

González-Díaz, Julio; Meléndez-Jiménez, Miguel A.

Working Paper

On the notion of perfect bayesian equilibrium

Discussion Paper, No. 1446

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: González-Díaz, Julio; Meléndez-Jiménez, Miguel A. (2011) : On the notion of perfect bayesian equilibrium, Discussion Paper, No. 1446, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/59628>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

On the Notion of Perfect Bayesian Equilibrium*

JULIO GONZÁLEZ-DÍAZ

Departamento de Estadística e Investigación Operativa, Universidad de Santiago de Compostela

MIGUEL A. MELÉNDEZ-JIMÉNEZ[†]

Departamento de Teoría e Historia Económica, Universidad de Málaga

Abstract

Often, perfect Bayesian equilibrium is loosely defined by stating that players should be sequentially rational given some beliefs in which Bayes rule is applied “whenever possible”. We argue that there are situations in which it is not clear what “whenever possible” means. Then, we provide an elementary definition of perfect Bayesian equilibrium for general extensive games that refines both weak perfect Bayesian equilibrium and subgame perfect equilibrium.

JEL classification: C72.

Keywords: non-cooperative game theory, equilibrium concepts, perfect Bayesian, Bayes rule.

1 Introduction

Perfect Bayesian equilibrium is profusely used to analyze the game theoretical models that are derived from a wide variety of economic situations. The common understanding is that a perfect Bayesian equilibrium must be sequentially rational given the beliefs of the players, which have to be computed using Bayes rule “whenever possible”. However, the literature lacks a formal and tractable definition of this equilibrium concept that applies to general extensive games and, hence, it is typically the case that perfect Bayesian equilibrium is used without providing a definition beyond the above common understanding.

The main goal of this paper is to make a concise conceptual and pedagogic contribution to the theory of equilibrium refinements. We first argue that there are games (indeed very simple ones) in which it is not clear what “whenever possible” is supposed to mean. Then, we introduce an elementary definition of perfect Bayesian equilibrium that works for all extensive games and that refines both subgame perfect equilibrium and weak perfect Bayesian equilibrium.

***Acknowledgements.** The authors are grateful to Johannes Hörner and Carlos Pimienta for helpful comments. Julio González-Díaz acknowledges financial support from the European Commission through a Marie Curie fellowship and from the Spanish Ministry for Science and Innovation and FEDER through project ECO2008-03484-C02-02. Miguel A. Meléndez-Jiménez acknowledges financial support from the Fulbright Program, from the Spanish Ministry of Education and Science (FU2005-1173 and ECO2008-04424), and from the *Junta de Andalucía* (SEJ2007-2986 and SEJ2009-4794). The authors thank the hospitality of Northwestern University (CMS-EMS), where part of this research was carried out.

[†]Corresponding author: Departamento de Teoría e Historia Económica, Universidad de Málaga. Plaza El Ejido s/n 29013 Málaga, Spain. e-mail: melendez@uma.es.

We call this equilibrium concept *simple perfect Bayesian equilibrium*. The main idea is to refine weak perfect Bayesian equilibrium in the same spirit in which subgame perfection refines Nash equilibrium, but to do so in such a way that it has bite also for imperfect information games. From our point of view, this new equilibrium concept provides a minimal requirement that should be imposed on equilibrium concepts that are based on Bayesian rationality. In addition, further requirements might be imposed depending on the specific characteristics of the games being analyzed. Finally, building up on the approach of Aumann (1964), we present a general model for games in which the players may have a continuum of actions and in which the number of information sets may be countably infinite. In this context, we extend the notion of Bayesian updating to what we call conditional updating, which is needed to explicitly work with beliefs in this general setting and, thus, to extend the notion of simple perfect Bayesian equilibrium.

When game theory started to analyze models with imperfect information, there was a need to refine the classic concepts of Nash equilibrium and subgame perfect equilibrium. Weak perfect Bayesian equilibrium can be regarded as the first step in this direction (though it does not even refine subgame perfection). This equilibrium concept was introduced by Myerson (1991) when preparing the ground for the definition of sequential equilibrium.¹ Also, Mas-Colell et al. (1995) follow this approach and introduce weak perfect Bayesian equilibrium as a bridge between the classic equilibrium concepts and the belief-based ones. Informally, a strategy profile is a weak perfect Bayesian equilibrium when it is sequentially rational given a system of beliefs that is consistent with Bayes rule on the path of the strategy (no restriction is imposed on the beliefs at information sets that are off-path). Though pedagogically useful, weak perfect Bayesian equilibrium has many shortcomings as an equilibrium concept; remarkably, it does not even imply subgame perfection. Thus, this equilibrium concept also has to be refined.

Sequential equilibrium is probably the most widely used equilibrium concept for games with imperfect information and, yet, there are classes of games where less demanding equilibrium concepts that are easier to handle select reasonable strategy profiles. Even setting aside the fact that it is often hard to deal with sequential equilibrium, there are natural economic settings in which it cannot even be defined. The notion of consistent beliefs cannot be (trivially) extended to games in which the players have a continuum of strategies; the models of auctions being the most outstanding example of the latter type of games.

Despite the common practice of using perfect Bayesian equilibrium without providing a formal definition, there have also been some exceptions. To the best of our knowledge, the first paper introducing a formal definition is Harris and Townsend (1981), in the context of mechanism design.² A more elaborate approach is taken in Fudenberg and Tirole (1991) for multistage games with observed actions. They impose some restrictions on how off-path beliefs can be formed and present a definition of perfect Bayesian equilibrium that is natural within the class of games to which their analysis is confined and, moreover, it is easier to study than sequential equilibrium. Yet, these definitions of perfect Bayesian equilibrium cannot be easily extended to general extensive games. Finally, Battigalli (1996) analyzes some natural restrictions on how off-path beliefs should be computed in general extensive games and derives several refinements of subgame perfection which, as he says, may be generically called perfect Bayesian equilibria. The main idea underlying these refinements is what Battigalli called strategic independence: when forming beliefs, the strategic choices of different players should

¹Originally, Myerson called this equilibrium concept weak sequential equilibrium.

²In their setting, all the uncertainty a player faces during the game is about the types of the other players. Once a player knows his type, he forms a prior over the types of the other players and this prior is updated using Bayes rule as the game unfolds. If Bayes rule cannot be applied, *i.e.*, after a history that is inconsistent with the type of the player and the equilibrium strategies, the beliefs of the player are set to coincide with his prior.

be regarded as independent events. Although this approach has been conceptually insightful, it is relatively hard to use in practice, since it requires the use of conditional probability systems on the set of strategy profiles.

Possibly because of the technical complications associated with a formal definition of perfect Bayesian equilibrium, there are many papers in the literature that carry out their analysis for equilibrium concepts that lie in between weak perfect Bayesian equilibrium and sequential equilibrium. They are generically referred to as perfect Bayesian equilibrium and, as mentioned above, they are typically defined as a strategy profile that is sequentially rational given a system of beliefs that is obtained using Bayes rule “whenever possible” (as opposed to do it only on the path as in weak perfect Bayesian equilibrium).

This paper is structured as follows. As we said above, we start by noting in Section 2 that one has to be careful when using the words “whenever possible” to refer to the way in which Bayes rule is applied. In Section 3 we present an elementary but general definition of perfect Bayesian equilibrium. Then, in Section 4 we extend the definition of simple perfect Bayesian equilibrium to a broader context in which players may have a continuum of actions. Finally, we conclude in Section 5.

2 The meaning of “whenever possible”

Since our contribution is primarily conceptual, we introduce as few notations as possible. We think of extensive games as modeled in Selten (1975) and Kreps and Wilson (1982) and refer the reader to those papers for the definitions of the concepts we discuss here.³ An *assessment* is defined as a pair (b, μ) where b is a (behavior) strategy profile and μ is a system of beliefs.

As discussed in the introduction, it is quite common to see papers in which perfect Bayesian equilibrium is defined as a sequentially rational assessment (b, μ) in which the beliefs are computed using Bayes rule “whenever possible”. Although this definition is clear for some classes of games, it is not precise enough for general extensive games. The reason, as illustrated in Example 1 below, is that it is not clear what “whenever possible” is supposed to mean in rigorous mathematical terms.

Example 1. Consider a game whose initial part corresponds with the one depicted in Figure 1. The strategy profile b stands for anyone in which player 1 plays D and player 2 plays d at his two decision nodes. The only nontrivial issue about beliefs has to do with the information set




Figure 1: “Whenever possible” is imprecise.

³Nonetheless, unlike the definition of sequential equilibrium, our approach can be naturally extended to more general settings, such as games in which players have a continuum of actions (see Section 4).

w_3 .⁴ As far as weak perfect Bayesian is concerned, since w_3 is not on the path of strategy b , there is no restriction for the beliefs at the nodes in w_3 , namely x_1 , x_2 , and x_3 . Yet, should we impose any restriction on the beliefs in w_3 if Bayes rule is to be applied “whenever possible”? On the one hand, it can be argued that Bayes rule cannot be applied at w_3 since, given b , w_3 is reached with probability 0 and we cannot condition on probability 0 events. On the other hand, given b and conditional on either x_1 or x_2 being reached, *i.e.*, conditional on player 1 having played U , any belief consistent with Bayes rule should put probability 0 at x_1 , since player 2 is playing d at $w_{2,1}$. Note that, even in the latter case, Bayes rule imposes no restriction on the relative probabilities of x_2 and x_3 . Hence, what approach should we take? We consider that the natural interpretation of “whenever possible” goes in the lines of the second reasoning. In any case, this illustrates that providing a formal definition of what is meant by “whenever possible” that can be applied to any extensive game is by no means a trivial exercise. Moreover, any such definition would probably be difficult to work with and, since any such equilibrium concept would always be refined by sequential equilibrium, its applicability would probably be very limited.⁵ \diamond

3 Simple Perfect Bayesian Equilibrium

Even after giving up on the objective of finding a compelling and useful definition of “whenever possible”, there is still some room for a definition that pushes Bayesian requirements further than weak perfect Bayesian equilibrium and that does it in a clean and practical way. On the other hand, one cannot expect the simple definition of perfect Bayesian equilibrium that we present below to imply (or coincide with) the perfect Bayesian equilibrium introduced in Fudenberg and Tirole (1991). The reason is that, in order to decide what restrictions on beliefs are reasonable, they use the specific structure of the games in the class to which they have restricted. Somehow, by formally disentangling the meaning of “whenever possible” in multistage games with observed actions and independent types, they get to something that is very close to sequential equilibrium. Since we want a definition that is valid for every extensive game and that is easy to deal with, we abstract from the implicit implications of Bayes rule such as the one presented in Example 1.

Let G be an extensive game and Γ its game tree. Given a node x , we denote by $w(x)$ the information set that contains x . We say that a node y comes after a node x if x is on the path from the root to y ; equivalently, we say that y is a *successor* of x . In particular, a node comes after itself. Similarly, a node comes after an information set if it comes after one of the nodes in the information set. The definitions below generalize the definitions of subtree and regular subtree in Selten (1975). Given an information set w , the *quasi-subtree* that begins at w , Γ_w , consists of all the nodes that come after w in Γ and all edges connecting them. The quasi-subtree Γ_w is *regular* if every information set of Γ that contains one node of Γ_w does not contain nodes outside Γ_w , *i.e.*, if an information set v has a node that comes after w , then all the nodes in v come after w . Note that, in the special case in which w is a singleton, the quasi-subtree Γ_w is indeed a subtree and, if Γ_w is regular, there is a well defined subgame of G that begins at w . If Γ_w is a regular quasi-subtree we say that w is itself a *regular* information set. Given an assessment (b, μ) and a regular information set w , we can associate a game with Γ_w in a very natural way. More specifically, let $G_w(\mu)$ be the game defined as follows. First, Nature moves and selects each node x in the information set w with probability $\mu(x)$. Second,

⁴We denote by $w_{i,k}$ the k -th information set of player i . For simplicity, if a player i has only one information set, we denote it by w_i .

⁵This example is similar to the example used in Battigalli (1996) to illustrate the notion of *strategic independence*.

the game unfolds in Γ_w and all remaining elements of the game are the restrictions of the corresponding ones in G .⁶

For the sake of completeness we present now the definition of weak perfect Bayesian equilibrium.

Definition 1. Let G be an extensive game. An assessment (b, μ) is a *weak perfect Bayesian equilibrium* if it is sequentially rational and, on the path of b , μ is derived from b by Bayes rule.

As we have already argued above, the main drawback of weak perfect Bayesian equilibrium is that it does not impose any restriction on the beliefs at off-path information sets and, hence, although it is an equilibrium concept stronger than Nash equilibrium, a weak perfect Bayesian equilibrium does not even need to be subgame perfect. Below we define a version of perfect Bayesian equilibrium that imposes some simple restrictions on the off-path beliefs in a natural way.

Definition 2. Let G be an extensive game. An assessment (b, μ) is a *simple perfect Bayesian equilibrium* if, for each regular information set w , the restriction of (b, μ) to $G_w(\mu)$ is sequentially rational and, on the path of b , the beliefs are derived from b by Bayes rule, *i.e.*, the restriction of (b, μ) to $G_w(\mu)$ is a weak perfect Bayesian equilibrium.

To some extent, simple perfect Bayesian equilibrium is to weak perfect Bayesian equilibrium in imperfect information games what subgame perfection is to Nash equilibrium in perfect information games. We present now a series of straightforward results and examples to illustrate the properties of this new equilibrium concept and its relationships with other equilibrium concepts.

- Proposition 1.**
- (i) *Every sequential equilibrium is a simple perfect Bayesian equilibrium.*
 - (ii) *Every simple perfect Bayesian equilibrium is a weak perfect Bayesian equilibrium.*
 - (iii) *Every simple perfect Bayesian equilibrium is subgame perfect.*
 - (iv) *Every extensive game with perfect recall has, at least, one simple perfect Bayesian equilibrium.*

Proof. (i) A sequential equilibrium is defined as a sequentially rational and consistent assessment. Consistency in the sense of Kreps and Wilson (1982) requires that the beliefs μ are the limit of the beliefs associated with a sequence of completely mixed behavior strategies converging to b . Clearly, this implies that Bayes rule must be applied on the path of all games associated with regular information sets. Hence, sequential equilibrium implies simple perfect Bayesian equilibrium.

(ii) The root node, r , is a regular information set. Then, since (b, μ) is a simple perfect Bayesian equilibrium, we have that the restriction of (b, μ) to $G_r(\mu)$ is sequentially rational and, on the path of b , the beliefs are derived from b by Bayes rule. Since $G_r(\mu) = G$ and the restriction of (b, μ) to G is (b, μ) , we have just stated the definition of weak perfect Bayesian equilibrium.

(iii) There is a one to one correspondence between subgames and games defined from regular information sets that are singletons. In each of these subgames, subgame perfect equilibrium requires that a Nash equilibrium is played, whereas simple perfect Bayesian equilibrium is more

⁶Note that if Γ_w is not a regular quasi-subtree, then it is not so clear how the game $G_w(\mu)$ should be defined. For instance, in the game in Figure 1, $\Gamma_{w_{2,1}}$ is not a regular quasi-subtree, since it does not contain x_3 . Nonetheless, what happens after x_3 might be important to define the game $G_{w_{2,1}}(\mu)$ since, conditional on w_3 being reached, player 3 might put positive probability at x_3 ($\mu(x_3) > 0$).

demanding, asking for a weak perfect Bayesian equilibrium. Hence, simple perfect Bayesian equilibrium implies subgame perfect equilibrium.

(iv) Every extensive game with perfect recall has, at least, one sequential equilibrium. Hence, the result follows from (i). \square

Consider the game in Figure 2. The strategy $b = ((D, d), D)$ is part of a weak perfect Bayesian equilibrium; it suffices to take the beliefs such that $\mu(x) = 1$; but it is not part of any simple perfect Bayesian equilibrium since it is not a weak perfect Bayesian equilibrium in the subgame that begins at $w_{1,2}$. Moreover, the unique simple perfect Bayesian equilibrium is $((U, d), U)$, which leads to the outcome $(2, 1)$. Take now the game in Figure 3. The strategy $b = (D, D)$ is a subgame perfect equilibrium but, since choice D of player 2 is strictly dominated in the information set Γ_{w_2} , it is not a simple perfect Bayesian equilibrium. Now, the unique simple perfect Bayesian equilibrium is (U, U) , which leads again to the outcome $(2, 1)$.




Figure 2: A weak perfect Bayesian equilibrium that is not a simple perfect Bayesian equilibrium.




Figure 3: A subgame perfect equilibrium that is not a simple perfect Bayesian equilibrium.

3.1 An Illustrating example

In section we present an example to illustrate the notion of simple perfect Bayesian equilibrium (SPBE) and its connections with sequential equilibrium (SE) and weak perfect Bayesian equilibrium (WPBE). Consider the three-player multistage game with incomplete information presented Figure 4. Only player 3 has a non-degenerate type space. In particular, with probability $p \in (0, 1)$, player 3 has type I and, with probability $1 - p$, player 3 has type II. Once the type of player 3 is realized, but without knowing it, player 1 chooses whether to continue the game (C) or to stop (S). If C is played, player 2 chooses between U and D , knowing player 1's action but not knowing player 3's type. Then, player 3 is called to move if he has type I, being only uninformed about player 2's action.

We start by looking at SE. Clearly, for any completely mixed behavior strategy of player 1, conditional on information set w_2 being reached, x_3 would be reached with probability p and x_4 with probability $1 - p$. Therefore, for each behavior strategy b , the consistency requirement on assessments imposes a unique possible system of beliefs, μ_b , where $\mu_b(x_1) = \mu_b(x_3) = p$ and the beliefs at w_3 are pin down by the probabilities with which player 2 chooses U and D . We now move to SPBE. The admissible system of beliefs is not necessarily unique. If b prescribes player 1 to play C , then the system of beliefs is necessarily given by μ_b (all the information sets are on the path of b). On the contrary, if b prescribes player 1 to play S , then SPBE does not impose any restriction on the beliefs at w_2 , so the set of admissible beliefs is larger. In any case,




Figure 4: An example illustrating the relations between simple perfect Bayesian equilibrium and the classic belief based equilibrium concepts.

since w_2 is a regular information set, SPBE still requires that the beliefs at w_3 are obtained by applying Bayes rule from w_2 , given b .⁷ Finally, in the case of WPBE, the set of admissible beliefs is even larger. If b prescribes player 1 to play C , as in the case of SE and SPBE, then the admissible system of beliefs is necessarily μ_b . However, if b prescribes player 1 to play C , then WPBE neither imposes no restriction on the beliefs at w_2 and w_3 . This is consistent with the inclusion relations in Proposition 1: the more demanding equilibrium concepts are those that are less permissive with beliefs.

We now characterize the pure strategy equilibrium profiles according to SE, SPBE, and WPBE. First of all, regardless of the value of p , (S, D, d) will never be a Nash equilibrium, since player 1 always gains by deviating to C . Further, in any strategy profile in which player 1 plays C , all the information sets are on the path and, hence, beliefs are uniquely pin down by Bayes rule. Therefore, for those strategy profiles in which player 1 plays C , the requirements of SE, SPBE, and WPBE coincide. Also, note that both SE and SPBE impose that player 3 plays u whenever player 2 plays U and that he plays d whenever player 2 plays D (regardless of player 1's action). As we show below, this is not the case for WPBE. Then, under SE and SPBE, the only equilibrium candidates are $\bar{b} = (S, U, u)$, $\tilde{b} = (C, U, u)$, $\hat{b} = (C, D, d)$.

SE: Recall that in every sequential equilibrium (b, μ) , we have $\mu(x_3) = p$ and $\mu(x_4) = 1 - p$. If $p > 1/2$, the unique SE is \tilde{b} ; if $p \in (1/3, 1/2)$, the unique SE is \bar{b} ; and, if $p < 1/3$, the unique SE is \hat{b} (if $p = 1/2$, both \tilde{b} and \bar{b} are sequential equilibria and if $p = 1/3$, both \bar{b} and \hat{b} are sequential equilibria).

SPBE: Now, there is no restriction on the beliefs at w_2 when S is played. Then, for all $p \in (0, 1/2]$, \bar{b} is a SPBE supported by any beliefs such that $\mu(x_3) \geq 1/3$ (Bayes rule on the “subgame” starting at w_2 imposes that $\mu(x_5) = 1$); in some sense, under a SPBE

⁷An equilibrium concept in the spirit of perfect Bayesian equilibrium (Fudenberg and Tirole, 1991) would also pin down the beliefs μ_b . The reason is that, on top of the requirements of SPBE, it imposes, among others, a “no signaling what you do not know” condition for the formation of beliefs. When applied to information set w_2 , this condition just says that player 2 cannot infer anything about the type of player 3 from the fact that player 1 has played C (since player 1 does not have extra information about this type). Thus, this equilibrium concept would coincide with SE in this example.

player 2 is allowed to think that there is some correlation between the realized type of player 3 and the action of player 1, whereas SE does not allow for it. When $p \geq 1/2$, the SE \tilde{b} is also a SPBE. When $p \leq 1/3$, the SE \hat{b} is also a SPBE. It can be easily verified that there are no other SPBE.

WPBE: When player 1 plays S , there is complete freedom on the beliefs at w_2 and w_3 . Then, for all $p \in (0, 1)$, the set of pure WPBE is constituted by the SPBE plus two additional equilibria: $b^* = (S, U, d)$ and $b^{**} = (S, D, u)$. The equilibrium b^* is supported by $\mu(x_3) \geq 1/3$ and $\mu(x_5) \leq 2/3$, whereas the equilibrium b^{**} is supported by $\mu(x_3) \leq 1/3$ and $\mu(x_5) \geq 2/3$. However, it is clear that the beliefs at w_3 in these two equilibria are quite unreasonable.

Therefore, we see that SE strictly refines SPBE when $p < 1/3$ and that the two equilibrium concepts select the same strategy profiles when $p \geq 1/3$. On the other hand, regardless of the probability p , SPBE strictly refines WPBE. More importantly, SPBE removes precisely the most unnatural WPBE, *i.e.*, those in which player 3 is not even best replying at what player 2 is doing at w_2 .

3.2 Connection with extended subgame perfect equilibrium

When discussing the properties of sequential equilibrium, Kreps and Wilson (1982) define what they call *extended subgame perfect equilibrium* and show that sequential equilibrium is a refinement of it. Interestingly, extended subgame perfect equilibrium and simple perfect Bayesian equilibrium build upon the same idea, but the former goes one step further and is a refinement of the latter. This is because extended subgame perfection imposes restrictions also at information sets that are not regular. The role played by regular information sets in our definition is played by subforms in the definition of extended subgame perfect equilibrium. Roughly speaking, a subform is a collection of information sets that is closed under *succession*; in particular, given a regular information set w , Γ_w is a subform. Then, similarly to what we did for regular information sets, one can associate a game to each subform. Yet, since a subform can have multiple “initial” information sets, one has to specify the relative probabilities with which they are chosen. Hence, to check that a given strategy profile is an extended subgame perfect equilibrium, it does not suffice to provide a system of beliefs, but probability measures over the “initial” information sets of the different subforms have to be specified as well.

To illustrate the above discussion, suppose that we have the three-player game described in Figure 5 (since payoffs are irrelevant for the discussion, we omit them). The game only has two regular quasi-subtrees: Γ_{w_1} and Γ_{w_3} . Hence, to check whether a strategy b is a simple perfect Bayesian equilibrium, it suffices to find a system of beliefs μ such that the restrictions of (b, μ) to $G_{w_1}(\mu)$ and $G_{w_3}(\mu)$ are sequentially rational and, on the path of b , μ is derived from b by Bayes rule. On the contrary, the game has four subforms as defined by Kreps and Wilson (1982). Specifically, in addition to Γ_{w_1} and Γ_{w_3} , there are two subforms with multiple initial information sets: a subform initiated in information sets $w_{2,1}$ and $w_{2,2}$, and a subform initiated in information sets $w_{2,1}$ and w_3 . Now, checking that a strategy profile is an extended subgame perfect equilibrium also requires to check sequential rationality in the corresponding “subgames”. In order to do this, it is necessary to know the relative probabilities of the different information sets, so that we know how Nature moves at the start of the different “subgames”. Clearly, this goes beyond the scope of an assessment (b, μ) , since a system of beliefs just pins down the probabilities of the different nodes in each information set, but contains no information about the relative likelihood of the different information sets. Further, these relative probabilities must be derived from b by Bayes rule on the path of the restriction of b to each “subgame”.




Figure 5: Comparing simple perfect Bayesian and extended subgame perfect equilibria.

4 Simple Perfect Bayesian Equilibrium in General Games

In this section we show that it is not hard to extend simple perfect Bayesian equilibrium to games in which the players may have a continuum of actions at some (possibly all) of their information sets and in which the number of information sets may be countably infinite. As we will see, the main difficulty lies on the definition of the game itself, not on the extension of our equilibrium concept.

4.1 General extensive games

To define general extensive games we build upon the approach in Aumann (1964). Aumann's approach delivers one of the most general settings under which Kuhn's theorem has been proved and, therefore, it seems a natural starting point.⁸ We now present a concise description of the ingredients needed to define a general extensive game; the reader is referred to Aumann (1964) for a much more complete exposition.

Let I be the unit interval endowed with Borel's σ -algebra. The main assumption in Aumann's approach is to restrict attention to what he calls *standard m -spaces*. This reduces to consider only the following measurable sets:

- Sets that are finite or countable and endowed with the discrete structure (all subsets are measurable).
- Sets that are isomorphic with I .⁹

This restriction seems quite natural and, at the same time, should be enough to model most real-life applications. In words of Aumann, "most measurable spaces that one 'encounters in practice' are standard; for example, any Borel subset of any Euclidean space or of Hilbert

⁸Alós-Ferrer and Ritzberger (2008) allow for more general games and characterize the class of game trees for which all pure strategy combinations induce unique outcomes. In their own words, "the paper addresses the question of what it takes to obtain a well-defined extensive form game". In particular, the main complication there is to allow not only for continuous strategy sets, but also for continuous time. In Alós-Ferrer and Ritzberger (2007) the same authors aim to characterize the game trees under which a version of Kuhn's theorem holds.

⁹An isomorphism is a one-one correspondence that is measurable in both directions.

space is standard". Let Ω be the the m -space I endowed with the Lebesgue measure, λ . The randomizations of the players and Nature are defined as random variables from the fixed sample space Ω to the action sets. For the sake of exposition, we assume that Nature moves once and for all at the beginning of the game; the information functions defined below will be the ones controlling what and when do the players learn about Nature's moves. We are ready to define general extensive games.¹⁰

Definition 3. A *general extensive game* GG with set of players $N = \{1, \dots, n\}$ consists of:

- (i) A (finite or infinite) sequence A_1, A_2, \dots of standard m -spaces called *action spaces*.
- (ii) Another sequence of the same cardinality P_1, P_2, \dots where each $P_k \in N$ indicates the player who has to choose an action from A_k .
- (iii) Another sequence of the same cardinality W_1, W_2, \dots of standard m -spaces called *information spaces* and whose elements are referred to as *information sets*.
- (iv) The set Nature's moves, given by the standard m -space A_0 , and a measurable function $n_0 : \Omega \rightarrow A_0$ (which induces a distribution of probability over Nature's moves).
- (v) A sequence of functions

$$g_k : A_0 \times A_1 \times A_2 \times \dots \times A_{k-1} \rightarrow W_k,$$

called *information functions* which, for each choice in A_0 , are measurable functions from $A_1 \times A_2 \times \dots, A_{k-1}$ into W_k .

- (vi) A function $h : A_0 \times A_1 \times A_2 \times \dots \rightarrow \mathbb{R}^n$ called the *payoff function* which, for each choice in A_0 , is a measurable function from $A_1 \times A_2 \times \dots$ into \mathbb{R}^n .

Informally, a general extensive game unfolds as follows. First, an element of the sample space $w \in \Omega$ is realized, which leads to a move by Nature: $n_0(w) = a_0$. Then, player P_1 gets information $g_1(a_0) \in W_1$ and makes a choice $a_1 \in A_1$. Next, player P_2 gets information $g_2(a_0, a_1) \in W_2$ and makes a choice $a_2 \in A_2$ and so on. The payoff will then be obtained as a function of the sequence a_0, a_1, a_2, \dots . We have to impose some conditions to ensure that the game has perfect recall.

Definition 4. A general extensive GG game has *perfect recall* if, for each $k < l$ such that $P_k = P_l$, there are two measurable functions

$$\begin{aligned} \alpha_k^l : W_l \rightarrow A_k \quad \text{and} \quad \beta_k^l : W_l \rightarrow W_k \quad \text{such that} \\ \alpha_k^l(g_l(a_0, a_1, a_2, \dots, a_{l-1})) = a_k \quad \text{and} \\ \beta_k^l(g_l(a_0, a_1, a_2, \dots, a_{l-1})) = g_k(a_0, a_1, a_2, \dots, a_{k-1}). \end{aligned}$$

Intuitively, a player can remember what he has previously done via the α functions and what he has previously known via the β functions. Since Aumann (1964) showed that Kuhn's theorem holds for general extensive games with perfect recall, it is justified to focus on behavior strategies.

¹⁰Despite following the same idea, our definition looks fairly different from the one in Aumann (1964). Aumann presented his definition from "an individual player's viewpoint", in the sense that the game is retained in extensive form for one player and normalized for the other players, "projecting" their strategy sets into a single set. This representation was sufficient (and very convenient) to prove Kuhn's theorem in his general setting. In our case, we leave explicit the extensive form for all players.

4.2 Behavior strategies and assessments in general extensive games

Definition 5. Let GG be a general extensive game.

- A *mixed strategy profile* in GG is a sequence $\mathbf{m} = (m_1, m_2 \dots)$ of measurable functions $m_k : \Omega \times W_k \rightarrow A_k$, where Ω is the (fixed) sample space. For each k and each information set $w_k \in W_k$, the random variables $m_k(\cdot, w_k)$ and n_0 are independent. Also, randomizations made by different players have to be independent, *i.e.*, for $k \neq l$ such that $P_l \neq P_k$, $m_k(\cdot, w_k)$ and $m_l(\cdot, w_l)$ are mutually independent random variables ($w_k \in W_k$ and $w_l \in W_l$ being arbitrary).
- A *behavior strategy profile* in GG is a mixed strategy profile \mathbf{b} in which also the randomizations made by a player in his different information sets are independent from each other, *i.e.*, for $k \neq l$, $b_k(\cdot, w_k)$ and $b_l(\cdot, w_l)$ are mutually independent random variables ($w_k \in W_k$ and $w_l \in W_l$ being arbitrary).

Definition 6. A *system of beliefs* in a general extensive game GG is given by a sequence $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots)$ such that, for each k and each $w_k \in W_k$, μ_k is a distribution of probability over Ω .

Then, an assessment will be a pair $(\mathbf{b}, \boldsymbol{\mu})$. Note that, combined with \mathbf{b} , μ_k induces a distribution of probability over A_0, \dots, A_{k-1} , *i.e.*, a distribution of probability about what may have happened previously in the game. As in the simple case of finite extensive games, we want the system of beliefs $\boldsymbol{\mu}$ to be obtained via “Bayes rule” along the “path” of \mathbf{b} . However, in general extensive games, it may be possible that some (possibly all) information sets of a given information space are reached with probability zero. Hence, Bayes rule will not be an appropriate way to update beliefs.¹¹ Instead, one has to use conditional probabilities. Conditional probabilities were already extensively used in Aumann (1964) and its need in connection with the use of Bayes rule for updating beliefs is discussed in Jung (2010).

Since each behavior strategy \mathbf{b} jointly with an element of the sample space $\omega \in \Omega$ uniquely pins down the realized actions a_0, a_1, \dots we can abuse notation and use $g_{\mathbf{b},k}(\omega)$ to represent $g_k(a_0, \dots, a_{k-1})$. Recall that λ represents the Lebesgue measure.

Definition 7. Given an assessment $(\mathbf{b}, \boldsymbol{\mu})$, we say that $\boldsymbol{\mu}$ is obtained by *conditional updating* from \mathbf{b} if, for each k , each $B \subset W_k$, and each $\Lambda \subset \Omega$,

$$\int_B \mu_k(\Lambda |_{g_{\mathbf{b},k}(\omega)=w_k}) d\lambda(g_{\mathbf{b},k}^{-1}(w_k)) = \lambda(g_{\mathbf{b},k}^{-1}(B) \cap \Lambda).$$

The above definition is simply saying that $\mu_k(\cdot |_{g_{\mathbf{b},k}(\omega)=w_k})$ has to be a version of the conditional probability $\text{cond } \lambda(\cdot |_{g_{\mathbf{b},k}(\omega)=w_k})$.¹² Therefore, if two systems of beliefs $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ are obtained by conditional updating from \mathbf{b} , they must agree on the information sets that are on the path of \mathbf{b} (even if the probability of reaching them is zero). When restricting attention to finite games, conditional updating is equivalent to the use of Bayes rule on the equilibrium path.

¹¹To illustrate, just think of a situation in which player 1 has to pick $x \in [-1, 1]$. Player 2 is then informed about $|x|$. Suppose player 1 strategy is to pick a number uniformly from $[0, 1]$ with probability $2/3$ and from $[-1, 0]$ with probability $1/3$. Then, if player 2 receives information 0.7 , he should deduce that $x = 0.7$ with probability $2/3$ and $x = -0.7$ with probability $1/3$. However, since the probability of $|x| = 0.7$ is zero, the above deduction cannot be made via Bayes rule: $P(x = 0.7 |_{|x|=0.7}) = P(x = 0.7 \cap |x| = 0.7) / P(|x| = 0.7) = 0/0$.

¹²The fact that any such conditional probability is a well defined probability measure crucially relies on the use of standard spaces.

4.3 Extending simple perfect Bayesian equilibrium

So far we have defined the ingredients we need to even think about defining a belief based equilibrium concept for general extensive games. The definition of sequential rationality of an assessment would now be straightforward: a player must be best replying at each information set given his beliefs. Now we can easily extend the definition of weak perfect Bayesian equilibrium to general extensive games.

Definition 8. Let GG be a general extensive game. An assessment $(\mathbf{b}, \boldsymbol{\mu})$ is a *weak perfect Bayesian equilibrium* if it is sequentially rational and $\boldsymbol{\mu}$ is obtained by *conditional updating* from \mathbf{b} .

In order to present the extension of simple perfect Bayesian equilibrium we first need to extend the notion of *regular information set*. Recall the idea of regular information set in finite extensive games: if an information set v has a node that comes after a regular information set u , then all the nodes in v come after u . Despite the fact that our definition of general extensive game hides the tree structure, the notion of regularity of an information set is very easy to generalize. First, given $k < l$, we say that an information set $w_l \in W_l$ comes after $w_k \in W_k$ if there is a realization of actions a_0, a_1, \dots such that $g_k(a_0, a_1, \dots, a_{k-1}) = w_k$ and $g_l(a_0, a_1, \dots, a_{l-1}) = w_l$, *i.e.*, according to the realized actions, the information w_l is given to player P_l after information w_k is given to player P_k .

Definition 9. Let GG be a general extensive game. An information set $w_k \in W_k$ is *regular* if, for each information set w_l such that w_l comes after w_k , then, for every realization a_0, a_1, \dots such that $g_l(a_0, a_1, \dots, a_{l-1}) = w_l$, we have $g_k(a_0, a_1, \dots, a_{k-1}) = w_k$, *i.e.*, we cannot reach information set w_l without reaching first information set w_k .

Now, as we did for finite games, given a general extensive game GG , an assessment $(\mathbf{b}, \boldsymbol{\mu})$ and a regular information set w_k we can easily define a general extensive game $GG_{w_k}(\mathbf{b}, \boldsymbol{\mu})$ as follows:

- (i) The sequence A_k, A_{k+1}, \dots
- (ii) The sequence P_k, P_{k+1}, \dots
- (iii) The sequence W_k, W_{k+1}, \dots
- (iv) Nature's moves are given by $A_0^{w_k} = A_0 \times \dots \times A_{k-1}$. The corresponding distribution of probability is the one induced by \mathbf{b} and $\boldsymbol{\mu}_k$.
- (v) The information functions remain unchanged: $g_l : A_0^{w_k} \times A_k \times A_{k+1} \times \dots \times A_{l-1} \rightarrow W_l$.
- (vi) The payoff function also remains unchanged: $h : A_0^{w_k} \times A_k \times A_{k+1} \dots \rightarrow \mathbb{R}^n$.

Definition 10. Let GG be a general extensive game. An assessment $(\mathbf{b}, \boldsymbol{\mu})$ is a *simple perfect Bayesian equilibrium* if, for each regular information set w_k , the restriction of $(\mathbf{b}, \boldsymbol{\mu})$ to $GG_{w_k}(\mathbf{b}, \boldsymbol{\mu})$ is sequentially rational and $\boldsymbol{\mu}$ is obtained by *conditional updating* from \mathbf{b} , *i.e.*, the restriction of $(\mathbf{b}, \boldsymbol{\mu})$ to $GG_{w_k}(\mathbf{b}, \boldsymbol{\mu})$ is a weak perfect Bayesian equilibrium.

As we have seen, despite of the complexity of the general setting under consideration, the notion of simple perfect Bayesian equilibrium can be easily extended. Although at this stage it would be desirable to provide some existence result for simple perfect Bayesian equilibrium of general extensive games, this is a far from trivial issue. Indeed, the literature has already struggled to find existence results for the weaker notion of subgame perfect equilibrium in

games with a continuum of actions,¹³ even after imposing more structure than the one we have here (such as perfect and almost perfect information, for instance). We refer the reader to Harris (1985), Harris et al. (1995, 2005), and Luttmer and Mariotti (2003) for references on the topic.

5 Conclusions

It is easy to find extensive games in which simple perfect Bayesian equilibria selects unreasonable strategy profiles. One such example might be constructed from Figure 1. For an assessment (b, μ) to be a simple perfect Bayesian equilibrium of a game starting with the game tree in Figure 1, conditional on w_3 being reached, there is complete freedom in the beliefs, *i.e.*, it is not necessary that $\mu(x_1) = 0$. Nonetheless, we consider that this equilibrium concept can be good enough in different situations and, because of its simplicity, it is much easier to deal with than sequential equilibrium. Indeed, also weak perfect Bayesian equilibrium is sometimes good enough and, because of this, it is often used in the literature.

Concerning the way in which equilibrium concepts are presented in most specialized books, there is a positive feature of simple perfect Bayesian equilibrium. Namely, it allows to restore the inclusion relation for the equilibrium concepts. Most books start by defining Nash equilibrium and subgame perfection comes immediately afterwards. Then, since subgame perfection does not perform well in imperfect information games, they introduce the systems of beliefs and, with them, weak perfect Bayesian equilibrium as a first step towards sequential equilibrium. Yet, we consider that it would be more pedagogic to use the concept of simple perfect Bayesian equilibrium instead of weak perfect Bayesian equilibrium.

References

- ALÓS-FERRER, C. AND K. RITZBERGER (2007): “Discrete Extensive Form Games,” Mimeo.
- (2008): “Trees and extensive forms,” *Journal of Economic Theory*, 143, 216–250.
- AUMANN, R. J. (1964): “Mixed and Behavior Strategies in Infinite Extensive Games,” in *Advances in Game Theory*, ed. by M. Dresher, L. S. Shapley, and A. Tucker, Princeton University Press, vol. 52 of *Annals of Mathematical Studies*, 443–476.
- BATTIGALLI, P. (1996): “Strategic Independence and Perfect Bayesian Equilibria,” *Journal of Economic Theory*, 70, 201–234.
- FUDENBERG, D. AND J. TIROLE (1991): “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53, 236–260.
- HARRIS, C. (1985): “Existence and Characterization of Perfect Equilibrium in Games of Perfect Information,” *Econometrica*, 53, 613–628.
- HARRIS, C., P. RENY, AND A. ROBSON (1995): “The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: A Case for Public Randomization,” *Econometrica*, 63, 507–544.

¹³In our context, a subgame is a game $GG_{w_k}(\mathbf{b}, \boldsymbol{\mu})$ where the regular information set w_k is such that the player is told exactly what has happened before in the game (μ_k pins down the realized history of play). Then, requiring Nash equilibrium in all these special regular information sets would clearly be less demanding than requiring weak perfect Bayesian equilibrium in all the regular information sets.

- HARRIS, C. J., M. B. STINCHCOMBE, AND W. R. ZAME (2005): “Nearly compact and continuous normal form games: characterizations and equilibrium existence,” *Games and Economic Behavior*, 50, 208–224.
- HARRIS, M. AND R. TOWNSEND (1981): “Resource Allocation under Asymmetric Information,” *Econometrica*, 49, 33–64.
- JUNG, H. M. (2010): “Perfect Regular Equilibrium,” Mimeo.
- KREPS, D. M. AND R. WILSON (1982): “Sequential Equilibria,” *Econometrica*, 50, 863–894.
- LUTTMER, E. G. T. AND T. MARIOTTI (2003): “The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: A Comment,” *Econometrica*, 71, 1909–1911.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*, Oxford University Press.
- MYERSON, R. B. (1991): *Game Theory: Analysis of Conflict*, Harvard University Press.
- SELTEN, R. (1975): “Reexamination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory*, 4, 25–55.