that

$$\begin{aligned}2 \left(\int_0^\infty f_1(s)g_1(s)ds \int_0^\infty f_2(s)g_2(s)ds - \int_0^\infty f_2(s)g_1(s)ds \int_0^\infty f_1(s)g_2(s)ds \right) \\ = \int_0^\infty \int_0^\infty F(x, y)G(x, y)dxdy,\end{aligned}$$

where $F(x, y) = f_1(x)f_2(y) - f_1(y)f_2(x)$ and $G(x, y) = g_1(x)g_2(y) - g_1(y)g_2(x)$. Further,

$$\left(\frac{f_1}{f_2} \right)'(s) = \frac{q(K + s)Pr(\Delta\zeta < -K + s) - Pr(\Delta\zeta < K + s)q(-K + s)}{(Pr(\Delta\zeta < -K + s))^2},$$

where q is the density of $\Delta\zeta$. Then $(\frac{f_1}{f_2})'(s) < 0$ since $0 < q(K + s) \leq q(-K + s)$ and $Pr(\Delta\zeta < -K + s) < Pr(\Delta\zeta < K + s)$ by Assumptions 3

and 3.¹⁷ Thus $\frac{f_1}{f_2}$ is strictly decreasing and hence $F(x, y) > 0$ if $x < y$ and $F(x, y) < 0$ if $y < x$ (of course, $F(x, y) = 0$ if $x = y$). By construction $G(x, y) > 0$ if $(x, y) \in]\delta, \infty[\times]0, \delta]$, $G(x, y) < 0$ if $(x, y) \in [0, \delta] \times]\delta, \infty[$, and $G(x, y) = 0$ otherwise. Hence $F(x, y)G(x, y) \leq 0$ for all $(x, y) \in \mathbb{R}_+ \times \mathbb{R}_+$ and $F(x, y)G(x, y) < 0$ for all $(x, y) \in]\delta, \infty[\times]0, \delta] \cup [0, \delta] \times]\delta, \infty[$. This implies that $\int_0^\infty \int_0^\infty F(x, y)G(x, y)dx dy < 0$ which proves the claim. \square

¹⁷The fact that $q(K+s) \leq q(-K+s)$ follows by unimodality. If $s \geq K$ then $q(-K+s) \leq q(K+s)$ since q is nonincreasing in the positive domain. If $s < K$ then $q(-K+s) > q(-K-s) = q(K+s)$ since q is nondecreasing in the negative domain and symmetric.

Appendix 4.B: Lotteries


Table 4.5: The lottery pairs.

Lottery pair	P-bet				S-bet			
	Prob	Outc	EV	StdDev	Prob	Outc	EV	StdDev
1	0.44	7	4.20	3.536	0.36	7	3.80	3.536
2	0.40	8	4.40	4.243	0.40	7	4.00	3.536
3	0.82	11	9.38	6.364	0.10	48	6.60	32.527
4	0.94	9	8.58	4.950	0.20	30	7.60	19.799
5	0.80	11	9.20	6.364	0.20	24	6.40	15.556
6	0.90	10	9.20	5.657	0.30	22	8.00	14.142
7	0.60	15	9.80	9.192	0.21	23	6.41	14.849
8	0.80	10	8.40	5.656	0.40	15	7.20	9.192
9	0.89	6	5.56	2.828	0.11	36	5.74	24.042
10	0.81	6	5.24	2.828	0.19	18	5.04	11.314
11	0.97	12	11.70	7.071	0.31	34	11.92	22.627
12	0.94	8	7.64	4.242	0.39	16	7.46	9.899
13	0.82	9	7.74	4.243	0.50	13	7.50	7.778
14	0.87	7	6.35	3.536	0.50	11	6.50	6.364
15	0.68	7	5.40	2.828	0.20	25	6.60	16.971
16	0.79	8	6.74	2.828	0.30	24	8.60	15.556
17	0.80	6	5.20	2.828	0.40	18	8.40	11.314
18	0.90	6	5.60	2.828	0.30	18	6.80	11.314
19	0.60	9	6.20	4.950	0.45	17	8.75	10.607
20	0.60	10	6.80	5.657	0.40	16	7.60	9.899

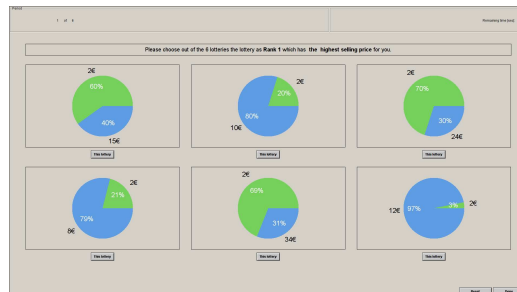
Note: All lotteries pay an amount of €2 with the corresponding converse probabilities. The table shows for each P-bet and S-bet within a pair the probability with which the outcome occurs, the expected value and the standard deviation. Lottery pairs 1 and 2, containing strictly dominated bets, were only used in Experiment 1 as a basic rationality check.

Appendix 4.C: Screenshots

(a) Pricing phase (Experiment 1)



(b) Ranking phase (Experiment 2)



(c) Choice phase (Experiments 1 and 2)




Figure 4.7: Screen displays.

References Chapter 4

- ACHTZIGER, A., AND C. ALÓS-FERRER (2013): “Fast or Rational? A Response-Times Study of Bayesian Updating,” *Management Science*, forthcoming.
- ALÓS-FERRER, C., D.-G. GRANIĆ, F. SHI, AND A. K. WAGNER (2012): “Choices and Preferences: Evidence from Implicit Choices and Response Times,” *Journal of Experimental Social Psychology*, 48(6), 1336–1342.
- ALÓS-FERRER, C., AND F. STRACK (2013): “From Dual Processes to Multiple Selves: Implications for Economic Behavior,” *Journal of Economic Psychology*, forthcoming.
- BALL, L. J., N. BARDSLEY, AND T. ORMEROD (2012): “Do preference reversals generalise? Results on ambiguity and loss aversion,” *Journal of Economic Psychology*, 33(1), 48–57.
- BARGH, J. A., AND T. L. CHARTRAND (2000): “The Mind in the Middle: A Practical Guide to Priming and Automaticity Research.,” in *Handbook of Research Methods in Social and Personality Psychology*, ed. by H. T. Reis, and C. M. Judd, pp. 253–285. Cambridge University Press, NY.
- BATEMAN, I., B. DAY, G. LOOMES, AND R. SUGDEN (2007): “Can ranking techniques elicit robust values?,” *Journal of Risk and Uncertainty*, 34(1), 49–66.
- BECKER, G. M., M. H. DEGROOT, AND J. MARSCHAK (1964): “Measuring Utility By A Single Response Sequential Method,” *Behavioral Science*, 9(3), 226–232.
- BEM, D. (1967): “Self-perception: An alternative interpretation of cognitive dissonance phenomena,” *Psychological Review*, 74(3), 183–200.
- BLAVATSKY, P. R. (2009): “Preference reversals and probabilistic decisions,” *Journal of Risk and Uncertainty*, 39(3), 237–250.

- BLEICHRODT, H., AND J. L. PINTO PRADES (1994): “New Evidence of Preference Reversals in Health Utility Measurement,” *Health Economics*, 18(6), 713–726.
- BREHM, J. W. (1956): “Postdecision changes in the desirability of alternatives,” *Journal of Abnormal and Social Psychology*, 52(3), 384–389.
- BUTLER, D. J., AND G. LOOMES (2007): “Imprecision as an Account of the Preference Reversal Phenomenon,” *The American Economic Review*, 97(1), 277–297.
- CASEY, J. T. (1991): “Reversal of the Preference Reversal Phenomenon,” *Organizational Behavior and Human Decision Processes*, 48(2), 224–251.
- (1994): “Buyers’ Pricing Behavior for Risky Alternatives: Encoding Processes and Preference Reversals,” *Management Science*, 40(6), 730–749.
- CHABRIS, C. F., D. I. LAIBSON, C. L. MORRIS, J. P. SCHULDT, AND D. TAUBINSKY (2009): “The allocation of time in decision-making,” *Journal of the European Economic Association*, 7(2), 682–637.
- CUBITT, R. P., A. MUNRO, AND C. STARMER (2004): “Testing Explanations of Preference Reversal,” *The Economic Journal*, 114(497), 709–726.
- DELQUIÉ, P. (1993): “Inconsistent Trade-Offs Between Attributes: New Evidence in Preference Assessment Biases,” *Management Science*, 39(11), 1382–1395.
- EVANS, J. S. B. T. (2008): “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition,” *Annual Review of Psychology*, 59(1), 255–278.
- FESTINGER, L. (1957): *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford.
- FISCHBACHER, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10(2), 171–178.

- FISCHER, G. W., Z. CARMON, D. ARIELY, AND G. ZAUBERMAN (1999): “Goal-Based Construction of Preferences: Task Goals and the Prominence Effect,” *Management Science*, 45(8), 1057–1075.
- GOLDSTEIN, W. M., AND H. J. EINHORN (1987): “Expression Theory and the Preference Reversal Phenomena,” *Psychological Review*, 94(2), 236–254.
- GREETHER, D. M., AND C. R. PLOTT (1979): “Theory of Choice and the Preference Reversal Phenomenon,” *The American Economic Review*, 69(4), 623–638.
- HERSHEY, J. C., AND P. J. H. SCHOEMAKER (1985): “Probability versus Certainty Equivalence Methods in Utility Measurement: Are They Equivalent?,” *Management Science*, 31(10), 1213–1231.
- HEY, J. D. (2001): “Does Repetition Improve Consistency?,” *Experimental Economics*, 4(1), 5–54.
- JARCHO, J. M., E. T. BERKMAN, AND M. D. LIEBERMAN (2011): “The neural basis of rationalization: Cognitive dissonance reduction during decision-making,” *Social Cognitive and Affective Neuroscience*, 6(4), 460–467.
- JOHNSON, E. J., AND D. A. SCHKADE (1989): “Bias in Utility Assessments: Further Evidence and Explanations,” *Management Science*, 35(4), 406–424.
- LINDMAN, H. R. (1971): “Inconsistent Preferences Among Gambles,” *Journal of Experimental Psychology*, 89(2), 390–397.
- MAAFI, H. (2011): “Preference Reversals Under Ambiguity,” *Management Science*, 57(11), 2054–2066.
- MOFFATT, P. G. (2005): “Stochastic Choice and the Allocation of Cognitive Effort,” *Experimental Economics*, 8(4), 369–388.

- OLIVER, A. (2013): “Testing Procedural Invariance in the Context of Health,” *Health Economics*, 22(3), 272–288.
- PIOVESAN, M., AND E. WENGSTRÖM (2009): “Fast or Fair? A Study of Response Times,” *Economics Letters*, 105(2), 193–196.
- PURKAYASTHA, S. (1998): “Simple proofs of two results on convolutions of unimodal distributions,” *Statistics & Probability Letters*, 39(2), 97–100.
- ROTTENSTREICH, Y., S. SOOD, AND L. BRENNER (2007): “Feeling and Thinking in Memory-Based versus Stimulus-Based Choices,” *Journal of Consumer Research*, 33(3), 461–469.
- RUBINSTEIN, A. (2007): “Instinctive and Cognitive Reasoning: A Study of Response Times,” *Economic Journal*, 117(523), 1243–1259.
- SCHKADE, D. A., AND E. J. JOHNSON (1989): “Cognitive Processes in Preference Reversals,” *Organizational Behavior and Human Decision Processes*, 44(2), 203–231.
- SCHMIDT, U., AND J. D. HEY (2004): “Are Preference Reversals Errors? An Experimental Investigation,” *The Journal of Risk and Uncertainty*, 29(3), 207–218.
- SCHNEIDER, W., AND R. M. SHIFFRIN (1977): “Controlled and Automatic Human Information Processing: I. Detection, Search, and Attention,” *Psychological Review*, 84(1), 1–66.
- SEIDL, C. (2002): “Preference Reversal,” *Journal of Economic Surveys*, 16(5), 621–655.
- SHAROT, T., B. DE-MARTINO, AND R. J. DOLAN (2009): “How Choice Reveals and Shapes Expected Hedonic Outcome,” *The Journal of Neuroscience*, 29(12), 3760–3765.
- SHULTZ, T. R., E. LÉVEILLÉ, AND M. R. LEPPER (1999): “Free choice and cognitive dissonance revisited: Choosing “lesser evils” versus “greater goods,”” *Personality and Social Psychology Bulletin*, 25(1), 40–48.

- SLOVIC, P., AND S. LICHTENSTEIN (1968): “Relative Importance of Probabilities and Payoffs in Risk Taking,” *Journal of Experimental Psychology Monograph*, 78(3, Part 2), 1–18.
- STALMEIER, P. F. M., P. P. WAKKER, AND T. G. G. BEZEMBINDER (1997): “Preference Reversals: Violations of Unidimensional Procedure Invariance,” *Journal of Experimental Psychology: Human Perception and Performance*, 23(4), 1196–1205.
- STRACK, F., AND R. DEUTSCH (2004): “Reflective and Impulsive Determinants of Social Behavior,” *Personality and Social Psychology Review*, 8(3), 220–247.
- TRAUTMANN, S. T., F. M. VIEIDER, AND P. P. WAKKER (2011): “Preference Reversals for Ambiguity Aversion,” *Management Science*, 57(7), 1320–1333.
- TVERSKY, A., S. SATTATH, AND P. SLOVIC (1988): “Contingent Weighting in Judgment and Choice,” *Psychological Review*, 95(3), 371–384.
- TVERSKY, A., P. SLOVIC, AND D. KAHNEMAN (1990): “The Causes of Preference Reversal,” *The American Economic Review*, 80(1), 204–217.
- WEBER, E. U., AND E. J. JOHNSON (2009): “Mindful Judgment and Decision Making,” *Annual Review of Psychology*, 60, 53–85.
- WIJSMAN, R. A. (1985): “A Useful Inequality on Ratios of Integrals, With Application to Maximum Likelihood Estimation,” *Journal of the American Statistical Association*, 80(390), 472–475.
- WILCOX, N. T. (1993): “Lottery Choice: Incentives, Complexity, and Decision Time,” *Economic Journal*, 103(421), 1397–1417.
- (1994): “On a Lottery Pricing Anomaly: Time Tells the Tale,” *Journal of Risk and Uncertainty*, 8(7), 311–324.