

**NORMAL AND EXTENSIVE FORM REFINEMENTS:
NOTES FOR AN ADVANCED SCHOOL IN SANTIAGO,
JANUARY 2017**

JOHN HILLAS

CONTENTS

1. Normal and Extensive Form Games	1
1.1. Normal Form Games	2
1.2. Extensive Form Games	3
1.3. Perfect Recall, Imperfect Recall, and Nonlinear Games	4
1.4. Defining Strategies	5
1.5. Kuhn Equivalence and Kuhn's Theorem	7
2. Admissible Equilibria and Normal Form Perfect Equilibria	7
3. Backward Induction	8
3.1. Subgame Perfect Equilibria	9
3.2. Extensive Form Perfect Equilibria	11
3.3. Quasi-Perfect Equilibria	12
3.4. Sequential Equilibria	14
Aside: Real Algebraic Geometry	16
3.5. Proper Equilibria	16
4. Invariance	19
5. The Need For Set-Valued Solutions	20
6. Forward Induction	22
7. An Introduction to Strategic Stability	27
References	30

1. NORMAL AND EXTENSIVE FORM GAMES

There are two prominent models of noncooperative games, the strategic or normal form and the extensive form. Both may be thought of as further abstractions from an incompletely defined abstract game or interactive decision problem. For the strategic form we consider the set of all complete plans of how to play the “game.” The instead of considering how the game actually unfolds we consider how the players

Date: 2017-01-10.

would, when they consider the game in advance choose their complete plan or strategy. One might argue, as several have done, that since a player when choosing his plan in advance can anticipate any future occurrence it should not matter whether they choose in advance or as the problem unfolds. In the strategic form of the game the players simultaneously choose their strategies and then the outcome implied by that profile of strategies is implemented.

1.1. Normal Form Games.

Definition 1. A *normal* or *strategic form game* consists of:

- (1) N , a finite set of players. (We abusively also use N to denote the number of players.)
- (2) For each player $n \in N$, a finite set of pure strategies S_n , with $S = \times_{n \in N} S_n$.
- (3) For each player $n \in N$, a payoff function $u_n : S \rightarrow \mathbb{R}$.

Definition 2. A mixed strategy for Player n in a normal form game is a probability distribution over the player's pure strategies. We denote the set of Player n 's mixed strategies Σ_n and call it Player n 's mixed strategy space and define $\Sigma = \times_{n \in N} \Sigma_n$ the mixed strategy space. We denote a typical element of Σ_n by σ_n and a typical element of Σ by $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$.

We extend the function u_n to Σ linearly, that is, by taking expectations, and indeed to $\times_n(S_n \cup \Sigma_n)$.

And we can now define an equilibrium.

Definition 3. Given a game (N, S, u) , a mixed strategy profile $(\sigma_1^*, \sigma_2^*, \dots, \sigma_N^*)$ is a *Nash equilibrium* if for each $n \in N$ and each $s_n \in S_n$,

$$u_n(\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*, \sigma_{n+1}^*, \dots, \sigma_N^*) \geq u_n(\sigma_1^*, \sigma_2^*, \dots, s_n, \sigma_{n+1}^*, \dots, \sigma_N^*).$$

Rather than speaking directly of utilities we can define what it means to be a best reply to what the other players are playing.

Definition 4. Given a game (N, S, u) , with Σ_n the mixed strategy space of Player n and Σ the space of mixed strategy profiles. Then the *best reply correspondence of Player n* , $BR_n : \Sigma \rightarrow \Sigma_n$ is defined as

$$\begin{aligned} BR_n(\sigma_1, \dots, \sigma_N) &= \{\tau_n \in \Sigma_n \mid u_n(\sigma_1, \dots, \tau_n, \dots, \sigma_N) \\ &\geq u_n(\sigma_1, \dots, s_n, \dots, \sigma_N) \text{ for any } s_n \in S_n\}. \end{aligned}$$

The BR_n are combined to give $BR : \Sigma \rightarrow \Sigma$

This allows us to restate the definition of a Nash equilibrium.

Definition 5. Given a game (N, S, u) , a mixed strategy profile $(\sigma_1^*, \sigma_2^*, \dots, \sigma_N^*)$ is a *Nash equilibrium* if for each $n \in N$,

$$\sigma_n^* \in BR_n(\sigma_1^*, \sigma_2^*, \dots, \sigma_N^*).$$

Theorem 1 (Nash (1950, 1951)). *Any finite normal form game (N, S, u) has at least one Nash equilibrium.*

1.2. Extensive Form Games. The strategic form is a convenient device for defining strategic equilibria: it enables us to think of the players as making single, simultaneous choices. However, actually to describe “the rules of the game,” it is more convenient to present the game sequentially. The first such formal definition of an extensive form game was given in the book by von Neumann and Morgenstern (1944). Somewhat later Kuhn (1953) gave a slightly more general definition, in the form of a game tree, that is essentially the definition almost universally used today..

The *extensive form* is a formal representation of the rules of the game. It consists of a rooted tree whose nodes represent decision points (an appropriate label identifies the relevant player), whose branches represent moves and whose endpoints represent outcomes. Each player’s decision nodes are partitioned into *information sets* indicating the player’s state of knowledge at the time he must make his move: the player can distinguish between points lying in different information sets but cannot distinguish between points lying in the same information set. Of course, the actions available at each node of an information set must be the same, or else the player could distinguish between the nodes according to the actions that were available. This means that the number of moves must be the same and that the labels associated with moves must be the same.

Random events are represented as nodes (usually denoted by open circles) at which the choices are made by Nature, with the probabilities of the alternative branches included in the description of the tree.

The details of the formal definition of an extensive form game are a bit cumbersome and we’ll skip it. We shall however briefly list the notation.

- A finite set of players, $N = \{1, 2, \dots, N\}$. We add an artificial Player 0 or Nature.
- A finite set of nodes, X , and X is a game tree, where $T \subset X$ is the set of terminal nodes and x_0 is the initial node.
- A set of actions, A . $\alpha(x) \in A$ is the action at the predecessor of x that leads to x . If x and x' are distinct and have the same predecessor then $\alpha(x) \neq \alpha(x')$.

- A collection of information sets, \mathcal{H} . For all nodes except the terminal nodes $H(x)$ is the information set containing x . \mathcal{H}_n are the information sets of Player n .
- We assume that all information sets in \mathcal{H}_0 are singletons and assign a probability $\rho(x)$ to each node that immediately follows such singleton information set of Nature.
- For each terminal node t and each Player n we have $u_n(t)$, the payoff or utility of Player n at terminal node t .

In what follows rather than listing all of the elements of an extensive form game we shall simply refer to the game Γ and understand that all of these elements are specified.

1.3. Perfect Recall, Imperfect Recall, and Nonlinear Games.

The information partition is said to have *perfect recall* (Kuhn, 1953) if the players remember whatever they knew previously, including their past choices of moves. In other words, all paths leading from the root of the tree to points in a single information set, say Player n 's, must intersect the same information sets of Player n and must display the same choices by Player n .

Selten (1975) gave the same definition of perfect recall as Kuhn, but his formal definition is a little more straightforward. We give the definition here in the way that Selten did.

Definition 6. A player is said to have perfect recall if whenever that player has an information set containing nodes x and y and there is a node x' of that player that precedes node x there is also a nodes y' in the same information set as x' that precedes node y and the action of the player at y' on the path to y is the same as the action of the player at x' on the path to x . If all players have perfect recall then we say the game has perfect recall.

Much of the literature on the refinement of equilibrium in extensive form games has restricted attention to games with perfect recall.

One implication of perfect recall is that each path from x_0 to a terminal node cuts each information set at most once. In games without perfect recall we distinguish between linear games, the games defined by Kuhn, and nonlinear games, an extension by Isbell (1957), and later under the name “repetitive games” by Alpern (1988), and more recently discussed under the name “absent-mindedness” by Piccione and Rubinstein (1997), and, following Piccione and Rubinstein, by a number of others. In linear games each play of the game reaches an information set at most once. In a nonlinear game we remove that restriction.

There is some issue as to how the moves of Nature are handled. Kuhn (1953) assumed that each of Nature's information sets was a singleton. Isbell (1957), following Dalkey (1953), relaxed this assumption, allowing Nature to have nontrivial information sets and commented that even more general structures were possible. To avoid notational complexity, we shall follow Kuhn's assumption, though extension to the more general cases is not problematic.

1.4. Defining Strategies. Recall that we introduced the idea of a strategy earlier saying that it was a player's "complete plan of how to play the game." When we discuss normal form games we treat strategies as primitives and so this is an intuitive justification rather than a definition. When we come to deal with extensive form games a strategy is not among the primitive components given in the definition of the game. Rather it is a derived concept defined in terms of the primitives.

Definition 7. A pure strategy in an extensive form game for Player n is a function that maps each of his information sets to one of the actions available at that information set. We denote the set of Player n 's pure strategies by S_n and the set of pure strategy profiles by $S = \times_{n \in N} S_n$.

This is a convenient manner in which to define a strategy. However there are some issues here. The definition as we have given it gives 4 pure strategies in the one player game of Figure Figure 1, namely, TL , TR , BL , and BR . This has some advantages. It defines a behaviour in parts of the tree that the player himself avoids and gives a cleaner embedding of the pure strategies in the set of behaviour strategies that we shall define below. However I'd say there is no meaningful difference between the strategies TL and TR . We shall below define what it means for strategies to be Kuhn-equivalent, and see that TL and TR are equivalent. We can then, if we wish, consider only one representative of each equivalence class of pure strategies.

Having now defined pure strategies we can associate to any extensive form game some associated normal form game. The player set is the same; we have just described the strategies, and for each profile of strategies we obtain a probability distribution over the terminal nodes—a probability distribution since there may be moves of nature; if there are no moves of nature then one terminal node will have probability 1—and hence an expected utility.

We define a mixed strategy precisely as we do for the normal form.

Definition 8. A mixed strategy in an extensive form game for Player n is a probability distribution over the player's pure strategies.




Figure 1

We denote the set of Player n 's mixed strategies by Σ_n and the set of mixed strategy profiles by $\Sigma = \times_{n \in N} \Sigma_n$. We say that the mixed strategy σ_n of Player n is completely mixed if $\sigma_n(s_n) > 0$ for all $s_n \in S_n$. We say that the profile $\sigma = (\sigma_1, \dots, \sigma_N)$ is completely mixed if σ_n is completely mixed for all $n \in N$.

Rather than having the player randomise over pure strategies we could have them randomise independently at each information set. We call such a randomisation a behaviour strategy.

Definition 9. A behaviour strategy in an extensive form game for Player n is a function that maps each of his information sets to a probability distribution on the actions available at that information set. We denote the set of Player n 's behaviour strategies by B_n and the set of behaviour strategy profiles by $B = \times_{n \in N} B_n$. We say that the behaviour strategy b_n of Player n is completely mixed if at each information set b_n assigns strictly positive probability to each of the actions available at that information set. We say that the profile $b = (b_1, \dots, b_N)$ is completely mixed if b_n is completely mixed for all $n \in N$.

Most of the extensive form refinements we consider were originally defined in terms of behaviour strategies, and defined only for games with perfect recall.

In nonlinear games we also need to consider randomisations over behaviour strategies. We *can* consider such strategies for linear games, but we do not need to do so.

Definition 10. A general strategy in an extensive form game for Player n is a probability distribution over the player's behaviour strategies. We denote the set of Player n 's general strategies by G_n and the set of general strategy profiles by $G = \times_{n \in N} G_n$.

If we wanted to define completely mixed general strategies there would be some difficulties and choices, but still a number of ways in which to do so. Since we won't be using them we shall not do so.

1.5. Kuhn Equivalence and Kuhn's Theorem. We now define what it means for two strategies to be equivalent. Simply put two strategies of a player are equivalent if, whatever the other players do, the two strategies induce the same probability distribution over the terminal nodes. We now state this a bit more formally.

Definition 11. Two strategies of Player n $x_n, y_n \in S_n \cup B_n \cup \Sigma_n \cup G_n$ are said to be Kuhn equivalent if for any strategies of the others $x_{-n} = (x_m)_{m \neq n}$ with $x_m \in S_m \cup B_m \cup \Sigma_m \cup G_m$ for all $m \neq n$ the profiles (x_n, x_{-n}) and (y_n, x_{-n}) induce the same probability distributions over the terminal nodes.

Kuhn (1953) showed that in linear games (the only games he considered) for any behaviour strategy there is always an equivalent mixed strategy and that if the player has perfect recall the converse is also true.

Theorem 2 (Kuhn (1953)). *In a linear game for any behaviour strategy b_n of Player n there exists a mixed strategy σ_n of Player n that is Kuhn equivalent to b_n . If in some extensive form game Player n has perfect recall then for any mixed strategy σ_n of Player n there exists a behaviour strategy b_n of Player n that is Kuhn equivalent to σ_n .*

2. ADMISSIBLE EQUILIBRIA AND NORMAL FORM PERFECT EQUILIBRIA

We turn now to the first refinements we shall consider. Here we are working with the normal form.

Definition 12. In a normal form game (N, S, u) a strategy $s_n \in S_n$ of Player n is *admissible* or *undominated* if there is no mixed strategy $\sigma_n \in \Sigma_n$ such that for all $s_{-n} \in \times_{m|m \neq n} S_m$

$$u_n(\sigma_n, s_{-n}) \geq u_n(s_n, s_{-n})$$

and for at least one $t_{-n} \in \times_{m|m \neq n} S_m$

$$u_n(\sigma_n, t_{-n}) > u_n(s_n, t_{-n}).$$

Definition 13. An admissible equilibrium is a Nash equilibrium σ such that for all n in N if $\sigma_n(s_n) > 0$ then s_n is admissible, that is a Nash equilibrium in which only admissible strategies are played with positive probability.

This seems a mild requirement. A slightly stronger refinement is that of normal form perfection.

Definition 14. A completely mixed strategy profile $\sigma \in \Sigma$ is a ε -perfect equilibrium if $\varepsilon > 0$ and for all $n \in N$ and all $s_n, t_n \in S_n$ if $u_n(s_n, \sigma_{-n}) < u_n(t_n, \sigma_{-n})$ then $\sigma_n(s_n) < \varepsilon$. A strategy profile $\sigma \in \Sigma$ is a normal form perfect equilibrium if there is a sequence of strategy profiles $\sigma^t \rightarrow \sigma$ and positive numbers $\varepsilon^t \rightarrow 0$ with σ^t a ε^t -perfect equilibrium.

Selten (1975) proved the following theorem.

Theorem 3 (Selten (1975)). *For any finite normal form game (N, S, u) there is at least one normal form perfect equilibrium.*

We give two results relating the two concepts we have just defined.

Theorem 4 (van Damme (1991)). *For any finite normal form game (N, S, u) if σ is a normal form perfect equilibrium then σ is admissible.*

Theorem 5 (van Damme (1991)). *For any finite two player normal form game (N, S, u) if σ is admissible then σ is a normal form perfect equilibrium.*

Thus, in two player games the two concepts coincide.

3. BACKWARD INDUCTION

I'd recommend Section 10 of Hillas and Kohlberg (2002) to be read in parallel with this section.¹ It is less formal but more detailed in its discussion of the underlying ideas. In that chapter we say

Selten (1965, 1975) proposed several ideas that may be summarized as the following *principle of backward induction*:

A self-enforcing assessment of the players' choices in a game tree must be consistent with a self-enforcing assessment of the choices from any node (or, more generally, information set) in the tree onwards.

The concepts defined by Selten were subgame perfect equilibrium and extensive form perfect equilibrium. We define these below. Kreps and Wilson (1982) defined sequential equilibrium and showed it was very closely related to perfect equilibrium. This proved a very influential solution concept with its explicit consideration of the beliefs of the players at each information set an intuitive and useful addition.

¹In fact, I'd recommend all of that chapter, but then again, I would.

Another solution closely related to perfect equilibrium, namely quasi-perfect equilibrium, was given by van Damme (1984) who showed that its requirements were implied by the requirements for proper equilibrium in the normal form of Myerson (1978).

3.1. Subgame Perfect Equilibria. Selten (1965) defined the notion that we now call subgame perfect equilibrium calling it, at the time, perfect equilibrium. ten years later, in Selten (1975) he pointed out that this concept did not imply all of what he thought a notion of perfect should and renamed the original concept subgame perfect equilibrium and introduced a new concept of perfect equilibrium, or extensive form perfect equilibrium, as we shall call it in these notes. Between 1965 and 1975 the concepts were not widely used in most areas of game theory or economics and so the new terminology was easily and widely adopted. The exception was in the field of repeated games where there were some results about what we are calling subgame perfect equilibrium. Thus in repeated games you will still find that when perfect equilibria are referred to the intended meaning is usually subgame perfect equilibria.

We shall now give a definition of subgame perfect equilibria for finite extensive form games.

Definition 15. Given an extensive form game Γ a subgame Γ' of Γ consists of some node of Γ and all nodes following it, together with those structures that pertain to those nodes such that each information set of Γ is either completely in Γ' or completely outside it

Since we may take the starting node of the subgame to be x_0 the initial node of the original game this means that for any extensive form game Γ the whole game Γ is one of the subgames of Γ . It may be the only subgame but often there are others.

Consider the game in Figure 2. This game has two proper subgames. Strictly speaking, the whole game is a subgame of itself, so we call subgames that are not the whole game *proper* subgames. The first is the game that begins from Player 2's decision node and the second is the game that begins from player 3's decision node.

However, it would be a mistake to think that every game has proper subgames. Consider the extensive form game in Figure 3, which is the extensive form version of matching pennies. In this game it doesn't make sense to say there is a subgame that begins at either of Player 2's decision nodes. This is because both of these nodes are in the same information set. If we said that a subgame started at either one of them, then we would be saying that Player 2 knew which node she is at, which violates the original structure of the game.




FIGURE 2. An extensive form game with three players.




FIGURE 3. The extensive form version of matching pennies. This game has no proper subgames.

As another example, due to Selten (1975), consider the extensive form game shown in Figure 4. In this game, the “game” starting from Player 2’s decision node is not a subgame, because it splits player 3’s information set.

Definition 16. In an extensive form game Γ with perfect recall a *subgame perfect equilibrium* is a profile of behaviour strategies b such that for every subgame the parts of b relevant to the subgame constitute an equilibrium of the subgame.




FIGURE 4. The game beginning at Player 2's decision node is not a subgame because it splits Player 3's information set.

The fact that the whole game is one of the subgames means that a subgame perfect equilibrium is an equilibrium. It is straightforward to see that every finite game with perfect recall has a subgame perfect equilibrium. It is also true that with the appropriate definition every finite extensive form game has a subgame perfect equilibrium.

The definition we have given is for games with perfect recall. There is nothing about the idea of subgame perfect equilibria that requires perfect recall, though, of course, such a concept cannot be defined in terms of behaviour strategies—in games without perfect recall there may be no equilibria in behaviour strategies. How to formulate such a definition is here left as an exercise for the reader.

3.2. Extensive Form Perfect Equilibria. There are a number of equivalent ways to define extensive form perfect equilibria in games with perfect recall. We shall give only one. We first make a few comments that will show that the details of the definition are well defined. If a behaviour strategy profile is completely mixed then every information set is reached with strictly positive probability and so the conditional probability distribution on the nodes of the information set is well defined. Moreover the assumption of perfect recall implies that

while the strategy of the player who owns the information set may affect the probability that the information set is reached it will not affect the conditional distribution on the nodes, given that the information set is reached. Thus, given a completely mixed strategy profile, at each information set each action of the player who owns that information set gives a conditional distribution over the terminal nodes, conditional on that information set being reached and all the players playing their parts of the behaviour strategy profile after the information set. And so there is a conditional expected utility associated to each of the actions available at the information set.

Definition 17. A completely mixed behaviour strategy profile b is an *extensive form perfect equilibrium* of Γ if there is a sequence of completely mixed behaviour strategy profiles b^t converging to b such that for each Player n and each information set of Player n the choices of Player n given by the strategy b_n are optimal given the conditional distribution on the information set implied by b^t and the behaviour of all the players given by b^t at the information sets following that information set.

3.3. Quasi-Perfect Equilibria. Quasi-perfect equilibria were defined by van Damme (1984).

Definition 18. A completely mixed behaviour strategy profile b is a *quasi-perfect equilibrium* of Γ if there is a sequence of completely mixed behaviour strategy profiles b^t converging to b such that for each Player n and each information set of Player n the choices of Player n given by the strategy b_n are optimal given the conditional distribution on the information set implied by b^t and the behaviour of the other players given by b^t and the behaviour of Player n given by b_n at the information sets following that information set.

If we compare this to the definition of extensive form perfect equilibria above we see that the only difference is that the “mistakes” of the player moving at an information set at later nodes are taken into account in the definition of extensive form perfect equilibria and are not taken into account in the definition of quasi-perfect equilibria.

The notion of quasi-perfect equilibrium has certain advantages over the notion of extensive form perfect equilibrium. We might see them as coming from quite different ideas. In his motivation Selten refers to the small possibility of the players not being completely rational, that we all make mistakes. In that case a player would do well to consider the possibility that he himself may make a mistake, and Selten’s definition makes some sense. (Though his model of irrationality is

somewhat limited.) While van Damme does not really give an intuitive justification of his concept, one might think of it as requiring a certain robustness to strategic uncertainty, that one is never completely sure of what the others will do. In this case there is no reason for a player to be uncertain as to what he himself will do in the future. This idea of robustness to uncertainty about the other players might also motivate, in the context of the normal form, the concept of normal form perfect equilibrium, that we defined earlier. And we have the following result.

Theorem 6 (van Damme (1984)). *A quasi-perfect equilibrium of an extensive form game is Kuhn equivalent to a normal form perfect equilibrium of the associated normal form game, and thus an admissible equilibrium.*

We say Kuhn equivalent since one is a profile of behaviour strategies and the other a profile of mixed strategies. Mertens (1995) argues that quasi-perfect equilibrium is precisely the right mixture of admissibility and backward induction.

Mertens (1995) offers the game shown in Figure Figure 5 in which the set of extensive form perfect equilibria and the set of admissible equilibria have an empty intersection and hence also the set of extensive form perfect equilibria and set of quasi-perfect equilibria.

The game may be thought of in the following way. Two players agree about how a certain social decision should be made. They have to decide who should make the decision and they do this by voting. If they agree on who should make the decision that player decides. If they each vote for the other then the good decision is taken automatically. If each votes for himself then a fair coin is tossed to decide who makes the decision. A player who makes the social decision is not told if this is so because the other player voted for him, or because the coin toss chose him. The payoffs are such that each player prefers the good outcome to the bad outcome. (In Mertens (1995) there is an added complication to the game. Each player does slightly worse if he chooses the bad outcome than if the other chooses it. However this additional complication is, as Mertens pointed out to me, totally unnecessary for the results.)

In this game the only admissible equilibrium, and hence the only quasi-perfect equilibrium, has both players voting for themselves and taking the right choice if they make the social decision. However, any perfect equilibrium must involve at least one of the players voting for the other with certainty. At least one of the players must be at least as likely as the other to make a mistake in the second stage. And such a player, against such mistakes, does better to vote for the other.




Figure 5

As an exercise, possibly after the School is over, you should calculate all the equilibria, all the extensive form perfect equilibria, and all the quasi-perfect equilibria of this game. Then represent them graphically. You will need to suppress some of the dimensions in which the strategy does not vary across the equilibria.

3.4. Sequential Equilibria. We have seen that extensive form perfect equilibrium and quasi-perfect equilibrium may be thought of as requiring robustness to the possibility of mistakes and a small degree of strategic uncertainty respectively. What if we asked only for a backward induction requirement, namely that players choose rationally at each information set, without saying anything about mistakes or uncertainty.

The problem is, of course, that unless players take all actions with strictly positive probability then some information sets may not be reached with strictly positive probability and thus the equilibrium strategy does not tell a player what he should believe about where he is in the information set. Kreps and Wilson (1982) get around this difficulty by making the beliefs of the players part of their definition of an equilibrium.

The concept of sequential equilibrium recognizes this by defining an equilibrium to be a pair consisting of a behaviour strategy and a system of beliefs.

Definition 19. A *system of beliefs* μ gives, for each information set, a probability distribution over the nodes of that information set. An assessment is a pair (b, μ) where b is a profile of behaviour strategies and μ a system of beliefs.

Definition 20. Given an assessment (b, μ) , the behaviour strategy b_n of Player n is said to be *sequentially rational* with respect to that assessment, if at every information set at which a player moves, it maximizes the conditional payoff of the player, given his beliefs at that information set and the strategies of the other players.

Definition 21. The assessment (b, μ) said to be *consistent* if there is a sequence of completely mixed behaviour strategy profiles b^t converging to b such the the beliefs μ^t obtained from b^t as conditional probabilities converge to μ .

Definition 22. An assessment (b, μ) is a *sequential equilibrium* if the strategy of each player is sequentially rational with respect to the assessment and the assessment is consistent.

If (b, μ) is a sequential equilibrium for some beliefs μ then we shall say that b is a sequential equilibrium strategy profile.

Sequential equilibrium is a weakening of both extensive form perfect equilibrium and of quasi-perfect equilibrium.

Theorem 7 (Kreps and Wilson (1982); van Damme (1984)). *Given an extensive form game with perfect recall if b is an extensive form perfect equilibrium then b is a sequential equilibrium strategy profile. Similarly if b is an quasi-perfect equilibrium then b is a sequential equilibrium strategy profile.*

But not much of a weakening. If we fix the extensive form, that is the game without the payoffs we can think of the space of games with this extensive form as some finite dimensional real space, \mathbb{R}^K for some K .²

Theorem 8 (Kreps and Wilson 1982; Blume and Zame 1994). *For any extensive form, except for a closed set of payoffs of lower dimension than the set of all possible payoffs, the sets of sequential equilibrium strategy profiles and extensive form perfect equilibrium strategy profiles coincide.*

² K will be, of course, the number of terminal nodes times the number of players.

And similarly.

Theorem 9 (Pimienta and Shen 2013; Hillas, Kao, and Schiff 2016). *For any extensive form, except for a closed set of payoffs of lower dimension than the set of all possible payoffs, the sets of sequential equilibrium strategy profiles and quasi-perfect equilibrium strategy profiles coincide.*

Aside: Real Algebraic Geometry. The last two results are proved using techniques from real algebraic geometry, the mathematics of the semi-algebraic sets that Sylvain spoke of yesterday. It is not obvious from the little that Sylvain said that the various sets that we define are semi-algebraic, that is are defined by a finite number of polynomial equalities and inequalities. In fact they don't seem to be, involving statements about the limit of infinite sequences. However the Tarski-Seidenberg Theorem (Tarski, 1951; Seidenberg, 1954) shows that they are.

A first-order formula is an expression involving variables and constants, the quantifiers \forall and \exists , the logical operators \wedge , \vee , and \neg , the operations $+$, $-$, \cdot , and $/$, and the relations $=$, $>$, and $<$. Variables in a first-order formula which are quantified are *bound*, while unbound variables are *free*. By definition, $X \subset \mathbb{R}^n$ is semi-algebraic if and only if it is defined by a first-order formula with n free variables and no bound variables. However, the Tarski-Seidenberg Theorem states that every first-order formula with n free variables is equivalent to a first-order formula with n free variables and no bound variables and hence all sets defined by first-order formulas are semi-algebraic.

An implication of the Tarski-Seidenberg Theorem is that all of the sets, functions, correspondences, and so on, that we have defined are in fact semi-algebraic.

3.5. Proper Equilibria. We turn again to normal form games, though we shall soon see the relevance to the extensive form. Myerson (1978) noticed what seemed to him the poor performance of the perfect equilibrium concept in some games and suggested a strengthening of normal form perfect equilibrium, which he called proper equilibrium.

Definition 23. An ε -*proper equilibrium* is a completely mixed strategy vector such that for each player if, given the strategies of the others, one strategy is strictly worse than another, then the first strategy is played with probability at most ε times the probability with which the second is played. In other words, more costly mistakes are made with lower frequency. A strategy profile is a *proper equilibrium* if it is the limit of a sequence of ε -proper equilibria as ε goes to 0.

Theorem 10 (Myerson (1978)). *Every finite game has at least one proper equilibrium of its normal form. Every proper equilibrium is normal form perfect, and hence admissible.*

While proper equilibrium is defined in terms of the normal form it does, in fact, have implications for the analysis of extensive form games.

Theorem 11 (van Damme (1984)). *For any normal form game and any extensive form game having that normal form, any proper equilibrium of the normal form game is Kuhn equivalent to a quasi-perfect equilibrium of the extensive form game.*

In the light of Theorem 7, this implies that any proper equilibrium is Kuhn equivalent to a sequential equilibrium strategy, a result proved independently by Kohlberg and Mertens (1982, 1986).

A partial result in the other direction also holds.

Theorem 12 (Mailath, Samuelson, and Swinkels (1997); Hillas (1997)). *An equilibrium σ of a normal form game G is proper if and only if there exists a sequence of completely mixed strategies $\{\sigma^t\}$ with limit σ such that for any extensive form game Γ having the normal form G , for any sufficiently small $\varepsilon > 0$, for sufficiently large t , some behaviour strategy corresponding to σ^t is an ε -quasi-perfect equilibrium of Γ .*

The full converse is false; it is not true that an equilibrium that is quasi-perfect in any extensive form game with a given normal form is necessarily proper in that normal form.

Consider the game given in Figure 6. The equilibrium (A, V) is not proper. To see this we argue as follows. Given that Player I plays A , Player II strictly prefers W to Y and X to Y . Thus in any ε -proper equilibrium Y is played with at most ε times the probability of W , and also at most ε times the probability of X . The fact that Y is less likely than W implies that Player I strictly prefers B to C , while the fact that Y is less likely than X implies that Player I strictly prefers B to D . Thus in an ε -proper equilibrium C and D are both played with at most ε times the probability of B . This in turn implies that Player II strictly prefers Z to V , and so there can be no ε -proper equilibrium in which V is played with probability close to 1. Thus (A, V) is not proper.

Nevertheless, in any extensive form game (with perfect recall) having this normal form there are perfect—and quasi-perfect—equilibria equivalent to (A, V) . The proof of this fact is a little tedious, but the idea is quite simple. It is that there is no way of arranging Player II's decisions in an extensive form so that the requirements for a quasi-perfect equilibrium will imply both that W is much more likely than Y

	<i>V</i>	<i>W</i>	<i>X</i>	<i>Y</i>	<i>Z</i>
<i>A</i>	1, 1	3, 1	3, 1	0, 0	1, 1
<i>B</i>	1, 1	2, 0	2, 0	0, 0	1, 2
<i>C</i>	1, 1	1, 0	2, 1	1, 0	1, 0
<i>D</i>	1, 1	2, 1	1, 0	1, 0	1, 0

Figure 6

and that X is much more likely than Y . Recall that this conjunction was needed in the argument that (A, V) is not proper.

Let us suppose that we have an extensive form game having the normal form of Figure Figure 6. We wish to show that there is a behaviour strategy equivalent to (A, V) that is quasi-perfect. To simplify things a bit let us first transform the game by expanding information sets so that neither player learns anything about what the other does. While this might change, for example, the set of sequential equilibria, it will not change the set of quasi-perfect equilibria—or, for that matter, the set of perfect equilibria. (The proof of this is straightforward, and is left as an exercise for the reader.)

Now what we have is essentially a one person extensive form for each player. And by the assumption of perfect recall these one person problems will have perfect information. Consider Player II's problem. For the moment remove all the branches leading to only V or Z . In this reduced problem consider one of the nodes at which Player II has a nontrivial decision and has, for the last time, the option to play Y . There are three possible cases: W is available, but not X ; X is available, but not W ; and both W and X are available. I shall make the argument for the first case; the second case is symmetric; and the third even easier.

We now construct the behaviour strategy in the following way. On the path to this node each branch is taken with probability at least ε . Action Y at any node following this one, is taken with probability ε . At this node, if W is "optimal" then Y is taken with probability ε ; if not, then W is taken with probability ε and Y with probability ε^2 . Every other "suboptimal" action is taken with probability ε^K , where K is some very large number, larger than the total number of nodes in the tree, say. Here "optimal" and "suboptimal" mean with respect to

the order $V \succ W \succ X \succ Z \succ Y$. For this order optimal actions are always unique so, since we have assigned probability to all suboptimal actions we have uniquely defined a behaviour strategy.

Now for the defined strategy we have the following: W is played with probability of order ε^{K_W} ; Y is played with probability of order ε^{K_Y} ; and X is played with probability of order ε^{K_X} , where $K_W < K_Y < K < K_X$. (Also, though we shall not need these facts since they do not affect Player I's preferences, V is played with probability of order 1 and Z is played with probability of order ε^{K_Z} , where $K < K_Z$.) Now this strategy for Player II can be seen to induce the following preferences for Player I: $A \succ D \succ B \succ C$. So, by the same argument as that given by van Damme for the proof of Theorem 1, there is a behaviour strategy satisfying the conditions for a quasi-perfect equilibrium equivalent to the mixed strategy $(1 - \varepsilon - \varepsilon^2 - \varepsilon^3, \varepsilon^2, \varepsilon^3, \varepsilon)$. And this strategy induces the preferences that we hypothesized for Player II. Thus this pair of strategies is an ε -quasi-perfect equilibrium, as required.

4. INVARIANCE

We have seen above that proper equilibria of a normal form game are quasi-perfect in all extensive forms having that normal form. And also that a partial converse is true. This suggests that the requirement of backward induction, at least as far as it is implied by the solution concept quasi-perfect equilibrium, to some extent depends only on the normal form. von Neumann and Morgenstern (1944) had already argued that the normal form encompassed all the relevant information about a game. Kohlberg and Mertens (1986), and later and more strongly, Mertens (2003, 1989, 1991b, 1992), have argued for even more. Kohlberg and Mertens argue that since the players can, in any case already play mixed strategies it should not matter if we add an existing mixed strategy as a new pure strategy. We would, in fact just be saying twice that the player could play this mixed strategy. Thus the solution should depend only the reduced normal form, the normal form when all strategies equivalent to mixtures of other strategies have been removed. This property is called reduced normal form invariance.

Mertens goes further adding that if two games have the same best reply correspondence then their solutions should be the same, and even that if the best reply correspondences of two games were the same on the interior of the strategy space—that is, on the admissible best reply correspondence—then the solutions of the game should be the same. He terms these properties ordinality.

I find this argument convincing, but shall not address it here. Rather, I'll just point out two things. First, that Nash equilibrium and normal form perfect equilibrium satisfy reduced normal form invariance and even ordinality defined in terms of the best reply correspondence, while normal form perfect equilibria also satisfy ordinality defined with the admissible best reply correspondence. And second that if we also require other properties such as backward induction that adding the requirement of invariance or ordinality can substantially increase the implications of those requirements.

5. THE NEED FOR SET-VALUED SOLUTIONS

[Large parts of this section are taken directly from Hillas and Kohlberg (2002).]

We are seeking an answer to the question: What are the self-enforcing behaviours in a game? As we indicated, the answer to this question should satisfy the various invariances we discussed above. We also require that the solution satisfy stronger forms of rationality than Nash equilibrium and normal form perfect equilibrium, the two equilibrium concepts that we have said do satisfy those invariances. In particular, we want our solution concept to satisfy admissibility that we defined in Section 2, and some form of the iterated dominance condition we shall define in the next section, the backward induction condition we discussed in Section 3, and the forward induction condition we shall also define in the next section. We also want our solution to give some answer for all games.

It is impossible for a single valued solution concept to satisfy these conditions. In fact, two separate subsets of the conditions are inconsistent for such solutions. Admissibility and iterated dominance are inconsistent, as are backward induction and invariance.

	<i>L</i>	<i>R</i>
<i>T</i>	3, 2	2, 2
<i>M</i>	1, 1	0, 0
<i>B</i>	0, 0	1, 1

Figure 7

To see the inconsistency of admissibility and iterated dominance consider the game in Figure 7 (taken from Kohlberg and Mertens (1986, p. 1015)). Strategy B is dominated (in fact, strictly dominated) so by the iterated dominance condition the solution should not change if B is deleted. But in the resulting game admissibility implies that (T, L) is the unique solution. Similarly, M is dominated so we can delete M and then (T, R) is the unique solution.

To see the inconsistency of ordinality and backward induction consider the game of Figure 8 (taken from Kohlberg and Mertens (1986, p. 1018)). Nature moves at the node denoted by the circle going left with probability α and up with probability $1 - \alpha$. Whatever the value of α the game has the reduced normal form of Figure Figure 8a. (Notice that strategy Y is just a mixture of T and M .) Since the reduced normal form does not depend on α , ordinality implies that the solution of this game should not depend on α . And yet, in the extensive game the unique sequential equilibrium has Player 2 playing L with probability $(4 - 3\alpha)/(8 - 4\alpha)$.




Figure 8

Thus it may be that elements of a solution satisfying the requirements we have discussed would be sets. However we would not want these sets to be too large. We are still thinking of each element of the

	L	R
T	$1, -1$	$1, -1$
M	$2, -2$	$-2, 2$
B	$-2, 2$	$2, -2$

Figure 8a

solution as, in some sense, a single pattern of behaviour. In generic extensive form games we might think of a single pattern of behaviour as being associated with a single equilibrium outcome, while not specifying exactly the out of equilibrium behaviour. One way to accomplish this is to consider only connected sets of equilibria. In the definition of Mertens (1989, 1991b) the connectedness requirement is strengthened in a way that corresponds, informally, to the idea that the particular equilibrium should depend continuously on the “beliefs” of the players. Without a better understanding of exactly what it means for a set of equilibria to be the solution we cannot say much more. However some form of connectedness seems to be required.

6. FORWARD INDUCTION

The concept of *forward induction* was introduced and discussed by Kohlberg and Mertens (1986). The precise status of this concept is not clear in their paper. They do not list “forward induction” as one of the requirements for a solution but it seems to be important in the motivation and stable sets do, in fact, satisfy their concept of forward induction. The property that Kohlberg and Mertens call forward induction is the following: “A stable set contains a stable set of any game obtained by the deletion of a strategy which is an inferior response in all the equilibria of the set.” ((Kohlberg and Mertens, 1986, p. 1029.)

This is obviously a strong property. For example, it means that if there are two strategies that are both inferior responses then when the more preferred of these strategies is deleted and the less preferred kept the solution should remain stable. Even before this full strength there are those who have argued against this kind of requirement. Cho and Kreps (1987) show how a series of stronger implementations of the forward induction like ideas refine the set of sequential equilibria in signaling games. They argue that the relatively strong implementations are quite unintuitive.

We next look at an example of a game that had a large role in the motivation of the idea of forward induction and show that a combination of forward induction and reduced normal form invariance is sufficient to eliminate the sets of equilibria that do not satisfy forward induction. The extensive form game is shown in Figure 9 and the associated normal form game in Figure 9a

In this game there are equilibria in which Player 1 plays T . The argument that has been given against such equilibria is as follows: If Player 2 is called on to move he does not see whether Player 1 has played M or B . However he does know that if Player 1 chooses M the most that Player 1 can obtain is 1 while Player 1 could have obtained 2 with certainty by choosing T . On the other hand if Player 1 can convince Player 2 that he has indeed played strategy B —so that Player 2 will choose L —then Player 1 gains by playing B . Thus when Player 2 sees that Player 1 has played either M or B Player 2 should believe that Player 1 has actually played B . This will lead Player 2 to play L and so Player 1 should not play T but rather B .




Figure 9

This reasoning and variants of it are what is meant by “forward induction”. One is either convinced by such reasoning or not. I myself find it fairly compelling. Here however I wish to examine how much of the results that follow from this kind of reasoning can be obtained from two other kinds of properties: invariance to “inessential” changes in the game; and backward induction. I shall become more precise about these properties as the paper progresses. For this example all

	<i>L</i>	<i>R</i>
<i>T</i>	2, 5	2, 5
<i>BU</i>	4, 1	0, 0
<i>BD</i>	0, 0	1, 4

Figure 9a

that is needed is invariance to the reduced normal form and that the solution contain a subgame perfect equilibrium.




Figure 10

The extensive form game shown in Figure Figure 10, with its normal form given in Figure Figure 10a, has the same reduced normal form as the original game of Figure Figure 9. Nevertheless one does not require any kind of forward induction argument to argue against the equilibria in which Player 1 plays *T* in this game. Subgame perfection suffices. It is straightforward to see that (B, L) is the unique equilibrium of the subgame, and so (B, L) is the unique subgame perfect—and hence unique sequential—equilibrium of the game.

	<i>L</i>	<i>R</i>
<i>T</i>	2, 5	2, 5
<i>BU</i>	4, 1	0, 0
<i>BD</i>	0, 0	1, 4
<i>X</i>	$\frac{5}{2}, 4$	$\frac{3}{2}, \frac{15}{4}$

Figure 10a

In a number of games that have been examined in the literature similar thing occur. At one point I had conjectured that some form of forward induction was implied by backward induction and various forms of invariance; I was wrong.

Let us now look at an example of a game in which it will be clear that we cannot get the full strength of forward induction from backward induction arguments.

In some ways this game resembles some of the signaling games that Cho and Kreps (1987) argue show the unreasonable strength of stability, and other strong forms of forward induction. It differs in that the game is unavoidably a three person game, and so the kind of techniques that work in signaling games cannot be applied.

The normal form perfect equilibria are (T, R, U) and $\{(T, L, (z, 1 - z)) \mid z \leq \frac{1}{3}\}$. Consider the second set. Since each player moves only once, and has only two actions, the proper, perfect, and quasi-perfect equilibria all coincide. And all equilibria in the second set are perfect/quasi-perfect/proper. However, even taken as a whole, this set does not satisfy what Kohlberg and Mertens call forward induction. Nowhere in this set is B a best response for Player 1. Thus forward induction requires that the set should remain stable if the strategy B is eliminated. If B is eliminated then D becomes weakly dominated for Player 3 and so Player 3 will not play D with positive probability in any stable set. However, in all of the equilibria in the set we are considering D is played with probability at least $\frac{2}{3}$.

Before going further consider a slight modification of the game

The only change is that Player 1's payoff following (B, L, U) has been changed from 2 to 0. Now B is weakly dominated for Player 1, and the argument can proceed as before, but now in terms of the elimination of weakly dominated strategies.




Figure 11




Figure 12

We consider a final example. This is a two person normal form game and the forward induction arguments are in terms of iterated deletion of weakly dominated strategies. I learned this example from Hari Govindan.

In this example we see an equilibrium that is eliminated by forward induction arguments. Nevertheless, the equilibrium is proper, and remains proper no matter what mixtures are added as new strategies.

	t_1	t_2	t_3
s_1	1, 1	0, 1	1, 1
s_2	1, 0	1, 1	0, -1
s_3	1, 1	-1, 0	1, 1

Figure 13

There are two perfect equilibria in this game, (s_1, t_1) and (s_2, t_2) . Both these equilibria are also proper. The equilibrium (s_1, t_1) does not satisfy forward induction. The strategies s_3 and t_3 are weakly dominated. And, if these strategies are eliminated, s_1 and t_1 become weakly dominated. Thus the only equilibrium that survives iterated deletion of weakly dominated strategies is (s_2, t_2) .

But the equilibrium (s_1, t_1) is proper and remains proper no matter what mixtures are added as new strategies. This is a little tedious to show, but the idea is clear. The relevant mixtures will be the ones that put large weight on the first strategy.

One simply observes that if t_3 is much more likely than t_2 or any of the added strategies, then s_1 and s_3 , which are best replies against t_3 will be better than s_2 or any of the added mixtures. And so, we can construct ε -proper equilibria with the first strategy played with large probability and the third strategy played with larger probability than any other strategies.

7. AN INTRODUCTION TO STRATEGIC STABILITY

Kohlberg and Mertens (1986) gave a list of requirements that a concept of strategic stability should satisfy. They showed that even quite weak versions of their requirements implied that the solution concept should assign sets of equilibria as solutions to the game. Thus a stability concept is a rule that assigns to each game in the domain of games under consideration a collection of subsets of the space of (mixed) strategy profiles of the game. Since the paper of Kohlberg and Mertens the list of requirements a concept of strategic stability should satisfy has been modified and expanded, particularly in the work of Mertens (2003, 1989, 1991b,a, 1992, 1995) The list presented here is a somewhat modified and expanded version of the original one.

- (1) Existence. *Every game has at least one stable set.*
- (2) Connectedness. *Stable sets are connected.*
- (3) Admissibility. *Consists only of normal form perfect equilibria.*
- (4) Backward Induction. *Contains a proper equilibrium.*
- (5) Independence of Inadmissible Strategies. *One form of the forward induction idea.*
- (6) Ordinality. *A stability concept is ordinal.*
- (7) The Small Worlds Axiom *Suppose that the players can be divided into insiders and outsiders, and that the payoffs of the insiders do not depend on the strategies of the outsiders. Then the stable sets of the game between the insiders are precisely the projection of the stable sets of the larger game.*

Kohlberg and Mertens (1986) consider a space of perturbations such as that defined by Selten (1975), in which each strategy of a player has attached to it a small probability. Lets assume that all these probabilities are no more than δ and call this set P_δ . Perturbed games are defined in a natural way. (There are a number of alternative methods.)

Kohlberg and Mertens then define a *stable set of equilibria* to be a set of Nash equilibria that is minimal with respect to the property that all sufficiently small perturbations of the game have equilibria close to the stable set.

They define a *hyperstable set of equilibria* to be a set of Nash equilibria that is minimal with respect to the property that for any game with the same reduced normal form and for any sufficiently small perturbation of the payoffs of that game the game has equilibria close to the stable set.

They also define a third concept, that of fully stable sets of equilibria, but we won't say anything further about that here.

There are some problems with two aspects of these definitions. The minimality requirement does not achieve exactly what was intended, and may be inconsistent with the ordinality requirement. The attempt to impose invariance in the definition of hyperstability is also, for reasons that we won't go into here, not completely satisfactory.

Thus, we'll define a *KM-stable set of equilibria* to be a connected set of normal form perfect equilibria such that all sufficiently small perturbations of the game have equilibria close to the stable set.

Kohlberg and Mertens point out that their original definition does not satisfy the backward induction property and express the hope that some modification of that definition will.)The slightly modified definition we gave above does not satisfy backward induction either.)

The paper of Kohlberg and Mertens gave one model for definitions of strategic stability. One defines a space of perturbations. Defines how each perturbation gives perturbed games, or at least how one can associate a set of “equilibria” to each perturbation and then require that the stable set be such that all small perturbations have nearby equilibria.

We give now a definition of stable equilibria that does satisfy backward induction. One defines BR-stability by considering directly perturbations to the best reply correspondence (together with a fairly fine topology on such perturbations). One can show that this definition is equivalent to a definition that looks at continuous functions from Σ to P_δ as the space of perturbations. (This is the main result of Hillas, Jansen, Potters, and Vermeulen (2001))

One can also make a definition based on the idea of making the minimal change to the definition of KM-stability to give the desired properties. This was one of the approaches I took in Hillas (1990). It seemed to me at the time that the more radical approach of perturbing the best reply correspondence was more promising. I’m no longer sure this is true. The work of De Stefano Grigis (2014, 2015) seeks to make the minimal modifications to the definition of Kohlberg and Mertens so that it satisfies the requirements.

Mertens gave a more fundamentally different way of redefining stability. Rather than keeping the form of the definition and changing the space of perturbations Mertens kept the same space of perturbations, P_δ , and changed the way of defining stability.

Consider a small neighbourhood of 0 in η . Let’s call this P_δ and call the boundary of this neighbourhood ∂P_δ . Consider also some part S of the graph of the equilibrium correspondence $E : \eta \rightarrow \Sigma$ and let the part above P_δ be called S_δ and the part above ∂P_δ be called ∂S_δ .

Mertens says that the part of S above zero is a stable set if the projection map from S_δ to P_δ is nontrivial (in some sense) for sufficiently small δ . Mertens gave a number of definitions involving different specifications of “nontrivial”. The “right” definition seems to be the one involving homology theory, which says that the projection should not be homologically trivial.

Since homology theory puts a fairly coarse structure on things it is relatively “easy” to be homologically trivial, that is, homologous to a map to ∂P_δ . Thus the definition in terms of homology is a strong one.

The easiest definition to understand simply says that the projection map from S_δ to P_δ should not be homotopic to a map from S_δ to ∂P_δ under a homotopy that leaves the map from ∂S_δ to ∂P_δ unchanged. We’ll call sets that satisfy this requirement *homotopy-stable sets*.




Figure 14

Hillas, Jansen, Potters, and Vermeulen (2001) show that every stable set in the sense of Mertens is a CKM-stable set. They also show that every CKM-stable set is a BR-stable set. So, every stable set in the sense of Mertens is a BR-stable set. These and some known relations between various stability concepts are displayed below. The relations are shown in Figure Figure 14 with those marked K&M were proved in Kohlberg and Mertens (1986); those marked H90 in Hillas (1990)); that marked McLennan in McLennan (1995); that marked Govindan in Govindan (1995); and the unmarked relations proved in Hillas, Jansen, Potters, and Vermeulen (2001).

REFERENCES

- STEVE ALPERN (1988): “Games with Repeated Decisions,” *SIAM Journal of Control and Optimization*, 26, 468–477.
- LAWRENCE E. BLUME AND WILLIAM R. ZAME (1994): “The Algebraic Geometry of Perfect and Sequential Equilibrium,” *Econometrica*, 62, 783–794.
- IN-KOO CHO AND DAVID M. KREPS (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102, 179–221.
- NORMAN DALKEY (1953): “Equivalence of Information Patterns and Essentially Determinate Games,” in *Contributions to the Theory of Games, Vol. 2*, edited by H. W. Kuhn and A. W. Tucker, 127–143, Princeton University Press, Princeton, NJ.
- FEDERICO DE STEFANO GRIGIS (2014): “Strategic Stability of Equilibria: the Missing Paragraph,” unpublished.

- FEDERICO DE STEFANO GRIGIS (2015): *Essays on Nash Equilibrium Refinements*, Ph.D. thesis, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.
- SRIHARI GOVINDAN (1995): “Every Stable Set Contains a Fully Stable Set,” *Econometrica*, 63, 191–193.
- JOHN HILLAS (1990): “On the Definition of the Strategic Stability of Equilibria,” *Econometrica*, 58, 1365–1390.
- JOHN HILLAS (1997): “On the Relation between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games,” unpublished.
- JOHN HILLAS, MATHIJS JANSEN, JOS POTTERS, AND DRIES VERMEULEN (2001): “On the Relation Between Some Definitions of Strategic Stability,” *Mathematics of Operations Research*, 28.
- JOHN HILLAS, TINA KAO, AND AARON SCHIFF (2016): “A Real Algebraic Proof of the Generic Equivalence of Quasi-Perfect and Sequential Equilibria,” unpublished.
- JOHN HILLAS AND ELON KOHLBERG (2002): “Foundations of Strategic Equilibrium,” in *Handbook of Game Theory with Economic Applications*, Vol. 3, edited by Robert J. Aumann and Sergiu Hart, Elsevier.
- J. ISBELL (1957): “Finitary Games,” in *Contributions to the Theory of Games, Volume III, Annals of Mathematics Study*, 39, edited by Melvin Drescher, Albert William Tucker, and Philip Wolfe.
- ELON KOHLBERG AND JEAN-FRANÇOIS MERTENS (1982): “On the Strategic Stability of Equilibria,” CORE Discussion Paper 8248, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.
- ELON KOHLBERG AND JEAN-FRANÇOIS MERTENS (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54, 1003–1038.
- DAVID M. KREPS AND ROBERT WILSON (1982): “Sequential Equilibria,” *Econometrica*, 50, 863–894.
- HAROLD W. KUHN (1953): “Extensive Games and the Problem of Information,” in *Contributions to the Theory of Games*, Vol. 2, 193–216, Princeton University Press, Princeton NJ.
- GEORGE J. MAILATH, LARRY SAMUELSON, AND JEROEN M. SWINKELS (1997): “How Proper is Sequential Equilibrium?” *Games and Economic Behavior*, 18, 193–218.
- ANDREW MCLENNAN (1995): “Invariance of Essential Sets of Nash Equilibria,” unpublished.
- JEAN-FRANÇOIS MERTENS (1989): “Stable Equilibria—A Reformulation, Part I: Definition and Basic Properties,” *Mathematics of Operations Research*, 14, 575–624.
- JEAN-FRANÇOIS MERTENS (1991a): “Equilibrium and Rationality:

- Context and History Dependence,” in *Issues in Contemporary Economics, Vol. 1: Markets and Welfare, Proceedings of the Ninth World Congress of the International Economic Association, I.E.A. Conference Vol. 98*, edited by Kenneth J. Arrow, 198–211, MacMillan Co., New York.
- JEAN-FRANÇOIS MERTENS (1991b): “Stable Equilibria—A Reformulation, Part II: Discussion of the Definition and Further Results,” *Mathematics of Operations Research*, 16, 694–753.
- JEAN-FRANÇOIS MERTENS (1992): “The Small Worlds Axiom for Stable Equilibria,” *Games and Economic Behavior*, 4, 553–564.
- JEAN-FRANÇOIS MERTENS (1995): “Two Examples of Strategic Equilibria,” *Games and Economic Behavior*, 8, 378–388.
- JEAN-FRANÇOIS MERTENS (2003): “Ordinality in Non-Cooperative Games,” *International Journal of Game Theory*, 32, 387–430.
- ROGER MYERSON (1978): “Refinement of the Nash Equilibrium Concept,” *International Journal of Game Theory*, 7, 73–80.
- JOHN NASH (1950): “Equilibrium Points in N -Person Games,” *Proceedings of the National Academy of Sciences*, 36, 48–49.
- JOHN NASH (1951): “Non-Cooperative Games,” *Annals of Mathematics*, 54, 286–295.
- MICHELE PICCIONE AND ARIEL RUBINSTEIN (1997): “On the Interpretation of Decision Problems with Imperfect Recall,” *Games and Economic Behavior*, 20, 3–24.
- CARLOS PIMIENTA AND JIANFEI SHEN (2013): “On the equivalence between (quasi-)perfect and sequential equilibria,” *International Journal of Game Theory*, 43, 1–8.
- ABRAHAM SEIDENBERG (1954): “A new decision method for elementary algebra,” *Annals of Mathematics*, 60, 365–374.
- REINHARD SELTEN (1965): “Spieltheoretische Behandlung Eines Oligopolmodells mit Nachfragertragheit,” *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–324, 667–689.
- REINHARD SELTEN (1975): “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 4, 25–55.
- ALFRED TARSKI (1951): *A Decision Method for Elementary Algebra and Geometry*, University of California Press, Berkeley, 2nd edition.
- ERIC VAN DAMME (1984): “A Relation between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games,” *International Journal of Game Theory*, 13, 1–13.
- ERIC VAN DAMME (1991): *Stability and Perfection of Nash Equilibria*, Springer-Verlag, Berlin, 2nd edition.
- JOHN VON NEUMANN AND OSKAR MORGENSTERN (1944): *Theory of*

Games and Economic Behavior, Princeton University Press, Princeton NJ.