Games and Economic Behavior 95 (2016) 178–181

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Reviews and Comments

With the intent of stimulating discussion, this section is reserved for book reviews, comments, and letters; your input is welcome. By nature, this material may be subjective, reflecting the opinions of the authors; your responses are therefore encouraged.

The Language of Game Theory: Putting Epistemics into the Mathematics of Games, by Adam Brandenburger.

When I received the request to write a review on Adam Brandenburger's book, we were just in the middle of giving a course on epistemic game theory at Maastricht University – a course that heavily builds on Adam Brandenburger's work in this area. Is this a coincidence? The timing may have been, but the latter part certainly not. Adam Brandenburger will always stand as one of the big pioneers and ambassadors of epistemic game theory, and his contributions to this field have been so fundamental that *any* course in epistemic game theory – now and in the future – will naturally be inspired by many of his works.

The book is a collection of some of his most important works in epistemic game theory, done in collaboration with his co-authors Robert Aumann, Lawrence Blume, Eddie Dekel, Amanda Friedenberg and Jerome Keisler. Many of these works are milestones in the development of the field, and for that reason have become real "classics" in the literature on epistemic game theory. This made my life as a reviewer relatively easy, as I knew each of the papers very well – as do most of the researchers in epistemic game theory I guess. For those people this book will be a wonderful bundling of works that have probably inspired much of their research, and shaped much of their views on the field. However, for students and researchers that are relatively new to the field, and want to learn about epistemic game theory, this book provides an excellent overview of some of the key ideas in the area.

Roughly speaking, the papers in the book can be organized around three themes: *language, common belief in rationality*, and *Nash equilibrium*. In what follows I will briefly discuss each of these themes, together with the contributions of Adam Brandenburger and his collaborators to that particular topic.

Language

Epistemic game theory describes, and investigates, what players in a game believe about the opponents before making a choice themselves. These beliefs do not only concern the opponents' possible choices, but also the beliefs that the opponents may hold about the choices of other players, the beliefs that the opponents may hold about the beliefs of other players about the choices of their opponents, and so on. Such constructs are called *belief hierarchies*, and in a sense they constitute the *language* of epistemic game theory.

But before we can make precise statements about belief hierarchies, we first need a *belief model* in which such belief hierarchies can be expressed. Most belief models used nowadays are inspired by the pioneering works of Harsanyi (1967–1968), Kripke (1963) and Aumann (1976). An interesting feature of these models is that they are *self-referring* – a type for a player in Harsanyi's model holds a belief about the opponents' types in this model, and at every state in the Kripke–Aumann model, every player holds a belief about the set of states in the same model. So, players within the model reason about the model itself.

We know from Russell's paradox that self-references may create problems in some specific settings – such as naive set theory – and hence a natural question is whether the self-referring belief models above lead to problems as well. This question is addressed by Adam Brandenburger and Jerome Keisler in Chapter 1. More specifically, Chapter 1 asks whether every possible belief sentence in the language can be expressed inside the belief model. The answer is "no" if beliefs are modeled by *possibility sets* – which are sets of states in the Kripke–Aumann model – instead of probability distributions over states. Hence, for every possible belief model that is based on possibility sets, there will be at least one belief sentence in the language that is not expressible within that model.

This result can be seen as an epistemic analogue to Russell's paradox, and is now known as the Brandenburger–Keisler paradox. The key message of this impossibility result is that there are limits to what we can express by a belief model that is based on possibility sets. This result is of fundamental importance for epistemic game theory, as it puts boundaries to its language.

http://dx.doi.org/10.1016/j.geb.2015.12.006

Chapter 2, on the other hand, shows that things look very different if beliefs are modeled by *probability distributions* rather than possibility sets. This is precisely the approach taken by Harsanyi (1967–1968). Adam Brandenburger and Eddie Dekel show that under certain topological conditions, we can construct a Harsanyi model with types in which every "well-behaved" belief hierarchy can be expressed. So we get a *possibility result* now, in contrast with the impossibility result in Chapter 1. A message to be learned from these two chapters is that it crucially matters, for the expressiveness of the belief model, which kind of beliefs we use.

In game theory it has long been argued that weakly dominated choices should not be made, since there is another choice (or a probability distribution over choices) which is always at least as good, and sometimes better. But in order to rule out weakly dominated choices on formal decision-theoretic grounds we must somehow assume that the player considers *each* of the opponents' choices possible. This may be in conflict, however, with the assumption that the opponents will not make weakly dominated choices either. Indeed, if a player believes that his opponents do not make weakly dominated choices, then this typically means that he must deem some opponents' choices impossible – contradicting our assumption above. This "inclusion–exclusion paradox" has been a real puzzle in game theory.

The deeper reason for this puzzle is that the traditional language for beliefs – in which beliefs are either represented by possibility sets or probability distributions – is not expressive enough to resolve this paradox. In the probabilistic setting, for instance, an opponent's choice is either deemed possible (receives positive probability) or impossible (receives probability zero) – there is nothing in between. And it is exactly this restriction that creates the puzzle.

In Chapter 6, Lawrence Blume, Adam Brandenburger and Eddie Dekel offer a beautiful resolution to this paradox by enriching the language of beliefs. Instead of modeling the belief of a player by a *single* probability distribution, they model it by a *finite sequence* of probability distributions, and call this a *lexicographic probability system*.

If a player must decide between two possible choices, he first uses the first probability distribution to decide which choice is better. However, if both choices are equally good under the first level, then he will turn to the second probability distribution to decide which choice is better, and so on. The interpretation of these different levels in the lexicographic probability system is that all opponent's choices that receive positive probability in the first level are deemed *infinitely more likely* than all opponent's choices that receive positive probability in some further level, but not in the first. However – and that is crucial – the opponent's choices that receive positive probability in some further level are still deemed *possible*. Similarly, all opponent's choices that receive positive probability in the second level are deemed infinitely more likely than all choices that receive positive probability in some further level, but not in the first or second, and so on. Hence, a player with a lexicographic probability system may deem an opponent's choice a infinitely more likely than another choice b, while still deeming b possible.

Let us now go back to the "inclusion–exclusion paradox". With the richer language of a lexicographic probability system at hand, we can now model a player who (a) deems all opponent's choices that are not weakly dominated infinitely more likely than all opponent's choices that are weakly dominated, and (b) still deems all opponent's choices (including the weakly dominated ones) possible. This resolves the paradox.

To me, the invention of lexicographic probability systems has been one of the most important milestones in the development of epistemic game theory. With this model at hand we could suddenly model states of mind in which a player deems some opponent's choices infinitely more likely than others, while not completely discarding any opponent's choice from consideration. I will come back to this issue later on.

Common belief in rationality

The concept of *common belief in rationality,* which states that all players choose rationally, all players believe that all players choose rationally, all players believe that all players believe that all players choose rationally, and so on, is certainly the most fundamental notion studied in epistemic game theory. Most other concepts in the area can be seen as refinements, or variations, of this idea.

A natural question is: What choices can a rational player make if he expresses common belief in rationality? Chapters 3, 4, 7 and 8 give different answers to this question. The reason is that the various chapters differ in the type of belief being used, and in the restrictions on these beliefs.

In Chapter 3, Adam Brandenburger and Eddie Dekel model the belief of a player by a single probability distribution. In the first part of the chapter, they investigate the case in which any kind of correlation in the beliefs is allowed – an issue which is only important if we have three players or more. Here, by correlation we mean that i's belief about j's choice may be correlated with i's belief about k's choice. They show that in this setting, the choices that can rationally be made under common belief in rationality are exactly the choices that survive *iterated elimination of strictly dominated choices*. The same result can be found in Tan and Werlang (1988). In the introduction of the book, Adam Brandenburger calls this the Fundamental Theorem of Epistemic Game Theory – for a good reason. Indeed, this result serves as a starting point for almost every survey or course in epistemic game theory, and is even discussed in some books and courses in classical game theory.

The second part of Chapter 3, and Chapter 4, investigate what happens to this result if we impose additional restrictions on the beliefs. In the second part of Chapter 3 *no* correlation is allowed, whereas in Chapter 4, Adam Brandenburger and Amanda Friedenberg study the case in which only *intrinsic* correlation is allowed. That is, correlation between the beliefs about the various opponents' choices must entirely be explained by correlation between the beliefs about the various opponents' beliefs. This resembles the common assumption in game theory that players choose *independently*. Note, however, that this "independent choice" assumption still allows for correlated beliefs, as long as the correlation is intrinsic.

It turns out that in each of the three cases mentioned above, the set of choices that can rationally be made under common belief in rationality may be different. So, a key lesson to be learned from Chapters 3 and 4 is that for the Fundamental Theorem of Epistemic Game Theory it not only matters whether we allow for correlation or not, but also which type of correlation we allow for.

Chapters 7 and 8 no longer model the beliefs by single probability distributions, but rather by lexicographic probability systems. The reason is that these chapters wish to model players that are *cautious* – that is, deem each of the opponents' choices possible. We have already seen above that lexicographic probability systems suit this purpose perfectly. In this new setting, the concept of common belief in rationality is replaced by an analogous notion – *common assumption of rationality*.

The key idea in this concept is *assumption of the opponents' rationality*, which roughly states that a player deems all "potentially optimal" choices for the opponent infinitely more likely than all choices that are "necessarily suboptimal". Here, "potentially optimal" and "necessarily suboptimal" are to be taken relative to the set of types that are present in the model. By appropriately iterating this condition we arrive at *common assumption of rationality*.

In Chapter 7, Adam Brandenburger, Amanda Friedenberg and Jerome Keisler show that the choices that can rationally be made under common assumption of rationality always constitute a *self-admissible set*. This result may thus be seen as a possible analogue to the Fundamental Theorem of Epistemic Game Theory in the world of lexicographic probability systems. It also shows that epistemic game theory is not only capable of formulating epistemic foundations for *existing* concepts, such as iterated elimination of strictly dominated choices, but also of generating *new* concepts, such as self-admissible sets. In Chapter 8, Adam Brandenburger and Amanda Friedenberg expand on the new notion of self-admissible sets and show what this concept implies for the behavior of players in some well-known games of interest.

The last part of Chapter 7 explores what happens if we choose a *complete* belief model, which contains *all* possible belief hierarchies. Two interesting things happen: First, there will in general be no type in the model that expresses *common assumption of rationality* – so we get another impossibility result here. On the other hand, if we only iterate the condition of *assumption of the opponents' rationality* k times – instead of infinitely often – then the choices that result are precisely the choices that survive the first $k+1$ rounds of iterated elimination of weakly dominated choices. So, in a sense, *iterated assumption of rationality* may be viewed as a foundation for the well-known procedure of iterated elimination of weakly dominated choices, provided we work with a *complete* belief model.

The latter result has by now become one of the classics in epistemic game theory. Part of the reason is that iterated elimination of weakly dominated choices has a long tradition in game theory, and for many years it remained unclear what type of reasoning by the players would lead to this procedure. Chapter 7 provides a brilliant answer to that question.

Nash equilibrium

For many decades, Nash equilibrium has been the central concept in game theory. Much of classical game theory is based on this idea, and its various refinements. An important and natural question is therefore: What conditions must be imposed on the reasoning of players, such that their beliefs – or choices – constitute a Nash equilibrium? Somewhat surprisingly perhaps, it took game theory a very long time before it could answer this question. In fact, Robert Aumann and Adam Brandenburger were the first, in 1995, to systematically address this question, and to provide a satisfactory answer. This paper appears as Chapter 5 in the book.

So, why did game theory have to wait until 1995 for an answer? The main reason is that before the rise of epistemic game theory, we simply did not have the appropriate language to formally define and investigate this question. And we know from experience that without a formal language to speak about certain issues, false statements are around the corner.

An important message that Robert Aumann and Adam Brandenburger reveal in Chapter 5 is that it takes more than just common belief in rationality to arrive at a Nash equilibrium. What we also need is that players believe that their opponents are *correct* about the beliefs they hold. In games with two players, this "correct beliefs assumption" – together with common belief in rationality – indeed leads to beliefs that constitute a Nash equilibrium. However – and that is another important lesson to learn from that chapter – we do not need *all* the layers in common belief in rationality once we impose the correct beliefs assumption: In a two-player game, imposing *mutual* belief in rationality, together with the correct beliefs assumption, is already enough to arrive at a Nash equilibrium. For games with more than two players, additional conditions are needed to obtain a Nash equilibrium.

These sufficient conditions for Nash equilibrium constitute yet another classic in epistemic game theory, and have inspired a whole body of literature that investigates the foundations of Nash equilibrium. The Aumann–Brandenburger conditions also enable us to have a transparent discussion about how reasonable a concept Nash equilibrium is, since the epistemic conditions for Nash equilibrium are now on the table. One of the conditions that seems problematic is the correct beliefs assumption. Indeed, why should a player necessarily believe that his opponents are correct about the beliefs he has? After all, his opponents cannot look inside his mind. Such a discussion would not have been possible without the fundamental work by Robert Aumann and Adam Brandenburger.

The lines above hopefully illustrate the enormous impact that Adam Brandenburger has had upon the development of epistemic game theory. The book – as a collection of some of his most important papers – constitutes a wonderful testimony of his many contributions to this blooming area. For the many people who already know his papers I can still recommend the beautiful introduction by Adam Brandenburger himself, written in the clear and transparent style which is so characteristic of both his writing and his research.

References

Aumann, R.J., 1976. Agreeing to disagree. Ann. Stat. 4, 1236–1239.
Harsanyi, J.C., 1967–1968. Games with incomplete information played by “Bayesian” players, I–III. Manage. Sci. 14, 159–182. 320–334, 486–502.
Kripke, S., 1963. A semantical analysis of modal logic I: normal modal propositional calculi. Z. Math. Log. Grundl. Math. 9, 67–96.
Tan, T., Werlang, S.R.C., 1988. The Bayesian foundations of solution concepts of games. J. Econ. Theory 45, 370–391.

Andrés Perea

7 July 2014
Available online 31 December 2015