Game Transformations and Game Equivalence

Boudewijn P. de Bruin*

Abstract

In 1952, Thompson defined four transformations on extensive games, and proved that they leave the reduced normal form intact (a so–called "adequacy" theorem). Kohlberg&Mertens, in 1986, proposed two additional game transformations. In this paper, I will state and prove adequacy results for all six transformations in a uniform manner. Then, a number of well–known solution concepts will be surveyed. Finally, I will examine the (in)variance of the solution concepts under the game transformations, so as to determine whether they induced a plausible notion of game equivalence.

1 Introduction

Game theory is often defined as the analysis of conflict—a description, however broad and misleading, that sets the stage for many of its applications in fields as diverse as the political sciences and evolutionary biology. More specifically, game theory is regarded as the study of solutions to conflicts. That is, a game theorist is expected to come up with a solution to particular games—this is the euphemism used to rid the discipline of its bellicose character—that are in some sense "optimal" to all combatants, where unsurprisingly this latter notion is rather opaque. Games and solution concepts being the primary objects of game theoretical studies, we ought to be amazed by the scant attention that has been paid to an issue that, not only at first sight, seems to be fundamental to the whole field, viz., the question what a game actually is, or stated differently: when are two games equivalent?

The motivation underlying this paper is that a possible answer to this question may come from logic. Yet, being a game theoretical rather than a logical paper, it is only a preliminary discussion, a survey, that is, of what

*I would like to thank Professor Johan van Benthem for his encouraging discussions, and for suggesting the topic to me.

concepts of game equivalence can be distilled from the existing game theoretical literature. Logic as such will not return until the last section. First, I will introduce a notion of equivalence that is induced by a set of game transformations that can be shown to respect the reduced normal form (rnf) of extensive games. I will discuss three theorems showing that rnf and game transformations are inextricably bound up with each other. Then, I will survey various solution concepts needed to evaluate the equivalence notions of the preceding section. Finally, we will see that none of the more advanced solution concepts is invariant under these equivalence notions. Nevertheless, a striking theorem will put this arguably hapless situation into the right perspective.

Let us fix some notation and terminology. The distinction between one–shot games, and games where players have the opportunity to (re)consider the game *while it is being played*, will be expressed by opposing "strategic games" to "extensive games." A strategic game will typically be denoted by $G = (N, (A_i)_i, (\succeq_i)_i)$, where N is a finite set of players, A_i the set of actions open to player i (from which he has to choose only one in this one–shot setting), and $\succeq_i$ is i's preference relation over the set of action profiles (N–tupels of actions taken by all players). When preference relations are represented by means of utility functions mapping strategy profiles to the reals, it will always be assumed that the Von Neumann–Morgenstern conditions be satisfied so as to enable us to reason sensibly about the *expected* utility. I will try to be as consistent as possible to distinguish mere outcomes (that what is the result of a particular play, e.g., a "win," or a candy), from payoffs, or utilities. That is, payoffs (utilities) are always assumed to satisfy the Von Neumann–Morgenstern conditions, thus expressing the players' divergent attitudes towards risk.

In general I will use Osborne&Rubinstein's way to describe extensive games, although at some points deviations will appear. So, an extensive game (of imperfect information, and with moves by nature) will be denoted $G = (N, H, \iota, f_c, (\mathcal{I}_i)_i, (\succeq_i)_i)$, where N is the set of players, H the set of histories, ι the player function assigning to each non–terminal history the player whose turn it is, f_c a probability distribution over the actions open to nature, $\mathcal{I}_i$ the set of information sets where it is i's turn, and $\succeq_i$ the preference relation of player i over terminal nodes. Under the same proviso, we will often use utility functions to represent the preference relations. (Osborne&Rubinstein define the preference relation on *lotteries* over the set of terminal nodes. I will assume that this is not accomplished until we have considered a utility function satisfying the Von Neumann and Morgenstern conditions.) Sometimes, instead of referring to the histories it will be more natural to refer to the nodes they consist of; then, the ordering implicit in

H will remain implicit.

2 Thompson's Transformation Rules

By some authors accepted as completely self–evident, by others rejected as lacking any theoretical or empirical justification, the "reduced normal form" is the most traditional game theoretical approach to game equivalence. In this section, I will survey how this essentially one–shot, or strategic, concept has fine extensive correspondents in sets of game transformations respecting reduced normal form. The study of such transformations has at least two reasons which have not always had the same priorities in the history of game theory. The first reason, to be found in Thompson's 1952 paper, is to find an effective (and, if possible, not too complex) procedure to obtain the rnf of any (large) extensive game—note the year of issue in relation to the history of automated computing. Some thirty years later a more theoretical interest arrose in game transformations. Indeed, in discussions about refinements of the Nash equilibrium concept (having started in the 1970's), game theorists often referred to particular game transformations in order to show that a particular refinement did, or did not, leave the so–called "strategic features" of a game untouched—implicitly assuming that the transformations were completely inessential. In 1986, Kohlberg&Mertens enlarged the set of four Thompson transformations with two new ones, which extended the domain of application to games with moves by nature. Whereas Thompson provided a sketch of the proof of the theorem identifying rnf–equivalence and (Thompson–)transformation–equivalence, Kohlberg&Mertens' paper focusses on other issues. In this section, I will present three such "adequacy" results in a uniform manner. However, since the main goal of this paper is to systematize game theoretical approaches to the notion of game equivalence, I will not go into too much detail.

As I said, game transformations have not always been used in a fully critical way to defend, or to refute, the plausibility of several solution concepts. A very interesting recent paper discussing the (in)acceptability of such arguments is Elmes&Reny's 1994 paper. Explicitly described as an attempt to "sharpen the focus" of the debate "whether or not the normal form is sufficient for all strategic considerations" [ER94, p. 2], credibility is given to the claim that the proposition that all four Thompson transformations are strategically inessential is inconsistent by arguing that one of them is strategically inessential only if another one is not.

To sum up, then, the adequacy theorems that will be proven below relate two essentially different approaches to game equivalence: the static one–shot view of the reduced normal form, and the dynamic extensive transformational view.

2.1 The Normal Form

Before exemplifying Thompson's transformation rules, I will give more details about the notion of "normal form" of an extensive game. The main idea of the normal form is to capture an extensive game G in a strategic format by taking the strategies open to each player (functions mapping each (!) of a player's decision nodes to a move open at that node) as *actions* of a strategic game, and to define preference relations in an obvious way. For reasons of convenience we will use utility functions to describe the extensive game.

Definition 1 (Normal Form) *Let $G = (N, H, \iota, f_c, (\mathcal{I}_i)_i, (u_i)_i)$ be an extensive game. Then, $nf(G) = (N, (S_i)_i, (\succeq_i)_i)$ is the normal form of G iff S_i is the set of strategies open to player i in G, and $\succeq_i$ is defined by*

$$s \succeq_i t \Leftrightarrow u_i(s) \geq u_i(t),$$

where s, and t are action profiles in game $nf(G)$ (that is, strategy profiles in game G).

As an illustration, consider extensive game G in figure 1, and its normal form $nf(G)$ in figure 2. Here a fairly standard representation of strategic games by means of matrices is applied (which works for games with less than 4 players). Note, that player I has *two* information sets in G. Also, note that the expressions used to denote the actions of the strategic game bear obvious resemblance to the strategies they originate from.

In many cases, the normal form of a game can be made smaller by removing redundant information thus obtaining the "reduced normal form." In figure 2, e.g., we would remove action LU, or (exclusive or!) action LD, resulting in the strategic game of figure 3. This, in fact, means that equivalence classes of actions are considered, which is reflected in the following definition.

Definition 2 (Reduced Normal Form) *Let $G = (N, H, \iota, f_c, (\mathcal{I}_i)_i, (u_i)_i)$ be an extensive game, and $nf(G) = (N, (S_i)_i, (\succeq_i)_i)$ its normal form. Define actions s_i and t_i of player i to be equivalent iff for all (!) i–deleted action profiles s_{-i} (i.e., profiles of the form $(s_0, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$) it is true that $(s_{-i}, s_i) \sim_i (s_{-i}, t_i)$. Then, let T_i be the set of equivalences classes of S_i modulo $\sim_i$, and let $\succeq_i'$ be player i's preference relation over profiles of such equivalence classes induced by $\succeq_i$. Then, $rnf(G) = (N, (T_i)_i, (\succeq_i')_i)$ is the reduced normal form of G.*

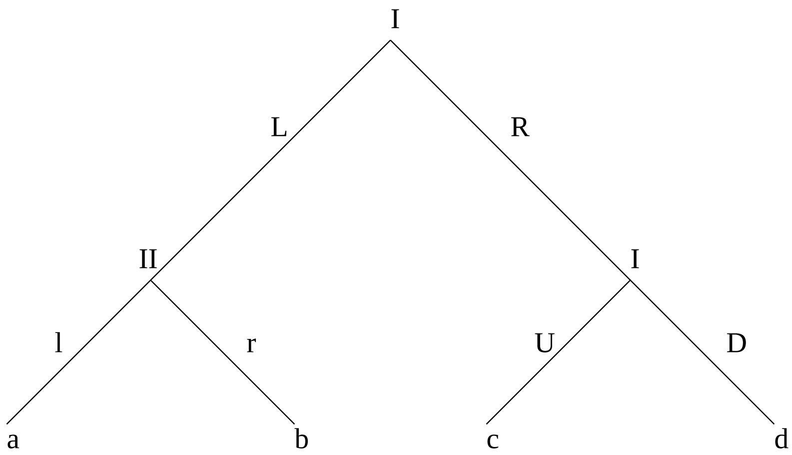

Figure 1: Game G

I \ II	l	r
LU	a	b
LD	a	b
RU	c	c
RD	d	d

Figure 2: The Normal Form of G

I \ II	l	r
L’	a	b
RU	c	c
RD	d	d

Figure 3: The Reduced Normal Form of G

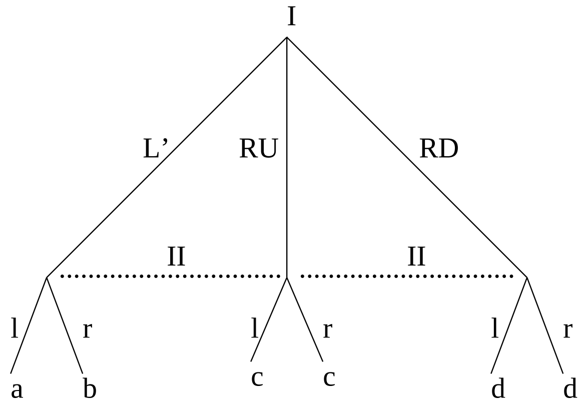

Figure 4: The Extensive Representation of the Reduced Normal Form of G

Notice that both the normal form and the reduced normal form can also be described in terms of utility functions. This more convenient perspective will be necessary in section 2.4.

Obviously, the matrix representation of strategic games will give us a fairly easy way of drawing the reduced normal form. However, in this paper it will be more convenient to display it by means of a graph. In a sense, then, the reduced normal form is depicted as if it were an extensive game; yet by means of a careful choice of information sets, the one–shot character of a strategic game can be preserved in the extensive setting. Figure 4, for instance, is the extensive representation of the reduced normal form of figure 3. Every player has been given exactly one information set so as to ensure that the players' choices of actions are independent of each other. More generally, the extensive representation of an N–player strategic game is an N–tiered tree generated by the information set of the first player. For every player it contains exactly one information set the number of branches departing from it being the number of actions open to the player in the strategic game. Furthermore, every information set has to be maximal in the sense that it intersects with every terminating path of the game tree.

2.2 Thompson's Transformations

Thompson's game transformations appear in different forms in the literature. Thompson himself gave formal definitions of each of them which he illustrated by means of simple examples. The examples in this paper (due to Osborne&Rubinstein) are more general, and hence more informative. I will briefly discuss the transformations separately. For all Thompson transformations, the original game is depicted in figure 5. Note that player I has three information sets, while II has only one.

By means of ID, the number of information sets of a player can be decreased. While in the original game, player I had perfect recall (that is, he was exhaustively informed about *his* previous actions), ID changes this into a game with imperfect recall. Although there are several reasons why this could worry us, it need not detain us here. The reader is referred to the above mentioned paper by Elmes&Reny. They show that if the definition of Thompson's ADD transformation is slightly amended, ID becomes redundant in the sense that every two games of *perfect* recall that are equivalent modulo Thompson's original transformations remain equivalent modulo the revised set of transformations (and *vice versa*).

The transformation ADD exploits the fact that in the original game player I has two successive choices of action by prying open the second one and filling it with a superfluous move of player II. Evidently, it will leave player I cold whatever action his opponent goes for at this added node. Player II in turn, one might say, is deluded with images of larger influence, since the very distribution of the outcomes (or payoffs) in the leftmost part of the changed game makes sure that the added moves are really superfluous. This also shows that the introduction of imperfect information on the part of player II is inconsequential. Remark that the number of information sets does not change.

COA is a second use of player I's consecutive moves. It plainly lumps them together eliminating one information set. Yet, it does not introduce or remove imperfect information.

Finally, INT takes advantage of the fact that in the original game player I is imperfectly informed about player II's moves U and D. Player I can force "c or e" (by playing l), and "d or f" (by r), and similarly player II can force "c or d" and "e or f." This holds for the transformed game as well. In other words, the *control* of the game both players have does not change.

One might wonder if Thompson's transformations do not rather transform game *forms* than games *proper*, since it looks as if the preference relation is not really used at all. This is not completely true. In fact, what we need to know of a game is whether outcomes of two different runs are equally preferable for all players. For instance, we saw that the *pointe* of ADD was that the outcomes in the leftmost part appear twice, rendering player II's choice of move at the added node completely irrelevant. It will be evident that different assignment of the outcomes would result in two games with different reduced normal forms. What we need, then, is more than the game form. Yet, we need not be fully informed about the relative preferences of the players; that is, in order to determine the reduced normal form it is sufficient to know which outcomes are equally preferable to all players. The corresponding notion is the "game structure," due to Thompson. In terms of utilities (so as to enable application to game structures with chance moves), two games have identical game structure iff at most the utility functions differ in such a way that if in one game tree two paths have the same utility, this is also true in the other game tree. The three adequacy results in this section can then be refined by substituting "games" by "game structures." This is, however, slightly far–fetched.

Before stating Thompson's adequacy theorem, let me note that the above "argumentation" for the transformations is completely irrelevant from a combinatorial perspective. If we set out to obtain a set of transformations that leave the reduced normal form of a game intact, the only proper justification of a particular transformation is that it does in fact not change the reduced normal form. As I noted above, however, the intuitive plausibility of the transformations (of which the above paragraphs give some impression) is often used to defend the reduced normal form. Altogether we can see, then, that Thompson's theorem (Theorem 3) is the pivot of both arguments: starting from the plausibility of the transformations one argues for the reduced normal form, and conversely, if one takes the reduced normal form as given, the transformations follow immediately. We can now state the theorem.

Theorem 3 (Thompson, 1952) *Let G and G' be two extensive games. Then the following two propositions are equivalent.*

(i) G and G' are equivalent modulo a finite sequence of transformations from: ADD, ID, INT, COA.

(ii) G and G' have the same rnf.

A note about my usage of "modulo." Condition (i) of the theorem requires that there be a finite sequence of extensive games $G_0, \ldots, G_n$, such that $G_0 = G$, $G_n = G'$, and for all $0 \leq i < n$ either G_i can be obtained from G_{i+1} by means of one application of one of Thompson's four transforms, or *vice versa*. So in other words the *direction* of the transformations is completely inconsequential.

Proof. For formal details of this proof, the reader is referred to Thompson's original paper, or to Elmes&Reny's 1994 paper. Continuing my pictorial exposition of Thompson's transformations my proof will lack some generality as it refers to the above examples.

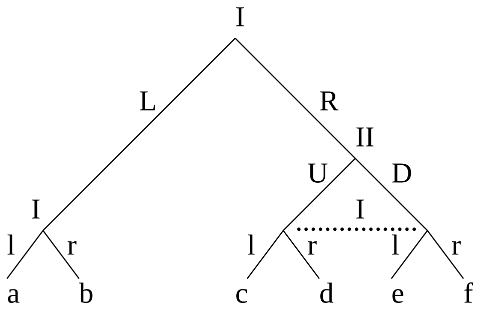

Figure 5: The Original to be Transformed

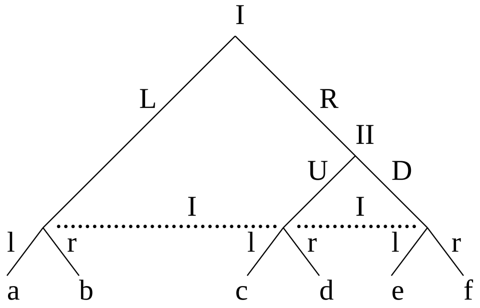

Figure 6: Inflation–Deflation (ID)

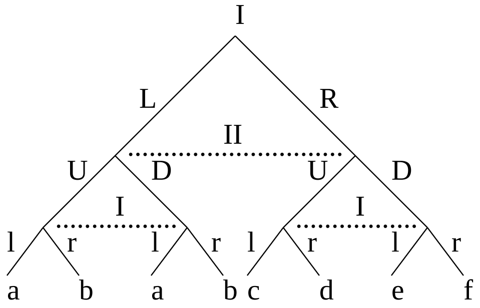

Figure 7: Addition–of–a–Superfluous–Move (ADD)

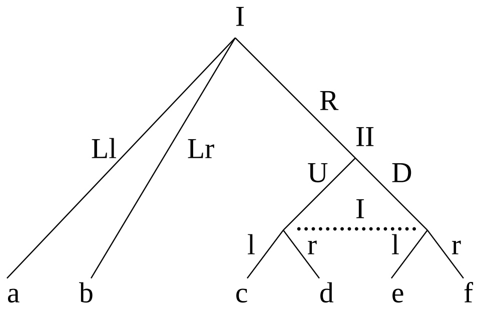

Figure 8: Coalescing–Moves (COA)

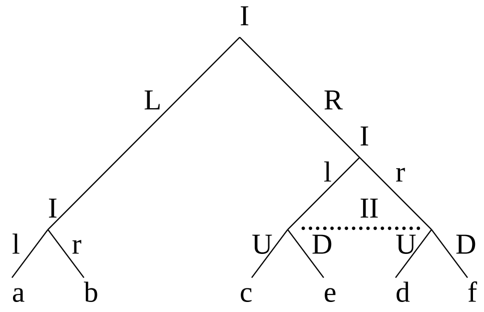

Figure 9: Interchanging–Moves (INT)

The "left–to–right" direction is proved by induction on the number of transformations. The underlying idea is that many actions specified by a strategy (profile) do not affect the outcome of a game.

Games that differ *qua* INT, or *qua* ADD, have exactly the same rnf since these transformations do not introduce any new strategies. COA and ID decrease the number of information sets by one, thus only renaming strategies. As a case in point, consider an application of ID as exemplified in the above figures. All strategies in the transformed game have analogs in the original game. For instance, the strategy for player I in the transformed game that prescribes L at his first information set, and r at his second, has Lrr as its pre–image in the original game. Hence, no strategies are *added* by transforming the game. But, strategies might be *lost*. Yet, this will not happen either because the strategies in the original game that do *not* mirror strategies in the transform, will be deleted in determining the rnf. This is so, because in the original game they specify two different actions at information sets that will fuse in the transformation, and one of those specified actions is not relevant for the outcome of the game.

For the other direction, I will sketch a procedure transforming an arbitrary extensive game into an extensive game that is the graph representation of an rnf according to the above definition.

- By using ADD and/or COA finitely often, we increase or decrease the number of branches of each decision node (except the starting point of the game) to two.[1]

[1]Thompson claims to be able to manage this by means of COA only. This is only true if we abstract from the *direction* in which the transformations are applied, i.e., if we consent to a sequence of games $G_0, \ldots, G_n$ for which either G_{i+1} is a transform of G_i, or *vice versa*. This rather unrefined "vice versa" can be eliminated if we allow for the application of ADD, too. This also emphasizes the differences as well as the similarities with Church–Rosser–like theorems in λ–calculus.

- Using ADD and/or ID every information set in the game is expanded to a maximal one. Induction on a particular ordering of the information sets gives this a neat algorithmic format.

- Applying COA to minimize the number of actions open to each player, and INT to obtain the right order of the N levels of the game, we obtain the rnf. (Note: COA and INT leave maximality of the information sets untouched.) $\dashv$

2.3 Kohlberg&Mertens' First Extension: COAc

In the beginning of this section we saw how to think of the reduced normal form of games with moves by nature in terms of expected utility: strategy profiles of the extensive game are mapped to the expected payoffs. However, although application of one or more of Thompson's transformations to games with moves by nature does not change the reduced normal form of the game, there is no way to actually arrive at the *extensive representation* of the reduced normal form, for to that purpose one would need to eliminate nature as a player. To see what that exactly means and how we can take care of the problem, consider the game in figure 10, which is obviously not a tree displaying a reduced normal form. The desired format being the game in figure 11, Kohlberg&Mertens simply define a game transformation COAc (my terminology) carrying the former game into the latter. Note that indeed both games share the reduced normal form. Notice, too, that this transformation only removes chance moves at the bottom of the game. If nature plays a role higher up in the game tree, the four Thompson transformations have to be used to transport nature to the bottom of the game. This will be important in the proof of the adequacy theorem below.

Theorem 4 *Let G and G' be two extensive games with chance moves. Then the following two propositions are equivalent.*

(i) G and G' are equivalent modulo a finite sequence of transformations taken from: ADD, ID, INT, COA, COAc.

(ii) G and G' have the same reduced normal form.

Proof. If two games with chance moves are equivalent modulo one or more of the first four transformations, no moves by nature have been removed; possibly, one or more chance moves have been transformed as if they were actions to regular players (where chance only has singleton information sets). From the preceding theorem, together with the definition of expected utility,

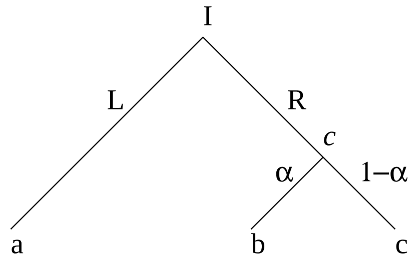

Figure 10: The Original to be Transformed

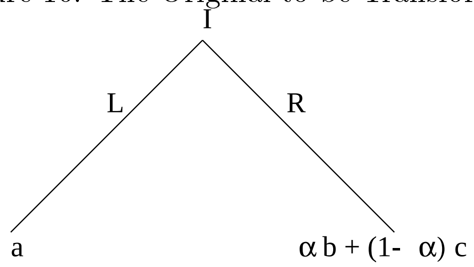

Figure 11: Coalescing–of–a–Chance–Move (COAc)

it follows that the games share the rnf. The particular tailoring of COAc (so as to fit the revision of the reduced normal form!) takes care of the other case.

Conversely, we only need to give a recipe for transforming a game with chance moves to its rnf. Dealing with nature c as a regular player (under the above proviso about information sets), we first relegate its moves to the bottom of the game by means of the techniques touched upon in the foregoing proof. Then (after a possible refurbishing of the top–part of the game), by COAc we delete the chance moves, and replace them by the expected payoffs. This finishes the procedure. ⊣

2.4 Kohlberg&Mertens' Second Extension: ADDc

I will now deal with the second transformation proposed by Kohlberg&Mertens (see figures 12–13). This transformation introduces (superfluous) moves by nature corresponding to "convex combinations" of pure strategies of a player of the game. Since this transformation is a bit more complicated *qua* structure then the first one, I will dwell upon it a little longer. Kohlberg&Mertens marshal the following motivation. For a particular extensive game G, a player i has particular pure strategies. However, in most game theoretical reasoning one prefers to consider the *mixed* strategies, i.e., probability distributions over pure strategies, since, for instance, whereas some games lack pure strategy

Nash equilibria, every game has at least one mixed strategy Nash equilibrium. Since for any profile of strategies (of which some may be mixed) the expected utility can be easily calculated, mixed strategies of an extensive game can be very naturally *incorporated* in its normal form. Indeed, what will be added to the normal form are simply convex combinations[2] (of the payoffs) of pure strategies, and if we take this to its extreme by allowing for the addition of all (infinitely many) convex combinations, any Nash equilibrium of this enlarged normal form will be a (possibly mixed strategy) Nash equilibrium of the original extensive game.

It is not difficult to see that this tinkering with the *matrix* representation of the normal form can be mirrored in the *graph* representation. In fact, for every convex combination of pure strategies of an extensive game G, we simply add to G a move by nature exactly reflecting it. Troubles only arise if we set out to get a graph representation of the above described infinite matrix. However, instead of expanding the normal form by adding convex combinations, the above observations may as well suggest a converse way of thinking, viz., to make the normal forms smaller. Indeed, we might reverse the process of translating mixed strategies into moves by nature and decide to delete in the normal form of an extensive game all those strategies the expected utility of which can also be obtained as a result of following a convex combination of pure strategies (i.e., a mixed strategy). This, then, is the revision of the definition of the reduced normal form.

Definition 5 (Reduced Normal Form, Revised) *Let $G = (N, H, \iota, f_c, (\mathcal{I}_i)_i, (u_i)_i)$ be an extensive game. Define an action a_i of player i to be a convex combination via scalars $\alpha_0, \ldots, \alpha_k$ of a finite set of actions $b_0, \ldots, b_k$ of player i iff for all i–deleted action profiles and for all players $j \in N$ (including player i) the expected utility $u_j(a_{-i}, a_i)$ is a convex combination via scalars $\alpha_0, \ldots, \alpha_k$ of the expected utilities of the same finite set of actions, i.e,*

$$u_j(a_{-i}, a_i) = \Sigma_{n=0}^{k} \alpha_n u_j(a_{-i}, b_n).$$

Then, the revision of the reduced normal form is defined to be the old version of the reduced normal form from which all convex combinations of actions have been deleted.

Since the revised reduced normal form depends on the old reduced normal form only, one direction of an analogous invariance result for ADD, ID, INT, COA, COAc can be readily obtained: if two games differ *qua* a finite sequence

[2]A convex combination over a set $\{x_0, \ldots, x_n\}$ is an expression $\Sigma_{i=0}^{n} \alpha_i x_i$ for which $\Sigma_{i=0}^{n} \alpha_i = 1$, and $\alpha_i \geq 0$.

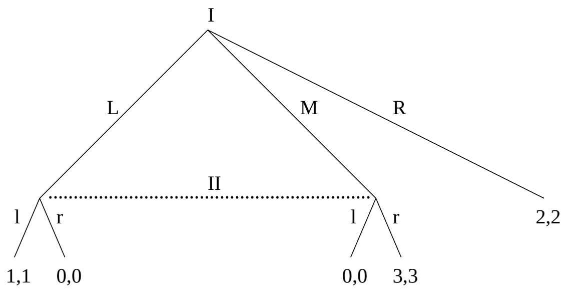

Figure 12: The Original to be Transformed

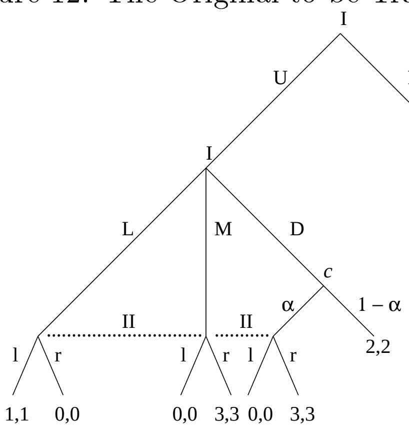

Figure 13: Addition–of–a–Superfluous–Chance–Move (ADDc)

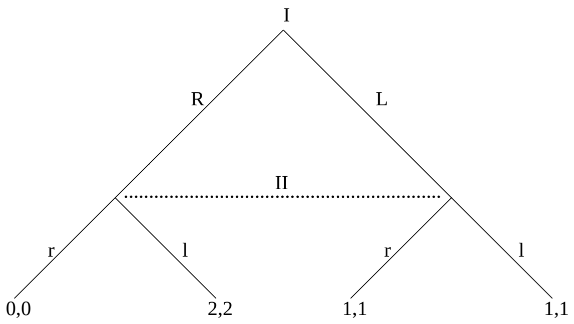

Figure 14: Game G_0

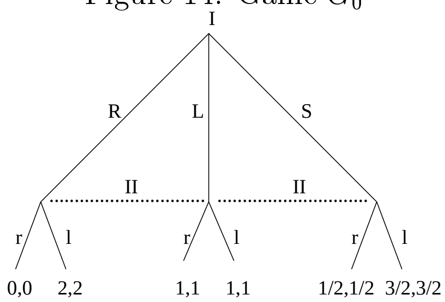

Figure 15: Game G_1

of such transformations, then they share the old reduced normal form, and hence they share the revised version, too. Yet, even for games without chance moves the converse is not true, for no convex combination of pure strategies can be introduced in a game by means of one or more applications of the above five transformations. For instance, since G_0 of figure 14 and G_1 of figure 15 have different old reduced normal forms, they are not indentical modulo the above transformations. However, they share the revised reduced normal form. To overcome this asymmetry, Kohlberg&Mertens introduce yet another transformation, which I will call ADD*c*. The following theorem states its adequacy. Note, that every $\alpha \in [0, 1]$ is allowed for in the transformation, since the goal is to add *any* convex combination. Yet, no infinite game trees will result.

Theorem 6 *For extensive games G and G' the following two statements are equivalent.*

(i) G and G' are equivalent modulo a finite sequence of transformations taken from: ADD, ID, INT, COA, COAc, and ADDc.

(ii) G and G' have the same reduced normal form (revised version).

Proof. Note that (i)⇒(ii) follows from the preceding remarks together with the observation that ADDc is an *exact* correspondent of the revision of the notion of reduced normal form.

Conversely, as we have seen above, the sixth transformation is indispensable. The recipe, then, is roughly as follows. For every convex combination of pure strategies we have the following finite procedure.

- One brings a game in the old reduced normal form by means of ADD, ID, INT, COA, COAc. Then, we are done if no convex combinations exist, since then we are in the revised reduced normal form as well.

- If, on the other hand, there are convex combinations, we apply ADD, ID, INT, COA, COAc, again to bring the game in a form such that ADDc can be applied (this may very well mean a detour, but for reasons of effectiveness this is unavoidable).

- Then, we apply ADDc. Finally, we use ADD, ID, INT, COA, COAc again to go to the old reduced normal form. If there are still convex combinations we are not in the revised reduced normal form, and we start over. Otherwise we are done. ⊣

3 Solution Concepts

In this section I will give an brief outline of the history of game theory along the lines of solution concepts. In almost all of the recent literature on game theory, Nash's solution concept is taken as the starting point of game theoretical investigations, being the coarsest concept, and, in some way, the top or bottom of the lattice (this is not intended as a precise statement!) of all possible solution concepts. Nash simply enumerated all possible strategy profiles of an extensive game, and stipulated that from this set all profiles satisfying a particular intuitively appealing condition had to be considered as rational solutions.

It will be enlightening to distinguish the solutions for strategic games from those for extensive games, although this dissimilarity is lop–sided in that any strategic solution concept can be translated in a solution for extensive games *via* the reduced normal form.

This anthology does not pretend to be balanced in any sense, except that only those concepts are mentioned that we could not do without. The consequence is that in particular the notion of properness will appear rather unmotivated.

3.1 Solution for Strategic Games

The most basic solution concept is the Nash equilibrium.

Definition 7 (Nash, 1951) *Let $G = (N, (A_i)_i, (\succeq_i)_i)$ be a strategic game, a an action profile. a is a Nash equilibrium of G iff for every player i in the game the following condition is satisfied for all actions b_i open to player i.*

$$O(a_{-i}, a_i) \succeq_i O(a_{-i}, b_i),$$

where a_{-i} denotes the i–deleted strategy profile.

After the introduction of the Nash equilibrium in the 1950's, it was attacked from different sides objecting against its too permissive character. One of the alternatives put forward—the only one we will need for strategic games—is Myerson's proper equilibrium. Myerson tried to account for the tendency of not fully rational agents to err, and to that purpose, the ϵ in the definition can be conceived of as a measure of the probability that the players fail to act rationally. To put it differently, in determining whether a strategy profile is an equilibrium, it is considered whether "sufficiently small" possibilities of error on the side of the opponents would force you to deviate.

This problem of erring agents will return in a disguised way in our discussion of the notion of sequential equilibrium for extensive games. There, too, the possibility of making mistakes is tried to be accounted for in that not only strategies *off* the equilibrium path (that is, paths that result from one or more players erring) are *considered* (which we also see in the notion of subgame perfect equilibrium for extensive games), but also related to the *probability* that such an off–path play would actually take place. Although the notions of properness and sequentiality seem only vaguely connected, a striking theorem by Kohlberg&Mertens shows a deep similarity between them. We will come to this later. Also, examples will be deferred until section 4.

Definition 8 (Myerson, 1978) *Let $G = (N, (A_i)_i, (\succeq_i)_i)$ be a strategic game. For $\epsilon > 0$, let $s(\epsilon)$ be a completely mixed strategy profile, i.e., $s(\epsilon) = (s(\epsilon)_1, \ldots, s(\epsilon)_N)$, where $s(\epsilon)_i$ is a probability measure over A_i. We call $s(\epsilon)$ an ϵ–proper equilibrium, iff the following implication holds for all players i, and all actions (!) a_0, and a_1.*

$$U_i(s(\epsilon)_{-i}, a_0) < U_i(s(\epsilon)_{-i}, a_1) \Rightarrow s(\epsilon)_i(a_0) \leq \epsilon s(\epsilon)_i(a_1),$$

where, to be precise, a_0, denotes the strategy that assigns probability 1 to a_0 (and a_1 is defined similarly).

A completely mixed strategy profile s is called a proper equilibrium of G iff it is a limit point of a sequence $(s(\epsilon))_{\epsilon\downarrow 0}$ of ϵ–proper equilibria.

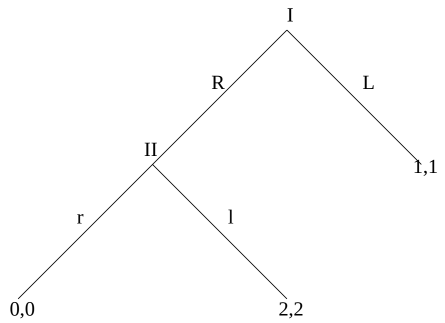

Figure 16: An Unreasonable Nash Equilibrium

A non–trivial theorem shows that every strategic game has at least one proper equilibrium (see [Mye78]).

3.2 Solutions for Extensive Games

As was mentioned before, *via* the normal form of an extensive game we easily define the extensive analog of the Nash equilibrium: take a Nash equilibrium of the normal form. However, in the extensive case similar criticism can be raised to the ineptitude of this notion in the light of erring players. In 1965, Selten proposed an alternative definition of equilibrium using the process of "backward induction" to arrive at strategy profiles that for "off–path" plays, too, would prescribe optimal strategic choices.[3]

3.2.1 Subgame Perfect Equilibrium

The following example illustrates the need for such an improvement (see fig. 16). This game has two Nash equilibria, viz., (L, r), and (R, l). The latter one seems to be the most rational: if both players set out to maximize payoff, it is most definitely in the interest of both to play (R, l). The former equilibrium, however, is less intuitively plausible. In fact, there are good reasons why any solution concept prescribing it should be rejected. The argument runs as follows. Suppose that, contrary to the state of affairs prescribed by equilibrium (L, r), player II has to play (i.e., player I has deviated from the prescription by playing R, instead of L). Now the equilibrium would *still* advise player II to play r, which evidently is not to his profit, since he would

[3]It should be noted that Selten's proposal was not exclusively directed at extensive games, and concerned strategic games as well. His notion of ϵ–perfectness for strategic games fits in between Nash and proper equilibria. Here, we restrict attention to perfectness for extensive games.

be better of playing l. Note, that this criticism cannot be obviated by referring player II in this situation to the second equilibrium since that would simply amount to advising a completely different strategy to him, viz., "play l." A strategy is conceived of as a prescription to a player what to do in all possible circumstances; the notion of circumstance, however, has to be taken *internal* to the game. The *external* circumstance of discussing which strategy *profiles* are equilibria is not part of the game, and consequently cannot be incorporated in the notion of equilibrium.[4] I will now state Selten's alternative definition.

Definition 9 (Selten, 1965) *Let G be an extensive game, and s a strategy profile. s is a subgame perfect equilibrium of G iff the following holds true for every player i, for all subgames H of G, and for all strategies t_i of player i*

$$O_H(s_{-i}\,|_H, s_i\,|_H) \succeq_i|_H O_H(s_{-i}\,|_H, t_i),$$

where $|_H$ is used to denote the obvious restriction of the various objects to subgame H.

In words, this definition merely states the requirement that a strategy profile be a Nash equilibrium of every so–called "proper" subgame (not a relative of proper equilibrium). This notion is somewhat vague, especially for games with imperfect information. Selten, roughly speaking, defines a proper subgame of a game G to be "generated" by one single node in G, i.e., the game that starts playing from that node. In games with perfect information every node generates a proper subgame. In games with imperfect information, however, the proper subgame is, very naturally, required to be closed under successors, and under information sets. It will be clear that this condition is quite a restriction of the number of subgames of games with imperfect information; in many games the only proper subgame will be the game itself.

As an example, it is straightforward how the undesired Nash equilibrium (L, r) is excluded by subgame perfectness, as it is not a Nash equilibrium of the proper subgame generated by player II's decision node. (The other Nash equilibrium does stand the test of subgame perfectness.)

[4] In some sense, this is an arbitrary restriction; but it is a very sensible one. Cf. [KM86, p. 1005n3]. Also, see the preface of [vD83]: "Since in a noncooperative game binding agreements are not possible, the solution of such a game has to be self–enforcing" (page iii).

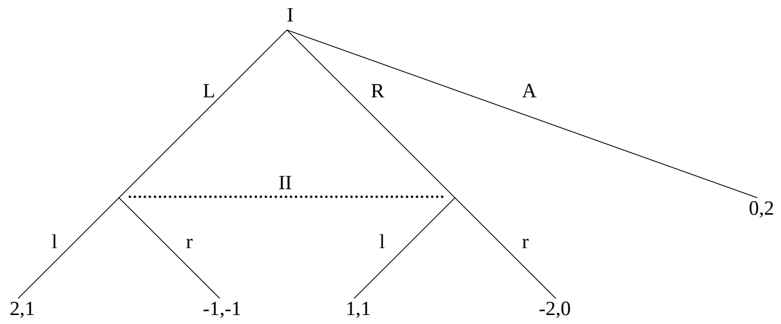

Figure 17: An Unreasonable Subgame Perfect Equilibrium

3.2.2 Sequential Equilibrium

Subsequently, consider an example of a game with imperfect information (fig. 17). This game has two Nash equilibria, viz., (L, l), and (A, r), and both are subgame perfect, since the only proper subgame is the game itself. However, we are in the same trouble as in fig. 16, for consider (A, r), and suppose that because player I has deviated from the prescription, player II has to move. This subgame perfect equilibrium (SPE) prescribes him to play r, which is doubtlessly to his disadvantage. Although player II does not know where he is, he knows that playing l will yield him a sure payoff of 1, while r yields him at most 0. This shows that the concept of SPE will not do for games with imperfect information, thus casting doubt, at least partly, on the possibililty of applying the procedure of backwards induction here.

One might be inclined to conclude from the above example that the notion of subgame is too narrow to be applied in any acceptable way, and that a redefinition is wished for. However, to my knowledge one does not find this approach in the literature, and that is certainly not without reason. A more penetrating analysis shows that the difficulties that arise in attempts to extend the notion of SPE to games with imperfect information are related to particular intrinsic features of such games rather than to a particular definition of the notion of subgame. Indeed, what seems to be more natural is to assume that a player, once he is in a non–singleton information set, has to make guesses about how he got there (i.e., what the other players, or, in case of imperfect recall, what he himself did) on the basis of which he will define a probability measure on the set of possible moves he can take. This idea was put forward by Kreps&Wilson in their 1982 paper as the notion of "sequential equilibrium." I will briefly summarize this.

First, a formalization of the above probabilistic notion of strategy has to be given. Let G be an extensive game, i a player, and $\mathcal{I}_i$ the set of his information sets. Kreps&Wilson define a "behavioral strategy" to be a function π_i that assigns to every information set $I \in \mathcal{I}_i$ a probalility measure on the set $\alpha(I)$ of actions open to i at I. Unless one uses the same name for actions open at different information sets, one can equivalently view π_i as a function mapping the set of all actions of player i to the closed interval $[0, 1]$ such that $\Sigma_{a\in\alpha(I)}\pi_i(a) = 1$ for all $I \in \mathcal{I}_i$.

Second, the players' beliefs about the past (i.e., the probabilities they ascribe to the ways they arrived at a particular information set) have to be grasped in a formal definition. To that end, Kreps&Wilson define a "system of beliefs" to be a function μ mapping the decision nodes of an extensive game G to the closed interval $[0, 1]$ such that for each information set I we have $\Sigma_{x\in I}\mu(x) = 1$.

These two definitions provide us with the components the new notion of equilibrium will consist of. While up to now an equilibrium was defined as a profile of strategies, the new candidate, called an "assessment," will add to this a system of beliefs. An assessment of an extensive game G, according to Kreps&Wilson, is a pair (π, μ) of a profile π of behavioral strategies, and a system of beliefs μ.

Not every assessment, of course, will serve as a reasonable equilibrium; hence, particular conditions will have to be fulfilled. The first condition will mirror the ideas expressed in the Nash equilibrium; that is, in some—as yet vague—way, it will be required that deviation from the strategy prescribed by the equilibrium will not be to a player's advantage unless another player deviates as well.[5] In order to formally describe that condition, a notion of "outcome" *relative to the notion of assessment* has to be defined first. Kreps&Wilson's approach is to consider the probability $P_{\pi,\mu}$ (calculated on the basis of an assessment (π, μ)) that a particular terminal node z will be reached *given* that a particular information set I has been reached.[6] Indeed, if z cannot be reached from I, unsurprisingly $P_{\pi,\mu}(z \mid I) = 0$. If, on the other hand, one can reach z from I, say, following nodes $x_0, x_1, \ldots, x_n$, where $x_0 \in I$, and $x_n = z$, then

$$P_{\pi,\mu}(z \mid I) = \mu(x_0) \cdot \Pi_{i=1}^{n}\pi(\alpha(x_i)).$$

Here, $\alpha(x)$ denotes the action taken to arrive at node x.

On the basis of probability measure $P_{\pi,\mu}$ one can obtain a conditional expectation $E_{\pi,\mu}$ in a straightforward fashion. The Nash–like condition touched upon above requires that an assessment (π, μ) be "sequentially rational,"

[5]The significance for the foundations of game theory of this principle of self–enforced collective non–deviation cannot be overestimated. The literature, however, is remarkably silent about it.

[6]The conditionality expressed in this sentence does not always fully coincide with the usual conception in probability theory.

which amounts to the statement that, for all information sets I it is true that, for all behavioral strategies $\rho_{\iota(I)}$ of player $\iota(I)$,

$$E_{(\pi_{-\iota(I)},\pi_{\iota(I)}),\mu}(u_{\iota(I)}(z)|I) \geq E_{(\pi_{-\iota(I)},\rho_{\iota(I)}),\mu}(u_{\iota(I)}(z) \mid I).$$

Here, u_i is the Von Neumann and Morgenstern utility function for player i.

Since this condition does not—in any significant way—involve the system of beliefs μ, we will need a second condition restricting the number of allowable systems of beliefs. This condition asks for the existence of a sequence $(\pi_n, \mu_n)_{n=0}^{\infty}$ converging to (π, μ) in Euclidean space, such that, first, every strategy profile π_n assigns positive probability to every action at every information set, and second, that for every system of beliefs, and every node x, the belief system is consistent with Bayes' rule for conditional probabilities, i.e., $\mu_n(x) = \frac{P_\pi(x)}{P_\pi(I)}$, where I is the information set containing x. An assessment satisfying this condition is called "consistent."[7] Finally, the new notion of equilibrium can be defined.

Definition 10 (Kreps&Wilson, 1982) *An assessment (π, μ) is called a sequential equilibrium iff it is sequentially rational and consistent.*

4 Equivalence: Solution Concepts and Transformations

In this section, I will discuss the relations between the solution concepts we encountered in the last section, and the notion of game equivalence of the penultimate one. This investigation is motivated by the idea that, for an *acceptable* notion of game equivalence, equivalent games should share equilibria, for an *acceptable* notion of equilibrium. (Equilibria and game equivalences have to be *pairwise* acceptable, so to speak.)

4.1 (In)variance of Equilibria under Transformation

Now, for the Nash equilibrium, as well as for the proper equilibrium, we obtain the following invariance result as a consequence of Theorem 3.

Corollary 11 *If two extensive games are equivalent modulo a finite number of applications taken from ADD, ID, INT, COA, then modulo renaming the strategies they have the same set of Nash equilibria, and the same set of proper equilibria.*

[7]Note that as a consequence of the Bayes conditions, μ can simply be abstracted from as there is a natural way to define it in terms of π. We will see this in Theorem 14.

Proof. While the transformations respect the (old) reduced normal form, the equilibrium notions are defined on the *non–reduced* normal form; hence the proposition is not entirely trivial. However, since the point of the old reduced normal form is to delete equally preferable strategies, and the Nash equilibrium is defined in terms of relative preferences only, the first half of the statement is unproblematic. For the proper equilibrium, it is slightly less straightforward. Recall that for an action profile to be a proper equilibrium, it need to be a limit point of a sequence of ϵ–proper equilibria. Now since in the reduced normal form the set of actions open to a particular player may be different from his action set in the original game, the probability distributions that the ϵ–proper equilibria actually are may be different, too. However, since equivalent actions in the original game are assigned equal probability, we can easily obtain a new sequence of ϵ–proper equilibria for the reduced normal form with the same limit point. From the fact that the converse can also be done it follows that the set of proper equilibria of a strategic game, and its reduced normal form coincide. $\dashv$

I will now show—in an admittedly rote way—that such invariance results cannot be given for several of the refinements of the Nash equilibrium. Note, that this impossibility is *not* the result of the fact that they "refine" the Nash equilibrium; rather, it shows that in a sense the refinements are not *uniform* with respect to the rnf.

Fact 12 *Games exist which are equivalent modulo a finite number of transformations taken from ADD, ID, INT, COA, that have different sets of subgame perfect equilibria.*

As a case in point, let us consider game G_0 (see figure 18), and its rnf, i.e., the extensive game G_1 (figure 19) obtained from G_0 following the recipe of the preceding theorem (one single application of ADD suffices). As was explained, the unique SPE of G_0 is (R, l). Obviously, this is an SPE of G_1, too; haplessly, it has to endure the presence of another SPE, viz., (L, r)[8] showing that Thompson transformations do not respect SPE's. Similar arguments take care of the other transformations. Only INT does not change the number of subgames, and that is why it is the sole transformation respecting SPE's.

Fact 13 *Games exist which are equivalent modulo a finite number of transformations taken from ADD, ID, INT, COA, that have different sets of sequential equilibria.*

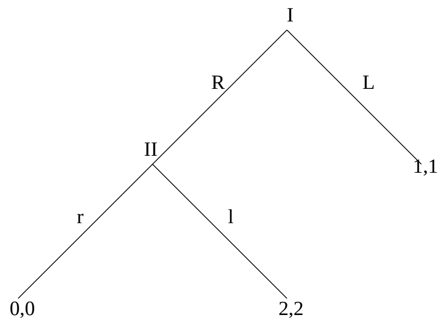

Figure 18: Game G_0

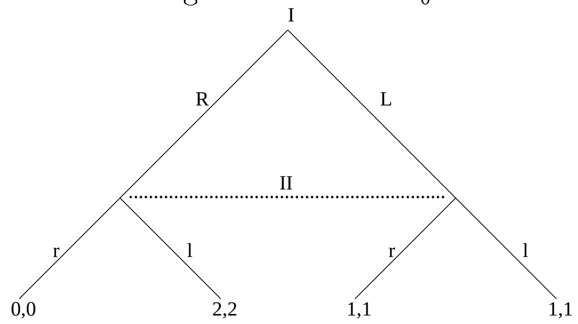

Figure 19: Game G_1: The Reduced Normal Form of G_0

As an instance I will show this for ADD. The other transformations can be done in similar ways. Consider game G_0 (figure 20) and its normal form G_1 (figure 21). First, we explain that the assessment (π, μ), where $\pi_1 = \pi_2 = L$, and $\mu(L|\{L, M, R\}) = 1$,[9] is a sequential equilibrium (SQE) of G_1. To show consistency we consider the sequences $\pi_1^\epsilon = (1 - \epsilon - \epsilon^2, \epsilon, \epsilon^2)$, and $\pi_2^\epsilon = (1 - \epsilon, \epsilon)$, which evidently converge to the desired limits. Furthermore, for $\mu^\epsilon(R|\{L, M, R\}) = \epsilon^2$, $\mu^\epsilon(M|\{L, M, R\}) = \epsilon$, and $\mu^\epsilon(L|\{L, M, R\}) = 1 - \epsilon - \epsilon^2$, the Bayes requirement is fulfilled as well. (It can be shown that this is the only consistent system of beliefs given the strategy profile π.) Sequential rationality follows from the observation that to compute the expected payoff (as defined with respect to SQE) the middle and the rightmost part of the game can be ignored. Indeed, a deviation of player II to R instead of L, would result in exactly the same payoff, since in calculating this expected

[8] Recall that according to the definition all Nash equilbria of "proper" subgames are SPE's; also notice that G_1 has only one proper subgame.

[9] The notation of the belief system is slightly at odds with the preceding section. "$L|\{L, M, R\}$" denotes the leftmost decision node (reached by player I's action L) in player II's information set (reachable by player I's actions L, M, and R). Note also that as, for each information set, μ is a probability measure on its decision nodes, there is no need to specify explicitly that R and M are assigned probability 0.

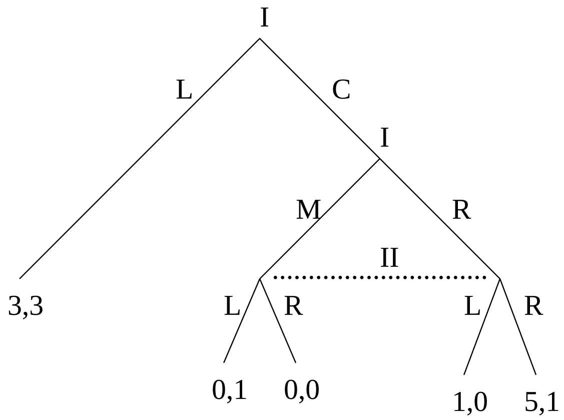

Figure 20: Game G_0

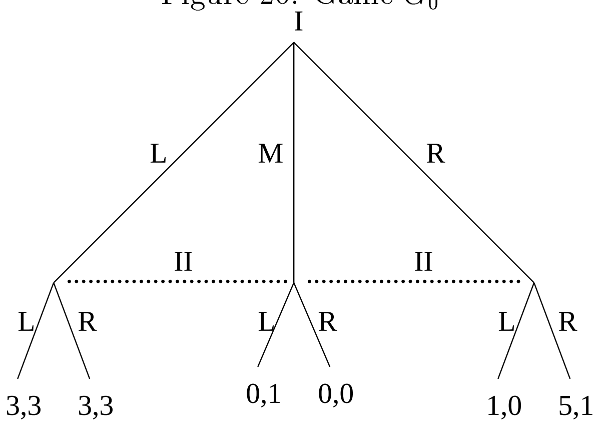

Figure 21: Game G_1: The Reduced Normal Form of G_0

value, the probability attached to the middle and rightmost decision nodes is 0.

The above line of reasoning shows that the payoff profile $(3,3)$ can be obtained in game G_1 as the result of a sequential equilibrium; the following argument shows that no SQE in game G_0 will have such payoff. Assume that it *is* indeed obtained from an SQE. Then player I's behavioral strategy will assign probability 1 to move L. Since move M for player I is stricly dominated by his move R, any equilibrium attaching positive probability to I playing M is ruled out. Consequently, the system of belief will satisfy $\mu((C,R)|\{(C,M),(C,R)\}) = 1$. Given such belief system, no sequential equilibrium would attach positive probability to player II playing L, since L is strictly dominated by R. The only remaining candidate for sequential equilibrium, then, is defined by $\pi_1(L) = \pi_1(R) = \pi_2(R) = 1$, and $\mu((C,R)|\{(C,M),(C,R)\}) = 1$.

It is not difficult to show that this equilibrium is consistent (again, no other definition of the belief system would be consistent given the strategy profile). Nevertheless, the candidate fails to be sequentially rational. To see that, we compute the expected payoff for player I at his first information set (i.e., the origin of the game). The candidate equilibrium would give him a payoff of 3. Deviation, though, would increase his score with 2, since playing a behavioral strategy σ_1 defined by $\sigma_1(C) = \sigma_1(R) = 1$, would provide him with a payoff 5—provided, as required by the Nash–like condition of the definition of SQE, that player II does not deviate.[10]

4.2 Properness and Sequentiality

Altogether, we have seen that the notion of game equivalence induced by the reduced normal form (or by the corresponding game transformations for that matter) is quite choosy with respect to what solution concepts it allows for. This feature, it was already noted before, plays an important role in arguments against, and in favor of, particular proposals for solution concepts. However, if we slightly change the perspective, a striking result by Kohlberg&Mertens (refined by Van Damme) shows that the situation is not as hopeless as we might think. Indeed, they show that under the notion of game equivalence induced by the reduced normal form, it will be

[10]This example is due to Osborne&Rubinstein. On page 226 of their textbook, they intend to show that in the unique SQE of game G_0 player I plays C and R with probability 1, player II plays R, and $\mu((C,R)|\{(C,M),(C,R)\}) = 1$. As far as I can see, however, though their argument shows that of any *consistent* equilibrium the system of beliefs has to satisfy this condition, it lacks the necessary observation that there is only one such assessment that is *sequentially rational* as well.

the case that any proper equilibrium of the reduced normal form of a game will be a sequential equilibrium of any equivalent game. Apart from this, the theorem is significant for another reason. I said before that both the proper equilibrium and the SQE try to account for mistaken players. This theorem, then, shows the connection between these two approaches to rationality and belief that are at first sight quite dissimilar.

Theorem 14 (Kohlberg&Mertens, 1982) *Let G be a graph representation of an rnf (old version). Suppose that a strategy profile s is a proper equilibrium of G. Then, if H is a game of which G is the rnf, there exists a profile π of behavioral strategies that are equivalent to the strategies from s such that π is an SQE of H.*

Proof. Since s is proper, there exists a sequence of ϵ–proper strategy profiles $(s(\epsilon))_\epsilon$ converging to s if ϵ tends to 0. Denote the strategy of player i by $s(\epsilon)_i$, and define a sequence of equivalent behavioral strategies $(\pi(\epsilon))_\epsilon$ by

$$\pi(\epsilon)_i(I)(a) = \frac{\sigma(\epsilon)_i(I,a)}{\sigma(\epsilon)_i(I)},$$

where $\sigma(\epsilon)_i$ is the sum of the probabilities that $\pi(\epsilon)_i$ assigns to all pure strategies that do not exclude x.[11] Furthermore, define a sequence $(\mu(\epsilon))_\epsilon$ of systems of belief satisfying Bayes' rule with respect to $(\pi(\epsilon))_\epsilon$. Now we claim that any ordered pair of limit points (π, μ) of the sequence $(\pi(\epsilon), \mu(\epsilon))_\epsilon$ is an SQE.

Consistency follows from the definition. Assume, then, that $(\pi(\epsilon), \mu(\epsilon))_\epsilon$ is *not* sequentially rational. Then, for some behavioral strategy π_i, and some "lowest" (i.e., "bottom–most") information set I for player i, and some actions L, and R open at I, it is true that $\pi_i(I)(L) > 0$, although the expected utility of applying π_i is less than the expected utility of applying ρ_i, where ρ_i coincides with π_i except that it attaches propability 1 to R at I (here the assumption that I be bottom–most is used). Now, translating that in the language of mixed strategies (paving the way to the level of the reduced normal form) gives that for any pure strategy $l = \bar{x}L\bar{y}$, and its "modification" $r = \bar{x}R\bar{y}$, that do not avoid information set I, we have, for the expected utility U_i of player i

$$U_i(s(\epsilon)_{-i}, l) < U_i(s(\epsilon)_{-i}, r).$$

Let m be the number of pure strategies that do not avoid I. Then from the properness of the equilibrium the following inequality can be deduced.

$$\pi(\epsilon)_i(I)(L) = \frac{\sigma(\epsilon)_i(I,L)}{\sigma(\epsilon)_i(I)} \leq \frac{m\epsilon s(\epsilon)_i(r)}{\sigma(\epsilon)_i(I)} \leq m\epsilon,$$

[11] This construction, due to Kuhn, has become standard in the literature.

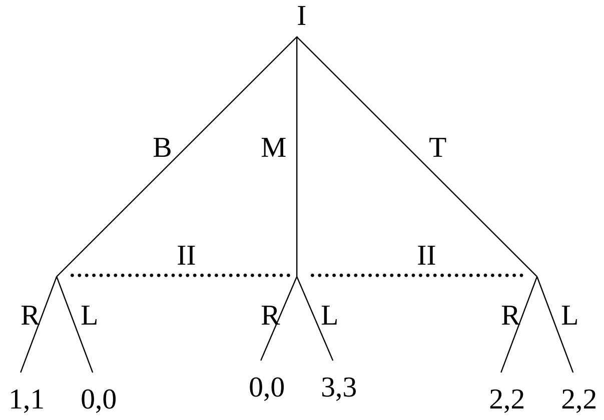

Figure 22: Game G

which shows that for ϵ tending to 0, the probability that $\pi(\epsilon)_i$ assigns to move L at information set I converges to 0. Hence, $\pi_i(I)(L) = 0$, which contradicts the assumption. This concludes the proof of the theorem. $\dashv$

As an illustration of the theorem as well as of the notions examined in this section, consider game G depicted in fig. 22. We show that the strategy profile (T, R) is proper. Indeed, let player I have a mixed strategy $(\epsilon^2, \epsilon^3, 1 - \epsilon^2 - \epsilon^3)$ (probabilities for B, M, and T, respectively), where player II plays $(1 - \epsilon^2, \epsilon^2)$ (R, and L, respectively). It is clear that they converge as desired. Also, the properness conditions are satisfied. For instance, for player I consider move B, which has lower expected payoff than move T; it is true that $\epsilon^2 \leq \epsilon(1 - \epsilon^2 - \epsilon^3)$ for ϵ sufficiently close to 0. In a similar fashion, it can be seen that (M, L) is proper, too. (One can verify that changing the payoff profile after (M, R) in 2,0 invalidates the above reasoning; indeed, player I would not play T, since it is (weakly) dominated by M. In (rough!) terms of properness: even if player II would mistakenly play R after I's move of M, I's payoff would still be 2.)

5 Conclusion

What is the upshot of our investigations? I believe that the adequacy theorems show that the reduced normal form (or the corresponding game transformations, for that matter) is the most natural notion of game equivalence from a game theoretical view point. Granted, the fact that many solution concepts are not invariant under game transformations constitutes a pretty predicament. But why not exploit this alleged drawback? Could not this fairly stable notion of game equivalence serve as a standard (in the sense of the standard meter, etc.) against which we measure all kinds of proposals for solution concepts? Well, perhaps indeed it can *not*. But if rejecting this use of game equivalence as a standard leads one to adopt a new notion of game equivalence that *is* invariant for a particular solution concept, the dilemma simply returns in different clothes, for there will be other solution concepts for which invariance cannot be proved.

I will not dwell upon this any longer. I would like to conclude by turning to the original motivation of this paper, that is, to logic. What, then, may a logical analysis of games provide us with? The first point I will mention hints at the relation between games and processes (in the sense in which the word is used in computer science). Indeed, a process being a directed generated graph the arrows of which represent particular actions, there is a straightforward analogy with extensive games, and since a quite elaborate theory of process equivalences has been developed, it is interesting to check whether these notions would lead to plausible alternatives to the game theoretical notion of game equivalence induced by the reduced normal form. For instance, in order to decide upon equivalence of two processes (and so of two games) their respective "traces" might be compared, where a trace can be thought of as a path in a game tree. Different requirements on the relations of the traces of two games would then result in different notions of game equivalence. Another notion of process equivalence leads us to logic proper, viz., the notion of simulation. Succinctly put, a logical language is developed for which the graph (representing the process, or the game) is a (Kripke) model. For two graphs to be equivalent, then, it is needed that particular *actions* taken in one graph can be mirrored in the other in such a way that the corresponding *nodes* satisfy the same propositions.[12]

However, even if among the notions of process equivalence we do not find a good alternative for the reduced normal form, the development of a logical language with games as models may still help us to obtain a deeper understanding of game equivalence. We might, then, compare the theories (in the sense of all the propositions true in a model) of two games with the same reduced normal form. If we take a predicate logical perspective, such investigations may give us information about what "global" properties of the games change upon game transformations, and in a trivial sense, such a predicate logical description is already provided in the formal definitions of the transformations. On the other hand, we can take a modal perspective,

[12]For detailed treatment of both trace and simulation equivalences, the reader is referred to Van Glabbeek's 1990 survey paper.

showing what happens "to the players," a "local" view point so to speak. In the literature, both approaches can be found[13] but to my knowledge no direct application to the problem of game equivalence has been given yet. Let that be our next challenge.

References

[AB95] R. J. Aumann and A. Brandenburger. Epistemic Conditions for Nash Equilibrium. *Econometrica*, 63:1161–1180, 1995.

[BK90] J. C. M. Barten and J. W. Klop, editors. *CONCUR '90. Theories of Concurrency: Classification and Extension.* Springer Verlag, Berlin, 1990.

[Bra92] A. Brandenburger. Knowledge and Equilibrium in Games. *Journal of Economic Perspectives*, 6:83–101, 1992.

[ER94] S. Elmes and P. Reny. On the Strategic Equivalence of Extensive Form Games. *Journal of Economic Theory*, 62:1–23, 1994.

[KM86] E. Kohlberg and J.–F. Mertens. On the Strategic Stability of Equilibria. *Econometrica*, 54:1003–1037, 1986.

[KN96] M. Kaneko and T. Nagashima. Game Logic and its Applications I. *Studia Logica*, 57:325–354, 1996.

[KT53] H. W. Kuhn and A. W. Tucker, editors. *Contributions to the Theory of Games*, volume II. Princeton University Press, Princeton, 1953.

[Kuh] H. W. Kuhn. Extensive Games and the Problem of Information. In *Contributions to the Theory of Games* [KT53], pages 193–216.

[KW82] D. M. Kreps and R. Wilson. Sequential Equilibria. *Econometrica*, 50:863–894, 1982.

[Mye78] R. B. Myerson. Refinements of the Nash Equilibrium Concept. *International Journal of Game Theory*, 7:73–80, 1978.

[Nas51] J. F. Nash. Non–Cooperative Games. *Annals of Mathematics*, 54:286–295, 1951.

[OR94] M. J. Osborne and A. Rubinstein. *A Course in Game Ttheory.* The MIT Press, Cambridge, MA; London, 1994.

[13] See for the former [KN96], and for the latter see, e.g., [Sta94].

[Sel65] R. Selten. Spieltheoretische Behandlung eines Oligomonopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 12:301–324, 667–689, 1965.

[Sta94] R. Stalnaker. On the Evaluation of Solution Concepts. *Theory and Decision*, 37:49–73, 1994.

[Tho52] F. B. Thompson. Equivalence of Games in Extensive Form. Research Memorandum RM–759, The RAND Corporation, Santa Monica, CA, 1952.

[vD83] E. E. C. van Damme. *Refinements of the Nash Equilibrium Concept.* Lecture Notes in Economics and Mathematical Systems. Springer–Verlag, Berlin, 1983.

[vG] R. J. van Glabbeek. The Linear Time–Branching Time Spectrum. In *CONCUR '90. Theories of Concurrency: Classification and Extension* [BK90], pages 278–297.