TRIMMED STRATEGIES: ACHIEVING SEQUENTIALLY RATIONAL EQUILIBRIA WITH ONLY PARTIALLY SPECIFIED STRATEGIES

by

Valeska Groenert

Faculty Sponsor: Myrna Wooders

Working Paper No. 08-WG01

January 2008

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

Trimmed strategies: achieving sequentially rational equilibria with only partially specified strategies[1]

Valeska Groenert
Vanderbilt University

January 8, 2008

[1] I would like to thank Rabah Amir, Andrew Daughety, Rosa Ferrer, Alper Nakkas, Jennifer Reinganum, Quan Wen, Myrna Wooders, and participants of the 2006 summer research presentations at Vanderbilt University and of the 2007 SAET conference for helpful comments and discussions.

Abstract

Both subgame-perfect equilibrium and weak-perfect Bayesian equilibrium impose rationality at information sets that are irrelevant for a sequentially rational outcome. In this paper, for each of these equilibrium concepts, I characterize the maximal set of such information sets. These information sets can be trimmed (or removed) from strategy sets without affecting the equilibrium outcome. A trimmed equilibrium is an equilibrium for the game where the irrelevant information sets have been ignored. It is shown that the trimmed version of an equilibrium concept, either subgame-perfect or weak perfect Bayesian equilibrium, is just sufficiently restrictive to ensure that equilibrium outcomes are consistent with the original concept, i.e., ignoring any additional information sets could change the equilibrium outcome. An example demonstrates that trimming irrelevant information sets can lead to existence of an outcome consistent with subgame-perfect rationality in cases where a subgame perfect equilibrium does not exist.

JEL Classification Number: C72 (Noncooperative Games)

Keywords: Nonexistence of equilibrium, partially specified games, subgame-perfect Nash equilibrium, weak-perfect Bayesian equilibrium

1 Introduction

Many refinements of Nash equilibrium (NE) incorporate some notion of sequential rationality. Such refinements serve to rule out outcomes sustained by irrational play, in particular by noncredible threats. For example, a subgame-perfect equilibrium (SPE) incorporates sequential rationality by requiring NE in each subgame, and a weak perfect Bayesian equilibrium (WPBE) incorporates sequential rationality by requiring best responses at each information set given the (Bayesian) belief at this information set, and given future play. This paper establishes that both the SPE and the WPBE can be more restrictive than necessary to ensure outcomes which are consistent with their particular notion of sequential rationality. More precisely, given an extensive form game and given a SPE (a WPBE), this paper characterizes the set of information sets off the equilibrium path which is maximal with respect to the property that choices within this set never affect the sequential rationality as imposed by a SPE (a WPBE) outside of the set.[1] A collection of information sets W' satisfies this property if for any pair of information sets w, w' with $w' \in W'$ and $w \notin W'$, whether the choice specified at w is sequentially rational does not depend on the payoff following w'. Hence, it is possible to relax the definitions of SPE and WPBE by dropping the requirement of sequential rationality in the maximally irrelevant set without losing the equilibrium property that the outcome is never *sustained* by sequentially irrational play. A "relaxed" SPE and WPBE, respectively, will be called trimmed SPE and trimmed WPBE. This terminology is chosen to emphasize that strategies in a trimmed equilibrium can be seen as "smaller" strategies, because the play at irrelevant sets does not even need to be specified. In that sense, in a trimmed equilibrium, players are not required to make complete contingent plans. It will be shown that a trimmed SPE (WPBE) can not be relaxed further. Its conditions are necessary and sufficient to ensure outcomes are not sustained by the respective notion of sequentially irrational play.

Knowing which parts of a game are irrelevant for a particular equilibrium proves helpful in cases when there is nonexistence of equilibrium in the whole game due to nonexistence of equilibria in a subset of subgames. How, in such cases, the results of this paper can help finding an outcome that is nonetheless consistent with sequential rationality is demonstrated in an example (Section 4). Theorems 6 and 13 and their corollaries show that a trimmed equilibrium, even though less restrictive than its original concept, does not give up on the degree of rationality imposed: provided the original equilibrium exists, the sets of outcome paths for the original concept and for its trimmed version are the same.[2]

For motivation, consider the classical Hotelling game. In the first stage, two sellers of a homogenous good simultaneously choose a location on an interval. In the second stage, they simultaneously set a price for their goods, which at that time are differentiated by location. Consumers, who demand a single unit of the good, are uniformly distributed over

[1] In principle the characterization of inessential game parts could be done for other equilibrium concepts as well, but is not examined here further.

[2] Beside the potential usefulness of the results in this paper, one can also interpret them as calling into question whether individuals actually always make "complete contingent plans", as standard game theory assumes they do. To be clear, this is different from bounded rationality. The point is not that the degree of computational difficulty might lead an individual to be unable to determine the best response, but to able to do so only at a high cost. Intuitively, one can think of a player asking himself "why bother about this problem right now? It is not going to happen if we all stick to our plans. *If* it arises there is still time to worry.

the interval and incur a cost of visiting a particular seller. This cost is linearly increasing in the distance to the seller. As d'Aspremont et al. (1979) show, pure strategy price equilibria in the second stage exist only for a subset of the sellers' location pairs (each location pair represents a subgame of the second stage).[3] Therefore, a SPE in pure strategies fails to exist.

A trimmed SPE, as it will be defined in this paper, requires existence of equilibrium only for a subset of subgames. Figure 1 illustrates how this set of relevant subgames could look like. Suppose we could find pure strategy pricing equilibria for the location pair (l_1^*, l_2^*), and for all subgames in $G(l_1^*, l_2^*) = \{(l_1, l_2) \in [0,1]^2 \subset \mathbb{R}^2 : l_1 = l_1^* \text{ or } l_2 = l_2^*\}$. Suppose further that, given that player j chooses location l_j^*, player i's payoff across all pricing equilibria of subgames in $\{(l_1, l_2) \in [0,1]^2 \subset \mathbb{R}^2 : l_j = l_j^*\}$ is maximized at l_i^*. Then, let each player's strategy be to choose, respectively, l_1^* and l_2^* in the first stage and to play the pricing equilibria in subgames in $G(l_1^*, l_2^*)$ in the second stage and to set an arbitrary price in all other subgames. Such a strategy profile would constitute a trimmed equilibrium.

Note the following two features of this strategy profile. First, assuming that neither of the players makes mistakes, there is no reason for any of them to deviate from their strategy. Albeit the play in some of the subgames might be noncredible (not a NE of the subgame), the play cannot be a noncredible *threat*: suppose firm 1 would like to avoid subgame $(l_1, l_2) \notin G(l_1^*, l_2^*)$. Given that firm 1 itself chooses l_1^* in the first stage, it would require a deviation by both firms to reach (l_1, l_2), i.e., firm 2 *cannot* reach (l_1, l_2) and therefore it is not possible that some announcement by firm 1 keeps it from deviating to there. Second, if along with these strategies we were able to specify pricing equilibria for all irrelevant subgames in $[0,1]^2 \diagdown G(l_1^*, l_2^*)$, the resulting strategy profile would constitute a SPE (Corollary 7 proves this claim). Thus, an equilibrium as described above is as reasonable as a SPE, ruling out all Nash equilibria that are possibly sustained by noncredible threats ("if you choose location xy, I am going to set my price equal to zero").

While in the Hotelling game determining the set that is maximally irrelevant with respect to SPE might be relatively straightforward, the task of characterizing maximally irrelevant sets for general games is not trivial. In the example, a subgame belonged to the maximally irrelevant set if it could only be reached by a deviation of both players in the first stage. What is the equivalence to this in games with multiple stages, or dynamic games, which cannot be divided into several "stages"? The answer requires a partitioning of information sets into equivalence classes so that two information sets belong to the same class if and only if the minimal subgame containing them is the same. Then an information set is irrelevant if any deviation leading to the information set requires, in at least one of the equivalence classes, the deviation of at least two players.

---

[3] Non-existence is due to the discontinuity of the payoff functions over the price space. Intuitively the discontinuity arises because, given a location pair and a price set by its competitior, for prices below a certain threshold the firm will be able to attract all consumers, while at and above the threshold, each firm attracts a positive share of consumers.

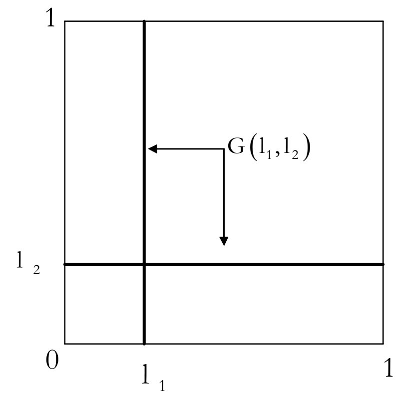

Figure 1: Set of subgames in the Hotelling game

Even if existence of equilibrium in subgames is not an issue, a trimmed equilibrium can be a useful tool for the analysis of a game. An indication of the extent to which such a technique can potentially simplify an analysis are the papers by Osborne and Pitchik (1987) and Kreps and Scheinkman (1983). Both papers deal with similar setups as in the Hotelling game described above. Pure strategy SPE do not exist, but mixed strategy SPE are known to exist due to a result in Dasgupta and Maskin (1982). At the cost of an intricate analysis, Kreps and Scheinkman derive the mixed strategy equilibrium, while Osborne and Pitchik derive approximations of an equilibrium, unable to provide a complete characterization.[4] The example provided in Section 4 takes up the game in Kreps and Scheinkman and demonstrates how the outcome they derive can be obtained in a greatly simplified way. In this connection, one drawback of a trimmed equilibrium should be pointed out. While a trimmed equilibrium, potentially, needs a significantly reduced amount of computations, the actual practicability of the concept relies on having a method of solving for a trimmed equilibrium that also uses fewer computations. I have not been able to find such a method. Until now, an application essentially relies on having an initial guess about the equilibrium outcome.

While there is no directly related prior literature, a somewhat related strand of prior research is concerned with the robustness of an equilibrium towards the game's specifications (such as number of players, or order of moves). Because often strategic situations are not completely specified, it is of interest to know which properties of an equilibrium guarantee its robustness towards the game's specifications. One part of the literature deals with robust equilibria in large games, where the number of (identical) players is uncertain (see for example Kalai 2004, 2005). Games with uncertain features are usually called "partially specified". Instead of determining equilibrium properties that guarantee robustness of equilibrium towards the specifics of a game, the present paper determines entire parts of a game that are irrelevant for the equilibrium path.

In another paper, Kalai and Neme (1992) introduce the concept of a p-subgame perfect equilibrium, which requires subgame perfection after histories with no more than p deviations from the equilibrium path. At first sight this concept appears to be similar to the approach in the present paper. However, their motivation is to rationalize behavior other than the one resulting from SPE play. They do so, arguing that after a certain number of deviations it is rational for players not to expect 'rational' behavior by other players. Accordingly, a

[4] The complete characeterization fails due to the computational complexity of the task. The authors proceed by identifying eight subgame subsets across which the kinds of mixed strategy equilibria vary. They compute equilibria for a grid of subgames, and use approximations for others.

p-subgame perfect equilibrium can result into outcomes not sustainable by a SPE.[5] This difference to a trimmed SPE also shows that the maximally irrelevant set of a game cannot be characterized by numbers of deviations.

Recently, and independently from my work, Briata, Garcia-Jurado, Gonzalez-Diaz, and Patrone have addressed similar questions. They identify what they call the "essential" collection of information sets - information sets that are sufficient for a particular equilibrium. However, our concepts do not coincide. In particular the complement of an "essential collection" is not equivalent to a maximally irrelevant set.[6] Also, while they focus on providing a general and unified framework for what they call essentializing equilibria, I focus on two widely used equilibrium concepts, for which irrelevant sets can be particularly large. I also provide results for infinite games and demonstrate the potential usefullness of the concept with an example.

To summarize, the main contributions of the paper are: (a) to demonstrate that SPE and WPBE can be unnecessarily restrictive; (b) to characterize maximally irrelevant sets; and (c) to provide modified definitions, shown to be just sufficiently restrictive to capture the original notion of sequential rationality, and (d) to demonstrate how these modified definitions can be useful in determining sequentially rational equilibrium outcomes. In addition, the last section of the paper briefly introduces the concept of a dominated subform, which can be seen as an alternative, though weaker, approach to eliminate unreasonable NE.

The rest of the paper is organized as follows. Section 2 introduces basic notation and informally discusses the reasoning behind the concepts to follow. Section 3 deals with the maximally irrelevant set for SPE and the trimmed SPE. Section 4 provides a detailed example. Section 5 parallels Section 3, but applies to WPBE. An additional tool to analyze games with uncertainties in subparts is introduced in Section 6. Section 7 concludes.

2 Preliminaries

This section introduces notation and provides a more informal discussion about the reasoning behind the results to follow. Because the information structure of a game will be important for all definitions and results the analysis is performed for extensive form games. Two simplifications should be pointed out: while games can be infinite, whenever a player has to move his choices are finite, and players have complete information. Intuitively, the concept should generally apply and the results established in the following sections should remain valid when these two assumptions are relaxed.[7] I also assume perfect recall. All terminology not introduced explicitly (for example, a path, a rooted tree etc.) is used in the standard game theoretic or graph theoretic sense.[8]

2.1 Notation

A reference for all notation used in this paper can also be found in Table 1 in the appendix.

[5] As long as p is not larger than the largest possible number of deviations in the game.

[6] Basically, this is because they require essential collections to be closed under $\leq$ (roughly meaning that if an information set is in the essential collection, so are all its predecessors). Another difference originates in their definition of essential collections for belief-based concepts. An information set belongs to the essential collection if it is relevant under *some* belief, while here irrelevant sets depend on a specific belief.

[7] Briata, Garcia-Jurado, Gonzalez-Diaz, and Patrone (2007) apply their concept to incomplete information games.

[8] The description of an extensive form game follows Selten (1975) and van Damme (1981).

Extensive form games

The following description of an extensive form game is only introduced to have a clearly defined notation for information sets, choices, and strategies. Other details of the description will not be used later on in the paper. An extensive form game is a quintuple $\Gamma = (T, \mathcal{P}, \mathcal{W}, \mathcal{C}, u)$, where

1. $T = (X, E)$ is a rooted tree with X being the set of vertices and E being a set of (unordered) pairs from X. The origin (root) of the tree is denoted by x_0. Furthermore, for every vertex x, its immediate predecessors and the set of its immediate successors is denoted by $succ(x)$ and $pre(x)$, respectively.[9] The set Z contains terminal nodes of T, i.e. nodes x with $succ(x) = \varnothing$.

2. $\mathcal{P} = (P_1, ..., P_n)$ is a partition of the set $X \setminus Z$ into n sets, one for each player $i \in I = \{1, ..., n\}$.

3. $\mathcal{W} = (W_1, ..., W_n)$ is an information partition, where W_i is a partition of P_i into information sets of player i, so that

   (a) every path from the origin intersects the information set at most once, and

   (b) nodes in the same information sets have the same number of immediate successors.

   Let $W = \cup W_i$.

4. $\mathcal{C} = \{\mathcal{C}_w\}_{w \in W}$ is a collection of partitions. Partition $\mathcal{C}_w$ divides nodes in $\cup_{x \in w} succ(x)$ into the choices available at information set w (which will be assumed to be finite), so that every choice contains exactly one element of $succ\,(x)$ for every $x \in w$. A generic choice at w (a member of the partition $\mathcal{C}_w$) is denoted by c_w and the set of choices at w is denoted by C_w.

5. $u = (u_1, ..., u_n)$ are n real-valued payoff functions, one for each player i, with domain $\bar{\mathcal{H}}$, the set of terminal histories. The payoff functions are assumed to be von Neumann-Morgenstern expected utility functions.

Behavior strategies

Let f_w denote a probability distribution over the members of the partition $\mathcal{C}_w$, let $\mathcal{F}_w$ be the set of all probability distributions over choices in w, and let $f_w\,(c_w)$ denote the probability that probability distribution f_w assigns to choice c_w. A (behavior) strategy for player i is a mapping $s_i : \{w\}_{w \in W_i} \longrightarrow \cup_{w \in W_i} \mathcal{F}_w$ such that $s_i\,(w) \in \mathcal{F}_w$. A profile of strategies is denoted by s. As system of beliefs is denoted by μ. Also, s^g is the strategy profile for subgame g induced by strategy profile s. It will also be useful to let $s(x, x')$ denote the probability that s attaches to a choice leading from x to x'. Given a set of information sets $\hat{W} \subset W$, and two strategy profiles s' and s'', $\left(s'_{\hat{W}}, s''_{W \backslash \hat{W}}\right)$ is the profile obtained by playing s' on $\hat{W}$, and s'' on $W \backslash \hat{W}$.

Notation concerning the structure of a game tree

$\iota(w)$ is the player who "owns" information set w, i.e., the player i such that $w \subset P_i$.
h is a history in Γ, i.e., h is a path of T starting at x_o.

[9] The order is naturally given by the distance to the origin.

$\mathcal{H}$ is the set of all histories in T.
$\bar{\mathcal{H}} \subset \mathcal{H}$ is the set of terminal histories in T. A history h is terminal if there is no other history $h' \in \mathcal{H}$ such that $h \subset h'$.
$H(w)$ is the set of histories ending at some $x \in w$.
$g(\Gamma)$ is a subgame of Γ.
$G(\Gamma)$ is the set of subgames of Γ (Γ will be omitted most of the times).
g^u is an "uncertain" subgame.
$h_{|x}$ is the subhistory of h starting at $x \in h$.
$\leq$ is an order, defined for both X and W, where (1) $x \leq x'$ if and only if x' is accessible from x meaning that there exists a history h such that $x, x' \in h$ and $x' \in h_{|x}$, and (2) $w \leq w'$ if and only if there exists a pair (x, x') with $x \in w$, and $x' \in w'$ such that $x \leq x'$.
$w(x)$ is the information set to which x belongs.
$A(s)$ is the set of information sets that have positive probability to occur under s.
$B(s)$ is the set of information sets that have zero probability to occur under s (i.e., $A(s)$ and $B(s)$ partition W for all s).
$o(s)$ is an outcome path corresponding to a strategy profile s (a probability distribution over W induced by s)

2.2 Discussion

For any game Γ, given an equilibrium concept and a strategy profile s, the goal is to characterize a set of contingencies that are irrelevant with respect to a specific notion of sequential rationality. Irrelevance of an information set w shall mean that for any set of terminal nodes Z' with $w \leq z$ for all $z \in Z'$, whatever payoffs follow at Z' it does not change whether or not s is sequentially rational at a relevant set. Thus the choice at the irrelevant set must be irrelevant. The notion of sequential rationality will depend on the equilibrium concept considered and be made precise later. Obviously, contingencies on the equilibrium path cannot be irrelevant.

The two equilibrium concepts I consider both incorporate some kind of sequential rationality. Furthermore, they require their particular notion of sequential rationality to be satisfied throughout the game. However, if one succeeds in finding a non-empty set of irrelevant contingencies, one has also shown that those concepts are overly restrictive. At some contingencies, requiring rational moves apparently "has no bite". In contrast, Nash equilibrium is an example for a concept that is not overly restrictive. Because it imposes rationality only on the equilibrium path, it cannot be more restrictive than necessary, even though its irrelevant set can be nonempty. For that reason, the focus of this paper is not on Nash equilibrium. Nevertheless, NE will be used for a first illustration of the concept.

Given a strategy profile, which contingencies need to be examined to determine whether the strategy profile is a Nash equilibrium? A strategy profile is a Nash equilibrium if no player can gain from a, unilateral, deviation. Hence, to calculate whether no such deviation exists one needs to know all choices at information sets that are reached by a unilateral deviation. Clearly, there is a possibility that not all information sets need to be considered, namely all information sets that only have a positive probability to be reached if at least two players deviate from their strategies. As an example, verify that in Figure 2 player 3's second choice from the left is irrelevant for the question whether the indicated strategy profile is a NE. Moreover, whatever payoff follows these information sets, s remains a NE.[10]

[10] In a sense, a NE is not restrictive enough, because it does not impose any rationality on information

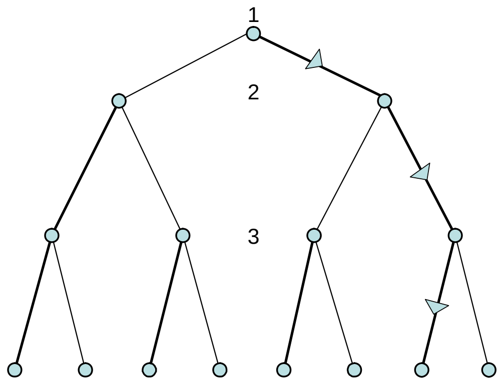

Figure 2: Irrelevant contingencies for NE

When conducting a similar thought experiment for equilibrium concepts which refine NE by incorporating some kind of sequential rationality, the reasoning is different. Instead of asking whether any player can gain from a deviation, one asks whether any player can gain from a deviation expecting sequentially rational play at future contingencies. With this kind of reasoning which contingencies might be irrelevant? Fix a strategy profile and answer the following question for each player using forward induction: Can player i gain from a deviation expecting sequentially rational play at all information sets reached by the deviation? Different from the NE-case, to answer this question, it is NOT sufficient to know the choices at all information sets reached by the deviation. One needs to know the sequentially rational choices at these information sets. But to know these choices information about the sequentially rational choices at further information sets is needed, which in turn requires more information and so forth. This higher order of reasoning is illustrated in Figure 3. The only sequentially irrational play occurs at player 3's middle information set, which cannot be reached by a unilateral deviation. However, player 2's play at his right information set is only sustained by 3 choosing this irrational move. If player 3 chose a' instead of b', player 2's sequentially rational choice would be B, in which case player 1's sequentially rational move was b. Hence, the outcome path indicated in Figure 3 is sustained by sequentially irrational play.

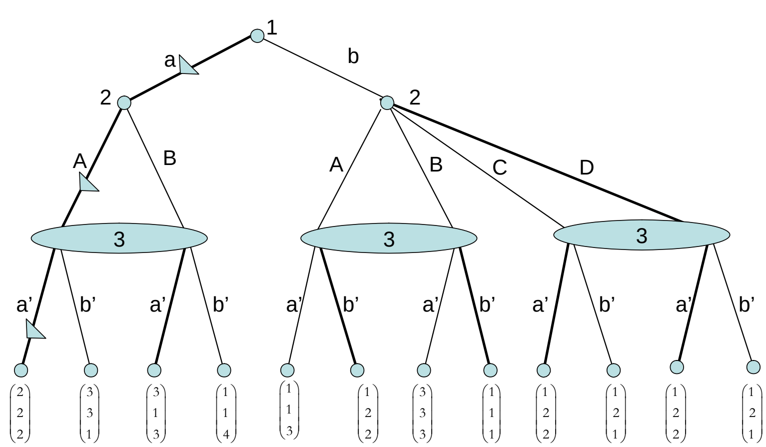

Figure 3: Not being reachable by a unilateral deviation is not sufficient to be irrelevant for a SPE

sets that matter for the decisions at information sets on the equilibrium path. However, I do not pursue this line of thinking here.

So, "where" might information sets be that are never considered under such a reasoning? Instead of getting to the answer to this question iteratively, let us approach the problem by requiring the set to satisfy a consistency property. Consider a game Γ and a strategy profile s for Γ. Fix an equilibrium concept and a corresponding notion of sequential rationality, call this notion ρ. Let $W_{irr}(s,\rho)$ denote a subset of $B(s)$ such that play at $w \in W_{irr}(s,\rho)$ does not affect the play outside of $W_{irr}(s,\rho)$. Hence, whenever $w \in W_{irr}(s,\rho)$ and $w' \in W \backslash W_{irr}(s,\rho)$ the choice at w does not affect the sequential rationality ρ at w'. Call $W_{irr}(s,\rho)$ an *irrelevant set*. Let $\mathcal{W}_{irr}(s,\rho)$ be the set of all irrelevant sets. $\mathcal{W}_{irr}(s,\rho)$ is a partially ordered set (under inclusion). Notice that it follows directly from the definition that $\mathcal{W}_{irr}(s,\rho)$ is closed under union. Hence, every chain in $\mathcal{W}_{irr}(s,\rho)$ has an upper bound, namely the union of all members of the chain. It follows from Zorn's Lemma that there exists a unique (possibly empty) maximally irrelevant set, which will be denoted by $\hat{W}_{irr}(s,\rho)$. I will consider SPE and WPBE separately and characterize $\hat{W}_{irr}(s,\rho)$ for both concepts. Then, irrelevant sets are examined more closely. In both cases the main result is whenever a strategy profile satisfies sequential rationality at all relevant information sets (whenever the profile is a trimmed equilibrium), one can find an equilibrium of the original concept that has the same equilibrium path, *provided that equilibrium exists for the game.*

3 Trimmed SPE

I will show that SPE-irrelevant contingencies are those that can *only* be reached from the equilibrium path if at least two players deviate in the "first stage" of some subgame (the emphasis on "only" is made because a non-singleton information set can be reached by several different deviations). The proof of this result proceeds by showing that the set containing all such contingencies is the maximally irrelevant set for a SPE. Returning to the reasoning in the discussion, these are the contingencies that do not play any role when a player considers a deviation while expecting NE in future subgames.

First, a precise definition for the "first stage" of a subgame is needed. For that purpose, partition the set of information sets W so that each member of the partition corresponds to one subgame of the game. Let the partitioning function be $\pi : W \to G$ where $\pi(w) = g$ is the minimal subgame containing w.[11] The resulting partition of information sets is denoted $\Pi = \{W_g\}_{g \in G}$ where $w \in W_g$ if and only if $\pi(w) = g$. Also, let s^{W_g} be the strategy s induces on the information sets in W_g.

Definition 1 *(information sets on and off the (unilateral) deviation path).*
Fix an extensive form game Γ and a strategy profile s. The set of information sets on the unilateral deviation path of s is denoted by $B_1(s)$. Information set $w \in B_1(s)$ if and only if
1. $w \in B(s)$, and
2. there exists s' such that
(a) $w \in A(s')$, and
(b) for each W_g, for at most one player i, $s_i^{W_g} \neq s_i'^{W_g}$ (i.e. $s_j^{W_g} = s_j'^{W_g}$ for all $j \neq i$).
The complement of $B_1(s)$ in $B(s)$, $B_2(s) \equiv B(s) \backslash B_1(s)$ is the set of information sets off the unilateral deviation path of s.

[11] That is, no subgame strictly contained in $\pi(w)$ contains w.

Furthermore, define $W_{SPE}(s) = A(s) \cup B_1(s)$ (so, $W \backslash W_{SPE}(s) = B_2(s)$). This set will be shown to contain all information sets that are possibly relevant for s being consistent with subgame-perfect rationality. The sets $A(s)$ and $B(s)$ partition W into information sets on the outcome path of s, and information sets off the outcome path of s. The criterion of a unilateral deviation path leads to a further partition of $B(s)$ into two sets, information sets on the unilateral deviation path, $B_1(s)$, and information sets off the unilateral deviation path, $B_2(s)$.

Let ρ_{SPE} stand for "sequential rationality, as required by SPE" and say that s satisfies ρ_{SPE} on $\hat{W} \subseteq W$ if it induces NE in all subgames with origin in $\hat{W}$. To economize on notation, I will omit the subscript SPE, and in this section ρ shall always mean ρ_{SPE}. I want to show that play at information sets in $B_2(s)$ does not affect whether s satisfies ρ on $W_{SPE}(s)$, and that $B_2(s)$ is the maximal set with this property. Notice the dependency of this statement on the strategy profile s. While a particular equilibrium might not be sensitive to specifics in certain parts of a game, the *set* of equilibria consistent with a certain concept, of course, might be sensitive to these specifics. If the Nash equilibrium of subgame $(l_1, l_2) = (0, 1)$ in the Hotelling game yielded a higher payoff than possible in any other subgame equilibrium, clearly there would exist a subgame perfect Nash equilibrium ending at this subgame.

Theorem 2 *For any s, $\hat{W}_{irr}(s, \rho) = B_2(s)$ (the set of information sets off the unilateral deviation path of s is the maximally irrelevant set for (s, ρ)).*

Proof.

1. $\hat{W}_{irr}(s, \rho) \subseteq B_2(s)$.

I will prove the contrapositive of the statement. Suppose that $w' \in W_{SPE}(s) = W \backslash B_2(s)$. Let $x_{og'}$ be the origin of g', the minimal subgame containing w'. The contrapositive is proven by showing that (1) $x_{og'} \in W_{SPE}(s)$ and (2) the play at w' matters for ρ at $x_{og'}$. Suppose that $x_{og'}$ is not in $W_{SPE}(s)$. Then $x_{og'}$ can only be reached if at least two players deviate from their strategies in the same first stage of some subgame. Because g' is the minimal subgame containing w', every history in $H(w')$ has to "pass through" $x_{og'}$. Hence, to reach w' from s, first $x_{og'}$ has to be reached, implying that w' can also only be reached if at least two players deviate in the same first stage of some subgame, contradicting $w' \in W_{SPE}(s)$. The fact that $w' \in W_{SPE}(s)$ also implies that can be reached from $x_{og'}$ by at most a unilateral deviation from s^g. Thus, w' is relevant for $s^{g'}$ being a NE or not.

2. $B_2(s) \subseteq \hat{W}_{irr}(s, \rho)$

Let $w' \in B_2(s)$. I will show that w' is irrelevant for ρ in any subgame with origin in $W_{SPE}(s)$. Let g be such a subgame.

Case 1) $w' \notin g$. In this case the irrelevance of w' is obvious.

Case 2) $w' \in g$.

First prove the following lemma.

Lemma. Let s be a strategy profile for game Γ. For any pair (w, g) such that $w \in B_2(s)$, g contains w, and $x_{og} \in W_{SPE}(s)$, $w \in B_2^g(s^g)$, i.e. w is off the (unilateral) deviation path that s^g induces on g.

Proof

Because $x_{og} \in W_{SPE}(s)$ it can be reached through a sequence of unilateral deviations from s. Because $w \in B_2(s)$, it can *only* be reached from s if there are at least two players deviating

at information sets in the same equivalence class. Because x_{og} is a singleton and w belongs to g, (a) only one history leads to x_{og} and (b) any history leading to w must contain x_{og}. Hence, $w \in B_2^g(s^g)$ for otherwise, the concatenation of the unilateral deviation path leading to x_{og} and the unilateral deviation path leading from x_{og} to w, is a unilateral deviation path as well, contradicting that $w \in B_2(s)$.■
Hence, we have $w' \in B_2^g(s^g)$, implying that $w' \notin W_{NE}(s^g)$. Whether or not s^g is a NE of g does not depend on the play at w'. ■

With this result we are ready to define a trimmed SPE.

Definition 3 *(trimmed SPE). Strategy profile s^* is a trimmed SPE in extensive form game Γ, if it induces a NE in all subgames g with origin $\{x_{0g}\} \in W_{SPE}(s^*)$.*

The following result relates trimmed SPE to SPE.

Theorem 4 *(1) Every SPE is a trimmed SPE, but the converse does not hold. (2) In a game of perfect information, $B_2(s) = \varnothing$ for all s.*

Proof.
(1) That a SPE is a trimmed SPE trivially from the definitions.
Figure 4 shows an example of a trimmed SPE that is not a SPE.

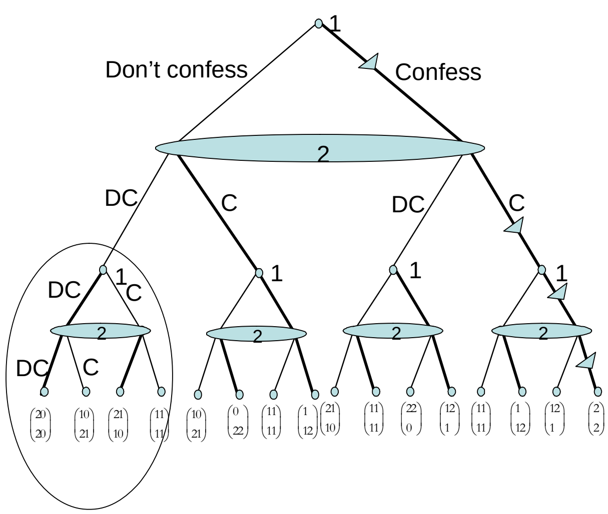

Figure 4: Trimmed equilibrium in a once repeated prisoner's dilemma

(2) Because every $x \in X$ is the origin of some subgame, the partition Π consists of singletons only, implying that every w can be reached by a sequence of unilateral deviations. ■

Corollary 5 *In a game of perfect information, every trimmed SPE is a SPE.*

Because in perfect information games every move is the first move of some subgame, every information set can be reached from an initial strategy profile through a sequence of unilateral deviations, at most one per "stage". And because a trimmed SPE requires optimal decisions at each such information set, subgame perfect play results.

The following theorem proves the claim made in the introduction. A trimmed SPE, while by definition less restrictive than a SPE, is just as "strong" as a SPE in the sense that it does not allow for any equilibrium outcome paths which cannot be supported by a SPE - provided NE exists for all subgames. Together with Theorem 2, this is the main result concerning SPE.

Theorem 6 *Assume that for the game Γ the set of subgame perfect equilibria $S_{spe} \neq \varnothing$. Let $S_{t.spe}$ be the set of trimmed subgame perfect Nash equilibria. Then the set of outcome paths induced by SPE denoted $O(S_{spe})$, and the set of outcome paths induced by trimmed SPE, denoted $O(S_{t.spe})$, are the same.*

Proof.
1. $O(S_{spe}) \subseteq O(S_{t.spe})$
This follows from the fact that $S_{spe} \subseteq S_{t.spe}$.
2. $O(S_{t.spe}) \subseteq O(S_{spe})$
Pick any $o \in O(S_{t.spe})$ and $s^* \in S_{t.spe}$ such that s^* induces o. If s^* happens to specify NE play in all subgames g with origin $w_g \in B_2(s^*)$, then it is a SPE and $s^* \in S_{spe}$, which shows the result. If not, choose any equilibrium in S_{spe}, say s_{spe}, and consider the strategy profile $s^{*spe} = \left(s^*_{W_{SPE}(s^*)}, s^{spe}_{B_2(s^*)}\right)$. Notice that s^{*spe} and s^* induce the same outcome paths because they only differ at information sets with zero probability to occur under s^*. Next, I will show that s^{*spe} is a subgame perfect equilibrium. First, consider subgame g with $x_{og} \in W_{SPE}(s^*)$. By Lemma 1, all changes made in subgame g when moving from s^* to s^{*spe} were made at information sets $w \in B_2^g(s^{g*})$. Hence these changes do not affect play on the (unilateral) deviation path of s^{g*} in g, and so s^{gspe*} is a NE of g. Second, consider any subgame g with $x_{og} \in B_2(s^*)$. Then, all information sets in g belong to $B_2\left(s^*\right)$ because they all follow x_{og}. Hence $s_g^{*spe} = s_g^{spe}$ and so s^{*spe} induces a NE in this subgame. Thus s^{*spne} induces NE in all subgames and therefore $s^{*spe} \in S_{spe}$, showing that the outcome path induced by s^* is also an outcome path of some subgame perfect equilibrium. ■

Corollary 7 *For any trimmed SPE s' and $s = \left(s'_{W_{SPE}(s')}, s''_{B_2(s')}\right)$, if $s''_{B_2(s')}$ induces Nash equilibrium in all subgames g with origin in $B_2(s')$, then s is a SPE.*

4 An Example

The following example demonstrates the usefulness of a trimmed SPE. Suppose two firms, i and j, produce the same homogenous good and compete against each other in a Bertrand-Edgeworth world: First they simultaneously build capacity, x_i and x_j, at a cost of K dollars per unit. Second, they simultaneously announce prices, p_i and p_j, and demand is realized. Market demand is given by $D = 20 - P$. Due to the limited capacity, it might be that a firm cannot serve everyone who demands to buy at the price it charges. Hence, a rationing rule is needed. With a surplus maximizing rationing rule (Levitan and Shubik (1972) and Shubik (1955)) the lower price firm serves the high demand consumers.[12] That is, if $p_i < p_j$,

[12] Suppose there is a mass of consumers of measure one, who all demand one unit of the good and whose willingness to pay is uniformly distributed on the interval $[0, 20]$.

demand for firm i is $20 - p_i$. For simplicity, assume variable production costs are zero. Firm i then produces $Min\{20 - p_i, x_i\}$. If the constraint binds for firm i, firm j might serve some consumers as well, but only up to its capacity, i.e. it produces $Min\{Max\{20 - p_j - x_i, 0\}, x_j\}$. If $p_i = p_j = p$, firm i's demand is given by $Min\{x_i, \frac{D(p)}{2} + Max\{0, \frac{D(p)}{2} - x_j\}\}$, and similarly for firm j, which means that if p is such that $D(p) = x_i + x_j$, firms simply produce respectively, x_i and x_j.

Suppose we allow only for pure strategies, and we are interested in the SPE of this game. Each capacity pair chosen in the first stage induces a pricing subgame in the second stage. As Kreps and Scheinkman (1983) show, pure strategy equilibria exist only for a subset of these subgames. More specifically, letting $r(q)$ be the best reply correspondence of the Cournot game with zero unit costs, pure strategy price equilibria exist only for the set $M = \{x \in \mathbb{R}_+^2 : x_i \leq r(x_j) \text{ and } x_j \leq r(x_i), \text{ or } x_i = x_j = 20\}$.[13] Hence, any attempt to specify a pure strategy SPE must fail. Allowing for mixed strategies, Kreps and Scheinkman derive an equilibrium, which leads to a complicated analysis due to the continuous strategy set. In the SPE they define, firms choose capacities equal to those of the Cournot equilibrium of a game when firms have unit costs of K, i.e. $x_i^{CE} = x_j^{CE} = \frac{20-K}{3}$. Except if $x_i = x_j = 20$, all pricing equilibria in M are such that both firms charge $D^{-1}(x_i + x_j)$. If $x_i = x_j = 20$, both firms charge zero.

Now, suppose that, in addition, the cost of capacity is discontinuous and capacity higher than some $\bar{x}$ is prohibitively expensive. Let $\bar{x}$ satisfy $\frac{20-K}{3} < \bar{x} \leq \frac{20}{3} + \frac{K}{6}$. Still, there remains a subset of subgames for which pure strategy equilibria do not exist (see Figure 5). However, given that firm j builds a capacity of $\frac{20-K}{3}$, all feasible capacity levels for firm i are in the set M. Hence, for all subgames reachable for firm i, conditional on firm j producing $\frac{20-K}{3}$, the expected pricing equilibrium will have both firms charging $P(x_i + x_j)$, and serving half of $20 - P(x_i + x_j)$. It is easy to verify that then a deviation for firm i is not worthwhile. Simply solve

$$\max_{x_i \leq \bar{x}} x_i(20 - \frac{20}{3} + \frac{K}{3} - x_i - K)$$

which yields $x_i^* = \frac{20-K}{3}$. By symmetry the same holds true for firm j.

Without having specified Nash equilibria for all subgames, we have found a reasonable prediction in form of a trimmed SPE, which rules out noncredible threats. However, while using a trimmed SPE to verify whether the above strategy is consistent with a SPE allows for a substantially simpler analysis than the one in Kreps and Scheinkman, their analysis also demonstrates the uniqueness of the derived equilibrium.

[13]For simplicity, assume that firms never build capacity beyond 20.

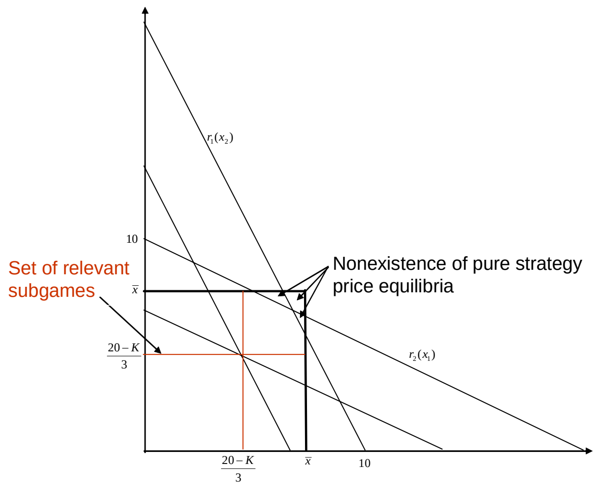

Figure 5: $W_{SPE}(s)$ for s with $x_i = x_j = \frac{20-K}{3}$

5 Trimmed WPBE

The task remains in principle the same as before. Given a strategy profile s and now also a system of beliefs μ, we are looking for the maximal subset of $B(s)$ such that the play inside this set does not affect the sequential rationality outside of the set. Here, sequential rationality means sequential rationality as required by a WPBE, which is denoted by ρ_{WPBE}. The pair (s, μ) is said to satisfy ρ_{WPBE} at w' if the play s specifies at w' is optimal given μ and future play induced by s. In addition, μ should satisfy Bayes' rule whenever possible. In this section, ρ stands for ρ_{WPBE}. Analogously to Section 3, $\hat{W}_{irr}(s, \mu, \rho)$ is defined to be the maximal subset of $B(s)$ such that if $w \in W \backslash \hat{W}_{irr}(s, \mu, \rho)$ and $w' \in \hat{W}_{irr}(s, \mu, \rho)$, then, given μ, ρ at w does not depend on the play at w'.

How can the set $\hat{W}_{irr}(s, \mu, \rho)$ be characterized? I suggest the notion of information sets "off the believed (unilateral) deviation path". The idea is that for any $w \notin \hat{W}_{irr}(s, \mu, \rho)$ and any history leading from w to $\hat{W}_{irr}(s, \mu, \rho)$, the player whose move it is does not believe that $\hat{W}_{irr}(s, \mu, \rho)$ can be reached - either due to his belief about past play or because s attaches zero probability to another player's future move in the history. In other words, players do not believe that $W_{irr}(s, \mu, \rho)$ can be reached by a unilateral deviation, but this belief does not need to be correct.

Definition 8 *(information sets on and off the believed (unilateral) deviation path). Given extensive form game Γ, strategy profile s, and system of beliefs μ, the set of information sets off the believed (unilateral) deviation path is denoted $B_2(s, \mu)$. Information set $w \in B_2(s, \mu)$ if and only if*
(i) $w \in B(s)$, and
(ii) $h = (x_o, x_1, ..., x_K) \in H(w)$ implies, for each $x_k \in h$
(a) $w(x_k) \in B_2(s, \mu)$ or
(b) $\mu(x_k) = 0$ or
(c) $\exists x_{k'} \in h$ with $k < k' < K$, $w(x_{k'}) \notin B_2(s, \mu)$, and $\iota(x_{k'}) \neq \iota(x_k)$ such that $s(x_{k'}, x_{k'+1}) = 0$.
Also, define $B_1(s, \mu) = B(s) \backslash B_2(s, \mu)$ and $W_{WPBE}(s, \mu) = A(s) \cup B_1(s, \mu)$.

The first result in this section parallels Theorem 2. It states that the set of information sets off the believed (unilateral) deviation path is indeed the maximally irrelevant set.

Theorem 9 *Fix* (s, μ) *and let* ρ *be sequential rationality as required by WPBE. Then,* $\hat{W}_{irr}(s, \mu, \rho) = B_2(s, \mu)$ *(the set of information sets off the believed unilateral deviation path of* (s, μ) *is the maximally irrelevant set for* (s, μ, ρ)*).*

Proof.
1. $\hat{W}_{irr}(s, \mu, \rho) \subseteq B_2(s, \mu)$.
To the contrary, suppose that $M \equiv \hat{W}_{irr}(s, \mu, \rho) \cap W_{WPBE}(s, \mu) \neq \varnothing$. Because $M \subseteq \hat{W}_{irr}(s, \mu, \rho) \subseteq B_2(s)$, we can pick $\hat{w} \in M$ such that for all $w \leq \hat{w}$, we have that $w(x) \notin M$ (because, surely, $\{x_o\} \notin M$). Because $M \subseteq W_{WPBE}(s, \mu) = W \backslash B_2(s, \mu)$, the information set $\hat{w} \notin B_2(s, \mu)$. However $\hat{w} \in B_2(s)$ and so there exists $\hat{h} \in H(\hat{w})$ for which none of the conditions in part (ii) of the definition of $B_2(s, \mu)$ holds. Write $\hat{h} = (x_0, x_1, ..., x_K)$ with $x_0 = x_o$ and $x_K \in \hat{w}$. By failure of (a), (b), and (c) there exists a node $x_{k'} \in h$ such that $w(x_{k'}) \notin B_2(s, \mu)$, $\mu(x_{k'}) > 0$ and $s(x_{k''}, x_{k''+1}) > 0$ for all $(x_{k''}, x_{k''+1}) \subseteq h$ for which $k'' > k'$ and $\iota(k'') \neq \iota(k')$. Therefore, the play at $\hat{w}$ matters at $x_{k'}$ (given the other players' strategies, $\iota(x_{k'})$ believes that he can choose a strategy so that $\hat{w}$ is reached). However, $w(x_{k'}) \notin B_2(s, \mu)$ and $w(x_{k'}) \notin M$ implies that $w(x_{k'}) \notin \hat{W}_{irr}(s, \mu, \rho)$, contradicting that $\hat{w} \in \hat{W}_{irr}(s, \mu, \rho)$.

2. $B_2(s, \mu) \subseteq \hat{W}_{irr}(s, \mu, \rho)$
I will show this by showing that $B_2(s, \mu)$ is an irrelevant set. Pick $w'' \in B_2(s, \mu)$, and $w' \notin B_2(s, \mu)$.
Case 1) not$(w' \leq w'')$.
At no node in w' there is a choice that leads to w''. Then, given μ, the choice at w'' does not matter for the sequential rationality at w'.
Case 2) $w' \leq w''$. Let $L \subseteq H(w'')$ such that $h \in L$ implies $h \cap w' \neq \varnothing$, i.e. L contains the histories that can access w'' from w'. Pick any $x_{k'} \in w' \cap L$. Because $w' \notin B_2(s, \mu)$, either (b) or (c) in part (ii) of the above definition must hold. Therefore, either $\iota(w')$ does not believe to be at $x_{k'}$ or believes that w'' cannot be reached from $x_{k'}$ because of the future play of other players. Therefore the play at w'' does not matter for the sequential rationality of play at w', showing that $B_2(s, \mu)$ is an irrelevant set. ■

Having shown that $B_2(s, \mu)$ is the maximally irrelevant set, a trimmed WPBE can be defined as follows.

Definition 10 *The pair* $\{s, \mu\}$ *is a trimmed WPBE if* s *satisfies* ρ *at all information sets* $w \in W_{WPBE}(s, \mu)$.

The following result compares trimmed WPBE to WPBE, and parallels Theorem 4.

Theorem 11 *(1) Every WPBE is a trimmed WPBE, but the converse does not hold. (2) In a game of perfect information,* $B_2(s, \mu) = \varnothing$ *for all* (s, μ).

Proof.
(1) That a WPBE is a trimmed WPBE follows trivially from the definitions.
Figure 6 shows an example of a trimmed WPBE that is not a WPBE.

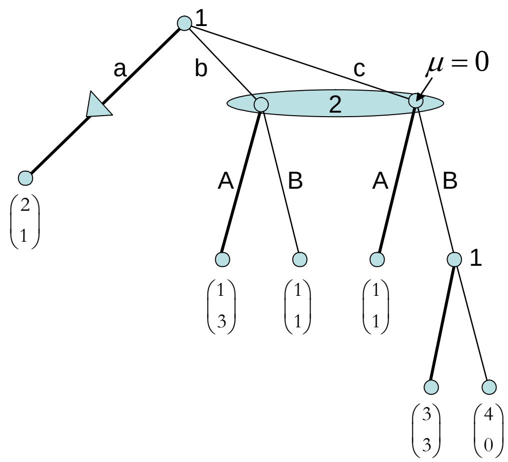

Figure 6: A trimmed WPBE that is not a WPBE

(2) Fix some (s,μ). Suppose $B_2(s,\mu) \neq \varnothing$. Pick $w' \in B_2(s,\mu)$ with $w \leq w'$ implying that $w \notin B_2(s,\mu)$. Let $x \in pre(w')$, so $w(x) \notin B_2(s,\mu)$, so (a) does not hold. By perfect information, $\mu(x) = 1$, so (b) does not hold. However, because x is an immediate predecessor of w' (c) cannot hold either, a contradiction. ■

Corollary 12 *In a game of perfect information, any trimmed WPBE is a WPBE.*

Paralleling Theorem 6, the next result establishes that, if a WPBE exists, every trimmed WPBE can be matched with a WPBE that has the same outcome path.

Theorem 13 *Assume that for Γ the set of weak perfect Bayesian equilibria $S_{wpbe} \neq \varnothing$ (the features of the game are such that a WPBE exists). Let $S_{t.wpbe}$ be the set of trimmed weak perfect Bayesian equilibria. Then the set of outcome paths induced by $s \in S_{wpbe}$, denoted $O(S_{wpbe})$, and the set of outcome paths induced by $s \in S_{t.wpbe}$, denoted $O(S_{t.wpbe})$, are the same.*

Proof.
1. $O(S_{wpbe}) \subseteq O(S_{t.wpbe})$.
This follows from the fact that $S_{wpbe} \subseteq S_{t.wpbe}$ as shown in Theorem 11.
2. $O(S_{t.wpbe}) \subseteq O(S_{wpbe})$
Consider any trimmed WPBE, $(s_{t.wpbe}, \mu_{t.wpbe})$. Construct a new strategy profile s' as follows.
(a) For each $w \in W_{WPBE}\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$, let $s'(w) = s_{t.wpbe}(w)$ and let $\mu'(w) = \mu_{t.wpbe}(w)$.
(b) For $B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$ construct a reduced game from Γ and $\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$ as follows. Let X_{WPBE} and X_{B_2} denote the sets of nodes such that $w(x)$ is an element of $W_{WPBE}\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$ and $B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$, respectively.
Step 1
Remove all $x \in X_{WPBE}$ such that $x' \leq x$ implies $x' \in X_{WPBE}$.

Step 2
For all $x \in X_{WPBE}$ such that $x \leq x'$ implies $x' \in X_{WPBE}$, if $pre(x) \in X_{B_2}$ replace x by the expected payoff induced by $s_{t.wpbe}$, and if $pre(x) \in X_{WPBE}$ delete x.

Step 3
For the remaining $x \in X_{WPBE}$, if $pre(x) \in X_{B_2}$, replace x by a move by nature as follows. All paths, which lead to final histories passing only through X_{WPBE}, are replaced by one final node with the expected payoff as induced by $s_{t.wpbe}$. Nature's probability to choose this final node is the probability that any of these terminal histories is reached conditional on x being reached. For each $x' \geq x$ such that $x' \in X_{B_2}$ and for which $x' \geq x'' \geq x$ implies $x'' \in X_{WPBE}$, replace the path leading to x' by a move by nature leading directly to x' with the probability as induced by $s_{t.wpbe}$ (which is well defined because all nodes on this path belong to X_{WPBE}).

It is not hard to verify that, after performing these three steps, all nodes in X_{WPBE} were either deleted, replaced by a final node, or replaced by a move by nature. Note that the resulting graph does not need to be connected and that there might be several initial nodes. Add an initial move by nature that attaches some positive probability to each initial node of the components obtained. Call the new game Γ^*. It is easy to verify that - apart from nature's moves -, the set of information sets for Γ^* is $B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$. Now, find a weak perfect equilibrium for Γ^*, denoted $\left(s^*_{wpbe}, \mu^*_{wpbe}\right)$ and complete the specification of s' by letting $s'(w) = s^*_{wpbe}(w)$ and $\mu'(w) = \mu^*_{wpbe}(w)$ for all $w \in B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$.

It remains to verify that (s', μ') is a weak perfect Bayesian equilibrium. For any $w \in W_{WPBE}\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$ we know that $s_{t.wpbe}(w)$ was sequentially rational given $s_{t.wpbe}$ and $\mu_{t.wpbe}(w)$. Because we only changed strategies and beliefs at information sets in $B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$, those changes could not affect the sequential rationality at w.

For any $w \in B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$, sequential rationality follows from the fact that $\left(s^*_{wpbe}, \mu^*_{wpbe}\right)$ is a weak perfect Bayesian equilibrium and that for each $x \in w$ and each choice at x, the expected payoff from that choice given s^*_{wpbe} is the same as the expected payoff given s'. Notice that, because $B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right) \subseteq B\left(s_{t.wpbe}\right)$, Bayes' rule does not need to be satisfied on $B_2\left(s_{t.wpbe}, \mu_{t.wpbe}\right)$. ■

Corollary 14 *For any trimmed WPBE* (s', μ'), $s = \left(s'_{W_{WPBE}(s',\mu')}, s''_{B_2(s',\mu')}\right)$, *and* $\mu = \left(\mu'_{W_{WPBE}(s',\mu')}, \mu''_{B_2(s',\mu')}\right)$, *if* $s''_{B_2(s',\mu')}$ *is sequentially rational at all* $w \in B_2(s', \mu')$, *given* s *and* $\mu''_{B_2(s',\mu')}$, *then* (s, μ) *is a WPBE.*

6 Dominated subforms

Even though the concept of a trimmed equilibrium can help in many cases of uncertainty in some game parts, there might also be games in which it fails to exist. Roughly, this will be the case if the structure of the game is such that players have incentives to choose a path "close" to the problematic subgames, where close means it is reached without a deviation, or can be reached by sequentially unilateral deviations. In this section, I introduce the concept of a dominated subform (which is, roughly, all parts of a game that follow some information sets). If one can assert that a subform is dominated, this part of the game can reasonably be considered as relatively irrelevant for the play of the rest of the game. This can provide a justification for a Nash equilibrium that is not a trimmed equilibrium. However, there still remain cases in which neither of the concepts in this paper helps to predict play that is consistent with sequential rationality.

To see that a trimmed SPE might not exist, consider the example in Figure 7.

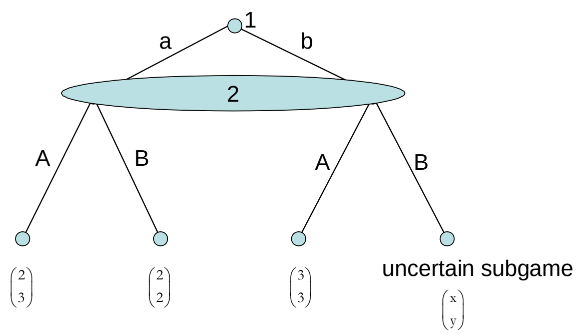

Figure 7: A trimmed SPE might not exist

Assume that a Nash equilibrium cannot be specified for the uncertain subgame. It can be easily verified that the only candidate for a trimmed SPE would involve the equilibrium path (a, A). This, however, cannot be part of a Nash equilibrium either because player 1 would want to deviate to b. All the "forces" in this game pull toward (b, B), which is too close to the uncertain subgame. Hence, the concept of a trimmed equilibrium does not help in dealing with this problem. Sometimes, however, one can still make a reasonable prediction. Suppose one knows that player 2's payoff y in the uncertain subgame does not exceed 2. Then, A, in a sense, strictly "dominates" B, and it is not reasonable to expect the uncertain subgame ever to be reached. Definition 15 makes this idea more precise. Let $\tilde{g}$ denote a subform of Γ, i.e. $\tilde{g}$ is defined on a set of nodes $\tilde{X}$ satisfying closure under succession and preservation of information sets, and $\tilde{g} = \left(\tilde{T}, \tilde{P}, \tilde{W}, \tilde{C}, \tilde{u}\right)$ is derived from Γ by restriction on $\tilde{X}$.[14] Let $\tilde{X}_o$ denote the set of initial nodes of $\tilde{g}$. See Figure 8 for an illustration of a subform. Also, for each $\tilde{g}$, a collection of information sets is called a separating set for $\tilde{g}$, denoted $W_{sep}(\tilde{g})$, if $\tilde{W} \cap W_{sep}(\tilde{g}) = \varnothing$ and for each h with $h \cap \tilde{X} \neq \varnothing$, there exists $x \in h$ such that $w(x) \in W_{sep}(\tilde{g})$.

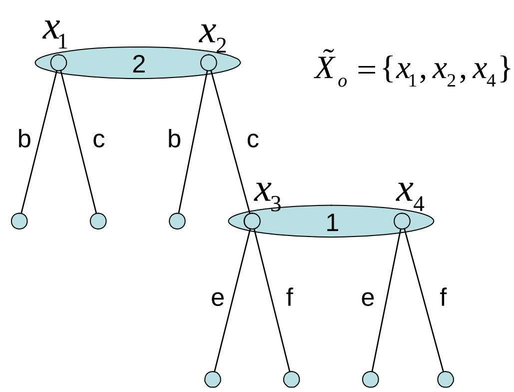

Figure 8: A subform

[14]This definition is taken from Kreps and Wilson (1982), p. 868.

Definition 15 *(dominated subform)*
The subform $\tilde{g}$ starting at $\tilde{X}_o$ is $\rho-$dominated at $w' \in W \backslash \tilde{W}$ with $w' \leq \tilde{W}$ in game Γ if there exists a probability distribution over the choices in w', $f_{w'}$, which assigns zero probability to choices leading to $\tilde{X}_o$, and is such that any *possible expected payoff for player $\iota(w')$ resulting from $f_{w'}$ under the expectation of $\rho-$play at future information sets, is not worse and sometimes better than* any *payoff resulting from any choice possibly leading to $\tilde{g}$, again under the supposition of* any *$\rho-$play, whenever such play exists. Subform $\tilde{g}$ is dominated if there exists a separating set $W_{sep}(\tilde{g})$ so that $\tilde{g}$ is dominated at each $w \in W_{sep}(\tilde{g})$.*

Remark: In contrast to irrelevant sets, domination of a subform does depend on payoffs, but does not depend on a particular strategy profile.

An example of a $\rho_{SPE}-$dominated game part is given by the game in Figure 9. Subgames, $g_1, ..., g_5$ have been replaced by their possible Nash equilibrium payoffs, and the uncertain subgames, g^{u1} and g^{u2}, have been replaced by their possible payoffs. Subform $\tilde{g}$ is dominated via the separating set containing only player 2's informations set.

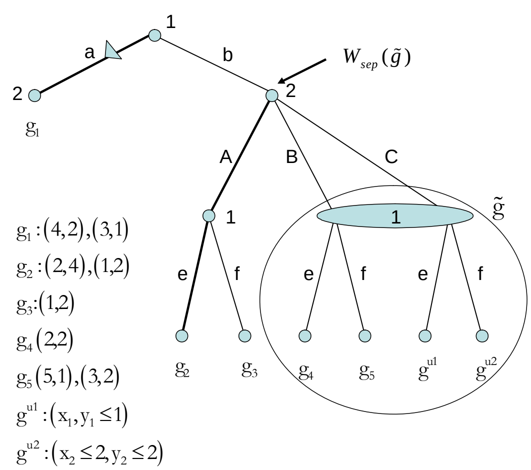

Figure 9: SPE-dominated game part

Because $\tilde{g}$ is dominated, player 1 should expect player 2 to pick B and therefore choose B herself. Even though this game does not have a trimmed SPE or WPBE (because it is not clear what would be a sequentially rational choice at player 1's information set in $\tilde{g}$), the uncertainties in subgames g^{u1} and g^{u2} still allow for a reasonable prediction of the game's outcome. Notice that the definition requires that player 2 prefers A to B and C, no matter which of the two Nash equilibria in g^5 he expects.

7 Conclusion

This paper introduces an equilibrium concept, called trimmed equilibrium, that potentially provides resolution if some parts of the game are difficult to predict due to non-existence of equilibrium, uncertainties about the game's specification in these parts, or computational difficulties. It is shown that a trimmed equilibrium is sufficiently restrictive to capture notions of sequential rationality of the original concept. At the same time, however, it is loose enough to dispose of entire parts of the game. By characterizing maximally irrelevant sets for both the SPE and WPBE, some insight is gained as to the kinds of sequential rationalities these concepts place on players. The analysis reveals redundancies in these concepts' definitions.

It is pointed out that, while a trimmed equilibrium might be invariant to the specification in certain parts of a game, the set of trimmed equilibria is not. Because the specification of the equilibrium remains open in parts of the game, it is not clear how to solve for such an equilibrium. A process of backward induction does not seem to be suitable if some of the subgames are uncertain. Of course, it is possible to determine the entire set of Nash equilibria and then check whether they are trimmed equilibria. This however might not be practical as the set of Nash equilibria in extensive form games can be very large. Constructing an algorithm that solves for trimmed equilibria - or which, at least, narrows down the set of candidates - remains an open task.

8 References

BRIATA, F., GARCIA-JURADO, I., GONZALEZ-DIAZ, J. and F. PATRONE (2007) Essentializing Equilibrium Concepts, mimeo.

DASGUPTA, P. and MASKIN, E. (1986) The existence of equilibrium in discontinous economic games, I: theory, *The Review of Economic Studies* 53 (1), 1-26.

D'ASPREMONT, C., GABSZEWICZ, J.J. and THISSE, J. (1979) On Hotelling's "Stability in competition", *Econometrica* 47 (5), 1145-1150.

FUDENBERG, D. and TIROLE, J. (1991) Game Theory, Cambridge, Mass.: MIT Press.

MAS-COLELL, A., WHINSTON, M.D. and GREEN, J.R. (1995) Microeconomic Theory, Oxford: Oxford University Press.

MYERSON, R. (1991) Game theory: analysis of conflict, Cambridge, Mass.: Harvard University Press.

NASH, J.F. (1951) Non-cooperative games, *Annals of Mathematics* 54, 155-162.

KALAI, E. and NEME, A. (1992) The strength of a little perfection, *Int. Journal of Game Theory* 20 (4), 335-355.

KALAI, E. (2004) Large robust games, *Econometrica* 72 (6), 1631-1666.

KALAI, E. (2005) Partially specified large games, Lecture notes in computer science 3828, 3-13.

KREPS, D.M. and R. WILSON (1982) Sequential Equilibria, Econometrica 50 (4), 863-894.

KREPS, D. M., and J. A. SCHEINKMAN (1983) "Quantity Precommitment and Bertrand Competition Yield Cournot Outcome", The Bell Journal of Economics, 14(2), 326-337.

LEVITAN, R., and M. SHUBIK (1972) "Price Duopoly and Capacity Constraints", International Economic Review, 13, 111-122.

OSBORNE, M.J. and C. PITCHIK (1987) Equilibrium in Hotelling's Model of spatial competition, *Econometrica* 55 (4), 911-922.

SELTEN, R. (1975) Reexamination of the perfectness concept for equilibrium points in extensive form games, *Int. Journal of Game Theory* 4, 25-55.

VAN DAMME, E. (1983) Refinements of the Nash equilibrium concept, Berlin: Springer-Verlag.

VIVES, X. (1999) "Oligopoly Pricing: Old Ideas and New Tools", MIT Press

9 Appendix

X	set of nodes
$succ(x)$	set of immediate successors of x
$pre(x)$	immediate predecessor of x
W	set of information sets
C_w	set of choices at w
f_w	probabiltiy distribution over w
F_w	set of probability distributions over w
$s_i(w) \in F_w$	i's strategy at w
g	subgame
$\tilde{g}$	subform
s^g	strategy for g induced by s
μ	system of beliefs
$\mu(w)$	belief at w
$\iota(w)$	player whose move it is at w
H	set of all histories
$H(w) \subset H$	set of histories ending at w
$h_{\|x} \subset h$	truncation of h by all nodes preceding x
$x \leq x'$	$\exists h$ such that $x' \in h_{\|x}$
$w \leq w'$	$\exists x \in w$ and $x' \in w'$ such that $x' \in h_{\|x}$
$w(x)$	information set to which x belongs
$A(s) \subset W$	information sets on the outcome path of s
$B(s) \subset W$	information sets off the outcome path of s
ρ	definition of sequential rationality
$W_{irr}(s,\rho) \subset B(s)$	set of information sets not affecting ρ at $w \in W \backslash W_{irr}(s,\rho)$
$\hat{W}_{irr}(s,\rho)$	maximal $W_{irr}(s,\rho)$
$W_{SPE}(s) \subset W$	information sets relevant for SPE under s
$W_{WPBE}(s,\mu) \subset W$	information sets relevant for $WPBE$ under (s,μ)
$\left(s'_{\hat{W}}, s''_{W\backslash\hat{W}}\right)$	strategy profile obtained from playing s' on $\hat{W}$, and s'' on $W\backslash\hat{W}$