



Individuating games

Michael Ridge¹

Received: 2 February 2019 / Accepted: 22 February 2020 / Published online: 3 March 2020
© The Author(s) 2020

Abstract

Games, which philosophers commonly invoke as models for diverse phenomena, are plausibly understood in terms of rules and goals, but this gives rise to two puzzles. The first concerns the identity of a single game over time. Intuitively one and the same game can undergo a change in rules, as when the rules of chess were modified so that a pawn could be moved two squares forward on its first move. Yet if games are individuated in terms of their constitutive rules and goals, this is incoherent—new rules mean a new game. The second concerns the individuation of games at a point in time. Intuitively, there can be different versions of a single game, where the versions differ in the details of their rules. I offer a solution to this problem that draws on an analogy with individuating languages. The resulting theory should illuminate the metaphysics of games more generally.

Keywords Game · Language · Meaning · Rule-following

Philosophers investigating wide variety of topics have found games to provide a useful analogy. Most famously, Wittgenstein argued that natural languages are usefully understood by way of analogy with games (Wittgenstein 1958).¹ Analogies with games are regularly invoked in normative ethics (Cf. Rawls 1955), meta-ethics (Cf. Katsafanas 2013; Enoch 2011) and applied ethics (Langton and West 1999), debates over particularism (McKeever and Ridge 2006), not to mention the philosophy of law (e.g. Hart 2012), speech-act theory (Boer 1974; Williamson 1996) and aesthetics (Walton 1990). The ubiquity of such analogies is aptly expressed in an anecdote from Mary Midgley: “Some time ago, an Innocent Bystander, after glancing through a copy of *Mind*, asked me, ‘Why do philosophers talk so much about Games? Do they play them a lot or something?’ (Midgley 1974: p. 321) Why *do* philosophers talk so much about games? Any monolithic answer to this question will oversimplify, but certain features stand out. The fact that games plausibly are (at least partly) *constituted* by certain rules and

¹ Perhaps most notably, informing the work of such heavyweights as David Lewis and Wilfrid Sellars. Cf. Lewis (1979) and Sellars (1954), though they are hardly alone in invoking the analogy.

✉ Michael Ridge
mridge@staffmail.ed.ac.uk

¹ University of Edinburgh, Edinburgh, Scotland, UK

goals is especially germane, since this provides an analogy with the role that rules and goals play in a variety of broadly “normative” domains.

Yet it is precisely this putative feature of games—their constitution by rules and goals—which gives rise to two puzzles. The first puzzle concerns the identity of a single game over time. Intuitively, one and the same game can undergo a change in rules. For example the rules of chess were modified so that a pawn could be moved two squares forward on its first move but it remained the same game (chess), albeit in a new form—a new “variant” of chess as it is usually put. Yet if games are individuated in terms of their constitutive rules and goals, this is incoherent—new rules trivially entail a new game, not a new “version of” or “variant on” the same game. The second puzzle concerns the individuation of games at a point in time. Intuitively, there can be different versions of a single game, where the versions differ in the details of their rules. For example, the rules of regulation five-a-side football differ from the rules of regulation FIFA football in several ways (https://en.wikipedia.org/wiki/Five-a-side_football). Yet it is commonplace to regard these as versions of the same game in at least one sense. To be clear, the puzzle is not that there is an intelligible and sometimes useful more fine-grained sense in which these are all different games; that is fine. Rather, what is puzzling is that there also seems to be an equally useful and pre-theoretically intuitive, but theoretically elusive, ordinary language sense of ‘game’ in which they are versions of the same game.

I begin by laying out the two puzzles (Sect. 1). I argue that these puzzles are strikingly analogous to certain problems facing orthodox accounts of the individuation of natural languages (Sect. 2). I then argue that one promising strategy for solving the problem in the case of natural language is also promising in the case of games (Sect. 3). In effect, I invert the more common methodology of learning more about language by reflecting on games, and try to learn more about games by reflecting on language. In particular, I argue that in one important sense of ‘game’ we should understand that term as denoting a historically embedded social practice. Games in the sense of abstract entities (rules/goal pairs) and games qua token contests bear important logical relations to games in this sense. I then discuss cases in which the changes to a game qua social practice mean we have an entirely new game qua social practice, though I lack the space to explore this important issue in full here (Sect. 4). Having defended this approach, I discuss some of the affinities of the account on offer with Fred D’Agostino’s work on the “ethos” of games and compare my approach with its main rival (Sect. 5).

1 Two puzzles about games

The most influential definition of ‘play a game’ comes from Bernard Suits, in his classic, *The Grasshopper*. Suits argues that ‘play a game’ should be defined in terms of four individually necessary and jointly sufficient conditions: (1) Engage in activity aimed at bringing about a specific state of affairs, (2) using only means permitted by rules, (3) where those rules prohibit more efficient means, and (4) where those rules are accepted just because they make the activity in question possible (Suits 2005: pp. 48–49). It is not hard to see the idea that games are constituted by certain kinds

of interrelated rules and goals as implicit in this definition.² Suits does not explicitly say very much about how to individuate games. The most obvious extension of his theory would be to hold that games are individuated by their rules and goals. Other theorists have explicitly endorsed a view of the individuation of games in this vicinity. Here, for example, is Gordon Reddiford: “We can say that a system of constitutive rules uniquely defines a particular game.” (Reddiford 1985: p. 42).

Suppose for the sake of argument that games *are* individuated in terms of their rules and goals. This quickly gives rise to two puzzles, which I shall call “the Diachronic Puzzle” and “the Synchronic Puzzle.” The Diachronic Puzzle arises out of the fact that in ordinary language we happily talk about how the rules of a single game have changed over time. I already gave the example of the way in which the rules of chess were changed to allow the pawn to move two squares on its first move, but this is hardly an isolated example. Indeed, in the case of chess alone there are many other examples of changes in the official rules of the game. For example, in the 1980’s the “50 move rule,” which holds that if 50 moves are made without any pawn moves or piece captures, the game is a draw, was modified in light of computer analysis which demonstrated that certain endgames could be won but only if the 50 move rule was violated. For a brief period, there were exceptional endgames were covered by a “75 move rule” but eventually in the interests of simplicity the 50 move was reinstated globally.³ Throughout the debate over this rule, all parties took themselves to be arguing about what the rules of chess, the game they had been playing all along, should be—they did not take themselves to be arguing about whether to quit playing chess and play some very similar but different game. Anyone familiar with the history of various popular games will know that countless such examples can be found across a wide spectrum of games.

Here, then, is the Diachronic Puzzle, which is just the obvious incompatibility of individuating games in terms of their constitutive rules and goals with the idea that games may persist through changes in their rules:

- (1) Games are individuated in terms of their constitutive rules and goals, such that different rules or different goals entail a different game.
- (2) Games can evolve over time, so that the rules or the goal of a given game can change, leading to a new version of the same game.

Absent an equivocation on ‘same game’ or some other key term, (1) and (2) entail a contradiction.

The Synchronic Puzzle arises out of reflection on the way in which there can at a single point in time be a variety of different versions of the same game. Above I

² Suits’ focus on the phrase ‘play a game’ raises another issue: what is the semantic relationship between ‘game’ and ‘play a game’? My own view is that ‘play a game’ gets its meaning compositionally—from the meaning of its parts and the way in which they are combined. This hypothesis is simple and fits with a broadly compositional view of semantics more generally. There is little evidence that ‘play a game’ is idiomatic, and much evidence that it is not. For example, it seems valid to infer from e.g. the claim that Jones was playing a game to conclusions that use ‘game’ but not ‘play’, (e.g. ‘there is at least one game’). A complication is that both ‘play’ and ‘game’ are multiply ambiguous. I do not have space to defend this broader commitment here, but I hope this is enough to motivate the view and suggest why I think theorists like Suits, whose explicit focus is on the phrase ‘play a game’ plausibly take on implicit commitments about the meaning of both ‘play’ and ‘game’.

³ https://en.wikipedia.org/wiki/Fifty-move_rule#History.

gave the example of different versions of football, but of course there are countless examples of this as well. Formal tournament games of chess include a “touch move” rule, according to which touching a piece with the intention of moving it means you must move that piece if it can legally be moved. Informal games of chess often do omit this rule. The official rules of Monopoly do not include any rules about paying certain fines into a slush fund in the middle of the board, such that whenever someone lands on the “Free Parking” square they collect it; however, in practice this is a common riff on the official rules. NFL football and touch football in the backyard are very different versions of American football, and so on. This gives rise to the Synchronic Puzzle:

- (1) Games are individuated in terms of their constitutive rules and goals, such that different rules or different goals entail a different game.
- (2) One and the same game can, at a given point in time, come in different versions, such that these versions vary in their rules or goals.

Once again, it is easy to see how, absent equivocation, these two these entail a contradiction.

A somewhat “bullet-biting” response to these puzzles would be to deny (2) in each case—that is deny that, strictly speaking, games can evolve or come in different versions at one and the same time. One might argue that we can accommodate ordinary talk of sameness of games in these cases in terms of people “speaking loosely,” whereby games that are strictly speaking distinct but sufficiently similar are treated as if they are instances of the same game. How similar two games must be for it to be sensible to treat them as instances of the same game, loosely speaking, will then be a pragmatic matter and will vary from context to context. By way of analogy, a critic might complain that Stephen King “writes the same book every year.” Obviously even the critic should admit that this is not strictly speaking true—*Carrie* and *The Shining* are not literally the same book. Still, the critic’s meaning will be relatively plain.

Although perhaps initially tempting, this reply is inadequate. First, it seems alien to the semantic intentions of ordinary speakers. Unlike the Stephen King example, ordinary speakers who say that the rules of chess changed when they introduced the “en passant” rule do not take themselves to be speaking loosely. They would typically insist that it is *strictly speaking true* that both before and after the rule change people were playing chess. Second, this is not a superficial feature of ordinary thought and talk. When people debate a proposed rule change to a beloved game, it is not at all uncommon for partisans to appeal to “the spirit of the game” and “the game’s rich tradition” as reasons for or against various proposed rule changes. Of course, this talk can be reinterpreted as loose talk, when what the speakers really intend is to be arguing that some new but very similar game would be better or worse than some older game. This seems ad hoc, and in the case of appeals to “tradition” it seems especially strained.

Both these objections have some force, but a third objection is in my view decisive. We can make sense of “speaking loosely” only against the backdrop of a plausible account of the strict and literal truth. What, then, are the strict and literal truth-conditions for “Smith and Jones played chess”? Is this true if they followed the rules of chess as encoded by FIDE (the international governing body of chess) in 2017? Or if they followed the official FIDE rules at some earlier time? Or if they followed the

rules of chess as understood before FIDE even existed? Choosing any of these (or any of countless other) options seems arbitrary, but given the fineness of grain of ‘same game’ we seem forced to choose one or give up on a determinate referent for ‘chess’. The fine-grained view of the individuation of games seems to entail massive semantic indeterminacy for natural language words, like ‘chess’, which are intuitively taken to be the names of games. This, though, is simply incredible.

Of course, one could bite this bullet too, and allow that strictly speaking most of the games we play do not have names, and that what we take to be the names of games actually name families of similar games. Although internally consistent, this move is desperate. Moreover, it does not avoid the need for further theorizing. For we now will need an account of what it means for two games to belong to the “same family” such that this can provide plausible truth conditions for words like ‘chess’ and ‘tag’. This is not progress so much as moving the bump in the rug.

Rather than tackle these two puzzles directly, I begin by discussing how a strikingly similar set of puzzles arise in the case of natural language. This sets the stage for my proposed solution in the case of games.

2 Two puzzles about natural language

The Diachronic Puzzle and the Synchronic Puzzle about games arose out of an otherwise tempting and orthodox conception of games. The conception might fairly be called “Platonic” because it effectively understands games as abstract entities defined formally by interrelated rules and goals.⁴ Intriguingly, an otherwise tempting and orthodox approach to natural languages is similarly Platonic, and leads to two strikingly analogous puzzles. A rich tradition in both linguistics and the philosophy of language identifies languages with such abstract objects as lists of expressions, grammatical rules, functions from contexts of utterance to semantic contents, and the like. The view was given a canonical defence by David Lewis (Lewis 1975). Donald Davidson was among the countless theorists heavily influenced by this conception, and he gives a nice statement of it:

A language may be viewed as a complex abstract object, defined by giving a finite list of expressions (words), rules for constructing meaningful concatenations of expressions (sentences), and a semantic interpretation of the meaningful expressions based on the semantic features of individual words. (Davidson 1992: p. 255)

It is easy to see how this conception of language leads very quickly to analogues of the Diachronic Puzzle and the Synchronic Puzzle. Once we define languages as abstract objects of the sort sketched here by Davidson, the individuation of languages just is a special case of the individuation of abstract objects consisting of words, rules, etc. Plausibly such abstract objects are individuated in terms of their constituents and their structure. In that case, any difference in constituents (e.g. adding a word) or structure (adding or revising a semantic or grammatical rule) entails a different abstract

⁴ Cf. D’Agostino (1981) who also refers to such views as “Platonistic” (D’Agostino 1981: 9). I discuss the relation of the positive view I develop in Sect. 3 to D’Agostino’s view in Sect. 4.

object. This trivially entails that, strictly speaking, languages can never evolve. What seems like the evolution of English is, in fact, just a matter of speakers switching from one language to another. This does not cohere with our ordinary criteria for sameness of languages, though. It is a commonplace that languages evolve over time. Once again we have a version of the Diachronic Puzzle. The same form of reasoning means that when two people seem to be speaking slightly different idiolects of a given natural language, like English, they are speaking two numerically distinct but very similar languages. This also does not cohere with our ordinary criteria for sameness of languages. We happily speak of different speakers of the same language using slightly different idiolects. Once again, we have a version of the Synchronic Puzzle.

Of course, one can once “bite the bullet” and deny that there is *any* useful conception of sameness of languages which can vindicate ordinary criteria of individuation. Moreover, one might partially vindicate the ordinary criteria by distinguishing what is true strictly speaking from what might be apt to say when speaking “loosely.” Indeed, in the case of languages, several prominent theorists at least flirt with this idea, although it is not always made entirely explicit. The rejection of any alternative sense of ‘same language’ which could vindicate ordinary thought and discourse when construed in a strict and literal manner may draw some plausibility from the idea that any clear line between a mere idiolect and a language will inevitably be arbitrary and historically has been driven more by politics than anything else. This idea is nicely captured by the quip, “A language is a dialect with an army and a navy.”⁵ Davidson, in the paper quoted above, is admirably clear about the upshot of his own conception of language and he does not suggest that another one is ready to hand: “...if we are precise about what constitutes a language, it is probably the case that no two people actually do speak the same language.” (Davidson 1992: p. 260) David Lewis thought literal sameness of language in actual practice is a “limiting case” which is never actually reached (Lewis 1969: p. 201). Noam Chomsky echoed these sentiments, remarking that, “no two people actually do speak the same language.” (Chomsky 1988: pp. 36–37)

Just as with the analogous bullet-biting response in the case of games, this is not very satisfying. Indeed, the passages just quoted seem deliberately provocative; they are so jarring to common sense that they wear their paradoxicality on their sleeve. Perhaps at the end of the day we must conclude that there is no coherent notion of language which can vindicate our ordinary criteria for individuating languages, but *ceteris paribus* we should prefer one which does. The same problems face the “loosely speaking” account here as well as in the case of games. Ordinary speakers do not take themselves to be merely speaking loosely when they discuss the evolution of a given natural language or characterize two people as speaking the same language. Moreover, they sometimes appeal to precedent and tradition to resist certain changes to their language in just the way that people invoke precedent and tradition in the face of proposed changes to the rules of a game; the “spirit of the game” perhaps never finds an

⁵ The quip is commonly attributed to the sociolinguist Max Weinreich, who himself claims he got the line from a comment from someone at one of his talks. https://en.wikipedia.org/wiki/A_language_is_a_dialect_with_an_army_and_navy.

analogue in the case of language, though even that is probably eminently debatable. Finally, the view entails massive semantic indeterminacy for our natural language names for the languages people speak. Characterizing any one of the many languages so-called “English speakers” speak as “really” being uniquely English would indeed be entirely arbitrary. Implausibly, sentences like “Churchill spoke English” will be semantically indeterminate.

Moreover, as Henry Jackman has argued in Jackman (1999), there is a further problem arising here in virtue of the plausibility of so-called “semantic externalism.” This is the problem that may undermine even theorist’s use of the more fine-grained conception of languages as abstract objects. This problem has no obvious analogue in the case of games, so the case against the Platonic conception of language seems even stronger in the case of language than in the case of games. Semantic externalism holds that the content of a word can depend on factors “outside the head” as it is often put. Semantic externalism comes in many forms, but on one of the more plausible conceptions the content of certain words—most notably natural kind words like ‘water’, ‘electron’, ‘negative charge’—is fixed in part by the underlying nature of the entity or property in question. Competent speakers may, then, have no ability to discriminate between the relevant kinds nor reliably pick out instances of them. How, then, do their words manage to have meanings rich enough to encode such distinctions? The most orthodox answer is given by Hilary Putnam—they do so through a “division of linguistic labour” (Putnam 1975). Putnam famously gave the example of the competence of ordinary speakers with ‘beech’ and ‘elm’ in spite of not being able to discriminate beeches from elms. Their words still encode the right distinction in virtue of their deference to experts who can in principle draw the relevant distinction. The importance of this sort of linguistic division of labour was usefully elaborated upon by Tyler Burge. Burge gave the example of someone who didn’t realize that *arthritis* was a disease of the joints, and who believed he had arthritis in his thigh (Burge 1979). Burge argues that if the wider linguistic community had used ‘arthritis’ to refer to a disease which *could* occur in one’s thigh (call it ‘tharthritis’) then Bert would have had a true belief. This would presumably depend on the semantic intentions of others (especially perhaps the relevant experts) in one’s linguistic community. Jackman argues that this whole picture depends crucially on the idea that speakers *share* a language, strictly speaking. He discusses several strategies for accommodating semantic externalism while still insisting, with Davidson, Lewis, Chomsky et al., that no two people ever really share a language. I cannot do his discussion of those replies justice here and anyway I would just be reinventing the wheel if I did. The sceptical reader should consult Jackman’s excellent discussion. I shall here take this general line of argument as providing yet another reason to try to make some sense of the idea that a natural language, like English, can literally be shared by many people despite variations in dialect, etc., and that natural languages so conceived can evolve over time.

3 Solving the puzzles

Jackman offers an attractive solution to the linguistic analogues of the two puzzles about games with which I began.⁶ He points out that it is very easy to end up thinking no two people could share a language if one begins with languages conceived as abstract *types* and then explains what it is for an individual to speak a language that is an instance of that type. Two people then share a language if they are related in the right way to the same abstract universal type. On this conception, sharing a language is like sharing a haircut. We can understand what it is for each of us to have a haircut (or speak a language) in isolation from anyone else. To see whether our haircuts (or languages) are the same, we see whether they are tokens of the same type.

This approach can seem very natural, if not inevitable, given a conception like David Lewis's. That conception, and the approach it drives, are not mandatory though. Instead of understanding sharing a language in terms of being appropriately related to the same abstract *type*, we can instead understand it in terms of sharing a single *token* language. The key to making this change in perspective is to think of token languages as *particulars*—more specifically, historically embedded social practices in which people cooperate in certain distinctive ways to communicate with one another. Here is Jackman:

I can be understood as sharing a language with other members of my community because I look for their help when deciding how my terms are correctly applied, and they do the same with me. It is in this cooperative sense of sharing that I can share a language with my fellow speakers in more than just a 'loose' sense...speakers are understood as speaking the same language not in virtue of speaking two type identical languages, but rather in virtue of both sharing a single (*token* identical) language...The languages we share are not abstract universals (as one might expect if they were 'functions from strings of sounds or of marks to sets of possible worlds') but rather *particular* practices. (Jackman 1999: p. 304)

Jackman draws a useful analogy between languages as historically embedded practices and biological species. He suggests that an animal which is molecularly indistinguishable from, but evolutionarily unrelated to, the animals we call 'tigers' would not be a tiger. In the same way, someone who spoke a language that instantiated the same abstract type as English, but which was historically unrelated to English, would not be an English speaker—though in another sense, of course (the one invoked by Lewis et al.) they do share a language. (Jackman 1999: p. 304; compare Richard 2019)

Much more could be said about Jackman's approach in the case of natural language. For example, one interesting question Jackman does not discuss is when there has been enough change over time that we should say the language spoken now is numerically distinct from some distant ancestor language even though they are all part of the

⁶ Although I focus on Jackman's approach in the text, which I find makes the analogy with games especially clear, the same basic move from abstract types to concrete socially embodied tokens can be found in many other discussions, perhaps most notably Kaplan (1990). Because this is not primarily a paper on the history of this strand of the philosophy of language, I shall not here canvass all of the ways in which this idea has arisen there. Thanks to an anonymous referee for drawing me out on this.

same historically embedded social practice. The analogy with evolutionary biology is not so helpful here, since in that context at least one standard criteria for sameness of species within a given lineage is whether reproduction between members of the two generations is possible.⁷ Clearly this approach has no application in the case of natural language. One *prima facie* plausible extension of the general approach would be that whether we have a new language within given lineage is whether current speakers generally take the practices of the earlier generation as broadly relevant to what counts as correct usage—where no clear shift in meaning has taken place, of course. However, I cannot investigate this and various other important accounts about Jackman’s approach in the case of natural language here; my primary focus is on the individuation of games, for which language simply provides a useful analogy.

The obvious parallel is to understand games as historically embedded social practices. We can then understand the evolution of games as an instance of the way in which social practices evolve over time.⁸ Similarly, we can understand different versions of a game at a given point in time as an instance of the way in which a rich social practice can contain a variety of sub-practices within it.⁹ When I play chess, I am in one sense playing the same game that François-André Danican Philidor played in the 1700’s, even though he played quite a different version of that game. In the version Philidor played, a pawn could be promoted only to a piece that one’s opponent had already captured, whereas the version I play allows me to promote to any piece (other than a king or pawn) regardless of what has already been captured. I am playing the same game as Philidor because I am taking part in the same historically embedded social practice. Again, my aim here is to articulate a sense of ‘game’ in which games can evolve and have different variants. I do not intend to argue that there is not also a legitimate but more fine-grained sense in which they are individuated entirely in terms of their constitutive rules and/or goals.

If this is right, then ‘game’ is context-sensitive. In some contexts, ‘game’ refers to a historically embedded social practice. In other contexts, ‘game’ refers to an individual token contest which is constituted by certain goals and rules. Although normally a game in the latter sense will be part of a game in the former sense, this is not essential. There could be a game in the latter sense which is not part of a historically embedded social practice. That is as it should be. It should be possible for me to invent a little game to pass the time while I wait to see the dentist but never share this game with anyone else. It is still a game in one useful sense even though it is neither identical to nor part of a game-constituting social practice. Still, there is an important and non-accidental connection between the two senses of ‘game’, since most token games are, in fact, part of larger social practices, and hence such games are parts of games in

⁷ This is not the only criteria for sameness of species. For a usefully thorough if somewhat dated discussion of some of the many ways of understanding biological species and their individuation, see Mayr 1988.

⁸ It would be interesting to see how Wittgenstein’s notion of a “form of life” might map onto the notion of a social practice in the present context, though a proper investigation of that issue would take me too far afield. See also Gert 1995b. Thanks to an anonymous referee for noting this connection.

⁹ For a similar view of games grounded more in an anthropologist’s perspective, see Malaby 2007. Malaby there argues that games are “grounded in (and constituted by) human practices and are therefore always in the process of becoming. This also means they are not reducible to their rules.” (Malaby 2007: p. 103) See also Consalvo 2009.

the first of my two senses. So, although we have a kind of ambiguity here, it is not a *brute* ambiguity. Plausibly, the two meanings of ‘game’ provide us an instance of what linguists call ‘polysemy’. Although the two meanings are distinct they are related in an important and interesting way. This ambiguity allows us to embrace both premises of the Diachronic Puzzle and the Synchronic Puzzle, but only if we read ‘game’ in *different senses* between the two premises. Given this equivocation, the two premises do not give rise to a contradiction, and the puzzles are therefore resolved.

Just to clarify the proposed solution, return to the canonical statement of the two puzzles, but now armed with this distinction. Let ‘game_{TC}’ refer to games qua token contests, and ‘game_{SP}’ refer to games qua social practices. To a first approximation, then we can then resolve the two puzzles as follows:

Diachronic Puzzle, Resolved:

- (1) Games_{TC} are individuated in terms of their constitutive rules and goals, such that different rules or different goals entail a different Games_{TC}.
- (2) Games_{SP} can evolve over time, so that the rules or the goal of a given game_{SP} can change, leading to a new version of the same game_{SP}.

Synchronic Puzzle, Resolved:

- (1) Games_{TC} are individuated in terms of their constitutive rules and goals, such that different rules or different goals entail a different Games_{TC}.
- (2) One and the same game_{SP} can, at a given point in time, come in different versions, such that these versions vary in their rules or goals.

This is only a first approximation, though, because I shall below draw a *threefold* distinction.

An instructive partial analogy is provided by ‘law’ in English. When we speak of ‘the law’ of a given country, we are referring to a historically embedded social practice. When we speak of ‘a law’, however, we are referring to a specific element of that practice which may hold only during some proper subpart of the practice’s temporal extension. The analogy with ‘game’ is not exact, since (as I just argued) you can have a game that exists outside of any game qua social practice, but you cannot have a law outside of the law qua social practice. Apart from this important difference, and one more I shall note below, the analogy holds up remarkably well. In particular, the use of the demonstrative article ‘the’ in both cases tends to be used to pick out the social practice. It is *the* game of chess which evolved over time, but it is *a* game of chess that I played at the club last night. To take an actual example, in the headline, “Arab Contributions to the game of Chess,” it is clear enough that ‘*the* game of chess’ refers to the historically embedded social practice.¹⁰ This form of words is often used to describe contributions to games qua social practices. Obviously, this depends on context, and you can use ‘the game’ to refer to a token contest rather than the social practice, as with ‘The game Kasparov played yesterday was brilliant’. When the context does not supply reference to some particular token contest, though, it is natural to take ‘the game’ to refer to the social practice.

¹⁰ <http://www.arabamerica.com/arab-contributions-game-chess/>.

However, in another respect the analogy between ‘game’ and ‘law’ is imperfect, but in an instructive way.¹¹ Consider the meaning of ‘game’ again. I suggested above that ‘game’ admits of a two-fold ambiguity, as between (a) historically extended social practices, and (b) token contests between particular players at a given time. In fact, we need at least a three-fold ambiguity because something very close to the fine-grained conception of games as constituted by rules and goals which gave rise to our paradox has at least some currency in ordinary language. Consider again the example of chess and the changes to its rules since the days of Philidor. Although Philidor and Kasparov both played chess, and so in that sense played the same game, it would also be natural to say e.g. that “Philidor was playing an importantly different game from the one played by Kasparov.” We can plausibly understand ‘game’ here as referring to the sort of abstract entity constituted by rules and goals that is suggested by the work of formalists like Suits. Players engage with these games qua abstract entities, though, only through social practices which bear the right relation to those abstract entities. When two people play a game of chess, their game_{TC} (qua token contest) is a proper part of the game_{SP} of chess (qua social practice). Their game_{TC} (qua token contest) is governed by certain rules R because those rules partly constitute the game_{AE} (qua abstract entity) they are playing, where that game_{AE} bears the right relations *both* to their activity (their playing of game_{TC}) and to the game_{SP} . In particular, they are playing that game_{AE} in virtue of engaging with the game_{SP} in the right way in a context in which that practice sanctions their following rules R .

Different theories will offer different accounts of what it is to engage with the social practice “in the right way,” with some theorists insisting that the players follow the rules and pursue the goal while intending to take part in the practice by doing so, and other theorists allowing for a less demanding relationship to those rules and goals, perhaps just undertaking a commitment to follow the rules and pursue the goal, “showing up” and doing enough in the way of taking part in the practice upon “showing up” (see below). I assume here that abstract entities can be indeterminate in various respects, and that this will allow the proposed account to accommodate the obvious fact that games as played by flesh and blood human beings often are indeterminate. We therefore need (at least) a threefold distinction in the case of ‘game’. Note that in one sense ‘game’ (game_{TC}) denotes the *activity* of playing a game, while in another sense ‘game’ denotes the *object* of that activity (game_{AE}), and in yet another sense ‘game’ refers to the broader practice which enable players to play it (game_{SP}). The close relation between these senses of ‘game’ suggest that it is polysemous rather than brutally ambiguous. Return again to our puzzles then:

Diachronic Puzzle, Resolved:

- (1) Game_{TC} and game_{AE} are individuated in terms of their constitutive rules and goals, such that different rules or different goals entail a different game_{TC} or a different game_{AE} .
- (2) Game_{SP} can evolve over time, so that the rules or the goal of the game_{AE} which are admissible forms of the game_{SP} can change over time, so that the game_{SP} comes to include new versions, where those versions can be understood as game_{AE} which bear the right relation to the game_{SP} at the relevant time.

¹¹ Thanks to an anonymous referee for drawing me out on this point.

Synchronic Puzzle, Resolved:

- (1) Games_{TC} and games_{AE} are individuated in terms of their constitutive rules and goals, such that different rules or different goals entail a different game_{TC} or a different game_{AE}.
- (2) One and the same game_{SP} can, at a given point in time, come in different versions, where those versions can be understood as games_{AE} which bear the right relation to the game_{SP} at the relevant time, such that these versions vary in their rules or goals.

In the case of ‘law’ we get something analogous to ‘game_{SP}’, namely “the law” as a broad social practice governing a jurisdiction over an extended period of time. We also get something analogous to ‘game_{AE}’, in that we can understand the law of the land at any given point in time as it applies to a given individual as an abstract entity constituted by rules (and associated sanctions), where that abstract entity maps onto “the law” qua social practice in a way analogous to the way in which the rules which constitute a game_{AE} are related to “the game_{SP}.” However, there is nothing especially analogous in the law to games_{TC} (games qua token contests), as “the game played by Kasparov and Karpov in game three of their second match.” Of course, we do have token legal contests in the sense that cases are argued in court before a judge or jury. However, those contests they are not relevantly analogous to token game contests. In legal contests, victory is constituted by convincing the relevant authority what the law entails in the case at hand. Given my analogy this would be like winning a game of chess by arguing more persuasively to an arbiter about some fine point in the application of the rules, which obviously isn’t the right analogy if we want the rules of chess to function analogously to the relevant laws.

Of course, even if it has the right structure to solve the two puzzles with which I opened, this general framework raises a lot of further questions that I cannot answer in detail here. Most obviously, not all social practices are games, so we need some account of what makes a given social practice constitute a game. We also need an account of what makes some token activity count as a game. To answer these questions, and show how the two answers provided cohere well together, is virtually to develop a fully free-standing philosophy of games, though, and I do not have the space to develop such a full-blown theory here. I will, though sketch one way in which such a broader theory might go just to provide some sense of how this theory of game individuation could be combined with a plausible broader philosophy of games.

One option would be to begin with a theory like Suits’ theory now understood as a theory of what it is to play games, where the *playing* of a game is itself a game_{TC} and the *game that is played* itself is a game_{AE}. We could then understand games qua social practices as social practices which originate with some the playing of some game_{AE} in Suits’ sense. Individuals take part in that practice by learning the rules of the game from others who are part of the practice. Participants in the practice could then suggest revisions to the rules or goal of the game. The merits of any proposed change in the rules can be debated by the participants who will give reasons for and against a given proposed change. Participants within the practice will characteristically appeal to what they take to be valuable about the game as it stands. Tradition, precedent and the “spirit” of the game will also serve as reasons, and so on. In addition to changes

to some implicitly canonical conception of the game, participants can also propose variants on the original game as an alternative to the original version of the game. Such variants need not be replacements of the older version(s) of the game, but supplements.

Insofar as Suits' theory independently provides a plausible account of games qua token contests, it would provide a nice foundation for a view which also posits games qua social practices in this way. In fact, I do not myself think Suits' theory is defensible—at least not as a fully general theory of playing games. In some respects, his theory does well enough for single-player non-competitive games, but it struggles with the arguably more paradigmatic competitive games. Suits' theory trivially entails that if you are not following the rules then you are not yet playing the game, and that if you are not trying to achieve the constitutive goal of the game then you are not yet playing the game. This generates several *prima facie* counter-examples. I already mentioned (in section one) the case of the novice player who we very naturally characterize as learning the rules *by* playing the game with others. This description is incoherent given Suits' definition. Suits' theory also entails that insofar as you cheat you are not playing the game, but we intuitively do not think that someone who cheats throughout a game has never played the game (think of someone secretly consulting the internet while playing Trivial Pursuit, say). Nor does common sense endorse the idea that someone cannot play a game without trying to win. Parents often play games with their children without making any serious effort to win, and someone who throws a boxing match or a chess game has still played the game. In formal contexts, this is even clearer. I could hardly complain that my rating should not go down after losing a tournament game of chess because I was not trying to win. If, though, Suits' theory was correct then my not trying to win entailed I did not play a game, and your rating only goes down if you played a game—and you cannot lose a game you never played! Even more clearly you can't win a game you never played, but someone playing a game of chess might win through sheer luck even if he was not trying to win at all—just making sure his moves were legal but otherwise choosing a move at random. Finally, Suits' theory seems overbroad. On Suits' account, parliamentary debate, seeking a profit in your business in a less efficient way instead of cheating someone out of a sense of moral duty, giving testimony in court and waging a just war all arguably come out as instance of playing games. In each case, an agent is pursuing some end (winning an argument, making money, finding the truth in a criminal case, defeating an enemy) in a less efficient way in virtue of accepting rules which prohibit more efficient means (the rules of parliamentary debate, the moral law, the rules governing the giving of testimony in court, and the rules of just war) where one accepts these rules just in order to engage in the activity in question (to engage in parliamentary debate, to be moral, to pursue criminal justice through legally sanctioned means).

Clearly much more would need to be said to make these objections stick. I mention them just to motivate another approach and illustrate how it too fits well with the framework I have developed here to solve the two puzzles with which I began. First, we can avoid the problem of overbreadth by “going functional.” The rough idea is that there is a range of functions we associate with the concept of playing a game. These include having fun, providing an opportunity for accomplishments (think of how we praise successful athletes), socializing in certain distinctive ways and other functions

as well.¹² Intuitively, these are what games as such are *for*: I take an etiological view of what functions are—either the intentions of those who invented the game, or the explanation of why the it proliferated over time, can fix its function. The analogy is with the way in which a maker’s intentions fixes the function of a chair as for sitting, while the evolutionary history of hearts explains why the function of a heart is to pump blood. With games, the relevant evolutionary history will be cultural rather than biological. Ideally, one would like a more precise account of what the relevant function or functions is (or are). My views here are not settled, but I am inclined to take a pluralist view on which a range of functions of the sort sketched above are all relevant, and having “enough” of these functions entails that a practice constitutes a game. In any event, a functional approach in this general vicinity can avoid the problem of overbreadth facing Suits’ theory. Parliamentary debate, giving testimony in a court of law and not cheating someone in business out of moral duty all lack the right functional profile to count as games.

Dealing with the possibility of playing games without following its constitutive rules or seeking its constitutive aim is more complex. Here I think the right approach is that in at least one useful sense of ‘play a game’, to play a game is to undertake a *commitment*, where this is a sort of speech-act, to follow the relevant rules and pursue the relevant goal and then to “show up” at the designated site of play so that one is thereby bound by the rules—and perhaps makes at least a few “moves” in the game, though how much one has to do on this front to count as playing the game can vary and depend on the details of the specific game in question.¹³ To be clear the relevant commitments are often tacit—just turning up at the ping pong table, raising a paddle and giving a knowing look to someone else can be a way of undertaking the commitment to follow the rules of ping pong in the right setting, e.g. One can do all of this without following the rules or seeking the goal, as with the novice who learns the rules as he plays, the cheater, and the parent who does not try to win the game (the “trifler” in Suits’ lingo). Insofar as one undertakes the relevant commitment and “shows up” and makes some moves (even illegal moves) one is a “player” and hence is playing the game. In the same way one can chair a meeting by being designated as the chair, showing up and invoking some of your rights as a chair even if you do not pursue the stated aims of the meeting, violate the rules governing the meeting, etc. One is the chair and one is thus chairing the meeting in one sense even if in another sense one isn’t “being much of a chair” because of one’s neglect of the duties of the station. The same distinction can be used in the proposed commitment based approach to ‘play a game’ to explain how in one sense, contra Suits, one can play a game without following its rules or pursuing its constitutive goal.

¹² For useful discussion of the role of having fun in game play as well as the nature of fun itself, see Bogost (2016). An interesting complication is that particular games can have additional functions, e.g. many sports have athletic fitness as one of their functions. The point here would be that not all functions of a given game play a role in explaining why it counts as a game. The recent enthusiasm for “gamification” is also relevant here, since the whole point of gamification is to motivate people to do things which are independently useful; see e.g. Burke (2014), Toppo (2015) and McGonigal (2011).

¹³ Compare Kreider (2011). I discuss the approach, which is in many ways very similar to Kreider’s, and explain why the theory I favour improves on Kreider’s in Ridge (2019).

A further complication is that this commitment based approach is plausible only for *multi-player games*. The idea that one must undertake a commitment qua speech-act, even a tacit one, to play a single-player game is simply not credible. However, we sometimes use ‘undertake a commitment’ in a sense which need not involve communicative acts at all. Here I have in mind the sense invoked when we say someone has undertaken a commitment *simply by setting themselves a goal*. Perhaps we should say that when one plays a single-player game that this consists, in part, in undertaking a commitment in this sense—a commitment qua intention. A commitment qua intention to do what, though? Plausibly, a commitment to achieve some goal which is intuitively the point of the game, but using only means allowed by the rules of the game. In the case of traditional forms of Solitaire, this would be a matter of producing four rows of cards in ascending order from Ace (as low) to King, where one attempts to achieve this goal only using means allowed by the rules of Solitaire. However, simply undertaking this commitment is intuitively not sufficient to play the game. One must also act on the commitment in pursuit of the relevant goal. The attentive reader will have noticed that I have effectively reinvented Suits’ definition, but only for single-player games. The only element of Suits’ definition I have not incorporated here is that the game player must accept the rules “just because” doing so makes the resulting sort of activity possible. I have left this out because on any natural interpretation this element of his theory is problematic—one can play a game without doing so for its own sake, as the jaded professional footballer illustrates. Admittedly, my approach involves more ambiguity than Suits’ theory, and *ceteris paribus* that gives us some reason to prefer Suits’ more elegant theory. However, *ceteris* is not *paribus*, since Suits’ theory is not well suited to provide a plausible conception of the individuation of games (among other problems). We should aim for theories that are only as complex as is necessary plausibly to explain the relevant data, but reject theories that are more elegant but unable to explain the data.

We can therefore improve on Suits’ account of the meaning of ‘game’ qua token activities by first “going functional” to avoid overbreadth, second distinguishing multi-player games from single-player games, and third understanding the former in terms of commitments qua speech-acts and the latter in terms of commitments qua intentions. Suppose we accept all of this. All that now remains is to show how this conception of games qua token activities can be combined with the proposed account of games qua historically embedded social practices. This is quite straightforward, though. We begin with a token game in one of the two senses just articulated. This activity then “catches on” and becomes a social practice into which new players are inducted by being taught the rules, goals, and perhaps the history of the game. As before, we can then understand how the game can evolve over time and how different versions of the game (which can be understood as abstract entities related to the social practice in the right way) can be played at the same time within the broader parameters of the relevant social practice. Moreover, the move to social practices fits especially well with the idea that we should “go functional,” at least insofar as functions are understood in terms of the cultural history of the game. Social practices are just the sort of thing which can fix functions over time in this way.

Although I have only discussed two ways of starting to fill in the details of my proposed solution of the two puzzles with which I began (Suits’ theory and my own),

the ease with which these approaches could be seamlessly combined with a conception of games as social practices is reassuring. It suggests, but does not prove, that any otherwise plausible conception of games qua token activities will be amenable to incorporation into a framework of games qua social practices. The basic approach to the puzzles here is *structural*. Any account of games qua token activities which adverts in some way to rules and goals (as any plausible theory will) can be integrated with this structure. This underscores a further advantage of the move away from a purely Platonic conception and to understanding games, at least in part, in terms of social practices.

An important recent trend in computer gaming is the increasing popularity of so-called “sandbox games.” *Minecraft* (at least in “creative mode”) is a good example of such a game. Sandbox games provide players with a sort of framework in which they can do various things—in *Minecraft*, for example, one can collect resources and build things in certain ways. While *Minecraft* and other sandbox games (*The Sims* is another good example) have rules in a certain sense—using “cheat codes” breaks the rules—they do not as such seem to have any constitutive *goals*. This is problematic on a purely Platonic conception of games, like Suits’, since having a constitutive goal is a necessary condition for something’s being a game. We could, of course, “bite the bullet” and say that *Minecraft* and many other so-called “sandbox games” are not, in fact, games. This, though seems like a significant cost.¹⁴

The approach developed here provides a very nice way of handling sandbox games. Sandbox games like *Minecraft* are, in and of themselves, simply games qua social practices. This is plausible on my account because such games plausibly have the right sort of *functional* profile. However, unlike more familiar games with associated social practices, *Minecraft* did not begin with any canonical goal (like checkmate). Rather, *Minecraft* has always been whatever the individual player makes of it. *Minecraft* and other sandbox games are thus games qua social practices which have, from the very beginning, admitted of an indefinitely wide plurality of versions, each of which has a different constitutive goal. This inverts the more normal order of explanation, in which games qua token activities with constitutive rules and goals come first, and only then does a practice evolve around those rules and goals. With sandbox games, we have a kind of social practice which grows up around a computer platform which is widely shared and discussed online, but which never privileged any particular goal as constitutive of the game. Of course, any *token* game of *Minecraft* will have some goal

¹⁴ Suits might argue that *Minecraft* is not a game, but rather a vehicle for what we might call play “full-stop.” A toddler splashing around in the tub or a child frolicking on a hill are playing, but they aren’t playing a game—or playing the trombone or a joke on someone, for that matter. They are just plain playing—hence my (not Suits’) terminology of “play full-stop.” *Minecraft* on this account is not a sandbox game, but it is rather like a literal sandbox. Sandboxes are not games, but they are vehicles for a certain sort of play full-stop—playing with sand; likewise *Minecraft* is not, on this construal a game, but a vehicle for a certain sort of play full-stop—playing now with certain kinds of pixels (or whatever), rather than with sand. Suits distinguishes play in this sense from playing a game in Suits (1988). I lack the space here for a full and proper reply to this interesting suggestion. It does seem evident to me, though that at least *ceteris paribus*, a theory which can vindicate the idea that so-called “sandbox games” do count as games in *some* useful sense is better than a theory which cannot. After all, it is completely standard in ordinary language to refer to these as games, so any reasonable principle of charitable interpretation would first try to make sense of why they are games before attributing a widespread, indeed systematic, semantic error to ordinary speakers in this way. Thanks to an anonymous referee for encouraging me to note this possible reply from Suits.

or other, as selected by the player, if only implicitly. Qua social practice, though, no such goal is privileged in this way. Thus, the social practice approach allows us to make sense of why sandbox games are properly thought of as games at all, whereas purely Platonic approaches like Suits will struggle to make sense of this data point. Quite independently of the need to solve the two puzzles, this is an advantage of proposed theory.¹⁵

4 Changes in social practices

While changes to a game-constituting practice over time can result in new versions of the same game, in other cases it seems more apt to describe such changes as generating new games qua social practices—new games_{SP}. To this extent, the theory on offer should have something to say about when a game-constituting practice undergoes changes which entail that those playing the game it constitutes after those changes are no longer playing the game it constituted before those changes. When, for example, did American Football stop being the same as Rugby, or when did Rugby stop being the same as British Football? In effect, in resolving the puzzles with which I began, my focus has been on the individuation of games qua token contests. I have, that is, tried to explain how two token contests can count, at least in one sense, as instances of the same game even though they may differ in their constitutive rules or goal. The idea was to appeal to games qua social practices and understand token contests as proper parts of these, such that being a part of the same game-constituting social practice is enough for two token contests to count as instances of the same game. That account of sameness of games qua token contests leaves open how we should understand sameness of games qua social practices.

An obvious consequence of the proposed theory is that we have different games qua social practices if the two games evolved in total isolation from one another. This is the analogue of Jackman’s point about Martian “tigers” not really being tigers and Martian “French” not literally being French. In the case of games, if a game with the same rules and goal as chess evolved on Mars it would not be the same games_{SP} as terrestrial chess. That is simply a consequence of the view. It may seem somewhat counterintuitive, but once we remember that we can still say that the Martians are playing the same game qua abstract entity—the same game_{AE} as many earthlings (the ones playing with identical rules/goal pairs)—it may not seem so counter-intuitive.

The harder problem is the diachronic one for a given social practice which evolves over time. We saw an analogue of this problem for Jackman’s account of the individuation of languages. At some point Latin evolved into French, and separately into Italian, Spanish, Portuguese and Romanian, for that matter. At what point in that history should we say that Latin ends and French begins? Recall that in that context I suggested that we might go for a rough analogue of a fairly orthodox view in the case of biological species. In particular, the criterion in the case of language could be whether current

¹⁵ The same sort of story may be plausible for some so-called “party games” like *Spin the Bottle*, which also doesn’t seem to have a constitutive aim as such, but which individual players will play with their own personal goals in mind. A proper discussion of these kinds of games would in this context be too much of a tangent from the main ideas of the paper, though.

speakers generally take the practices of the earlier generation as broadly relevant to what counts as correct usage—at least, where no clear shift in meaning has taken place. An analogue of this in the case of games would be whether those taking part in a later stage of the social practice take themselves to be bound, *ceteris paribus*, to respect precedents and values (implicit or otherwise) associated with an earlier stage of that practice. In the case of Rugby and British Football, this seems to sort the cases properly. Rugby enthusiasts presumably would not care about the practices or values implicit in British Football when deciding, e.g., how to resolve some internal dispute. By contrast, contemporary chess players would, I suspect, generally at least think it relevant, though not decisive, when thinking about the 50 move rule today and whether to retain it, that Ruy Lopez introduced a version of the 50 move rule in his classic manual on the game in the 1500's. In this and many other ways, contemporary chess players would think the rich history of the game matters to how the game should be played, the values it should manifest and embody and the ways in which it deserves to be valorised even today. This remains true even though the rules have evolved considerably since e.g. Ruy Lopez's day, when one could promote a pawn only to a piece that had already been captured to use an example of a rule change already mentioned.

Other examples are more clear-cut. Ultimate Frisbee apparently originated as a kind of “Frisbee Football,” (<http://thesportjournal.org/article/the-origins-and-development-of-ultimate-frisbee/>), that is, as a “version of American football played with a flying disc.” A distinctive and interesting aspect of the Ultimate Frisbee culture is its categorical rejection of any sort of refereeing, with players policing one another and relying on a kind of honour code. This strongly held value has been an impediment to the professionalization of the sport, so it obviously is a dearly held value (<https://www.nytimes.com/2016/06/17/sports/ultimate-frisbee-debates-a-role-for-referees.html>). I presume that if someone pointed out that they use referees in American football, Ultimate Frisbee players would not consider the slightest reason to change their practice. Once again, the proposal on offer seems to sort the cases correctly. Or to take another example, rugby enthusiasts would presumably be unmoved by the fact that the kind of behaviour routinely tolerated in the scrum would be considered totally out of bounds in football (soccer). Another useful test is whether players of the contemporary game would find it compelling to think of players of the ancestral game as “heroes” of their own cherished game. Papineau makes the point well:

The problem that faces any new sport is its lack of tradition. History is an essential component of sport. All established sports can tell tales of past heroes and famous victories...It is one thing to be good at hitting a little ball...with long sticks. It's another to follow in the footsteps of Old and Young Tom Morris, of Ben Hogan and Tom Snead, of Steve Ballesteros and Tom Watson. (Papineau 2017: p. 258)

Baseball and cricket probably both evolved from a common ancestor sport of some kind in Britain, but baseball enthusiasts view cricket as boring and alien, while cricket enthusiasts tend to view baseball as “nothing but vulgar razzmatazz” (Papineau 2017: p. 266). Neither would view precedents from the one as relevant to their own practices even if their common ancestry were somehow proved beyond a doubt.

The proposal on offer adverts to the idea that “whether those taking part in a later stage of the social practice take themselves to be bound, *ceteris paribus*, to respect precedents and values (implicit or otherwise) associated with an earlier stage of that practice.” *How many* of those taking part must endorse the precedents and values of the previous stage, though?¹⁶ If we say “all” then the proposed theory will be too restrictive, but if we just say “some” then it will be too permissive. We therefore seem forced to rely on the deliberately vague thesis that current participants “generally” take this attitude to the relevant previous stage of the social practice. This will, of course, result in considerable indeterminacy in when two stages of a social practice should count as the same game_{SP}. That, though, is not obviously objectionable. After all, it *does* seem quite indeterminate in many cases whether we have the same game or a different one in this sense. In fact, it is a virtue of a theory if it posits indeterminacy in a concept which roughly matches the indeterminacy in folk applications of that concept.

Actually, it might be useful to distinguish an “internal” and an “external” perspective on this question, in the spirit of H.L.A. Hart’s famous distinction between internal and external perspectives on the law (Hart 2012). Hart distinguished between a more sociological “external” perspective on law that one might make qua external observer from the “internal” perspective of someone bound by the law and who accepts it as providing a standard of conduct. This perspective is normative, in that it involves accepting certain norms to govern one’s conduct. The account of the individuation of games provided above would provide an analogue of Hart’s external perspective, here on sameness of game rather than content of law. From the point of view of an active enthusiast of some game, though, this sociological perspective may seem beside the point. Think of the famous remark, “it is just not cricket.” While the phrase is now used much more widely to denote any dishonest or unsavoury action, it was most famously used to describe the “under-arm bowling incident of 1981” in which the bowler delivered the ball to the batsman “under-arm,” along the ground. Although technically legal at the time (it was soon after banned), it was widely seen as not in the spirit of the game. The fact that many found it so natural to describe it as “just not cricket” provides a nice example of the internal point of view in the sense I have in mind. To say that some construal of the rules is “not cricket” in this sense is, in effect, to make a *normative* claim—to claim, roughly, that the values which make the game worth playing decisively count against that construal and that a game played with the rules so understood would “not be cricket”—that is, would constitute a different game_{SP} in my sense. Interestingly, this internal perspective on the individuation of game_{SP} can come apart from the external one. An enthusiastic participant of some gaming practice may concede that his/her peers generally (perhaps even universally apart from himself/herself) view a given precedent or value as binding but reject their conception as a perversion of what makes the game worth playing. Charles Stevenson’s famous discussion of “persuasive definitions” (Stevenson 1938) is also relevant here, since we can see such a maverick as offering a Stevensonian persuasive definition of “cricket” (or “chess,” or whatever). Although I lack the space to elaborate on this here,

¹⁶ Thanks to an anonymous referee for pressing me on this point.

I think both the internal and external perspective on individuating games_{SP} has a basis in ordinary discourse and practice.

Returning to the external perspective on individuating games, there is a puzzle worth at least briefly discussing, though it too raises broader issues I cannot fully address here.¹⁷ The proposed criterion for the identity of games_{SP} can naively seem to generate a paradox, given the transitivity of identity. Suppose generation 10 mostly respects how generation 9 played the game and treats their values as relevant to how to go on. Suppose this is also true of generation 9 in relation to generation 8, and so on all the way back to generation 1. Further assume that the changes made between generation 1 and generation 10 were, cumulatively, sufficient that generation 10 does not respect the values inherent in the practices of generation 1, and view their precedents as largely irrelevant to how to go on. On the one hand, given that each generation's game is identical to the game_{SP} of the immediately previous generation, the transitivity of identity seems to entail that the game_{SP} played by generation 10 is numerically identical to the game_{SP} played by generation 1. On the other hand, the canonical criterion for sameness of game_{SP} trivially entails that the game_{SP} played by generation 10 is *not* identical to the one played by generation 10. The proposed criterion therefore leads inexorably to contradiction and must be abandoned.

Before explaining how this objection can be disarmed, it is worth noticing that this phenomenon is not merely hypothetical. Furthermore, it can take a “branching form.” For example, modern forms of chess derive from the Indian game *chaturanga*. However, *chaturanga* is also the common ancestor of *shogi*, *sittuyin* and *makruk*. As it happens, in this case, these games are all often characterized as “variants of chess,” but one can easily imagine a case with this structure in which *shogi* might come out on the proposed account as identical to *chaturanga*, and chess might come out as identical to *chaturanga*, yet chess might still come out as not identical to (that is, not the same game_{SP} as) *chaturanga*. Once again, contradiction looms.

It is not hard to see how this paradox could equally be thought to arise for the fairly orthodox conceptions of the individuation of biological species and of natural languages which provided the model for the account offered in the case of games here. To that extent, in facing this objection, my proposed theory is at least in good company. Fortunately, though, the objection can be disarmed, or so I shall argue. The first point is a logical one. The account developed here entails in the generations 1-10 example above that there is *some* game_{SP} such that generations 10 and 9 are both playing that game, and that there is *some* game_{SP} such that generations 9 and 8 are both playing that game, and so on back to generation 1. The account does *not* entail that there is just one game_{SP} throughout this historic chain. So while generations 10 and 9 might both be playing chess, generation 9 might also (and in the same token activities) be playing chess* which generation 10 is *not* playing at all. Since generation 8 might be playing chess* but not chess, it would then follow that generations 10 and 9 participate in the same game practice, generations 9 and 8 participate in the same game practice, but there is no game practice such that generations 10 and 8 both participate in it.

For this resolution of the alleged paradox to work, it must be the case that one generation is, in virtue of precisely the same game playing activity, simultaneously

¹⁷ Thanks to an anonymous referee for drawing me out on this point.

playing two distinct games [chess and chess* in my toy example]. That might itself seem paradoxical, or at least counter-intuitive. In my view, though, this seems counter-intuitive only because it is hard to track the different senses of ‘game’ when thinking about the case. In the sense of ‘game’ qua abstract entity, of course, the relevant generation could and presumably would be playing just one game_{AE}. For simplicity we may assume that there are not multiple variations of the game in currency in that generation. Even if there are multiple variations in currency, it will remain true that for any given token contest, the participants will typically be playing just one game_{AE}.¹⁸ This one game_{AE} is, though, itself related in the right way to two distinct games_{SP}, and because the members of the relevant generation are simultaneously taking part in both of those practices, their token contest also counts as an instance of both games_{SP}. Because we do not, in ordinary language, explicitly track these different senses of ‘game’ it is easy to confuse them, and it is just this sort of confusion which in my view makes the proposed resolution seem problematic.

It might be useful to work through a concrete case to test the plausibility of this resolution. The history of *Monopoly* provides a nice test case.¹⁹ Although commonly considered the brainchild of Charles Darrow, the game’s origins go back to Elizabeth (“Lizzie”) Magie. Magie developed what she called *The Landlord’s Game* to illustrate the ideas of the economist/politician Henry George. George was a critic of unbridled capitalism as leading to problematic monopolies. He advocated what he called a “single tax,” where this was a tax on land, construed broadly to include natural resources and natural opportunities. George argued that such an approach was efficient (here drawing on ideas from Adam Smith), more egalitarian than other forms of taxation and uniquely just. Magie’s *Landlord’s Game* was similar in many ways to modern *Monopoly*, and was innovative in many ways, including that game play does not involve moving one’s piece to a fixed end-point which terminates the game, but instead involves going around the board over and over again. However, Magie’s version of the game was different from modern versions of the game in one especially salient way. At any point during the game, two or more of the players could, by voting, change the rules, switching the game into “Single Tax” mode. Single Tax mode effectively transformed the game into one which illustrates the advantages of George’s Land Tax. Perhaps most notably, under these rules rent on owned properties was *not* paid to the owner of that property when another player landed on it, but was instead paid into the “public treasury” to be used for public improvements. The fees on railroads are still paid to their owners, but only until they are nationalized when they become free. Players can, moreover, improve any property they like with a house, whether their piece is on the corresponding space or not. Eventually Free Colleges are introduced into play. Magie’s intention in developing the game was, of course, polemical. She wanted a game which illustrated

¹⁸ Actually, it is logically possible for one and the same activity simultaneously to constitute the playing of two distinct games_{AE}. One illustration of this logical possibility is found in “meta-games” in which one game is contained as a proper part of some other game. E.g. when the parties in a role-playing game like *Dungeons and Dragons* have their characters play a game of poker, and where the DM has them do that by actually playing poker with real playing cards, etc. Computer games with this structure are also not hard to find. For purposes of simplicity I put this possibility to one side in the text, as it is a complication that would not change any of the main lines of my proposed theory.

¹⁹ My discussion in the text draws upon Pilon (2015), which provides a rich and detailed history of *Monopoly*. Numerous online sources summarize some of the main contours of the history too, though.

both the perils of unfettered capitalism, most especially private ownership of land, and the advantages of George's Land Tax. The first part of the game, which is most similar to *Monopoly* as played today, was meant to illustrate the ruthless and immoral nature of unconstrained capitalism, while the latter stage was meant to illustrate the more humane and economically efficient advantages of George's Land Tax. Sadly, Magie's attempt to sell her game (for which she had a patent) to *Parker Brothers* failed, who passed on it twice on the grounds that it was too complex and too polemical.

However, this is far from the end of the story. The game proved popular in the northeast, especially with University professors and students, and homemade, bootleg versions of the game proliferated. Eventually, a version of the game was taken up by Scott Nearing who was a professor of economics at Wharton, and who picked up the game from the residents of Arden, a progressive utopian society based on George's single tax philosophy. Nearing, and eventually his student, Rexford Tugwell used their own versions of the game to teach their students economics at Wharton. Quite separately, bootleg versions of the game reached Atlantic City, where Charles Darrow learned about the game when he played it with some of his friends. Poignantly illustrating the force of Magie's original game, Darrow shamelessly *stole* the game, modified its rules into something much closer to the modern version by dropping out the second stage in which players implement the Land Tax altogether and eventually sold it to *Parker Brothers*. Darrow's version effectively inverted the values and message of Magie's original by valorising the cutthroat competition of unbridled competition and celebrating the sadistic joy of crushing one's opponents by acquiring monopolies and bleeding them dry. Darrow claimed the game was entirely his own invention and made a fortune as the game of course became extremely popular. Magie eventually sued *Parker Brothers* but received a very paltry settlement.

This rich history was all uncovered as the result of a lawsuit which arose over yet *another* game derived from Magie's original, namely *Anti-Monopoly*, which was invented by the San Francisco economist Ralph Anspach. The point of this game, which came in more than one version, was far closer to Magie's original vision. Anspach wanted a game which would illustrate why unfettered capitalism generated monopolies that were highly problematic, and thereby illustrate the need for robust anti-trust legislation. In one version of the game, each player chooses either to be a "monopolist" or a "competitor." Competitors charge lower rents (the idea being that they do not gouge people by acquiring monopolies), whereas monopolists charge higher rents. However, monopolists must own at least two properties of a given set before they can make improvements (houses and hotels), whereas competitors can start improving a given property immediately upon acquiring it. *Parker Brothers* unsurprisingly sued Anspach. In defending himself, Anspach uncovered the history of the game going back to Magie and argued that Magie had put the game into the public domain, thus insulating Anspach from *Parker Brothers'* lawsuit.

I have laid out this history in some detail because it provides a very nice test case for my proposed solution to the paradox. Clearly in a broad sense we have a social practice here which stretches from Magie's original game up to the present day, including *Parker Brother's* classic version of the game as well as Anspach's *Anti-Monopoly*. Moreover, we very plausibly have a case which richly illustrates the alleged failure of transitivity of identity, at least on some ways of filling in some of the details of

the history (here I make no claim to full historical accuracy). Plausibly the version of the game played at Wharton to teach economics is identical to Magie's original *The Landlord Game*, and Darrow's modern form of *Monopoly* is identical to the Wharton version of game, yet Darrow's game is not identical to Magie's game. The last point is easy enough to derive from the account on offer here, since the values of Darrow's game, namely the valorization of cutthroat capitalist competition, are roughly the polar opposite of Magie's. Enthusiasts of Darrow's game would, therefore quite rightly not consider the values and ideas of Magie's game as providing precedents or values they must respect. On the other hand, since the Wharton version of the game was less polemical without actively valorising unfettered capitalism, its enthusiasts might well consider the values and precedents of Magie's original game to have at least some relevance to their practice. Likewise, enthusiasts of Darrow's game might consider the Wharton less polemical version to provide relevant values and precedents. Finally, of course, *Anti-Monopoly* enthusiasts might well consider Magie's game to provide germane values and precedents but not care about the values of the Darrow/*Parker Brothers* game. This last point illustrates that social practices can have "gaps" so long as they are causally connected in the right way, as with the gap in anti-capitalist versions of the game as between Magie's *The Landlord Game* and Anspach's *Anti-Monopoly*.

So, at least on some ways of filling in the details of this story, we should say that the Wharton game, Magie's game and *Anti-Monopoly* are all proper parts of one game_{SP}, while the Wharton game and the Darrow/*Parker Brothers* games are all proper parts of a numerically distinct game_{SP}. These two game_{SP} do not have names, though we could name them. However, that should not provide any kind of knock-down objection, after all, it is intuitively perfectly possible for someone to invent a game and play it without yet having come up with a name for it; there is no contradiction in the idea of a game with no name. More interestingly, those playing the Wharton version of the game would, on this account, be playing two game_{SP}, since that version of the game is a proper part of both social practices. Once we understand what that means, though, it should not seem problematic. For all that really means is that those playing the Wharton game were, without realizing it, taking part in two distinct but overlapping social practices. Moreover, we can, of course also say that they were playing their own specific game_{AE}, and it is the game in this sense which fixed the rules and goal of the game they were playing.

To further minimize any residual sense of paradox, consider this analogy. My department cross lists some of our undergraduate courses with masters courses. This means that my token activity of giving a particular lecture will simultaneously constitute two lectures in one sense—the lecture for the masters course and the lecture for the undergraduate course. In another, perfectly familiar sense, though, there is just the one lecture. This is an apt analogy because the one activity constitutes two lectures only in virtue of being a proper part of two distinct social practices—our undergraduate programme and our masters programme. In another sense, it is a single lecture, where that could be understood as something broadly analogous to a token contest, though here of course it is simply a lecture rather than a contest. It could even be seen to bear salient relations to an abstract type where that type is constituted by the words and semantic values associated with them which make up my lecture.

Obviously much more could and should be said about this proposal about when a change in social practice constitutes a change in game_{SP} , its broader implications and my resolution of the alleged paradox. However, the proposal seems to do reasonably well at sorting cases in an intuitively plausible way. Unfortunately, I do not have the space to delve into the many cases one might consider when further testing this proposal, much less to consider rival approaches on this front.

5 Rival approaches

In this section I briefly consider some rivals theories. Probably the most obvious point of comparison is Fred D'Agostino's influential work on the "ethos" of a game (D'Agostino 1981). Like me, D'Agostino rails against "Platonistic" conceptions of games. He argues that we should distinguish the formal rules of a game from its "ethos" in a given context. He gives the nice example of basketball, which "on paper" is a non-contact sport but which in practice is not treated as a non-contact sport. The "ethos" of the game differs from its "official rules" in being more tolerant of certain forms of physical contact between the players. Like my account, D'Agostino's approach can make sense of different versions of the same game at a given point in time, since games governed by the same canonical rules can have a different "ethos" in different contexts. It might therefore seem that my approach simply reinvents D'Agostino's wheel.

In fact, our accounts are quite different. Granted, we both reject a purely Platonistic approach, but for rather different reasons. D'Agostino is keen to make sense of how players can violate the rules but still count as playing the game. I agree with him that a plausible account needs to do this, as my critique of Suits in the previous section should make clear. However, I do not think his own account goes far enough. On D'Agostino's account, violating the constitutive rules of the game still means you simply are not playing the game unless the violation is one for which the local "ethos" fixes some specific penalty. On my account, someone can still count as playing a game even if they violate constitutive rules which for which no specific penalty is attached by the local ethos. In any event, my argument for understanding 'game' as sometimes denoting a social practice is that doing so can solve the two puzzles with which I began, and is not an argument from making sense of how one can count as playing a game while breaking its rules.

Moreover, my approach differs from D'Agostino's not only in its motivation but in its content. D'Agostino's approach is to take some canonical set of rules as still constitutive of a given game, but then to add that these constitutive rules can be interpreted differently in different contexts. My approach, by contrast, allows that the constitutive rules themselves (and even the constitutive goals) can vary as between different versions of a given game, and that these can change over time as well as regards any canonical rules or goals. This is an important difference in the content in our accounts. Furthermore, his account simply does not make use of the idea that 'game' can be used to denote a historically embodied social practice. For all these reasons, our accounts are quite different. At the same time, the accounts do not need to be seen as rivals. At least, a suitably modified version of D'Agostino's account, which allows that you can play the game while violating even constitutive rules not

associated with specific penalties by the local ethos of the game, could be combined with the main ideas developed here.

How else might one try to solve the two puzzles with which I began? The only other *prima facie* promising approach which occurs to me is one picking up on Wittgenstein's idea that 'game' is a family resemblance term. Those working on the theory of concepts under the influence of Wittgenstein, but also drawing on interesting empirical evidence, have argued that many concepts do not have the classical "necessary and sufficient conditions" structure which my approach seems to endorse. Alternatives include "cluster concepts," "prototype concepts" and "exemplar concepts." A cluster concept is one which is fixed by a set of conditions, perhaps with different weights, such that a competent user of the concept will be more willing to apply the concept in direct proportion to how many of the conditions are true of it—adjusting for any differential weightings. None of these conditions is strictly *necessary* for the application of the concept, though. A prototype concept holds that people build a kind of implicit "prototype" of the relevant kind based on their beliefs about paradigm instances of the kind, and then compare new putative instances to this prototype. A prototype is a purely cognitive construction, though, and may have no actual instances as such—for example, perhaps the paradigm for 'water' is pure H₂O even if pure H₂O never occurs in nature. The greater the similarity to the prototype, the more confidently one should apply the concept to a new type. Finally, an exemplar concept does work via some token paradigm instances, with which a competent user of the concept compares new putative instances to decide whether the concept applies. For example, one might compare olives to apples and oranges and decide that they are not similar enough to count as fruits, given those exemplars. I have chosen an unfavourable example here (olives are fruits, at least in the biologist's sense) to illustrate a problem with the exemplar approach that I take seriously, but put that to one side. How might these various approaches, when applied to the concept of a game, help with the Diachronic Puzzle and the Synchronic Puzzle?²⁰

These approaches seem reasonably well suited to solving the Synchronic Puzzle. Given a prototype or exemplar of a specific game, like chess, we can classify variants on that paradigm, prototype or exemplar as being forms of chess insofar as it is similar enough to the prototype or exemplar. Alternatively, if 'chess' is a kind of cluster concept then we could understand different versions of chess as counting as chess in virtue of satisfying different of the relevant criteria in the cluster. However, none of these approaches seem very well suited to solving the Diachronic Puzzle. Once we have the concept of a game of chess (e.g.), it is hard to see how the game *itself* changes over time. Of course, certain versions of the game might be more prominent at one time rather than another, but that is different from the game itself changing. Only when we understand games (in one sense) as being historically embedded social practices, can we make good sense of the idea that games literally change their rules over time.

There is, though, another interesting if unorthodox way of drawing on Wittgenstein's idea of a family resemblance term worth discussing here. Here I have in mind Michael Pelczar's approach (Pelczar 2000). On Pelczar's account, family resemblance

²⁰ For a useful summary of these three approaches, as well as a critique of each, see Gert (1995a); Gert supplies ample references to prominent defenders of each approach as well.

terms/phrases have three defining features: (1) they are context-sensitive, with their content depending on the context of utterance, (2) they are predominantly topical, so that the topic of conversation tends to select a content (though other features of context can override this), and (3) they are semantically open. The third of these features is perhaps the most distinctive and least familiar, so I shall elaborate on it. Very roughly, for a term to be semantically open is for speakers have a kind of discretion in their use of that term which implies an ability to set semantic precedents. There is scope for talk of semantic rules in this framework, but the rules will not entirely fix the standards for correct usage of the term. Pelczar usefully compares semantic openness with the kind of openness we find in other familiar practices, e.g. in having a conversation, in playing jazz, and in the common law. The extension of a family resemblance term is a matter for argument and negotiation. Although I lack the space to argue for this here, in my view Pelczar's approach actually captures more of what is truly distinctive in Wittgenstein's remark than the three more dominant approaches (cluster concept theory, prototype theory and exemplar theory) do.²¹

The account of games qua historically embedded social practices fits very well with the idea that 'game' is a family resemblance term in Pelczar's sense. The context of utterance can fix different referents for e.g. 'game of chess' since different contexts will select different versions of the game, so we have context-sensitivity. Topicality is harder to establish, but a case can be made for this as well. If we are discussing a FIFA game of football then that naturally suggests one content for 'game' (qua token activity), whereas if we are talking about a pick-up game of football as played by some kids after school, another content will be triggered. Finally, semantic openness falls out of the account developed here *very* naturally. Games qua social practices are very much like the common law, relying on precedent and shared values to negotiate new cases and to argue for or against extending or retracting the practice in various ways. Interestingly, if this is right then 'game' (full-stop) is not a family resemblance term, but the words for specific games like 'chess' and 'basketball' are family resemblance words.²² That is a surprising and interesting result, but one I cannot explore here.

6 Conclusion

I have here laid out two puzzles about the individuation of games and argued that these puzzles are analogous to two puzzles facing a dominant way of thinking about natural languages. In both cases, we should reject a purely Platonic conception of the phenomena. Following Jackman's lead in the case of languages, I have argued that we should conceive of 'game' in one sense as denoting a historically embedded social practice. I have argued that this approach can be combined with a wide variety of ways

²¹ One advantage of Pelczar's view which he emphasizes is its ability to avoid the collapse into a notational variant on a classical style "necessary and sufficient conditions" reductive definition—a collapse which arguably is a big problem for e.g. "cluster concept," "prototype theory" and "exemplar theory." For useful discussion, see Gert (1995a).

²² Moreover, 'play a game' may also be a family resemblance phrase even though 'game' is not itself a family resemblance term. This is because what it is to *play* a game may itself depend on the kind of game it is. I discuss this further nuance in other work.

of understanding games qua token activities, and that the structure of the approach easily resolves the two puzzles with which I began. I have compared my approach with what I take to be its main potential rivals and argued that the account developed here need not be seen as a rival to those approaches, but that insofar as it is seen in that light it is superior to them. Finally, I have explored the idea that given the approach on offer here, in one interesting sense the words for *specific* games, like ‘chess’, may be family resemblance terms (in Pelczar’s sense) even though ‘game’ itself may not be—or if it is, this is less immediately compelling than it is in the case of words for specific games like ‘chess’. If this is right then it would mark an interesting contrast with standard ways of thinking in this area, which tend to treat ‘game’ as a family resemblance term but words for specific games (like ‘chess’) as amenable to a reductive definition of some kind.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Boer, S. (1974). Speech-acts and constitutive rules. *Journal of Philosophy*, 79, 169–174.
- Bogost, I. (2016). *Play anything*. New York: Basic Books.
- Burge, T. (1979). Individualism and the Mental. In U. French & Wettseti (Eds.), *Midwest studies in philosophy IV: Studies in metaphysics* (pp. 73–121). Minneapolis: University of Minnesota Press.
- Burke, B. (2014). *Gamify*. Brookline: Bibliomotion.
- Chomsky, N. (1988). *Language and the problem of knowledge*. Cambridge: Cambridge University Press.
- Consalvo, M. (2009). There is no magic circle. *Games and Culture*, 4, 408–417.
- D’Agostino, F. (1981). The ethos of games. *Journal of the Philosophy of Sport*, 8, 7–18.
- Davidson, D. (1992). The second person. *Midwest Studies in Philosophy*, XVII, 255–267.
- Gert, H. (1995a). Alternative analyses. *Southern Journal of Philosophy*, 28, 31–37.
- Gert, H. (1995b). Family resemblances and criteria. *Synthese*, 105, 177–190.
- Hart, H. L. A. (2012). *The concept of law*. Oxford: Oxford University Press.
- Jackman, H. (1999). Convention and language. *Synthese*, 117, 295–312.
- Kaplan, D. (1990). Words. *Proceedings of the Aristotelian Society*, 64, 93–113.
- Katsafanas, P. (2013). *Agency and the foundations of ethics*. Oxford: Oxford University Press.
- Kreider, A. J. (2011). Game-playing without rule-following. *Journal of the Philosophy of Sport*, 38, 55–73.
- Langton, R., & West, C. (1999). Scorekeeping in a pornographic language game. *Australasian Journal of Philosophy*, 77, 303–319.
- Lewis, D. (1969). *Convention*. Cambridge, MA: Harvard University Press.
- Lewis D. (1975). Languages and language. In K. Gundersons (Ed.) *Language, mind and knowledge*. Minneapolis.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 87, 339–359.
- Malaby, T. (2007). Beyond play: A new approach to games. *Games and Culture*, 2, 95–113.
- Mayr, E. (1988). *Toward a new philosophy of biology*. Cambridge: Harvard University Press.
- McGonigal, J. (2011). *Reality is broken*. Great Britain: Random House.
- McKeever, S., & Ridge, M. (2006). *Principled ethics: Generalism as a regulative ideal*. Oxford: Oxford University Press.
- Midgley, M. (1974). The game game. *Philosophy*, 49, 231–253.

- Papineau, D. (2017). *Knowing the score*. London: Constable.
- Pelczar, M. (2000). Wittgensteinian semantics. *Nous.*, 34, 483–516.
- Pilon, M. (2015). *The monopolists: Obsession*. Bloomsbury: Fury and the Scandal Behind the World's Favorite Board Game.
- Putnam H. (1975). The meaning of meaning. Reprinted in Putnam's *mind, language and reality*. New York: Cambridge University Press.
- Rawls, J. (1955). Two concepts of rules. *Philosophical Review.*, 64, 3–32.
- Reddiford, G. (1985). Constitutions, institutions, and games. *Journal of the Philosophy of Sport.*, XII, 41–51.
- Richard, M. (2019). *Meaning as species*. Cambridge: Harvard University Press.
- Ridge, M. (2019). How to play well with others. In T. Hurka (Ed.), *Games, sports and play*. Oxford: Oxford University Press.
- Sellars, W. (1954). Some reflections on language games. *Philosophy of Science*, 21, 204–228.
- Stevenson, C. (1938). Persuasive definitions. *Mind.*, 47, 331–350.
- Suits, B. (1988). Tricky triad: Games, play and sport. *Journal of the Philosophy of Sport.*, 15, 1–9.
- Suits, B. (2005). *The grasshopper: Games*. Life and Utopia: Broadview Press.
- Toppo, G. (2015). *The game believes in you*. London: Palgrave Macmillan.
- Walton, K. (1990). *Mimesis as make-believe*. Cambridge: Harvard University Press.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review.*, 105, 489–523.
- Wittgenstein, L. (1958). *Philosophical investigations*. New York: MacMillan.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.