CESifo WORKING PAPERS

7715
2019

June 2019

Strict Incentives and Strategic Uncertainty

Hans Carlsson, Philipp C. Wichardt

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7715*
Category 12: Empirical and Theoretical Methods

Strict Incentives and Strategic Uncertainty

Abstract

This paper proposes a comprehensive perspective on the question of self-enforcing solutions for normal form games. While this question has been widely discussed in the literature, the focus is usually either on strict incentives for players to stay within the proposed solution or on strategic uncertainty, i.e. robustness to trembles. The present approach combines both requirements in proposing the concept of robust sets, i.e. sets of strategy profiles which satisfy both strict incentives and robustness to strategic uncertainty. The result is a set valued solution, a variant of which is shown to exist for all finite normal form games.

JEL-Codes: C720.

Keywords: game theory, self-enforcing solution, strict incentives, strategic uncertainty.

*Hans Carlsson**
Department of Economics
Lund University
Box 7082
Sweden – 22007 Lund
hans.carlsson@nek.lu.se

Philipp C. Wichardt
Department of Economics
University of Rostock
Ulmenstr. 69
Germany – 18057 Rostock
philipp.wichardt@uni-rostock.de

*corresponding author

This Version: June 10, 2019
Wichardt thanks the Department of Economics at Lund University for its hospitality. Financial support from the Bank of Sweden Tercentenary Foundation, the Swedish Competition Authority (Carlsson) and the Arne Ryde Foundation (Wichardt) is gratefully acknowledged. The usual disclaimer applies.

1 Introduction

What does it mean for a proposed solution to a game to be self-enforcing? Almost ever since Nash (1950) proved the existence of equilibrium points in general n-person games, the game theoretic literature has been abound with discussions of this question. The most prominent branch in this, perhaps, is the refinement literature focusing on Nash equilibria which are robust to some form of error or uncertainty about the behaviour of others (e.g. Selten, 1975, Myerson, 1978, Kreps and Wilson, 1982, Kalai and Samet, 1984, Kohlberg and Mertens, 1986). In their seminal 1986 paper, Kohlberg and Mertens indeed motivate their analysis by saying that they intend to establish "...which Nash equilibria are strategically stable, i.e. self-enforcing,..." (Kohlberg and Mertens, 1986, p. 1003).

An alternative perspective on the matter is due to Basu and Weibull (1991) who propose a set-valued extension of the notion of strict equilibrium, named CURB (mnemonic for "closed under rational best replies"). In contrast to the refinement literature, their focus is not on robustness against potential errors but on the requirement of strict incentives to stay within the solution. Accordingly, they consider sets of strategies containing all best replies to the opponents' strategies. Referring to Kohlberg and Mertens, they write (Basu and Weibull, 1991, p. 142) "...while they select sets of Nash equilibria, i.e. sets contained in their own best replies, we here select sets containing all their own best replies, a 'dual' approach which can be viewed as a set-theoretic coarsening of the notion of strict Nash equilibrium while [Kohlberg and Mertens's] approach is a set-theoretic refinement of the Nash equilibrium concept."

In the present paper, we bring together both strands of the discussion considering strict incentives and strategic uncertainty at the same time, with a focus on static games. For the purposes of our analysis, we conceive of players as being in a context where a particular solution is proposed, this candidate solution always being a selection (i.e. a product set) of pure strategies.[1] The restriction to pure strategies here reflects the idea to identify all best responses rather than a particular equilibrium distribution.

In order to determine criteria for robust recommendations, it is useful to first make explicit the dual nature of game-theoretic solutions often left implicit in the literature (for an exception, though in the context of dynamic games, see Rubinstein, 1991): on the one hand the solution to a game is a set of recommendations about strategies that players should use, on the other, a set of conjectures that each player can entertain about the other players' behaviour. Eventually, these two aspects of a solution, of course, should be aligned: recommendations should contain strategies that can be rationalised given some justifiable conjecture, and conjectures about a player's behaviour should closely reflect the recommended strategies for this player. Robustness of a solution, therefore, should cover both aspects: strict incentives of players to remain within the solution and immunity regarding strategic uncertainty when it comes to the expected behaviour of others.

For our solution concept, we require players to have strict incentives to conform to the solution, i.e. each player should be recommended to play a given strategy if and only if it is rational under some conjecture justifiable under the candidate solution. Notice that this implies both *external* and *internal stability* with respect to the candidate solution: no rational strategy is excluded (external stability) and no irrational one included (internal stability). We view this requirement as a prerequisite for a truly self-enforcing solution: unless the recommended strategies coincide with the set of possibly rational choices, the corresponding conjectures will be inappropriate and

[1]This context is similar to the one proposed by Kohlberg and Mertens (1986, fn. 3) although in their set-up the proposed solution is a mixed strategy profile while in our approach each player is recommended a set of pure strategies. However it contrasts strongly with solution concepts – such as rationalizability (e.g. Bernheim, 1984; Pearce, 1984) and iterated dominance (e.g. Asheim and Dufwenberg, 2003; Brandenburger, Friedenberg and Keisler, 2008; Brandenburger and Friedenberg, 2010) – which only rely on each player reasoning deductively on the basis of his knowledge about the game and the other players.

will have to be revised, thus upsetting the solution. Sufficiently refined equilibrium concepts typically satisfy internal, but not external stability. To meet both requirements a set-valued approach is called for. Notably, the notions of tight curb set (Basu and Weibull, 1991) and tight σ-curb set (Balkenborg et al., 2013) both satisfy the strict-incentives condition.

In addition and more in line with the the refinement literature, we require that the players' conjectures reflect some amount of strategic uncertainty: although each player strongly believes that the others will act rationally, he is not absolutely certain. Accordingly, we assume that players only consider conjectures which assign a positive probability to every strategy of a given player.[2] Moreover, we again make an external-internal distinction: external uncertainty concerns strategies outside the solution while internal uncertainty refers to strategies within. However, we assume that the latter has a much larger magnitude: any strategy included in the proposed solution is considered much more likely than any strategy not included. This key assumption will be labelled *prevalence of internal uncertainty*.

Our main solution concept, which we name *robust set*, then, is derived from the above assumptions by means of a natural limit procedure. Formally, we describe it as a fixed point of the robust best-reply function $B : \mathbb{T} \to \mathbb{T}$ which maps the set of all selections $\mathbb{T}$ into itself. For any selection $T \in \mathbb{T}$, $B(T)$ contains the pure strategies that a rational player can be expected to play given T and our assumptions about strategic uncertainty. Intuitively, a robust set consists of recommendations from which no rational player has any incentive to deviate even in the face of some uncertainty about the behaviour of others.

As it turns out, there are games in which there is no robust set and, hence, no solution that satisfies all our desiderata. To amend for this, we propose the notion

[2]This assumption embodies a certain measure of realism: in a real-life setting, it is virtually impossible for a player to exclude events which could induce an opponent to make an irrational move. Our approach attempts to account for such unmodelled events in a practicable way.

of *quasi-robust set*. Such a solution is the outcome of an iterative procedure based on the robust best-reply function B; as we show, quasi-robust sets exist for every finite normal form game. In a sense, quasi-robust sets can be viewed as a natural weakening of robust sets which follows our goal to identify all justifiable behaviour within a solution. Hence, it is no surprise that a robust set is always quasi-robust and, if a selection T is quasi-robust but not robust, then $B(T)$ is a proper subset of T. Thus, what is lost in weakening robustness is that quasi-robust sets sometimes do not satisfy internal stability.

In order to illustrate the role of the key assumption of our main model – the prevalence of internal uncertainty – we also consider a variant where this condition is not met, a model with unrestricted strategic uncertainty. Solutions of this model are called *semi-robust sets*. As we will see, semi-robust sets always exist in finite games but, unlike robust sets, they sometimes contain strategies that are relatively dominated, i.e. dominated in the constrained game where players are restricted to their solution strategies.

Before we proceed to illustrate the the main features of our approach with some examples, we want to briefly return to some related papers. For the sake of brevity, we focus on those two concepts which we think are closest related to the subsequent discussion. A recent, concise and more comprehensive review of the related literature can be found in Myerson and Weibull (2015, introduction).

An interesting recent approach which bears some resemblance to the spirit of (semi-)robust sets is due to Myerson and Weibull (2015). In their paper, Myerson and Weibull consider selections of strategies – blocks in their terminology – which no player would want to deviate from once almost all others adhere to it. Similar to the present discussion, they also use a notion of robustness which requires behaviour outside the proposed solution to be considered much less likely. Unlike the present approach, however, they start by assuming that average behaviour within the selection constitutes a Nash Equilibrium of the perturbed game. Deviations, then, are judged against this equilibrium and not, as in the present context, against each single strategy in the proposed solution. Accordingly, Myerson and Weibull reach finer solutions, yet at the cost of already presuming some equilibrium coordination to begin with.[3]

Finally, a concept which is (technically) related to the present approach, albeit coming from a very different (evolutionary) direction, is the *most refined best-reply correspondence* σ defined by Balkenborg et al. (2013, 2015). More specifically, Balkenborg et al. (2013, p.169) "are interested [...] in identifying and characterizing the smallest faces [of the polyhedron of mixed strategy profiles] that are evolutionarily stable under some reasonable dynamics (appropriate for highly rational and highly informed human beings)" – a motivation very different from the present attempt to find those strategy sets which are robust to both strategic uncertainty and strict incentives to stay within the solution. From their notion of a best-response correspondence, they then define the notion of a σ-curb set which is a refinement of curb sets (Basu and Weibull, 1991). As it turns out, one version of this notion, namely that of a tight σ-curb set, is technically equivalent to our notion of a semi-robust set to be defined below. Given the difference in motivation and the fact that Balkenborg et al. are mainly interested in the dynamics generated by the correspondence σ, however, we want to emphasise that the criticism of semi-robust sets expressed in the present paper need not apply to σ-curb sets in the context where this notion is used.

The rest of this paper is structured as follows: In Section 2, we discuss some illustrating example before moving on to the formal modelling in Section 3. Section 4 concludes with some final comments. All proofs are gathered in the Appendix.

[3] Myerson and Weibull (2015) suggest to interpret (tenable) blocks as possible norms for a certain interaction. From that point of view, a strong reliance on equilibrium ideas, i.e. presuming some amount of equilibrium coordination to begin with, is very natural. The focus of the present discussion, however, is a totally different one, namely to clarify what can be argued for as rational in case without any existing form of equilibrium coordination.

2 Modelling Strategic Uncertainty: Motivation and Examples

In this section, we provide an informal account of robust (and quasi-robust) sets and illustrate their main features through two examples.

Model (informal)

As we have argued in the introduction, we are interested in selections of strategies which contain all best responses to strategies within the selection and which – if proposed to the players as a solution – are robust to some uncertainty regarding the behaviour of others; for our main solution concept, we will assume, though, that the uncertainty still focuses on strategies within the proposed solution.

To clarify the role of the assumption of prevalence of internal uncertainty, we use as a comparison the model with unrestricted strategic uncertainty and the associated notion of semi-robust set. Both models consider a situation where a selection T – a product set of pure strategies – is proposed as a solution to a given game. Each player, then, forms conjectures corresponding to completely mixed strategies that are concentrated on T. However, while the semi-robust model contains no further restriction, the one underlying robust sets is based on the additional assumption that any strategy within T is much more likely than any other in the conjectures of the players.

More specifically, the two models are generated by best-reply functions $B : \mathbb{T} \to \mathbb{T}$ and $A : \mathbb{T} \to \mathbb{T}$ to be defined formally later. For the two-person games to be analysed here, it is useful to consider the types of conjectures, to be called ε-conjectures and δ-conjectures respectively, by means of which these functions can be approximated. Given a selection T as the proposed solution, a player entertaining ε-conjectures, for some small $\varepsilon > 0$, will assume that the probability $c_j(s_j) > 0$ of any pure strategy s_j

for the other player *not* in T_j is less than $\varepsilon \cdot c_j(t_j)$, for any pure strategy t_j *within* T_j. If, on the other hand, he has δ-conjectures for some small $\delta > 0$, he will simply assume that the other player uses a completely mixed strategy such that the probability of some pure strategy not in T_j being played is less than δ. Thus, ε-conjectures generally prioritise strategies within T while δ-conjectures do not.

The pure strategies included in $B(T)$ and $A(T)$, then, are the ones that, for each player, maximise payoffs for some ε-conjectures and δ-conjectures with arbitrarily small ε and δ, respectively.[4]

To begin with, we make two observations: First we note that, for any T, $B(T) \subseteq A(T)$ so that B is potentially more precise than A. This follows from the fact that for any ε-conjecture c_j one can find $\delta^* > 0$ such that c_j is also a δ-conjecture for $\delta \in (0, \delta^*)$ while the converse does not hold. If, however, T is a singleton, i.e. contains a single pure strategy profile, the converse does hold. So, in this case B and A coincide.

Secondly, we note that the semi-robust best-reply A function has the following monotonicity property: for any selections $T$and U in $\mathbb{T}$, if $T \subseteq U$, then $A(T) \subseteq A(U)$. This ensues from the fact that, for any δ, if c_j is a δ-conjecture with respect to $T$and $T \subseteq U$, then c_j is also a δ-conjecture with respect to U. By contrast the robust best-reply function B does not have this property as will become clear from the examples below. We also note that monotonicity can be used to show existence of semi-robust set: the property implies that if we iterate A starting from the whole strategy space S, we obtain a decreasing sequence which will end up in a fixed point in a finite number of steps.

Examples

In the remainder of this section, we present some illustrating examples, including a version of the famous Beer-Quiche game (Cho and Kreps, 1982); cf. Example 3.

[4]This characterization works for the present examples but is not correct in general as will be seen when the best-reply functions are defined formally in Section 3.

	E	F	G	H
A	3,1	3,1	0,0	0,0
B	0,0	0,0	1,3	1,3
C	3,1	0,0	3,1	0,0
D	0,0	1,3	0,0	1,3

Figure 1: Game G_1.

Example 1: Prevalence of Internal Uncertainty and Non-Monotonicity of B

The first example, game G_1 (cf. Figure 1), illustrates that the non-monotonicity of B is closely linked to the prevalence of internal uncertainty.[5]

To see this, note first that game G_1 has no weakly dominated strategies. As a result, iterated weak dominance and, a fortiori, rationalizability have no bite. Moreover the whole game is a robust set and also a semi-robust set. One may also note that the game has eight pure-strategy Nash equilibria none of which satisfies external stability with respect to B or A.

What we are looking for now are minimal robust and semi-robust sets (a robust set – semi-robust set – is minimal if it does not contain a smaller robust set – semi-robust set). As can be seen, there is a unique minimal semi-robust set, namely $\{A, C\} \times \{E, F, G\}$. This selection, however, is not very plausible in a context where solution candidates are being recommended to the players in a pre-play phase: if player 2 expects player 1 to play either A or C, playing E seems like the only rational choice. – Note that, under the proposed selection, strategies F and G are relatively dominated. – The selection $\{A, C\} \times \{E, F, G\}$ is not a robust set, though: under ε-

[5]This game can be seen as a two-stage cheap-talk game in which player 1 first sends one out of two available signals to player 2 after which the players simultaneously choose actions in a 2×2 game with two strict Nash equilibria with payoffs $(3, 1)$ and $(1, 3)$. Here, however, we disregard this dynamic interpretation and assume that the game is completely static.

conjectures player 2 will consider both A and C much more likely than B and D and, hence, will strictly prefer to play E (as F and G are relatively weakly dominated). The selection $\{A, C\} \times \{E\}$, by contrast, is a robust set and makes good sense also for player 1: when player 2 plays E with high probability, only A and C can be optimal for player 1. The reason why $\{A, C\} \times \{E\}$ is not a semi-robust set is that, under that notion, player 2 is allowed to attach more weight to deviations than to some solution strategy. For instance, if player 2 has δ-conjecture $(1 - [2 \cdot 2\delta + 3\delta], 2\delta, 3\delta, 2\delta)$ on $S_1 = \{A, B, C, D\}$, F will be optimal for δ small enough.

It is instructive to point out that the robust set $\{A, C\} \times \{E\}$ relies on a violation of monotonicity of the best-reply correspondence B. To see this, notice first that B and A, obviously, agree on the singletons $\{A\} \times \{E\}$ and $\{C\} \times \{E\}$:

$$B(\{A\} \times \{E\}) = A(\{A\} \times \{E\}) = \{A, C\} \times \{E, F\}$$

and

$$B(\{C\} \times \{E\}) = A(\{C\} \times \{E\}) = \{A, C\} \times \{E, G\}.$$

Nevertheless, as $\{A, C\} \times \{E\}$ is a robust set, we have

$$B(\{A, C\} \times \{E\}) = \{A, C\} \times \{E\}$$

while monotonicity would have required $B(\{A, C\} \times \{E\})$ to include both F and G. (Indeed, we have $A(\{A, C\} \times \{E\}) = \{A, C\} \times \{E, F, G\}$.) This shows that a robust sets may have more external stability than its component strategy profiles, a feature which clearly derives from the prevalence of internal uncertainty.

It should also be noted that, in addition to $\{A, C\} \times \{E\}$, there are three more minimal robust sets, viz. $\{A\} \times \{E, F\}, \{C\} \times \{E, G\}$, and $\{B, D\} \times \{H\}$. These solutions have similar justifications.[6]

A number of features can be observed in the example. First of all, in the present case the unique minimal semi-robust set contains smaller robust sets, and the minimal robust sets are more numerous than the minimal semi-robust sets. Yet, as semi-robust sets always exist while robust sets sometimes do not, it should be clear that these two features do not hold in general. Substituting quasi-robust for robust, we will be able to prove (weak) versions of these observations in the following, though. Moreover, notice that minimal robust sets can overlap: for instance, $\{A, C\} \times \{E\}$ and $\{A\} \times \{E, F\}$ in Game G_1. This, again, derives from the prevalence of internal uncertainty. By contrast, such overlapping can never occur with minimal semi-robust sets.[7]

	L	R
T	1,1	1,1
M	0,1	2,0
B	1,0	0,1

Figure 2: Game G_2

[6] A tedious, but safe way to find all strict sets is to derive the best reply for each player to any selection using ε-conjectures and then look for fixed points. In Game G_1 this gives the following result (fixed points are underlined):

$B_1 : \underline{S_2 \to S_1}$,
$\{E, F, G\} \to \{A, C\}, \{E, F, H\} \to \{A, D\}, \{E, G, H\} \to \{B, C\}, \{F, G, H\} \to S_1$,
$\underline{\{E, F\} \to \{A\}}, \underline{\{E, G\} \to \{C\}}, \{E, H\} \to S_1, \{F, G\} \to \{A, C\}, \{F, H\} \to \{A, D\}$,
$\{G, H\} \to \{B, C\}, \underline{\{E\} \to \{A, C\}}, \{F\} \to \{A\}, \{G\} \to \{C\}, \underline{\{H\} \to \{B, D\}}$.
$B_2 : \underline{S_1 \to S_2}$,
$\{A, B, C\} \to \{E, G\}, \{A, B, D\} \to \{F, H\}, \{A, C, D\} \to \{E, F\}, \{B, C, D\} \to \{G, H\}$,
$\{A, B\} \to S_2, \underline{\{A, C\} \to \{E\}}, \{A, D\} \to \{F\}, \{B, C\} \to \{G\}, \underline{\{B, D\} \to \{H\}}$,
$\{C, D\} \to S_2, \underline{\{A\} \to \{E, F\}}, \{B\} \to \{G, H\}, \underline{\{C\} \to \{E, G\}}, \{D\} \to \{F, H\}$.

Thus, in addition to the four minimal robust sets, we also have the (non-minimal) robust set $S_1 \times S_2$. Semi-robust sets can be derived in a similar way.

[7] Proof: Assuming V is the non-empty intersection of two such sets, it is clear that $A(V) \subseteq V$. Iterating A from $A(V)$ on, one obtains – exploiting monotonicity – a decreasing sequence which necessarily ends up in a fixed point, i.e. V contains a semi-robust set, a contradiction.

Example 2: Possible Non-Existence of Robust Sets

Next, consider Game G_2, depicted in Figure 2, which illustrates the possible non-existence of robust sets. To see that there is no robust set in this game, notice first that strategy B is weakly dominated and, thus, cannot be part of any robust set. Secondly, assume that, for example, $\{T\} \times \{L\}$ is proposed as a solution. Although this would work for player 1, T being his only admissible best reply to L, strategy R has to be added for player 2 as he may consider player 1 more likely to deviate to B than to M. However, once R is added to the solution, we also need to add M which is player 1's unique best reply to R. Following a similar line of argument, we are then forced to move from $\{T, M\} \times \{L, R\}$ to $\{T, M\} \times \{L\}$ and, eventually, back to $\{T\} \times \{L\}$. Thus, what we obtain is a cycle, and it is easily checked that, whatever starting point we choose, we will end up in the same cycle.[8] In other words, the best-reply function B has no fixed point and, hence, there is no robust set.

In order to handle such situations, we subsequently propose the notion of a quasi-robust set as a possible remedy. Formally, this is defined as a selection of strategies which is part of a cycle generated by the best-reply function B and, for each player, consists of the strategies that appear at some point in the cycle. Accordingly, in game G_2, $\{T, M\} \times \{L, R\}$ is a quasi-robust set – actually, the only one. As we show below, quasi-strict sets always exist.

Example 3: The Beer-Quiche Game

As our last introductory example, the famous Beer-Quiche game introduced by Cho and Kreps (1987) provides another illustration of non-existence of robust sets. In the original (dynamic) version of this two-person game, Nature first chooses a type, Weak

[8]The non-monotonicty of B is crucial for this phenomenon: we have, for instance,

$$B(\{T\} \times \{L\}) = \{T\} \times \{L, R\}$$

and

$$B(\{T, M\} \times \{L\}) = \{T\} \times \{L\}$$

while monotonicity would have required $B(\{T, M\} \times \{L\})$ to include R.

or Strong, for player 1 with probabilities 0.1 and 0.9, respectively. Next, player 1, being informed of his type, sends either a "strong" (s) or a "weak" (w) signal. A true signal is costless, whereas a false signal costs 1 unit. Finally, player 2, who is only informed about the signal, decides whether to fight (f) or to retreat (r). If he fights, he will win or lose 1 unit depending on whether his opponent was Weak or Strong. Player 1 loses 2 whenever there is a fight. Game G_3 provides the normal form representation of this interaction (cf. Figure 3), which we will focus on for our discussion.[9]

	ff	fr	rf	rr
ss	0.9, -0.8	0.9, -0.8	**2.9, 0**	**2.9, 0**
sw	1, -0.8	1.2, -0.9	**2.8, 0.1**	**3, 0**
ws	0, -0.8	1.8, 0.1	0.2, -0.9	2, 0
ww	0.1, -0.8	2.1, 0	0.1, -0.8	2.1, 0

Figure 3: G_3: The Beer-Quiche Game, strategic form; the unique quasi-robust selection being indicated by bold corresponding payoffs. A strategy of the form xy for player 1 indicates x as his message if he is "strong" and y if he is "weak." For player 2, a strategy of the form xy indicates action x after getting message s and action y after getting message w.

[9]For the purpose of reference, we here provide the corresponding extensive form for the Beer-

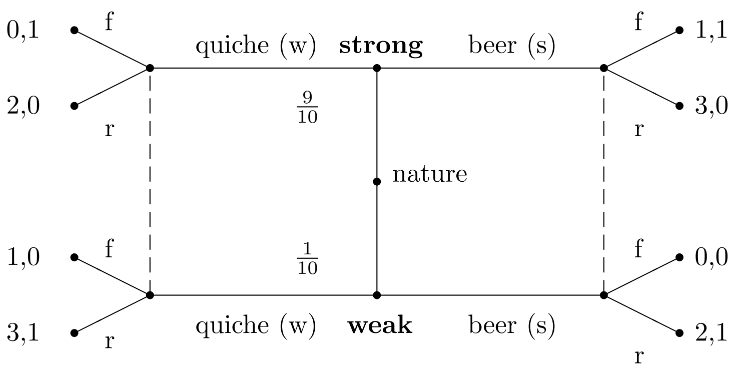

Quiche game.

The first thing to note is that the set of Nash equilibria for game G_3 consists of two connected components:

1. Both types of player 1 send the strong signal (s); player 2 retreats (r) if the signal is strong, and fights (f) with probability greater than or equal to 0.5 if the signal is weak (w).

2. Both types of player 1 send the weak signal; player 2 retreats if the signal is weak, and fights with probability greater than or equal to 0.5 if the signal is strong.

Arguably, equilibria in the first component make more sense than those in the second component where, counterintuitively, player 2's posterior belief about his opponent being strong after hearing a strong signal needs to be lower (at most 0.5) than his prior (0.9). Still many solution concepts fail to distinguish between these two types of Nash equilibria all of which are sequential, perfect, and proper. As Cho and Kreps (1987) note, however, only the first type of equilibrium satisfies strategic stability (Kohlberg and Mertens, 1986). The present approach provides partial support for this conclusion.

Applying the robust best-reply B function to game G_3, what is reached – irrespective of the specific starting point – is the following cycle

$$\{ss\}\times\{rf\} \to \{ss\}\times\{rf, rr\} \to \{ss, sw\}\times\{rf, rr\} \to \{ss, sw\}\times\{rf\} \to \{ss\}\times\{rf\}$$

. Hence, the game has no robust but a unique quasi-robust set, namely $\{ss, sw\} \times \{rf, rr\}$. Note that this set contains all Nash equilibria of the first component of equilibria described above and no other.

To see how this solution can be rationalised, consider a situation where the quasi-robust set $\{ss, sw\}\times\{rf, rr\}$ is recommended to the players. As in the game restricted to this set rr is dominated by rf, the latter strategy appears to be the only rational choice for player 2. Accordingly, considering this, player 1 should play ss. (Indeed, $\{ss\} \times \{rf\}$ is the unique robust set of the restricted game.) These conclusions about rational strategies, however, have been derived on the assumption that conjectures are concentrated on the quasi-robust set $\{ss, sw\} \times \{rf, rr\}$. In this sense, the quasi-robust set is part of the solution.

3 The Solution Concepts

Next, we provide the formal presentation of our model.

3.1 Preliminaries

For our analysis, we consider finite normal form games $G = \{I; S; u_i, i \in I\}$ with player set $I = \{1, ..., n\}$, set of pure strategy profiles $S = \times_{i \in I} S_i$ and payoff functions $u_i : S \to \mathbb{R}$ extended to mixed strategies in the usual way. $\mathbb{T}$ denotes the collection of all selections, i.e. all sets of the form $T = \times_{i \in I} T_i$ where each T_i is a non-empty subset of S_i. Moreover, for any finite set V, $\Delta(V)$ denotes the set of probability distributions on V.

A conjecture about player i's behaviour is an element c_i of $\Delta(S_i)$; we let $c_i(s_i)$ denote the probability associated with the pure strategy $s_i \in S_i$. A conjecture for player i about the other players' behaviour is an element c_{-i} of $\times_{j \neq i} \Delta(S_j)$ and a set of conjectures for i is a subset C_{-i} of $\times_{j \neq i} \Delta(S_j)$.[10] For a conjecture c_{-i}, we let $BR_i(c_{-i})$ denote the set of pure best replies for i to c_{-i}:

$$BR_i(c_{-i}) := \text{argmax}_{s_i \in S_i} u_i(s_i, c_{-i});$$

[10]Although, we do not allow correlated conjectures, this is more for pragmatic than ideological reasons. Dealing exclusively with uncorrelated conjectures will simplify both analysis and presentation.

for $c \equiv (c_1, ..., c_n) \in \times_{i \in I} \Delta(S_i)$, we write $BR(c) := \times_{i \in I} BR_i(c_{-i})$; moreover

$$BR_i(C_{-i}) := \bigcup_{c_{-i} \in C_{-i}} BR_i(c_{-i}).$$

Furthermore, the *stability set* $St(s_i)$ of a pure strategy s_i is defined as the set of conjectures for which s_i is a best reply:

$$St(s_i) = \{c_{-i} \in \times_{j \neq i} \Delta(S_j) : s_i \in BR_i(c_{-i})\}.$$

With these preliminary remarks, we now proceed to present our two models of strategic uncertainty. In both models, we assume that the opponents of player i consider some subset C_{-i} of $\times_{i \neq j} \Delta(S_j)$ to be the set of appropriate conjectures for player i. Given this set C_{-i}, the other players form conjectures about i's behaviour by determining which pure strategies i can be expected to use if he acts rationally.

3.2 δ-Conjectures and Semi-Robust Sets

In our first model conjectures are concentrated on candidate solution strategies but strategic uncertainty is otherwise unrestricted. Formally, players use δ-conjectures:

Definition 1 *Given $\delta \in (0,1)$ and $T_i \subseteq S_i$, $c_i \in \Delta(S_i)$ is a δ-conjecture about player i if*

(a) $c_i(s_i) > 0$ for any $s_i \in S_i$ and

(b) $c_i(T_i) > 1 - \delta$.

The set of all δ-conjectures about player i given T_i is denoted by $C_i^\delta(T_i)$.

Given a selection $T = T_1 \times \ldots \times T_n$, we use the notation $C^\delta(T) = \times_{i \in I} C_i^\delta(T_i)$ and $C_{-i}^\delta(T) = \times_{j \neq i} C_j^\delta(T_j)$. Clearly, for any $\delta \in (0,1)$ and T_i, $C_i^\delta(T_i)$ is non-empty.

To complete the model, we have to determine which pure strategies player i can be expected to use if he acts rationally given some set of appropriate conjectures C^{δ}_{-i}; these will be denoted by $T_i(C^{\delta}_{-i})$. Noting that $C^{\delta}_{-i}(T)$ is always a subset of positive measure in $\times_{j\neq i}\Delta(S_j)$, we let a strategy s_i be included in $T_i(C^{\delta}_{-i})$ if and only if the intersection of its stability set $St(s_i)$ with $C^{\delta}_{-i}(T)$ has positive measure.

Remark 1 *Note that the preceding assumption excludes not only strategies which are not best replies for any $c_{-i} \in C^{\delta}_{-i}(T)$ from $T_i(C^{\delta}_{-i})$, but also strategies which are best replies only on a subset of measure zero of $C^{\delta}_{-i}(T)$. The rationale for excluding the latter strategies is that if the other players' ideas about player i's conjectures are reasonably blurred, they should consider it most unlikely that player i will hold the conjectures required for choosing such strategies.*

	D	E
A	2,2	0,0
B	1,1	1,1
C	0,0	2,2

Figure 4: Game G_4

The situation is illustrated in game G_4 (cf. Figure 4): Assuming player 1 thinks both D and E are rational for player 2, his set of δ-conjectures about player 2 is given by $C^{\delta}_2(\{D,E\}) = \{c_2(D) : c_2(D) \in (0,1)\}$. Strategy B is only optimal – and then not uniquely so – under the conjecture $c_2(D) = 0.5$ which has measure zero in $C^{\delta}_2(\{D,E\})$. In this case, we assume that player 2 considers the event of player 1 having such a conjecture sufficiently unlikely to be ignored. Thus, player 2 will expect a rational player 1 to play A or C and player 1, even if he were to hold the conjecture $c_2(D) = 0.5$, will be aware of this fact.

Starting from a solution candidate T, we let $A_i^\delta(T)$ denote the pure strategies that player i can be expected to play if he acts rationally given appropriate δ-conjectures:

$$A_i^\delta(T) := T_i(C_{-i}^\delta(T)),$$

and write $A^\delta(T) := \times_{i\in I} A_i^\delta(T)$. The assumption of strict incentives requires that the solution candidate generates sets of conjectures which in turn generate sets of expected rational play which coincide with the solution candidate, i.e. what we are looking for is selections $T \in \mathbb{T}$ such that

$$A^\delta(T) = T.$$

A selection satisfying this condition will be called a *δ-robust set*. Semi-robust sets, then, are derived by taking limits. Specifically, we let A denote the limit of A^δ as δ goes to zero:

$$s_i \in A_i(T) \text{ if and only if, for any } \delta^* > 0, \text{ there is a } \delta \in (0, \delta^*) \text{ such that } s_i \in A_i^\delta(T).$$

Thus, we obtain the following definition:

Definition 2 *A selection T is a* semi-robust set *if $A(T) = T$.*

In the following, the mapping A will be considered as a function $A : \mathbb{T} \to \mathbb{T}$ and called the *semi-robust best-reply function*.

To characterise semi-robust and, later, robust sets, we need some basic definitions. Given a game G and a selection T, let $G(T)$ denote the game where players are restricted to strategies in T and inherit payoff functions from G. Thus, $G(S_i, T_{-i})$ is the game where i can use all his strategies in G while the other players are restricted to the selection T. Given a game G, a pure strategy s_i is *inferior* if its stability set $St(s_i)$ has Lebesgue measure zero in $\times_{j\neq i}\Delta(S_j)$; s_i is (*weakly*) *dominated* if there exists $\sigma_i \in \Delta(S_i)$ such that $u_i(\sigma_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$ for every totally mixed σ_{-i} in $\times_{j\neq i}\Delta(S_j)$. Clearly, dominated strategies are always inferior. Given a selection T, s_i is *relatively inferior* (*relatively dominated*) if s_i is inferior (dominated) in $G(S_i, T_{-i})$.

The following Proposition summarises basic properties of semi-robust sets.

Proposition 1

(a) Every finite strategic form game has a semi-robust set.

(b) Given a semi-robust set T, for each player i, if $s_i \in T_i$, then s_i is not inferior best reply to some $\sigma_{-i} \in \times_{j\neq i}\Delta(T_j)$.

(c) Given a semi-robust set T in a two-person game, for each player i, $s_i \in T_i$ if and only if s_i is not inferior and a best reply to some $\sigma_{-i} \in T_{-i}$.

(d) A semi-robust set may contain relatively dominated strategies

Proof. See Appendix. ■

To conclude this subsection, we return to the aforementioned equivalence between semi-robust sets and tight σ-curb sets. The notion of σ-curb sets as defined by Balkenborg et al. (2013, 2015) is based on what the authors call the most refined best-reply correspondence, which is defined as follows:[11] In a first stop, they define $s_i \in S_i$ to be a semi-robust best reply to $x_{-i} \in \times_{j\neq i}\Delta(T_j)$ if, for every neighbourhood U of x_{-i}, $St(s_i) \cap U$ contains an open set, and let $\Sigma_i(x_{-i})$ denote the set of semi-robust best replies to x_{-i}. Then, the most refined best-reply correspondence σ_i is given by $\sigma_i(x_i) := \Delta(\Sigma_i(x_{-i}))$.[12] Finally, σ-curb sets are defined using the function $\Sigma : \mathbb{T} \to \mathbb{T}$,

[11]Note that Balkenborg et al. (2013, 2015) restrict attention to a certain class of games – G^* in their paper – which is not necessary in the present context.

[12]The term σ here and in σ-curb sets does not refer to mixed strategies as in the present setting but to a best-response correspondence (cf. Balkenborg et al., 2013, 2015); we keep their notation as we do not see any hazard of confusion.

where $\Sigma(T) := \times_{i\in I}\Sigma_i(T)$, which maps any selection T on the selection which for each player i contains his semi-robust best responses to the strategies in $\times_{j\neq i}\Delta(T_j)$. A selection T is a σ-curb set if $\Sigma(T) \subseteq T$ and a tight σ-curb set if $\Sigma(T) = T$.

Proposition 2 *For any $T \in \mathbb{T}$, $\Sigma(T) = A(T)$.*

Proof. See Appendix. ■

Corollary 1 *A selection T is a tight σ-curb set if and only if it is a semi-robust set.*

3.3 ε-Conjectures and Robust Sets

Our second model is based on the assumption of prevalence of internal uncertainty: an appropriate conjecture about player i's behaviour should assign positive probability to all strategies in S_i but any strategy in T_i should be considered much more likely than any strategy not in T_i. This idea is formalised by means of the notion of ε-conjectures:

Definition 3 *Given $\varepsilon > 0$ and $T_i \subseteq S_i$, $c_i \in \Delta(S_i)$ is an* ε-conjecture *about i if*

(a) $c_i(s_i) > 0$ for any $s_i \in S_i$ and

(b) $c_i(s_i)/c_i(t_i) < \varepsilon$ whenever $s_i \in S_i \setminus Ti$ and $t_i \in T_i$.

The set of all ε-conjectures about player i given T_i is denoted by $C_i^\varepsilon(T_i)$.

Given a selection $T = T_1 \times ... \times T_n$, we define $C^\varepsilon(T) := \times_{i\in I}C_i^\varepsilon(T_i)$ and $C_{-i}^\varepsilon(T) = \times_{j\neq i}C_j^\varepsilon(T_j)$. Clearly, for any ε and T_i, $C_i^\varepsilon(T_i)$ is non-empty.

From here on, we proceed in the same way as in the above derivation of semi-robust sets, simply substituting ε-conjectures for δ-conjectures. Thus, if player i has conjectures C_{-i}^ε, the pure strategies he can be expected to use if he acts rationally, denoted by $T_i(C_{-i}^\varepsilon)$, are those whose stability sets have intersections of positive measure with C_{-i}^ε. Moreover, given a solution candidate T, we let $B_i^\varepsilon(T)$ denote the pure strategies that player i can be expected to use if he acts rationally given appropriate ε-conjectures, i.e.

$$B_i^\varepsilon(T) := T_i(C_{-i}^\varepsilon(T)),$$

and write $B^\varepsilon(T) := \times_{i\in I} B_i^\varepsilon(T)$. A selection T satisfying the strict-incentive condition $B^\varepsilon(T) = T$, then, is called an *ε-robust set*. Finally, we take limits letting B denote the limit of B^ε as ε goes to zero. From this, we obtain the following definition:

Definition 4 *A selection T is a* robust set *if $B(T) = T$.*

Proposition 3 below summarises some basic properties of robust sets.

Proposition 3

(a) Given a robust set T, for each player i, if $s_i \in T_i$, then s_i is neither inferior nor relatively dominated.

(b) Given a robust set T in a two-person game, for each player i, $s_i \in T_i$ if and only if s_i is neither inferior nor relatively dominated.

Proof. See Appendix. ∎

3.4 Quasi-Robust Sets

As observed earlier, there are normal form games for which no robust sets exist (cf. Example 2, Section 2). For these cases, we suggest an iterative procedure to find a solution which, in essence, still captures the spirit of our approach that all strategies be justifiable from an outside perspective given what is proposed.

Starting with an arbitrary selection T as a candidate solution, players compute their best replies under the function B and, thus, get $B(T)$ as the new candidate solution. Continuing in this way, they eventually end up in a cycle, i.e. a collection of selections $\mathbb{C} = \{T^1, ..., T^k\}$, such that $T^{j+1} = B(T^j)$ for $j = 1, ..., k-1$ and $T^1 = B(T^k)$. Once such a cycle has been reached, we take the next solution candidate to contain all strategies that appear in the cycle. Hence, for each player i, the next selection T^{k+1} obtained by the procedure is given by $T^{k+1} = \bigcup_{j=1,...,k} T_i^j$.

At this point, there are two possibilities depending on whether T^{k+1} is an element of $\mathbb{C}$ or not. If $T^{k+1} \in \mathbb{C}$, we consider T^{k+1} as a solution: iterating B from T^{k+1} on will just reproduce the cycle $\mathbb{C}$. If, on the other hand, $T^{k+1} \notin \mathbb{C}$, we suggest to proceed by iterating B starting from T^{k+1} until a new cycle is reached. The respective procedure is modelled below.

For a collection of selections $\mathbb{C} = \{T_1, ..., T_k\}$, $k \geq 1$, we use the following notation

$$U(\mathbb{C}) = \times_{i \in I} \bigcup_{j=1,...,k} T_i^j.$$

Each $T \in \mathbb{T}$ will be associated with an infinite path $P(T) = (P^0(T), P^1(T), ...)$ of elements from T as follows

(a) $P^0(T) = T$,

(b) $P^1(T) = B(T)$,

(c) for $j \geq 2$,

(i) if $P^{j-1}(T) \notin \{P^0(T), ..., P^{j-2}(T)\}$, then $P^j(T) = B(P^{j-1}(T))$, and

(ii) if $P^{j-1}(T) \in \{P^0(T), ..., P^{j-2}(T)\}$, then

$$P^j(T) = U(\{P^{k^*}(T), ..., P^{j-2}(T)\})$$

where $k^* := \max\{k : P^k(T) \in \{P^0(T), ..., P^{j-2}(T)\}\}$.

Remark 2 *In the above formulation, whenever $P(T)$ reaches an element $P^{j-1}(T) = V$ that has already occurred, we adopt the rule that the next element will be deter-*

mined by the elements between the penultimate occurrence of V and $P^{j-1}(T)$. It is easily shown that a rule which instead started at some other occurrence of V would be equivalent. For an informal proof, suppose V appears at least three times in $P(T)$. After the second occurrence of V the next element, say Y, will contain all the strategies that were used between its first and second occurrence. After the third occurrence of V, by definition, the next element contains all the strategies that were used between its first and third occurrence. However, as Y contains all the strategies that were used between its first and second occurrence and itself appears between the second and third occurrence of V, the strategies that appear between the second and third occurrence of V coincide with the strategies used between its first and third occurrence. Clearly, this argument can be generalised to any number of occurrences of V.

We say that a path $P(T)$ comes to a stop if there exist $Q(T) \in \mathbb{T}$ and $K \in \mathbb{N}$ such that $P^j(T) = Q(T)$ for any $j \geq K$. $Q(T)$ is then the solution reached starting from T. Any such solution will be called a *quasi-robust set*.

Definition 5 *Any limit point of a path $P(T)$ as defined above is called a* quasi-robust set.

Notice that, if T is a robust set, $P(T)$ comes to a stop and $Q(T) = T$. Thus, we obtain the following result.

Proposition 4 *Every robust set is quasi-robust.*

Clearly, we would like the path for any T to come to a stop and, thus, to be associated with a quasi-robust set. Another desirable property would be that if V is a quasi-robust set, i.e. $V = Q(T)$ for some T, then V is also the solution obtained starting from V itself, i.e. $V = Q(V)$. Yet another attractive property would be that all selections which appear on the same path end up in the same quasi-robust set, i.e.

if $V \in \{P(T)\}$, then $Q(V) = Q(T)$. The following Theorem implies that all these properties obtain.

Theorem 1

(a) For every $T \in \mathbb{T}$, $P(T)$ comes to a stop.

(b) There exist a partition $\{\mathbb{T}^1, ..., \mathbb{T}^k\}$ of $\mathbb{T}$ and a collection $\{Q^1, ..., Q^k\}$ of quasi-robust sets such that for each $j \in \{1, ..., k\}$ and $T \in \mathbb{T}_j$

(i) $Q(T) = Q_j$ and

(ii) $\{P(T)\} \subseteq \mathbb{T}_j$.

Proof. See Appendix. ■

Corollary 2 *Every game has a quasi-robust set.*

Moreover, quasi-robust sets exhibit the following properties.

Corollary 3 *If T is a quasi-robust set, then:*

(i) $P(T)$ comes to a stop at the second occurrence of T (the first occurrence being the first element of $P(T)$).

(ii) $T = U(\{P(T)\})$

(iii) Either T is a robust set or T is relatively dominated and $B(T)$ is a proper subset of T.

Proposition 5 *Every semi-robust set contains a quasi-robust set.*

Proof. See Appendix. ■

As the intersection of any two minimal semi-robust sets is empty – due to the monotonicity of $A(.)$ discussed in Section 2 – the following corollary is immediate.

Corollary 4 *In any game the number of minimal semi-robust sets is smaller than or equal to the number of minimal quasi-robust sets.*

Finally, any quasi-robust set T viewed as a restricted game obviously has a Nash equilibrium. By definition of quasi-robustness, the respective strategy combination also constitutes a Nash equilibrium of the unrestricted game.

Proposition 6 *Every quasi-robust set T of a normal form game G contains a Nash equilibrium of G.*

4 Discussion

Having provided the formal details of our approach, we proceed with some discussion about potential extensions and further illustrating examples.

4.1 Some Comments on Equilibrium Notions vs. Robust Sets

The notion of robust sets as defined above proposes a formal answer to the question for which combinations of strategies there is substantial reason to believe that rational players are going follow them once they are suggested as a solution - provided all relevant information about preferences is reflected in the payoffs.[13] Thus, contrary to the common aim of the refinement literature, the intention is not to identify particular equilibria that are more (or most) likely to be played in some strategic interaction given some specific line of reasoning. The reason for this is that equilibria often entail some implicit idea of repetition or average population behaviour, especially when mixed strategies are involved.[14] In single interactions, however, actual randomisation

[13]Note that if potential other-regarding concerns of the players as proposed, for example, by Rabin (1993), Fehr and Schmidt (1999) or Bolton and Ockenfels (2000) are not yet captured in the payoffs our analysis - as any argument in the refinement literature - loses empirical content.

[14]Myerson and Weibull, for example, explicitly argue in terms of a population average interpretation when providing intuition for the selection criteria underlying their solution concepts.

is difficult to discern and actual equilibrium play often almost impossible to identify. What can be said, though, is whether two (or more) pure strategies imply equal expected payoffs given some beliefs about others. Equilibrium arguments, then, might add valuable information regarding expected frequencies of specific strategies when observing multiple similar interactions. But for a single interaction all relevant information effectively concerns pure strategies, which is why the selections of strategies considered here only refer to those.

To exemplify this point by means of a simple example, consider the matching pennies game depicted in Figure 5. Obviously, the only Nash equilibrium of the game is the one in which both players play H and T with equal probability. Likewise the only robust set of the game is the full set of possible pure strategy profiles $S = \{H,T\} \times \{H,T\}$.

	H	T
H	-1,1	1,-1
T	1,-1	-1,1

Figure 5: G_5: Matching Pennies.

Put differently, while in the long run it is very reasonable to expect equal play of H and T, the only thing we can say for a single interaction is that both H and T are reasonable to be played. And this is what robust sets are intended to capture.

Of course, once some robust set is established as a solution candidate, further ideas from the refinement literature may be used to assess which of the strategies entailed in the proposed solution are more (or less) likely to be played if the interaction were to be observed more than once. Similarly, refinement arguments may be used to differentiate between robust sets. Consider, for example, the stag hunt game depicted in Figure 6. In this game, there are 2 strict (pure strategy) equilibria ((T,T) and

(B, B)) and a mixed equilibrium; the robust sets, in turn, are given by $\{T\} \times \{T\}$, $\{B\} \times \{B\}$ and $\{T, B\} \times \{T, B\}$.

	T	B
T	5,5	0,4
B	4,0	3,3

Figure 6: G_6: Stag Hunt

A standard equilibrium selection argument for this game would be in terms of risk dominance (Harsanyi, Selten, 1988), favouring (B, B). In a similar vein, one may argue for $\{B\} \times \{B\}$ as the more likely (or risk dominant) robust set. Yet again, such an argument, while certainly interesting to explore, would implicitly leave the idea of self-enforcing solutions for single interactions and thereby go beyond the scope of this paper.

4.2 Evaluating Conjectures

Related to the preceding discussion, a further question to ask is whether conjectures about opponents' play should differentiate only between strategies within the proposed solution and outside of it.

	T	B
T	1,1	0,0
B	0,0	1,1

Figure 7: G_7: A coordination game.

Consider, for example, the coordination game G_7 depicted in Figure 7. Note that G_7 has two strict Nash equilibria, the corresponding strategy sets of which also constitute minimal robust sets. As pointed out by Myerson and Weibull (2012, p. 949f), who discuss an extensive form version of this game, proposed solutions for this game tend to change once the outcome of a miscoordination, say (T, B), is replaced by a matching pennies type game with a unique (mixed strategy) equilibrium leading to zero payoffs in expected term (cf. Figure 8).[15] Obviously, what is a robust set for this type of game also depends on whether we analyse G_7 or G_{7a}.

	TL	TR	BL	BR
TL	1,1	1,1	2,-2	-2,2
TR	1,1	1,1	-2,2	2,-2
BL	0,0	0,0	1,1	1,1
BR	0,0	0,0	1,1	1,1

Figure 8: G_{7a}: Game G_7 with the cell (T, B) being replaced by a matching pennies game in which player 1 (rows) earns 2 in case of matching choices and loses 2 otherwise.

In our view, the fact that what is considered to be part of the solution (in the present setting) depends on the modelling of the interaction is not problematic, though. It rather emphasises the fact that once we ask "What would be justifiable ways to behave in a certain situation?" – which is not the question Myerson and Weibull try to answer – we have to be careful in choosing the model. Recall that with respect to applications almost any game is a drastic simplification of the corresponding "real life" situation to be commented on. That the choices made in the course of the simplification affect what can be viewed as justifiable, then, seems rather natural.

This notwithstanding, we believe that the dependence of solutions on the details of the modelling indicates an important point, namely which conjectures are to be

[15]Myerson and Weibull (2012) emphasise that their solution concept is independent of such changes; this is a consequence of them (different from us) focusing on particular equilibria to begin with.

thought of as relevant – a point that, in our view, should be taken into account in the abstract modelling of the real life situation. For the present coordination game example G_7/G_{7a}, we would, for instance, argue that a representation as in G_7 reflects the case where actual players have some experience in the matching pennies (sub-) game and, hence, can be assumed (in the modelling) to expect equilibrium behaviour there. A representation as in G_{7a}, by contrast, would rather reflect a case in which it is not reasonable to consider players as having any relevant experience in the matching pennies part.

An possible way to address such issues, of course, would be to introduce further requirements on eligible conjectures, which may or may not reflect rather informal assessments of the players' experience by the analyst.[16] However, while certainly a valuable avenue to consider, this would again have to be related to some external factors which are not captured by the purely technical description of the game, which is the basis for the arguments presented in this paper.

Finally, before leaving this discussion, we want to emphasise that the question which conjectures (or beliefs) about the likely play of opponents are reasonable to hold is one that is omnipresent in the refinement literature, with some requiring robustness of solutions to some trembles (and corresponding beliefs, e.g. Selten, 1975; Myerson, 1978; Kreps and Wilson, 1982) and others generally requiring robustness to any possible perturbations (e.g. Kohlberg and Mertens). The only requirement commonly made is that beliefs – inasmuch as considered – be consistent with actual strategies.[17] In the present context, we retain this assumption, albeit allowing for a broader basis of strategies to be considered – adding a 'strict incentives' condition (to remain within the proposed solution) to requiring robustness to 'strategic uncertainty.'

[16]The role of the analyst in choosing between technically eligible solutions is, for example, also emphasised by Kreps and Wilson (1982, p. 864f)

[17]See Carlsson and Wichardt (2012) for a discussion of how, in extensive games, existence of equilibria may rely on appropriate off equilibrium path choices of beliefs which follow no other logic than to support the existence of an equilibrium.

5 Concluding Remarks

In this paper, we have proposed the notion of a robust set as an answer to the question what would be a self-enforcing solution for a finite normal for game. As we have argued, the notion essentially captures both the core idea of the game theoretic refinement literature – namely robustness to strategic uncertainty – as well the key requirement of the various CURB-type notions – namely strict incentives for players to stay within the proposed solution.

While robust sets leave the realm of common equilibrium notions, they still rely on some form of coordination as the respective set of strategies is thought of as a proposal being made to the players of the game. Accordingly, strategies are not evaluated unconditionally but against a certain range of possible behaviours of others. The result, we believe, offers some kind of (weak) boundary to what might count as a self-enforcing solution.

Moreover, although robust sets are designed for the analysis of static games, they also indicate some possible avenues for dynamic analyses. More specifically, in a repeated interaction, robust sets of the static game could be seen as the a basis on which players might try different (justifiable) behaviours; recall that quasi-robust sets were defined based on an iterative procedure following a stepwise best response dynamic.

Thus, while we see robust sets as coming rather at the end of a discussion of static solution concepts, we believe that the approach offers various aspects that might be instructive to explore deeper also in the context of dynamic games.

Appendix

In the sequel, we provide the proofs for the results of our analysis.

Proof of Proposition 1

Parts (a) and (d) can be inferred from the examples in Section 2. Part (b) holds by definition. Part (c) can be seen as follows: In view of (b), we only need to show that if s_i is not inferior and a best reply to some $\sigma_{-i} \in \Delta(T_{-i})$, then $s_i \in T_i$. Now, for any $\delta > 0$, we can find $\eta > 0$ such that every totally mixed c_{-i} in the open set

$$C^\eta := \{c_{-i} \in \Delta^o(S_{-i}) : ||c_{-i} - \sigma_{-i}|| < \eta\}$$

is a δ-conjecture; $\Delta^o(S_{-i})$ denotes the interior of $\Delta(S_{-i})$. As s_i is not inferior, its stability set is a convex set with a non-empty interior. Moreover, $\sigma_{-i} \in St(s_i)$ so that $St(si) \cap C^\eta$ contains a non-empty open set. Therefore, $s_i \in T_i$. q.e.d.

Proof of Proposition 2

If $s_i \in A_i(T)$, then, for arbitrarily small δ, $St(s_i) \cap C^\delta_{-i}(T)$ has positive measure which implies that it contains an open set. Hence, there exists sequences $\{c_i^t\}_{t=1\ldots\infty}$ converging to some $\sigma_{-i} \in \times_{j\neq i}\Delta(T_j)$ and $\{?^t\}_{t=1\ldots\infty}$ converging to zero such that each c_i^t belongs to an open set in $St(s_i) \cap C^{\delta^t}_{-i}(T)$. Thus, s_i is a semi-robust best reply to σ_{-i}. Hence, $s_i \in \Sigma_i(T)$.

If $s_i \in \Sigma_i(T)$, then, there exists $x_{-i} \in \times_{j\neq i}\Delta(T_j)$ such that for every neighbourhood U of x_{-i}, $St(s_i) \cap U$ contains an open set. Clearly, for any δ, one can find a neighbourhood U of x_{-i} such that $U \subseteq C^\delta_{-i}(T)$. Thus, for any δ, $St(s_i) \cap C^\delta_{-i}(T)$ contains an open set and, consequently, has positive measure. Hence, $s_i \in A^\delta_i(T)$ for any δ and, as a result, $s_i \in A_i(T)$. q.e.d.

Proof of Proposition 3

(a) By definition, an inferior strategy cannot be part of $B^\varepsilon_i(T)$ for any ε. Assume some $s_i \in T_i$ is relatively dominated. Then, there exists $\sigma_i \in \Delta(S_i)$ such that $u_i(\sigma_i, s_{-i}) \geq$

$u_i(s_i, s_{-i})$ for any $s_{-i} \in T_{-i}$ with strict inequality for some $s'_{-i} \in T_{-i}$. In an ε-conjecture, strategy combinations such as s'_{-i} will be much more likely – in the limit infinitely more likely – than strategies not in T_{-i}. Thus, for sufficiently small ε, s_i cannot be a best reply to any ε-conjecture.

(b) In view of (a), we only need to show the "if" part. By Myerson (1991, Theorem 1.7) it is clear that if s_i is not relatively dominated, then s_i is a best reply to some totally mixed σ_{-i} in $\Delta(T_{-i})$. For any $\varepsilon > 0$, we can find $\eta > 0$ such that every totally mixed c_{-i} in the open set

$$C^\eta := \{c_{-i} \in \Delta^o(S_{-i}) : ||c_{-i} - \sigma_{-i}|| < \eta\}$$

is an ε-conjecture. ($\Delta^o(S_{-i})$ denotes the interior of $\Delta(S_{-i})$.)

As s_i is not inferior, its stability set is a convex set with a non-empty interior. Moreover, $\sigma_{-i} \in St(s_i)$ so $St(s_i) \cap C^\eta$ contains a non-empty open set. Therefore, $s_i \in T_i$.

q.e.d.

Proof of Theorem 1

Preliminaries

Given any game, consider the function $B : \mathbb{T} \to \mathbb{T}$. Iterating B starting from any T gives a sequence $B^\infty(T) = (T, B(T), B(B(T)) \ldots)$. Eventually this sequence will end up in a cycle $\mathbb{C}$. (We include the case of a degenerate cycle containing only one element, say V; in such a case $B(V) = V$ and V is a robust set.) We let $\mathbb{C}(T)$ denote the cycle in which $B^\infty(T)$ ends up. Clearly, we have:

- There exists a non-empty, finite collection $\mathcal{C}$ of cycles $\mathcal{C} = \{\mathbb{C}_1, \ldots, \mathbb{C}_K\}$ for some $K \geq 1$.

- There exists a partition $\mathcal{X} = \mathbb{X}_1, \ldots, \mathbb{X}_K$ of T such that $T \in \mathbb{X}_j$ iff $\mathbb{C}(T) = \mathbb{C}_j$.

The elements of $\mathcal{X}$ will be called cells. We let $\mathbb{X}(T)$ denote the cell to which T belongs.

The proof of Theorem 1 uses induction and is based on four lemmas. Lemma 1 to Lemma 3 mainly serve to show the proposition for the special case where $\mathcal{X}$ consists of a single cell while Lemma 4, then, is used in the inductive step of the proof.

Given a path $P(T) = (P_0(T), P_1(T), \ldots)$, as defined in Section 3.4, we define a sub-path $P_j^k(T)$ as follows $P_j^k(T) = (P_j(T), \ldots, P_k(T))$, where $j \in \mathbb{N}$, $k \in \mathbb{N} \cup \{\infty\}$, and $j \leq k$. Moreover $\{P_j^k(T)\} = \{P_j(T), \ldots, P_k(T)\}$.

The following claim will be used frequently in the sequel. It says that if a path $P(T)$ at a certain stage comes to a selection V, then $P(T)$ will mimick $P(V)$ from V on as long as $P(V)$ does not encounter any element that appeared in $P(T)$ before it came to V. The claim is immediate and stated without proof.

Claim 1 *If for some T, $V \in T$, $j \in \mathbb{N}$, and $k \in \mathbb{N} \cup \{\infty\}$,*

(a) $P_j(T) = P_0(V)$, and

(b) $\{P_0^k(V)\} \cap \{P_0^{j-1}(T)\} = \emptyset$,

then $P_j^{j+k+1}(T) = P_0^{k+1}(V)$.

A collection of selections $\mathbb{H} = \{T^1, \ldots, T^m\}$ is a *hypercycle* if, for each $T_j \in \mathbb{H}$, there exists $h(j) \in \mathbb{N}$ such that $\{P_0^{h(j)}(T_j)\} = \mathbb{H}$ and $P_{h(j)+1}(Tj) = U(\mathbb{H})$, i.e. $P_{h(j)+1}(Tj)$ is the union of all elements of the hypercycle. Clearly, every cycle is a hypercycle and every hypercycle contains a cycle. Notice that hypercycles share a crucial feature with cycles: if a path starts within a hypercycle, then it visits all its elemnts before, possibly, exiting.

The following lemmas show that every cell $\mathbb{X}$ contains a hierarchy of hypercycles which end up in a particular hypercycle $\mathbb{H}$ such that either $U(\mathbb{H}) \in \mathbb{H}$ or $U(\mathbb{H}) \notin \mathbb{X}$. The latter is the result in Lemma 2 while Lemma 1 provides the for the hierarchy.

Lemma 1 *If $\mathbb{H}$ is hypercycle such that $\mathbb{H} \subseteq \mathbb{X}_j$ for some $\mathbb{X}_j \in \mathcal{X}$ and $U(\mathbb{H}) \in \mathbb{X}_j \setminus \mathbb{H}$, then the collection $\{P_0^{m-1}(U(\mathbb{H}))\} \cup \mathbb{H}$, where*

$$m := min\{n : P_n(U(\mathbb{H})) \in \mathbb{H}\}$$

is also a hypercycle.

Remark 3 *As $U(\mathbb{H}) \in \mathbb{X}_j$ $\mathbb{H}$ and $\mathbb{C}_j \subseteq \mathbb{H}$, the existence of the set $P_0^{m-1}(U(\mathbb{H}))$ is guaranteed by the fact that the elements of $P(U(\mathbb{H}))$ until it reaches $\mathbb{C}_j$ are obtained by iterating B from $(U(\mathbb{H}))$ on.*

Proof. We define $\mathbb{H}' := \{P_0^{m-1}(U(\mathbb{H}))\} \cup \mathbb{H}$ and $V = P_m(U(\mathbb{H}))$. We need to show that, for every $T \in \mathbb{H}'$, $P(T)$ first visits all the elements of $\mathbb{H}'$ and no other and then moves to $(U(\mathbb{H}'))$.

If $T \in \mathbb{H}$, $P(T)$ first visits all the elements of $\mathbb{H}$ and then moves to $(U(\mathbb{H}))$. By Claim 1, $P(T)$ will then mimick $P(U(\mathbb{H}))$ until $V = P_m(U(\mathbb{H}))$ is reached. As V has already occurred in $P(T)$ its next element will now be determined by the sequence of selections having appeared since the last occurrence of V. Since $(U(\mathbb{H}))$ belongs to this sequence, the next element will be $U(\{P_0^{m-1}(U(\mathbb{H}))\} \cup \mathbb{H}) = U(\mathbb{H}')$ as desired.

If, instead, $T \in P_0^{m-1}(U(\mathbb{H}))$, say $T = P_j(U(\mathbb{H}))$, $P(T)$ will start with the subpath $P_j^m(U(\mathbb{H}))$ at which point V is reached. By Claim 1, $P(T)$ will then mimick $P(V)$ until $V = U(\mathbb{H})$ after which it will mimick $P(U(\mathbb{H}))$ until it returns to T. At this point it has visited all the elements of $\{P_0^{m-1}(U(\mathbb{H}))\} \cup \mathbb{H}$ and, thus, will move to $U(\{P_0^{m-1}(U(\mathbb{H}))\} \cup \mathbb{H}) = U(\mathbb{H}')$. ■

Lemma 2 *Every cell $\mathbb{X}_j$ contains a unique hypercycle $\mathbb{H}(\mathbb{X}_j) = \mathbb{H}_j$ such that either $U(\mathbb{H}_j) \in \mathbb{H}_j$ or $U(\mathbb{H}_j) \notin \mathbb{X}_j$.*

Proof. The result is obtained by constructing an increasing sequence $(\mathbb{H}_0, \mathbb{H}_1, \ldots)$ of hypercycles within $\mathbb{X}_j$ starting at $\mathbb{H}_0 = \mathbb{C}_j$. If $U(\mathbb{H}_0) \in \mathbb{H}_j$ or $U(\mathbb{H}_0) \notin \mathbb{X}_j$, we are done. Otherwise we note that, by Lemma 1, $\mathbb{H}_1 = \{P_0^{m-1}(U(\mathbb{H}_0))\} \cup \mathbb{H}_0$, for some $m \in \mathbb{N}$, is a hypercycle and, moreover, that $\mathbb{H}_0 \subset \mathbb{H}_1$. The increasing sequence is obtained by iterating this argument. As $\mathbb{X}_j$ is finite a hypercycle with the desired property will be reached in a finite number of steps. ■

The next Lemma shows that if $U(\mathbb{H}_j) \in \mathbb{H}_j$, then $U(\mathbb{H}_j)$is a quasi-robust set and every selection in $\mathbb{X}_j$ ends up in $U(\mathbb{H}_j)$.

Lemma 3 *If $U(\mathbb{H}_j) \in \mathbb{H}_j$ where $\mathbb{H}_j = \mathbb{H}(\mathbb{X}_j)$ and $T \in \mathbb{X}_j$, then $Q(T) = U(\mathbb{H}_j)$. Moreover $\{P(T)\} \subseteq \mathbb{X}_j$ and, if $T \in \mathbb{X}_j$, $\{P(T)\} \subseteq \mathbb{H}_j$.*

Proof. If $T \in \mathbb{H}_j$, $P(T)$ will first visit all the elements of Hj including $U(\mathbb{H}_j)$ and then move to $U(\mathbb{H}_j)$. The next element will contain all the strategies that have appeared since the last occurrence of $U(\mathbb{H}_j)$ and will, thus, clearly be $U(\mathbb{H}_j)$ once more. As a result $P(T)$ comes to a stop at $U(\mathbb{H}_j)$ on. As $P(T) = \mathbb{H}_j$ and $\mathbb{H}_j \subseteq \mathbb{X}_j$, clearly $\{P(T)\} \subseteq \mathbb{X}_j$.

If, instead, $T \notin \mathbb{H}_j$, $P(T)$ will start by iterating B from T on until it reaches some element, say V, of $\mathbb{H}_j$. At this point, by Claim 1, $P(T)$ will mimick $P(V)$ until it comes to a stop at $U(\mathbb{H}_j)$. Again, it is clear that $\{P(T)\} \subseteq \mathbb{X}_j$. ■

Finally, Lemma 4 below considers a modified best-reply function B^* which for elements in a particular hypercycle $\mathbb{H}_j$ moves directly to $U(\mathbb{H}_j)$ instead of first visiting all of $\mathbb{H}_j$. The lemma shows that the essential properties of the original best-reply function B will be preserved under this modification. This result will be used in the inductive step of the proof of Theorem 1 where its full significance will become clear.

Lemma 4 *Let $\mathbb{H}_j$ be a particular hypercycle as described in Lemma 2. If $B^* : \mathbb{T} \to \mathbb{T}$ is a modified best-reply function such that $P^*(T)$ comes to a stop for any T, $B^*(T) = U(\mathbb{H}_j)$ for $T \in \mathbb{H}_j$, and $B^*(T) = B(T)$ for $T \notin \mathbb{H}_j$, then, for any T,*

(a) $Q(T) = Q^*(T)$ and

(b) $\{P(T)\} \subseteq \{P*(T)\} \cup \mathbb{H}_j$

where $Q^*(T)$ *is the quasi-robust set reached from* T *and* $P^*(T)$ *is the path of* T *under* B^*.

Proof. If $\{P^*(T)\} \cap \mathbb{H}_j = \emptyset$, then $P(T)$ and $P^*(T)$ coincide so the lemma holds for such T. Now consider T such that $\{P^*(T)\} \cap \mathbb{H}_j \neq \emptyset$ and let V be the first element of $P^*(T)$ which lies in $\mathbb{H}_j$. By construction, $P(T)$ coincides with $P^*(T)$ up to the point where they both reach V. After this point $P^*(T)$ will move directly to $U(\mathbb{H}_j)$ while $P(T)$ will visit all the elements of $\mathbb{H}_j$ and then move to $U(\mathbb{H}_j)$. The continuation after this point depends on whether or not $U(\mathbb{H}_j) \in \mathbb{H}_j$.

If $U(\mathbb{H}_j) \in \mathbb{H}_j$ and $U(\mathbb{H}_j) = V$, $P^*(T)$ comes to a stop at $U(\mathbb{H}_j)$. If $U(H_j) \in \mathbb{H}_j$ and $U(\mathbb{H}_j) \neq V$, $P^*(T)$ next moves to $U(\mathbb{H}_j)$ and, thus, again comes to a stop at this selection. In order to see that $P(T)$ also comes to a stop $U(\mathbb{H}_j)$, first notice that, by Lemma 3, $P(V)$ comes to a stop at $U(\mathbb{H}_j)$ after visiting all the elements of $\mathbb{H}_j$. The result then follows from Claim 1: from V on, $P(T)$ will mimick $P(V)$. It follows that $\{P(T)\} = \{P^*(T)\} \cup \mathbb{H}_j$.

Now consider the case $U(\mathbb{H}_j) \notin \mathbb{H}_j$ and assume that both $P^*(T)$ and $P(T)$ encountered V, the first element in $\mathbb{H}_j$, at stage m. The first occurrence of $U(\mathbb{H}_j)$ in $P^*(T)$ will then happen at stage $m+1$ while $P(T)$ will encounter $U(\mathbb{H}_j)$ at some stage $n > m+1$ after visiting all the elements of $\mathbb{H}_j$. After their first encounter with $U(\mathbb{H}_j)$ both $P(T)$ and $P^*(T)$ will mimick $P(U(\mathbb{H}_j))$ as long as this path does not come across any element in $\{P_0^{n-1}(T)\}$ by Claim 1. If this never happens, clearly, $Q(T) = Q^*(T) = Q^*(U(\mathbb{H}_j))$. Again, it follows that $\{P(T)\} = \{P^*(T)\} \cup \mathbb{H}_j$.

If, instead, $\{P(U(\mathbb{H}_j))\} \cap \{P_0^{n-1}(T)\} \neq \emptyset$, we let Y be the first element of $P(U(\mathbb{H}_j))$ which lies in $\{P_0^{n-1}(T)\}$. First, we consider the case where Y also lies in $\{P_0^{*,m}(T)\}$. To see where $P(T)$ and $P^*(T)$ move after their second encounter with Y, let A denote the collection of selections occurring between the first and second appearance of Y in

$P(T)$ and let $\mathbb{A}^*$ be the corresponding collection for $P^*(T)$. We know that $\mathbb{A}$ consists of all elements of A^* and $\mathbb{H}_j$. As, however, $\mathbb{A}^*$ contains $U(\mathbb{H}_j)$, the strategies contained in $\mathbb{A}$ and $\mathbb{A}^*$ are the same, i.e $U(\mathbb{A}) = U(\mathbb{A}^*)$. This in turn implies that both $P(T)$ and $P^*(T)$ move to $U(\mathbb{A})$.

Next consider the case where Y does not belong to $P_0^{*,m}(T)$. Clearly this can only happen if $Y \in \mathbb{H}_j$ which means that $P^*(T)$ after its encounter with Y next moves to $U(\mathbb{H}_j)$ for a second time. Now, let B denote the collection of selections occurring between the first and second appearance of Y in $P(T)$, and let B^* denote the selections occurring between the first and second appearance of $U(\mathbb{H}_j)$ in $P^*(T)$. Again it is easily seen that $U(B) = U(B^*)$ so that both $P(T)$ and $P^*(T)$ next move to $U(B)$.

Whether $P(T)$ and $P^*(T)$ next move to $U(A)$ or $U(B)$, the proof can be completed by iterating the above argument, mutatis mutandi. The fact that $P^*(T)$ comes to a stop ensures that a common quasi-strict set will be reached in a finite number of steps.

■

Proof of Theorem 1

We now prove the claim of the theorem by induction on $\#\mathcal{X}$, the number of cells in $\mathcal{X}$.

If $\#\mathcal{X} = 1$, Theorem 1 holds by virtue of Lemmas 2 and 3: all selections come to a stop at $U(\mathbb{H}_1)$ which, thus, is the unique quasi-robust set.

Now assume the result holds for $\#\mathcal{X} = n$ and consider the case $\#\mathcal{X} = n + 1$ so that $\mathcal{X} = (\mathbb{X}_1, \ldots, \mathbb{X}_{n+1})$. If $U(\mathbb{H}_j) \in \mathbb{H}_j$ for every particular hypercycle $\mathbb{H}_j$, then any $T \in \mathbb{X}_j$ comes to a stop at $U(\mathbb{H}_j)$,$j = 1, \ldots, n + 1$, so Theorem 1 holds and the game has $n + 1$ quasi-robust sets.

If, instead, $U(\mathbb{H}_j) \notin \mathbb{X}_j$ for some j, we may assume this is the case for $U(\mathbb{H}_1)$ and, moreover, that $U(\mathbb{H}_1) \in \mathbb{X}_2$. We then consider a modified best-reply function $B^* : T \to T$ such that for $T \in H_1$, $B{*}(T) = U(\mathbb{H}_1)$ and for all other T, $B^*(T) = B(T)$. Iterating B^* to find the cycles under this function shows that the elements belonging to $\mathbb{X}_1$ will now end up in the cycle of $\mathbb{X}_2$. All other selections end up in the same cycles as under B. Thus, $\mathbb{X}_1 \cup \mathbb{X}_2$ will be a single cell under B^* and all the other cells will be as under B. As a result, there are now only n cells so the claim of Theorem 1 holds by assumption.

Item (a) of Theorem 1 now follows from the fact that, by Lemma 4 (a), $Q(T) = Q^*(T)$ for every T.

Regarding item (b) we need to show that, for every T, $Q(V) = Q(T)$ for any $V \in \{P(T)\}$. If $V \notin \mathbb{H}_1$, we know that $V \in \{P^*(T)\}$ from Lemma 4 (b). Thus, $Q^*(V) = Q * (T)$ as the proposition holds for B^* and $Q(V) = Q(T)$ by Lemma 4 (a). If, instead, $V \in \mathbb{H}_1$, we know that $\{P^*(T)\} \cap \mathbb{H}_1 \neq \emptyset$ as $P(T)$ and $P^*(T)$ obviously coincide until they reach an element of $\mathbb{H}_1$. Hence, by the definition of B^*, $U(\mathbb{H}_1) \in \{P^*(T)\}$. Similarly, we get $B^*(V) = U(\mathbb{H}_1)$. As a result $Q^*(V) = Q^*(U(\mathbb{H}_1)) = Q^*(T)$ and, thus, again $Q(V) = Q(T)$ by Lemma 4 (a).

Note that although $\{P_0^n(T)\}$ contains some elements that are not in $\{P_0^{*,m+1}(T)\}$, all these elements belong to $\mathbb{H}_j$. So we have $U(\{P_0^n(T)\}) = U(\{P_0^{*,m+1}(T)\}) = U(\{P_0^{*,m-1}(T)\} \cup \mathbb{H}_j)$. q.e.d.

Proof of Proposition 5

The proof of Proposition 5 again uses three initial steps before the proposition itself is proven.

Lemma 5 *For any T, $B(T) \subseteq A(T)$.*

Proof. First note that for any $\delta > 0$ there exists $\varepsilon(\delta)$ such that, for any $\varepsilon \in (0, \varepsilon(\delta)]$, every ε-conjecture is also a δ-conjecture. This can be seen by setting $\varepsilon(\delta) = \delta/M$, where M is the largest number of pure strategies for any player. As a result $C_{-i}^{\varepsilon}(T) \subseteq C_{-i}^{\delta}(T)$ for any T and $\varepsilon \in (0, \varepsilon(\delta)]$. This, in turn, implies $B^{\varepsilon}(T) \subseteq A^{\delta}(T)$ and, thus, $B(T) \subseteq A(T)$ for any T. ■

Before moving on to the second step, we need some further terminology. Given any function $F : \mathbb{T} \to \mathbb{T}$, say that a selection T is closed under F if, for any selection V such that $V \subseteq T$, we have $F(V) \subseteq T$.

Lemma 6 *If T is a semi-robust set, then T is closed under A.*

Proof. By definition, $A(T) \subseteq T$. The result, then, follows from the monotonicity property that for any selections $T$and V in T, if $V \subseteq T$, then $A(V) \subseteq A(T)$. ■

The next Lemma is an immediate consequence of the two previous ones.

Lemma 7 *If T is a semi-robust set, then T is closed under B.*

Finally, we can prove Proposition 5: From Proposition 3, we know that the iterative procedure starting at any selection$V$ comes to a stop at a quasi-robust set $Q(V)$. Thus, what we need to show is that if T is a semi-robust set and $V \subseteq T$, then $Q(V) \subseteq T$.

Consider the path $P(V)$ and assume that $P_h(V) \subseteq T$ for any $h < k$ and $kin\mathbb{N}$. Now consider $P_k(V)$. If $P_k(V) \notin P_0^{k-1}(V)$, then $P_k(V) = B(P_{k-1}(V))$. As $P_{k-1}(V) \subseteq T$, so is $P_k(V)$ by Lemma 7. On the other hand, if $P_k(V) \in P_0^{k-1}(V)$, then $P_k(V) = U(\mathbb{C})$, where $\mathbb{C}$ is some collection of selections from $P_0^{k-1}(V)$. Hence $U(\mathbb{C}) \subseteq T$. q.e.d.

References

Asheim, G., Dufwenberg, M., 2003. Admissibility and Common Belief. *Games and Economic Behavior* 42, pp. 208-234.

Balkenborg, D., Hofbauer, J., Kuzmics, C., 2013. Refined Best Reply Correspondence and Dynamics. *Theoretical Economics* 8, pp. 165-192.

Balkenborg, D., Hofbauer, J., Kuzmics, C., 2015. Refined Best-Response Correspondence in Normal Form Games. *International Journal of Game Theory* 44, pp. 165-193.

Basu, K., Weibull, J., 1991. Strategy Subsets Closed under Rational Behavior. *Economics Letters* 36, pp. 141-146.

Bernheim, D., 1984. Rationalizable Strategic Behavior. *Econometrica* 52, pp. 1007-1027.

Bolton, D., Ockenfels, A., 2000. ERC: A Theory of Equity, Reciprocity, and Competition. *American Econonomic Review* 82, pp. 166-193.

Brandenburger, A., Friedenberg, A., Keisler, H.J., 2008. Admissibility in Games. *Econometrica* 76, pp. 307-352.

Brandenburger, A., Friedenberg, A., 2010. Self-Admissible Sets. *Journal of Economic Theory* 145, pp. 785-811.

Carlsson, H., Wichardt, P., 2012. On the Belief (In-)Dependence of Sequential Equilibria. *Economics Letters* 115, pp. 504-507.

Cho, I.-K., Kreps, D., 1987. Signaling Games and Stable Equilibria. *Quarterly Journal of Economics* 102, pp. 179-221.

Fehr, E., Schmidt, K., 1999. A theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics* 114, pp. 817-868.

Harsanyi, J., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games.* MIT Press

Kalai, E., Samet, D., 1984. Persistent Equilibria in Strategic Games. *International Journal of Game Theory* 13, pp. 129-144.

Kohlberg, E., Mertens, J.-F., 1986. On the Strategic Stability of Equilibria. *Econometrica* 54, pp. 1003-1037.

Kreps, D, Wilson, R., 1982. Sequential Equilibria. *Econometrica* 50, pp. 863-894.

Myerson, R., 1978. Refinements of the Nash Equilibrium Concept. *International Journal of Game Theory* 15, pp. 133-154.

Myerson, R., Weibull, J., 2015. Tenable Strategy Blocks and Settled Equilibria. *Econometrica* 83, pp. 943-976.

Nash, J., 1950. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences* 36, pp. 48-49.

Pearce, D., 1984. Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* 52, pp. 1029-1050.

Rabin, M., 1993. Incorporating fairness into game theory and economics. *American Economic Review*, 83, pp. 1281-1302.

Rubinstein, A., 1991. Comments on the Interpretation of Game Theory. *Econometrica* 59, pp. 909-924.

Selten, R., 1975. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* 4, pp. 25-55.