

Cheating and Trifling¹
Michael Ridge
2015

The main questions addressed in this paper are motivated by two related paradoxes, which I call the “paradox of cheating” and the “paradox of trifling.” Both arise out of the need to distinguish forms of behaviour which can superficially seem like game playing from the genuine article. I begin by laying out these paradoxes (section I). I then discuss a variety of tempting strategies for resolving these paradoxes, and argue that none of them work (sections II-IV). I develop my own positive account of what it is to play a (competitive) game and argue that this account is both independently plausible and elegantly resolves the two paradoxes – or, rather, does so once a further important distinction is drawn (section V). I conclude by discussing how my account can be situated in a broader semantic framework according to which ‘play a game’ and cognate phrases are family resemblance phrases, at least in one useful sense of ‘family resemblance’ (section VI).

I. Two Paradoxes

What I call the “paradox of cheating” arises naturally out of reflection on how best to distinguish playing a game from merely engaging in behaviour which happens to accord with the rules of that game. Plausibly, someone who accidentally accords with the rule of a given game is not thereby playing. For example, consider two people who do not know any of the rules of chess, but take turns moving the pieces (from the starting position) at random. Miraculously, they end up always moving them in accordance with the rules of chess until they get bored - which happens to be when a checkmate has appeared. Intuitively, these two people have not thereby played a game of chess.

Some might balk at that example because it is so unrealistic, though I think it makes the point effectively enough. No matter; there are more realistic cases which also make the point. If two children who live across the street from one another leave their homes at the same time, but are running late and so run as fast as they can along the same route to school then they are acting in accordance with the rules of the game, “footrace to school.” Yet, intuitively, they are not thereby having a race. A child who hides from his parent to avoid getting in trouble may be acting in accordance with the rules of the game “hide and seek,” but they are not thereby playing a game of hide and seek, and so on.

The most obvious and tempting way to avoid counting such cases as instance of game playing is to define playing a game in terms of *following*, and not merely behaving in accordance with, the rules of the relevant game. In each of the preceding examples, what seems to be driving our intuitions is the fact that it is a sheer *coincidence* that the people involved are acting in accordance with the rules. The implicit contrast would seem to be with the paradigmatic way in which game players’ accordance with the rules is *not* accidental. Plausibly, their accordance with the rules is not accidental precisely because they know the rules and they deliberately guide their behaviour by the rules.

This has led many theorists to endorse definitions of ‘play a game at time t’ on which a necessary condition on one’s playing a given game at some specific time t is that one is *following* the rules of that game at t (most notably Suits 2005, but see also Dellatre 1976: 136, Fezell 1988, and Ross

¹ Thanks to Sebastian Koehler, Geoff Sayre-McCord and the members of my 2014 class, “Philosophy: Fun and Games,” at the University of Edinburgh, for very helpful discussion of the ideas developed in this paper.

1968: 25). Since the two people in my first example above do not even know the rules of chess, they were never *following* them, as opposed to merely accidentally according with them. So we have a nice, elegant explanation of why they were not playing chess. This approach also provides a ready explanation of why you have to know the rules of a game to play that game, which is plausibly a welcome corollary.

Unfortunately, any definition which makes rule-following a necessary condition on playing a game at a given time trivially entails that anyone who cheats at that game at time t does not play the game at t . This follows simply because cheating entails deliberately breaking the rules. Cheating may well involve more than deliberately breaking the rules, but it least involves that. For example, it would not seem right to call someone who deliberately fouls another player in a game of basketball to stop the clock a “cheater,” since this sort of rule violation is an accepted strategic part of the game (Cf. D’Agostino 1981), with certain characteristic penalties associated with it. Perhaps cheating necessarily involves an element of deception, or an attempt to gain an unfair advantage, as well; it may also presuppose that the rules are themselves fair and fairly enforced (compare Green 2003). All that matters here is that it is a *necessary* condition on one’s cheating that one deliberately violates the rules; that is enough to generate the paradox discussed below. I therefore put to one side the independently interesting question of what other necessary conditions must be satisfied for a piece of behaviour to count as cheating.

The implication that anyone who cheats thereby is not playing the game is paradoxical because it also is an extremely plausible that someone cannot be guilty of cheating at a given game at a given time unless she is playing that game at that time. To get a feel for why this is so plausible, consider the following statement: “Jones cheated us in our game of poker last night, but Jones has never played poker.” This certainly sounds incoherent, yet if there were no necessary connection between cheating at a game and playing that game, it is hard to see why it would sound so contradictory.

In fairness, there is a weaker necessary connection between playing a given game and cheating at that game which could accommodate the oddness of such statements. One might hold that one can count as playing a game even if one cheats at that game, but that one is not playing the game *in* cheating at it. So, for example, one might hold that up until the point at which Jones began cheating (when he started dealing from the bottom of the deck, say), he was playing the game.

There is a problem with this move, though. It only works in cases in which the cheater only starts cheating midway through the game. The same move clearly will not work for the *inveterate* cheater. The inveterate cheater cheats throughout the game. If, though, to play the game one must follow the rules, then the inveterate cheater never played the game to begin with. This is an implausible conclusion. Some readers may not find this obvious; I elaborate on the point and argue for it at greater length in section II below. For now, let us assume to see that the possibility of the inveterate cheater means that we should try to make sense of the idea that someone can cheat at a game and play that game at one and the same time.

With this assumption in place, we are faced with a paradox. In order to distinguish playing a game from accidentally according with its rules, we insist that following the rules at a given time is necessary for playing the game at that time. This, though, plus platitude that cheating minimally requires deliberate rule-breaking, entails that insofar as someone is cheating at a given time she is not playing the game at that time. I have suggested, however, that cheating at a given time actually *presupposes* that the person is playing the game at that time, for the reasons given above. From this it follows that if someone cheats at a given game at a given time that she both follows the rules

and deliberately breaks the rules of that game at that time. This is clearly impossible, but cheating clearly is possible, so we now have a contradiction. Indeed, we seem to have arrived from very compelling premises through impeccable reasoning to a logical contradiction. This is what I call the “paradox of cheating”:

1. If someone plays a game at time *t* then she follows the rules of that game at *t*. [otherwise we cannot distinguish merely accidentally according with the rules from proper game playing]
2. If someone cheats at a game at time *t* then she violates the rules of that game at *t*. [platitude]
3. If someone cheats at a game at time *t* then she is playing that game at time *t*. [otherwise we cannot make sense of someone cheating throughout a game she is nonetheless playing]
4. Cheating is possible.

Each of these theses has a lot to recommend it when taken in isolation. However, together they entail a contradiction. Something must give.

Interestingly, the paradox of cheating provides a fascinating and very close parallel with a more well-known paradox facing traditional Kantian theories in moral philosophy. Kantians hold that for an action to be autonomous, the agent must be following the moral law. So when an agent acts immorally, they lack autonomy. Yet one might have thought that for an action to be immoral (as opposed to merely non-moral, e.g. the actions of an infant), the action must in some sense be done autonomously. So unless (as some, including Kant, argue) there is some principled equivocation on ‘autonomy’, immoral action looks to be impossible. This problem plagues Christine Korsgaard’s recent neo-Kantian account. Korsgaard herself suggests that Platonic ethics faces a similar problem. Here is Korsgaard: “...it looks at first as if for Kant *nothing exactly counts as bad action.*” (Korsgaard 2009: 160) The main point for present purposes is to note how *precisely* analogous this paradox is to the one sketched above. Here the upshot is that “nothing exactly counts as cheating.” Just replace ‘plays a game *G* at time *t*’ with ‘performs an autonomous action at *t*’, replace ‘the rules of *G*’ with ‘the moral law’, and replace ‘cheats at a competitive game *G* at time *t*’ with ‘acts immorally at time *t*’ throughout, to see how striking this analogy is. I intend to return to this analogous paradox future work, and to see what fresh light might be shed on it by analogy with the paradox I discuss here. My hope is that my discussion of the seemingly more down to earth concept of a game in this paper will provide a more tractable domain for resolving a paradox with this form than the more high-flown and controversial context of morality.

The second paradox I discuss in this paper I call the “Paradox of Trifling.”² This paradox is structurally similar to the paradox of cheating. In each case, we seem driven to embrace a necessary condition on playing a game to explain why certain behaviour which seems like game playing is not, and in each case this necessary condition, when combined with other plausible theses, quickly leads to paradoxical conclusions.

Suppose two people move the chess pieces deliberately in accordance with the rules of chess from the starting position. However, they are trying to produce pretty patterns on the board, rather

² My use of the term ‘trifling’ here is obviously indebted to Bernard Suits’ discussion, though equally obviously I disagree with him about trifling! See Suits 2005: chapter 4.

than trying to deliver checkmate. Indeed, they may not even have the concept of checkmate – they may only know the rules governing how the pieces move. While these two people may well count as playing *some* game or other, they hardly should count as playing a game of chess, or indeed as playing any *competitive* game at all. Reflection on cases like this have led theorists to endorse definitions which entail that it is impossible to play a given game at a given time without trying to win at that time. Bernard Suits, for example, famously endorses something very close to this (Suits 2005), since he defines playing a game, in part, in terms of trying to achieve what he calls the game’s “lusory goal”. I elaborate on the details of Suits’ view below (in section II), but for now suffice it to that, at least for competitive games, to achieve the lusory goal of the game is to win the game.

However, it also seems like a platitude that one *can* play a game, even a competitive game (one which admits of winners and losers, that is) without trying to win. Parents routinely play such games with their children without trying to win in order to encourage them or simply to pass the time. Yet intuitively they are playing the game all the same. Admittedly, someone who discovers her opponent is not trying to win might intelligibly express her frustration by exclaiming, “Play the damn game!” Indeed, in the end I will agree that there may be *a* useful sense in which someone who is not trying is thereby not playing the game (see sections VI below). Nonetheless, there is also a useful and in my view more central sense of ‘play a game’ in which someone who is not trying to win can still count as playing the game. Therefore in order to block the looming paradox, an adequate theory must specify at least one useful sense of ‘play a game’ in which playing a game does not entail trying to win.

Some readers may balk at my example of parents playing games with their children. Perhaps in that case it is not so implausible to deny that the parent is really playing a game, as opposed to going through the motions for the sake of their child. Perhaps, though I myself find this interpretation strained. Fortunately, formal competitive contexts provide another useful example in which people can count as playing a game without trying to win.

Suppose I enter and turn up at a chess tournament, sit down at the board and make a series of legal moves but make *no* attempt to win. When I then find myself in checkmate, I complain to the arbiter that this should not count as a proper loss, but as a default, on the grounds that if I did not try to win then no game was played, and you cannot lose unless you played. Again; the distinction is not merely verbal – a loss will negatively impact a player’s rating; a default will not. Such a complaint would quite rightly be heaped with scorn. By the conventions of tournament chess, a game was played even if one player (or both, for that matter) did not try to win.

Once again, we face a paradox (the paradox of trifling):

1. If someone plays a competitive game at time *t* then she is trying to win that game at *t*.
[otherwise we cannot explain why e.g. two people moving the pieces in accordance with the rules as a purely aesthetic exercise are not playing chess]
2. It is possible to lose a game without ever trying to win. [as the example of the perverse chess player demonstrates]
3. You can only lose a game if you played it. [platitude]

Actually the paradox can be formulated in terms of winning, rather than losing, and this might be an even more forceful version of the paradox. Although I consider it to be a platitude that you can only lose a game if you played it, I can imagine some ways of pushing back on this point. The thesis that you can only win a game if you played it seems even more compelling, though, and this is enough to generate a version of the paradox:

1. If someone plays a competitive game at time t then she is trying to win that game at t . [otherwise we cannot explain why e.g. two people moving the pieces in accordance with the rules as a purely aesthetic exercise are not playing chess]
2. It is possible to win a game without ever trying to win. [a lucky trifler might win against a weak opponent]
3. You can only win a game if you played it. [platitude]

The paradox of trifling of course applies only in the case of games which admit of winners and losers – competitive games. Plausibly, there are games which do not admit of winners and losers even in principle, such as the computer game, *The Sims*. For those games, the idea that one must try to win to count as playing gets no purchase. I do not think this scope restriction ultimately does much to undermine the philosophical significance of the paradox, but to explain why this is the case I need to get more of my own theory in place. I therefore return to this qualification below (in section VI). In any event, with these paradoxes on the table, I now consider some of the main ways in which one might try to avoid the looming contradictions associated with each of them.

II. Rule-following and the inveterate cheat.

In this section I discuss what the most obvious way to respond to these paradoxes. In each case one might simply argue that cheaters and triflers are not playing the game. At least, one might argue that insofar as they are cheating and trifling they are not playing. This is in fact precisely the view taken by Bernard Suits, though he does not himself seem to consider these conclusions to be even *prima facie* paradoxical. Suits is certainly not the only theorist to endorse this view. Edwin Delattre, for example, declares, “A person may cheat at a game or compete at it, but it is impossible for him to do both. To cheat is to cease to compete.” (Delattre 1976: 136; similar remarks can be found at Ross 1968: 25 and Ganz 1971: 73, e.g.). However, I focus on Suits’ account because it is very clear and well worked out.

In his underappreciated masterpiece, *The Grasshopper*, Suits argues that ‘play a game’ should be reductively defined as follows:

To play a game is to attempt to achieve a specific state of affairs [prelusory goal], using only means permitted by rules [lusory means], where the rules prohibit use of more efficient in favour of less efficient means [constitutive rules], and where the rules are accepted just because they make possible such activity [lusory attitude].
(Suits 1995: 54-55)

A “prelusory goal” is a goal which can be understood independently of the relevant game. In golf, the prelusory goal is getting the ball in the cup. However, to play golf you must aim at achieving this goal only in ways permitted by the constitutive rules – the “lusory means.” These rules must impose some inefficiency on the accomplishment of the prelusory goal. Here, for example, these

rules exclude walking over and dropping the ball in the cup, and instead requires the characteristic use of clubs, etc. The “lusory goal” is just the prelusory goal achieved only through lusory means.

Suits argues that this definition has great explanatory power; I cannot begin to do justice to his subtle treatment here. An obvious consequence of this definition, though, is that insofar as one is cheating or trifling, one is not playing the game. Suits does not shy away from this implication. He openly embraces it (in his chapter, “Triflers, Cheats and Spoilsports”):

Perhaps we can say of the trifler that he is not playing chess because of a deficiency of zeal in seeking to achieve the prelusory goal of chess. If so, then perhaps we can say of the cheat that he is not playing chess because of an *excess* of zeal in seeking to achieve the prelusory goal. (Suits 1995: 59)

Here Suits discusses the prelusory goal of chess, but at this stage one might reasonably wonder just what the *prelusory* goal of an abstract game like chess could be. Unlike golf, there is no obvious way of characterizing the goal of chess which does not make reference (at least implicitly) to the rules. This is because chess, unlike golf, is a game in which the pieces and the moves can be understood functionally. So long as the bishop is understood as the piece which moves diagonally on a board with 64 squares, it makes no difference whether the bishop is made of wood, steel, glass, a digital representation (computer chess), or indeed whether there even is a physical realization of the bishop at all (think of blindfold chess).³

Suits’ reply to this worry is instructive in the present context, as in developing this reply he makes a distinction which he then uses to make his take on the cheat and the trifler seem less implausible. He distinguishes the *institution* of a game like chess from the game itself. Someone can take part in the institution of chess without playing a game of chess. For example, I can take out the pieces and show someone how the knight moves, or I might analyse a recent game or simply set up a position in which one side can force checkmate as a puzzle to solve. In all of these cases there is a sense in which I am taking part in the institution of chess without playing a game of chess. My behaviour and conception of what I am doing can only be understood against the backdrop of my understanding the rules governing the movement of the pieces.

Crucially, the concept of checkmate can be understood in the first instance relative to the institution of chess. Indeed, someone could have the concept of checkmate without realizing that checkmate is the goal of chess qua game. In this sense, checkmate is indeed the prelusory goal of chess, as it can be understood without yet (fully, anyway) understanding the game of chess. The lusory goal, then, is to achieve checkmate only through the means allowed by the game from the designated starting position. Since I can set up a checkmate position just to illustrate a mating pattern, I can achieve the prelusory goal without thereby having achieved the lusory goal, which is all Suits needs to meet the objection that he cannot distinguish lusory from prelusory goals for a game as abstractly defined as chess. I will return to Suits’ useful distinction between institutions and games below. First, though, it is worth pausing to be very precise about what problematic conclusions I am arguing Suits must accept, since both Suits himself and indeed some of his critics are not always as clear about this as they might be.

Suits sometimes writes as if insofar as someone is cheating or trifling at a putative token game, they simply have not played the game at all. He says, for example, that although he is “engaged in something chess-like, playing chess is not what he is engaged in.” (Suits 2005: 59) Similarly, he says of the cheat that he is “not really playing the game.” (Suits 2005: 59) In these remarks, Suits does

³ Compare Sellars 1963 and Lance 1998.

not explicitly relativize his claims about not playing the game to those times at which the person is trifling or cheating. It might therefore seem that he is saying that insofar as you are trifling or cheating, you are just not playing the game at all. This, though, would be an uncharitable reading of Suits' view. Clearly, his definition commits him to saying that insofar as the cheat and the trifler were *previously* pursuing the lusory goal through permissible means (and doing so just in order to take part in that activity), the cheat and the trifler were indeed playing the game *right up until the moment at which they started cheating/trifling*. In fact, I am sure that this is what Suits intended to convey. Because his book is written as an extended dialogue, in the style of the Socratic dialogues, I suspect he sometimes does not include all the caveats and qualifications (such as the relativization to points in time) one would typically expect of a philosophical argument. Because the book aims at (and in my view achieves) lofty literary aims as well as lofty philosophical ones, these slightly incautious formulations are eminently forgivable. Even so, we must be careful to avoid misunderstanding his view.

Nor is this exegetical point entirely hypothetical, as some of Suits' critics do misunderstand him on precisely this front. Fred D'Agostino characterizes Suits as a "formalist" and then defines formalism as entailing that "no activity is an instance of some particular game G if any rule of G is violated *during that activity*." (D'Agostino 1981: 9) On this reading, given that one of the players at some point cheats, Suits must infer that no game was played. It should by now be clear that this is just a mistake. So long as the cheater has not been cheating throughout the putative game, Suits need not and should not deny that a game took place. Rather he should simply deny that once the cheater starts cheating, she has stopped playing. What we need, then is to play close heed to the distinction between (a) whether a game G has been played by A at all, and (b) whether A's behaviour at some specific point in time should count as playing that game. Because he only works with the more coarse-grained question raised by (a), D'Agostino does not give the formalist a fair hearing.

As I explained in section I above, the real problem for formalists like Suits who hold that insofar as you cheat (or trifle) you are not playing the game is raised by the problem of the *inveterate* cheat and the *inveterate* trifler. It is now time to redeem the promissory note from section I, where I indicated I would return to the case of the inveterate cheat and the inveterate trifler and argue in more detail that such individuals are still playing the game.

Gaylord Perry infamously threw "spitballs" throughout his entire career.⁴ These are thrown when the pitcher coats the ball in some substance (typically Vaseline) which makes the ball move in an atypical manner by sliding out of the pitcher's hand without the usual spin. Such pitches are strictly prohibited by the rules of baseball. Perry was only "busted" quite late in his career, but it was long known he was probably throwing them. The problem was that nobody could prove it. Perry is in the Baseball Hall of Fame in spite of this misconduct.

I presume that not *every* pitch Perry actually threw in his career was a spitball. However, we can imagine a hypothetical version of Perry who *did* throw a spitball every time he was at the mound. Most people would not infer from this that our hypothetical Perry never played a game of professional baseball. Yet if cheating at time t entails that you are not playing at time t, it would seem to follow that this modified version of Perry never played. Moreover, there may well have been an entire game in which all of Perry's pitches were illegal. Again, intuitively, this does not seem to entail he did not play a game of baseball.

⁴ I learned about this nice example from A.J. Kreider's work, which I discuss in detail below. See Kreider 2011: 56. The example also figures in Craig Lehman's discussion (see Lehman 1981: 42). For a contrary view, see Feezell 1988, who argues (unconvincingly, in my view) that Perry was not cheating.

Other examples are easy to generate. Suppose someone uses a marked deck throughout a game of poker to gain an advantage. Plausibly, he was cheating throughout the game. However, this seems to imply that he never played if we accept that you are not playing at time t if you are deliberately violating the rules at t . Or suppose someone is playing a game of Trivial Pursuit while secretly consulting the internet on every turn through an earpiece and a confidant. Again, it seems very plausible that such a person played a game of Trivial Pursuit, yet he was cheating throughout.

Here is an example I find convincing. Suppose someone is playing a game of tournament chess and her first two moves are both illegal. This is very unlikely under normal conditions, but occasionally players have arrived at the board drunk, so it could happen. Under the current FIDE rules (FIDE is the international governing body for chess), if you make two illegal moves then you thereby lose the game (your opponent has to call you on these rule violations at the time for this rule to be enforced). Note that you *lose* the game. It is not a default, as when the player just does not show up at the board, e.g., as Bobby Fischer famously did in his bout with Spassky.

Nor is the distinction between a loss and a default a nominal one. If you lose a game of chess then your rating goes down; if you default then your rating is not adjusted. Unless we want to say you can lose a game you never played, it seems to follow from the fact that he lost the game that my hypothetical drunk chess player did indeed play a game of chess. Yet *both* of his two moves were illegal, so he *never* followed the rules. It seems, therefore, that we had better not say that you are playing a game at time t only if you are following the rules of that game at t .

Admittedly, my hypothetical drunk chess player presumably was not cheating, since he was not deliberately making illegal moves to gain some illicit advantage. In a way, that is irrelevant, though. We could simply rename our paradox the “paradox of illegal moves.” I do not think this is necessary, though. We could add that the chess player was so drunk he thought his opponent would not notice he had moved his pawns three squares down the board, say, and did so to try to gain an advantage. Even on this admittedly very eccentric version of the example, he had better count as playing the game since he will correctly be scored with a loss – and you can only lose a game you played. We should therefore conclude that if someone is cheating at a given game G at time t , then it follows that the person is playing a game of G at time t . This, though, Suits cannot accommodate. His account therefore still has counterintuitive consequences.

Of course, it is a contingent fact about tournament chess that it draws a robust distinction between a default and a proper loss.⁵ The conventions of tournament chess could be changed so that a loss by default would still impact your rating (and your opponent’s rating). In fact, there are other games in which there is no system of ratings and which a default is treated just like a loss (*Magic: The Gathering* apparently is like this). My appeal to this contingent fact about these conventions was really just a way of making vivid what I take to be pre-theoretically very plausible anyway, though, namely that you cannot lose a game you never played. Indeed, even if the conventions of chess did change in this way, I would still be inclined to say that there is a difference between a loss and a default. The point would be that for various pragmatic reasons, tournament directors and FIDE would treat a default just like a loss, but of course there can be very good reasons to treat something just like an X even when it manifestly is not an X . I would be inclined to say the same about *Magic* too.

Actually, all I really need here is the weaker thesis that in at least *one* useful sense of ‘play a game’ and ‘lose’, it is a platitude that you can only lose a game if you played it. I am going to argue below

⁵ Thanks to Sebastian Kohler for drawing me out on this point.

that ‘play a game’ is a context-sensitive phrase. I will not go into this complication here, but it seems very likely that if ‘play a game’ is context-sensitive then ‘lost a game’ will be context-sensitive as well. I can therefore allow that there may well be senses of ‘played a game’ and ‘lost a game’ in which it can be true that someone lost a game he never played. We do, after all, sometimes say someone “lost by default,” after all.

It is enough for my purposes that there is *also* a sense of these phrases on which you can only lose a game if you played it. I happen to think the senses in which this is a platitude are more central. If someone were to say that Fischer lost the first game of the match to Spassky, it would be quite natural to reply that, “Well, he didn’t *really* lose; he didn’t even play.” To my ear, at least, the use of ‘really’ to trigger a content of ‘play a game’ on which this is a platitude is entirely apt. If that is right then it suggests that this content (whatever it is) is somehow more at the core of our concept of what it is to play a game, or at least what it is to play a competitive game. This is not essential, though. As long as it is allowed that there are such senses, no matter how central they may be, we need some account of the meanings of these phrases which can vindicate this, and no existing theory seems capable of doing this. Moreover, it seems very plausible that such a content is very much in play in the context of tournament chess as it is currently practiced, too, even if this is not the case for other games. In fact, this suggests the intriguing possibility that which content for ‘play a game’ is most naturally invoked may also depend on the nature of the game in question (note the difference between chess and *Magic: The Gathering* noted above, e.g.) – I discuss this more below in section VI.

It seems equally plausible that someone could be trifling throughout a given game, yet still have played a game for all that. I have in effect already argued for this point in section I. Recall my discussion of the parent who does not try to win yet intuitively plays a game with her child, and the example of the tournament chess player who complains that he should only be scored with a default, rather than a loss, because he never played the game in virtue of not having tried to deliver checkmate. These sorts of examples put enormous pressure on the idea that the inveterate trifler has not played a game at all.

Suits is not without a further rejoinder. This is where his distinction between games and their associated institutions becomes relevant. Suits emphasizes that although the cheat and the trifler are not (when cheating and trifling, anyway) playing a game of chess (e.g.), they are still taking part in the institution of chess. Perhaps the subtlety of the distinction between a given game and the institution associated with that game means that our confidence that the cheat and the trifler are still playing is not well founded. Perhaps we are implicitly confusing taking part in the institution with playing the game, which case the intuitions behind my objection have been debunked.

Perhaps, but this seems very unlikely for reasons that should by now be apparent. Both our intuitions and actual practices seem to presuppose that inveterate cheats deserve to *lose* their games. Plausibly, though, it is a platitude that someone can only lose a game she has played (defaulting is another matter). The same point applies in the case of inveterate triflers. In fact, in the case of trifling we can appeal to the even more platitudinous thesis that you can only win a game you have played. Given that these presuppositions plausibly are built into our existing practices (again, consider the example of a game of tournament chess with an inveterate cheat), it is a cost of any theory of those practices that it entails that these presuppositions are analytically false. Methodologically, we should at least be reluctant to endorse theories of our practices which entail that they are somehow internally incoherent. We therefore have reason to see if we cannot find a theory which does better on these fronts. Fortunately, as I shall argue below (in the final two

sections), we can. First, though, I need to consider other strategies for disarming the two paradoxes.

III. Constitutive rules, regulative rules and the “ethos” of a game.

One tempting strategy for resolving our paradoxes invokes a distinction between different kinds of rules. A distinction between so-called “constitutive rules” of a game and rules of skill is very commonly drawn in discussions of games. On this way of thinking, “A knight on the rim is dim” is a rule of skill in chess, whereas “Bishops move diagonally” is a constitutive rule. The former is meant as a piece of advice; it is supposed to indicate how one plays *well*. The latter, though, is in some sense meant to be constitutive of the game.

In my view, this thesis about the rules being constitutive of the game must be handled very carefully. Here it is important to distinguish the rules constitutive of some specific *token* game (e.g. the final game of the World Series in 2010) from the rules allegedly constitutive of some game *type*. The very idea of constitutive rules in the latter sense is, in my view, problematic. In ordinary language, we are happy to refer to two different token games as instances of the same type even though the apparently constitutive rules of those tokens are different. For example, for a long time the pawns could not move two squares on their first move, but we still think of game played before that time as games of chess – the same game that is played now when pawns are allowed to move two squares on their first move. One could insist on a more fine-grained way of individuating games on which these are different games, but in my view this would very quickly lead to an account which deviates sharply from ordinary thought and discourse about games. However, I put these interesting complexities to one side for now, as they are not essential to the current debate.

In addition to the distinction between constitutive rules and rules of skill, some theorists distinguish both of these from what they call “regulative rules.” A standard way of defining regulative rules is as rules one can violate while still playing the game, but which are such that whenever one violates them, one is subject to certain penalties (see, e.g. D’Agostino 1981). Committing a foul in basketball is a clear example of a regulative rule in this sense.

Assuming some such distinction is viable, it suggests a strategy for disarming the paradox of cheating. One might hold that one only has to follow the constitutive rules of the game to count as playing, whereas one can cheat by violating the merely regulative rules of the game. This would imply that the original paradox may rely on some sort of equivocation.

Admittedly, players do often violate the regulative rules of a game without cheating. The fact that committing a foul in basketball is a paradigm case of violating a regulative rule should make this clear, since deliberately fouling someone to stop the clock is in some sense an entirely acceptable strategic move. Still, plausibly one *can* cheat by violating the regulative rules. For example, if a basketball player *covertly* commits a foul to get a strategic advantage, with the intention of having her foul go undetected, this begins to look more like a case of cheating.

However, for this to help resolve the paradox of cheating, it is not enough that one *can* cheat by violating the regulative rules. It would have to be the case that one can cheat *only* by violating the regulative rules, and that one cannot by violating the constitutive rules. Otherwise, the problem of the inveterate cheater would still loom, so long as our inveterate cheater is always cheating *by* violating constitutive rules of the game. Unfortunately, this strong thesis (that you cannot cheat by violating a constitutive rule) is simply not very plausible.

On any plausible construal of the constitutive rules of a typical token game of chess (as it is played in the early 21st century, anyway), the constitutive rules will include rules like “bishops can only move diagonally.” However, chess hustlers will sometimes try to cheat their opponents in blitz chess by having their bishop “phase shift” midway across the board from the dark squares to the light squares or vice-versa. This is an uncontroversial case of cheating, yet it violates what one might consider to be the paradigm of a constitutive rule.

Moreover, this strategy relies on the problematic idea that to be a regulative rule just is to be a rule with a specific penalty attached to its violation. If a player makes an illegal move in a tournament game of chess then the opponent gets extra time added onto her clock in the first instance, which is surely a penalty. Yet making an illegal move must be the paradigm of the violation of a constitutive rule – moving the knight like a bishop, or whatever.

Before concluding that this approach is entirely misguided, we should consider D’Agostino’s own account. Having argued against what he calls “formalist” accounts which simply draw a distinction between constitutive and regulative rules, D’Agostino argues that we instead need to draw on what he calls the “ethos” of a given game in a given concrete context of play. The ethos of a game is fixed by local conventions, conventions which sometimes involve systematically overlooking certain violations of the formal rules of the game. The way in which referees in professional basketball systematically overlook forms of physical contact which technically are fouls (given the formal rules) is one of D’Agostino’s nice example of how the ethos of a game might function to make certain violations of the formal rules come out as acceptable in some broader sense.

D’Agostino’s discussion suggests another strategy for disarming the paradox of cheating. The ethos of a given game in a given context will, he argues, allow us to distinguish (a) behaviour that is permissible, (b) behaviour that is impermissible but acceptable, and (c) behaviour that is unacceptable. Permissible behaviour will include *both* behaviour which accords with the formal rules *and* behaviour which violates those rules but which the ethos dictates does not warrant the imposition of a penalty. Behaviour which is impermissible and acceptable violates the rules in a way in which the ethos dictates does warrant the imposition of a penalty. Finally, unacceptable behaviour “violates the rules of the game in a way which, according to the ethos of the game, disqualifies its perpetrator as a player of the game.” (D’Agostino 1981: 15)

This account suggests a strategy for disarming the paradox of cheating which is structurally parallel to the strategy deployed by the defender of the distinction between constitutive and regulative rules. This is not surprising, since the ethos of a game in effect provides something functionally equivalent to that distinction.⁶ D’Agostino could allow that cheaters are still playing the game insofar as their cheating consists in covertly violating those rules whose violation is acceptable when not covert.

One problem with D’Agostino’s approach is that he defines playing a game in terms of one’s behaviour being acceptable, but acceptability in turn is defined in terms of whether the ethos of the game dictates that one’s behaviour disqualifies one *as a player of the game*. This means that the definition of what it is to play a game is circular, for we now need an account of the content ‘player of a game’ as it figures in the definition’s claim about the dictates of the ethos of the game.

⁶ D’Agostino prefers this approach to the formalist approach because in his view it provides a better account of the individuation of games.

There may be a way around this, though. D'Agostino's idea seems to be that the rules which do not have specific penalties attached to their violation are the ones which are such that their violation means one is not playing the game. This provides a way to avoid circularity, since we can now define playing a game in terms of following the rules whose violation does not dictate some specific penalty.

Unfortunately, this threatens to render the account of what it is to play a game vacuous in some cases. As I pointed out above, in the context of a tournament game of chess, the rules one would naturally invoke to explain what it is to play the game are *all* ones such that their violation warrants specific penalties according to the ethos of tournament chess. So the only obvious solution to the worry about circularity would in effect entail that many games do not have any constitutive rules – which hardly makes it plausible that what is to play games in general is to follow their constitutive rules!

I therefore conclude that neither the distinction between constitutive and regulative rules, nor D'Agostino's appeal to the "ethos" of a game can help resolve the paradox of cheating. Nor, for that matter, are they well suited to resolve the paradox of trifling, though I have not discussed that here. Perhaps, though, Wittgenstein's famous suggestion that 'game' is a family resemblance term can help avoid these paradoxes; I turn to this suggestion next.

IV. Wittgenstein and 'game' as a family resemblance term.

Both the paradox of cheating and the paradox of trifling arise out of an attempt to articulate a necessary condition for some activity to count as playing a game – a condition which explains why behaviour merely according with the rules of a given game is not sufficient to count as playing a game. However, on one way of understanding Wittgenstein's famous thesis that 'game' is a family resemblance term, this would entail that there are no strictly necessary conditions on some activity's counting as playing a game. Rather than finding necessary conditions, if we look closely at the things we call games Wittgenstein famously claims we will find nothing common to all (apart from their all being games, presumably!), but rather a "complicated network of similarities overlapping and criss-crossing." (Wittgenstein 1958: 32) If Wittgenstein is right about this, then, neither paradox can get off the ground, since both presume non-trivial necessary conditions on some activity's counting as playing a game.

Unfortunately, the very idea of a family resemblance term is a controversial one, and I cannot begin to do justice to the vast literature on this topic in the short space of this section.⁷ Instead I dogmatically lay out what I take to be a familiar and important problem with some orthodox ways of understanding the idea. I underscore this problem by juxtaposing it with my account of how the two paradoxes naturally arise. I then explain how even if those problems with more orthodox understandings could be overcome, there is a way to finesse the idea that 'game' is a family resemblance term which would still allow slightly modified versions of each paradox to arise. Finally, I outline what I take to be the most plausible understanding of family resemblance terms (one developed by Michael Pelczar) and explain why on that understanding, the thesis that 'game' is a family resemblance term simply provides a useful overarching framework for the two paradoxes, rather than dissolving them. I return to this point below (in section VI) where I argue that my positive account can be incorporated into a family resemblance framework – at least, given Pelczar's conception of such terms.

⁷ I intend to discuss this issue in more detail elsewhere in future work.

A standard way of unpacking the idea of a family resemblance term is as follows. Although there are no reductive necessary and sufficient conditions which define a family resemblance term, such terms are not primitive, either. There are certain characteristic features semantically associated with such terms, and insofar as a given object satisfies “enough” of these features, then it falls under the term. What counts as “enough,” though, is vague and on some readings highly context-sensitive as well. The crucial point is that no one of these features is necessary for the application of the term in question.⁸

An equally standard (and in my view, fatal) objection to this conception of family resemblance terms is that the sort of analysis it offers turns out to be logically equivalent to a complex classical analysis couched in terms of necessary and sufficient conditions. There will, after all, be *sufficient* conditions for the application of such terms (leave the issue of vagueness to one side for now). Suppose the four (reductively specified) marks of a given family resemblance term ‘t’ are A,B,C and D. Now suppose that satisfying any three of these is sufficient for falling under ‘t’, and that anything short of that is not sufficient. In that case, we can trivially articulate reductive necessary and sufficient conditions for the application of ‘t’ as follows: In that case, something x falls within the extension of ‘t’ if and only if x is *either* (i) A, B and C, or (ii) B, C and D, or (iii) A, C and D, or (iv) B, C and D, or (v) A, B and D. This objection has been raised more than once.⁹

Admittedly, this schematic counter-example does not fit well with the idea that family resemblance terms are vague, admitting of borderline cases. However, there must be more to being a family resemblance term than vagueness. After all, there is no reason for supposing that the classical approach to analysis (via necessary and sufficient conditions) cannot offer analyses of vague terms which deploy more basic vague terms. Indeed, this is probably commonplace.¹⁰ For example, a definition of ‘animal’ will probably use the concept of being alive, but this concept plausibly admits of borderline cases (e.g. viruses).

Putting these concerns to one side for the moment, it is also not at all obvious that ‘game’ is a family resemblance term in precisely this sense anyway. Quite apart from the move to disjunctive definitions, my original exposition of the paradox of cheating and the paradox of trifling itself set out *arguments* that there are non-trivial necessary conditions one must meet to count as playing a game – namely, following the rules and trying to win. Of course, this did then lead to paradox, but the paradoxes are based on theses which *prima facie* have a lot to recommend them. At any rate, the family resemblance theorist who insists that ‘game’ does not admit of any non-trivial necessary conditions would seem to owe us some account of why the examples discussed in previous sections (e.g. of people accidentally moving the chess pieces in accordance with the rules, or people who know the rules but not the goal of the game following the rules) should not count as playing the game.

Moreover, Wittgenstein’s own discussion was focused on the plurality of different *kinds* of games. The point there seemed to be that no “one size fits all” definition would work for all of the various activities we characterize as playing games. In that case, though, we can easily fashion a more narrowly tailored version of our paradoxes which is entirely consistent with the point. Rather than formulate the paradoxes in terms of games *in general*, we can pick some specific game which plausibly is well understood in terms of its constitutive rules and goal – e.g. chess or draughts.

⁸ This reading is clearly put forward in Bambrough, e.g. See Bambrough 1960: 209-210.

⁹ See Pelczar 1990 and Gert

¹⁰ See Gert 1995 for elaboration of this point. Gert also very usefully discusses how this sort of objection undermines modern approaches inspired by Wittgenstein’s discussion – prototype theory, cluster concept theory and paradigm theory approaches to concept possession.

Even if there are no non-trivial necessary conditions on playing a game *in general*, there might still be non-trivial necessary conditions on playing chess (e.g.). Indeed, it seems pretty clear that there are at least some such conditions. That the game is played with pawns, knights, bishops and rooks, e.g. seems one such plausible condition. If the paradoxes are reformulated in these ways, then it would then be irrelevant that no analogous paradox would tempt us for other games – games which perhaps are not so well understood in terms of constitutive rules and defining goals.

My own view is that there is a much more interesting way to understand family resemblance terms, but one which allows for non-trivial necessary conditions to play a role in the broader theory. On Michel Pelczar’s powerful account (to which I cannot do justice here), family resemblance terms 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; the first two are at least familiar from other discussions.

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.

For present purposes, the main point is that it may be that on Pelczar’s account in a given context of utterance a given family resemblance term may well have non-trivial necessary and sufficient conditions of application. For example, when the topic of conversation is chess, the phrase ‘play a game’ may have a determinate content which can be specified reductively in some way, even though that phrase will have a different content if we are talking about playing Dungeons and Dragons or solitaire, say. There will be discretion in certain cases, too, but that does not mean that once we have fixed the way in which that discretion is used that we do not have scope for necessary conditions. Moreover, some necessary conditions may hold for some terms no matter how the relevant discretion is used – just as in the common law there are limits on discretion, too. I therefore tentatively conclude that the Pelczar approach to understanding family resemblance terms does not preclude specifying necessary conditions for the application of ‘game’ *given* a context of utterance which specifies a topic and a role for discretion. This, though, is in principle all we need to have the logical space to set out the two paradoxes. The appeal to family resemblance theory therefore does not allow us to side-step the two paradoxes – even if we accept the (rightly controversial) thesis that ‘game’ *is* in some useful sense a family resemblance term in the first place. I return to Pelczar’s account below (in section VI).

V. Paradoxes Resolved: Commitments and Activities Defined by Normative Roles

A more promising approach to avoiding the two paradoxes is to move from defining ‘play a game’ in terms of actually following the rules and trying to win to defining it in terms of undertaking a *commitment* to follow the rules and try to win. After all, someone can undertake such commitments without following through on them, so depending on what else is required for a piece of behaviour to count as playing a given game, we may be able to make sense of both cheating and trifling while playing. Moreover, the fact that someone has undertaken such commitments can also elegantly explain why we consider cheating and trifling to warrant criticism. In general we are warranted in criticism when people fail to live up to their commitments.

To be plausible, the notion of a commitment will have to be understood in a way which does not require any highly formal or explicit speech-act to be performed. The commitments constitutive of playing games are very often going to be implicit, but this is not obviously a serious problem. We do, after all, have some intuitive grip on how people can implicitly undertake commitments to one another.¹¹ In any event, I do not have the space here for a detailed account of how we should understand implicit commitments, but shall have to rely on the reader's pre-theoretical conception.

The only account of which I am aware which takes this form is due to A.J. Kreider (see Kreider 2011), who emphasizes his account's ability to explain why the cheater is still playing the game as one of its main virtues. Kreider's implementation of this idea is open to decisive objections, but the basic idea of defining 'play a game' in terms of corresponding commitments is very promising.

Kreider effectively maintains that *all* one must do to play a competitive game is to undertake a *commitment* to follow its rules in an attempt to win the game. Kreider celebrates the fact that on his account undertaking such a commitment is sufficient for playing the game, and not just necessary. Here is Kreider:

What is new in the account presented here is that entering into such commitments is both necessary and sufficient for determining game-playing. Fraleigh may believe that such commitments are necessary...but it is clear that he believes that such agreements are not sufficient...This, however, is a mistake. (Kreider 2011: 61-62)

Kreider's definition elegantly avoids the two paradoxes, simply because one can undertake the relevant commitments without fulfilling them. The cheater violates the commitment to follow the rules, whereas the trifle violates the commitment to try to win. No matter; so long as they both take on these commitments, they will still count as playing the game on this account.

Unfortunately, Kreider's definition is clearly too broad. If I formally undertake a commitment to play a game of chess but do not even show up at the board, much less make a move, then I have not played a game, but merely defaulted. It is not as if a promise to play someone a game is literally unbreakable as a matter of the logic of 'play a game'.¹²

This objection might seem churlish. All Kreider needs to do, one might object, is to specify which forms of behaviour must follow from one's commitments in order for one to count as playing a

¹¹ Another closely related issue discussed by Kreider (Kreider 2011: 63) is how this account fits with the idea that young children can play games. Although this is an interesting question, it is not clear that it poses a serious problem for an account which requires commitments to count as playing a game, especially once it is allowed that 'play a game' is a family resemblance notion, as I shall discuss below. First, once we allow for implicit commitments, it may be that children are able to undertake commitments at a very young age. Second, insofar as young children cannot undertake commitments yet, it may well simply be unclear whether we should say they are *really* playing games yet, as opposed to learning how to play games by imitating those who properly play the game in question. It is perhaps instructive here to consider Wilfrid Sellars' discussion (in Sellars 1963). In that context, Sellars argues that such an intermediate stage of development is crucial to cashing out the idea that learning a first language is well modelled by learning the rules of a game in a way that avoids an infinite regress of rule deployment. Third, even if we insist that children who cannot yet undertake commitments can literally count as playing games, they may not count as playing games in precisely the same sense that competent adults count as playing games. Anyone who insisted that such children must count as playing games in *precisely* the same sense as adults is in the grip of a theory; our pre-theoretical intuitions are simply not anywhere near this fine-grained. I discuss this second suggestion in more detail in section VI below.

¹² Compare Royce 2013, who presses this objection and others in the same spirit at some length against Kreider.

game. Undertaking the relevant commitment is just a necessary condition, but not a sufficient one. Let us grant that Kreider himself massively overreached in claiming that the relevant commitments are sufficient for playing a game. Still, one might insist, his basic idea of making commitments central to our conception of game playing is on the right track.

In fact, I think this is basically correct, but putting it in this way dramatically understates how far we are from a complete theory. In particular, this rather large concession means that much of the hard work of explaining what it is to play a game is yet to be done, and it is not obvious from Kreider's approach that it *can* be done in a way that avoids our two paradoxes. For, unfortunately, there is reason to worry that whatever other necessary conditions we add to Kreider's account will reintroduce the two paradoxes.

Indeed, it might seem like we will be *driven* to embrace necessary conditions which generate these paradoxes for just the reasons given for finding the paradoxes compelling in the first place – the reasons given in section I. Suppose I have undertaken a commitment to play a given game. As we have seen, the problem is that not all of the behaviour which follows from this commitment should count as game playing. Driving to the venue is not game playing, nor is having dinner the night before, even if this is done to make sure I have enough energy for the game. Crucially, the most obvious and *prima facie* tempting way to distinguish behaviour which constitutes game playing from behaviour which does not would seem to be to require that I am following the relevant rules and trying to win. This, though, brings us directly back to the two paradoxes.

However, an account couched in terms of such commitments ultimately can help us develop an analysis which evades both paradoxes. Before I can lay out my positive account, though, I need to draw some further distinctions. First, for any game G, we need to distinguish between the content of 'played G' and 'played a game of G'. Plausibly, these can come apart. Someone, e.g., can count as having played chess without having played *a game of chess* if the game is cut short, for example, by external circumstances and never resumed. A vivid example would be one in which one player dies of a heart attack during the game, but there are of course many ways in which this can happen. In general, for any given game G, it seems that 'played G' picks out an activity which can be complete in itself at any given moment in which game-playing activity is taking place. By contrast, the occurrence of a count noun ('a game') in the phrase 'played a game of G' seems to apply to an activity only if somehow reaches its intended conclusion.¹³ I begin my analysis here with a focus on the more demanding concept ('played a game of G') and then consider how to extend the analysis to the less demanding concept ('played G').

Second, we must distinguish between what it takes for a game to have been played over some span of time, on the one hand, and what must be true of some specific individual at a given moment in time for him to count as playing at that specific moment. Being clear about this distinction is essential to taking our two paradoxes seriously. Ideally, we need some account which can vindicate not only the claim that a game was played by a cheater, but one which can vindicate the claim that someone can in at least some useful sense count as playing the game *while he is cheating*. Rival theories (like Suits') can accommodate the former, but not the latter. For reasons that shall become apparent below, I begin with an account of what it takes for a game to be played over some span of time and then return to the crucial question of what makes it true of some individual at some specific point in time that he is playing at that moment.

¹³ Compare Kenny 1963: 171- 186, who draws a canonical distinction between stative and non-stative verbs – a distinction which in turn draws on ideas from Aristotle. See also Chrisman 2012.

Third, we need to distinguish competitive from non-competitive games. The two paradoxes both presume that games are the sorts of thing one can win or lose, but for non-competitive games (e.g. Hide and Seek and Charades) this assumption is not warranted. Moreover, it may well be that the content of ‘play a game’ shifts depending on whether the topic of conversation is a competitive or a non-competitive game – a hypothesis I pursue below in section VI (mixed contexts will of course raise further complications). I shall therefore here focus my attention on how we might understand the content of ‘play a game’ in contexts in which it is clear that only competitive games are conversationally salient.

My focus on competitive games, in turn, leads me to offer an analysis which takes as its basic unit *groups* of people. After all, competitive games as I conceive them are always played by more than one person. Existing analyses in the literature tend to be uncritically couched in individualistic terms, but in my view this is methodologically dubious. Once we have an account of what it takes for some group of people to count as playing a game together at a given time, it is not hard to see how to derive from this that each of those individuals is playing that game at that time – this will follow trivially. Whereas starting with an analysis of what it is for some specific individual to play a game at a given time without at the same time defining what it is for the other players to play with him or her may make it hard indeed to explain what it takes for some group of people to count as playing a game with each other. Suits’ definition provides a good illustration of this individualistic bias, and it is not entirely trivial to work out how his account bears on when two people are playing a game together. The focus on groups is also natural given the substance of my account, which puts the undertaking of commitments to the other players at its centre.

What, though, must also be true of some group of people who have undertaken the right mutual commitments for them to count as having played a game? At a bare minimum, they must all in some sense “show up” for the game. What counts as “showing up” will vary from game to game, and the existence of online games makes it clear that being at some specific place will not always be essential. However, the nature of the game and the implicit terms of the agreement should fix some fairly clear sense of what is required for the players to count as showing up.

Even if all of the players show up, a game will not be played if they are not quorum. What counts as quorum of course also depends on the game and plausibly also depends on the locally specified ethos of the game, in roughly D’Agostino’s sense, too. For a proper game of American football, both sides must field eleven players, for example, but this is clearly not essential for a pick-up game of American football played in some less formal setting. Although working out in any given context just what it takes for the players present to be quorum, this is often clear (board games often explicitly say things like “for between 4 and 8 players,” e.g.) and is also not a trivial condition.

What more must happen for a game to be played for a given stretch of time by a given set of people? Having undertaken the relevant mutual commitments, showing up and being quorum for a game plausibly is not sufficient for a game to have been played. Here we must be careful, since the strong temptation will now be to add that the players follow the rules while trying to win throughout this period of time. This, though, would land us directly back in our two paradoxes.

At this point it might be tempting to insist that instead of actually following the rules, one must simply *present oneself as if* following them.¹⁴ After all, this is what the cheat does, and it is the cheat’s behaviour we intend to accommodate. Unfortunately, on any normal interpretation of ‘present oneself as if following the rules’ one need not present oneself in this way to count as playing. Most game players are not self-conscious enough to be thinking about how they are

¹⁴ Thanks to Nadia Mehdi for useful discussion of this proposal.

presenting themselves as they play (unless that is part of the game, as with poker, e.g.). Instead they lose themselves in the game itself. Their attention is focused on actually following the rules and winning the game, and not how they are seen by others.

Moreover, the proposal is open to direct counter-examples. In particular, someone intuitively can play a game while presenting herself as if *violating* the rules. This may be surprising, but consider the following hypothetical example. Suppose Miley and Taylor and have agreed to play a game of Scrabble. Taylor secretly hates Miley's company, but does not have the courage to confront her about this. However, she secretly plots to get caught cheating in their game of Scrabble, thereby making Miley despise her so she will no longer have to spend time with her – but without having to confront her directly, tell her how annoying she is, etc. When the time comes, though, Miley is too honest and cannot bring herself to cheat. So she instead *pretends* to cheat, in the hopes that Miley will think she is cheating. She pretends to consult the Scrabble Player's Dictionary during the game when she actually is just flipping through it at random. However, Taylor is too focused on her own tiles to notice this, or perhaps she does not even realize it is against the rules. Either way, Miley's behaviour goes undetected. This goes on until they have all played all their tiles and Miley is declared the winner. Intuitively, they have played a game of Scrabble, yet Miley certainly did not present herself as if she were following the rules of the game throughout – indeed, she did just the opposite.

So simply moving from a definition couched in terms of actually following the rules to one couched in terms of merely presenting oneself as if one is following them is not promising. What, then, must one do apart from “showing up” in order to play a competitive game you have agreed to play with others?

My answer is that there must have been “enough” rule-following activity by the players, where what counts as “enough” depends on the context of play. In particular, it will depend on the game being played and what D'Agostino would call the “ethos” of the game in the context in which it is played. In certain contexts, very little rule-following behaviour will be needed for a game to count as having been played. Indeed, as I shall argue below, in some contexts, perhaps no rule-following behaviour whatsoever may be required.

To be clear, then, my proposal is that ‘play a game’ is context-sensitive in at least two ways. First, the content of the phrase will vary depending on whether the game in question is a competitive or a non-competitive one; I return to what we should say about non-competitive games below (in section VI). Second, how much rule-following behaviour is required for a game to be played will depend on the topic – that is, on what game and what context of play (what “ethos”) is under discussion.¹⁵

Consider again my examples of inveterate cheats. In the case of Gaylord Perry, there was still quite a lot of rule-following – even by Perry himself. Perry did, after all, pitch from the designated pitcher's mound, he waited until there was a batter at home plate before pitching, etc. Moreover, of course, the other players were presumably all following the rules. How far could the players have deviated from the rules before we would say no game of baseball had been played? I do not know enough about the rules and ethos of Major League baseball to say with confidence, but I expect that there is some minimum amount of rule-following activity that needs to take place before a game has been played. For all I know, that minimum might be very low, though.

¹⁵ A further complication is that we can use phrases like ‘played several games’ to quantify over very different sorts of games – competitive and non-competitive alike, and very different competitive games as well. I discuss this problem (which is a problem for everyone, in my view) in the following section.

Furthermore, because one team can forfeit after the game has started, the real question is how much rule-following must take place before we can say a game was played (with a winner and a loser) rather than the game having effectively been cancelled due to lack of play. In my view, this will be dictated by the rules of the game and the locally specified ethos. There may, of course, be considerable indeterminacy about this, too. After all, our conventions are generally only as determinate as they need to be, and cases in which both sides show up but fail to engage in normal play are with good reason very rare indeed.

Consider again my inveterate poker cheat, who uses a marked deck throughout. To make the example more interesting, we can add that it is a two player game and that the opponent is also cheating – by using a hidden camera which transmits the opponent’s hand to a confederate who then communicates with the player through a hidden earpiece. In this case, we have rule-breaking throughout the game, yet a game has been played. Here, too, though, although we have rule-breaking throughout, we also have considerable rule-following too. After all, the players do deal the correct number of cards to each other, bet when and only when they are supposed to bet, declare who wins the pot in accordance with the rules, etc. If none of these conditions were met then it would be less clear that a game of poker was played. If, for example, the two people showed up and, having agreed to play poker, actually followed the rules of Blackjack throughout we would correctly judge that they had not in fact played a game of poker.¹⁶

One thing these examples bring out is that one and the same piece of behaviour can count as rule-following and rule-violating at one and the same time, given that games have multiple constitutive rules such that you can follow one while breaking another. This is important to bear in mind when thinking about the inveterate cheat. We can, in fact, now distinguish at least two notions of ‘inveterate cheat’. First, there is the sense in which someone is an inveterate cheat if and only if his every move in the game violates some rule. Second, there is the much stronger notion in which someone is an inveterate cheat only if his every move both violates some rule and does not follow *any* of the rules of the game. Once this distinction is made explicit, it should be clear that it is only in the weaker of these two senses that it is clear that an inveterate cheat can count as playing a game.

Perhaps my example of the drunken chess player (from section II above) who makes two illegal moves and thereby loses is an instance of an inveterate cheat in the stronger sense who nonetheless counts as having played a game. Both of the cheat’s moves are illegal – moving a pawn three squares on the first move, say. He is clearly violating the rules of chess and it is not at all clear that he is following any of them. Yet given that two illegal moves entail a loss, it would seem that a game has been played.

Actually, the example is more complicated than I originally presented it as being. The official rules of chess as promulgated by FIDE seem not to give a clear verdict in this sort of case. On the one hand, the rules clearly state that two illegal moves by one player mean that the player loses the game. On the other hand, the rules also clearly state that a game of chess has not been played unless each player has played at least one move – and in my hypothetical scenario this is not the case. In actual practice, there would be room for discretion by the arbiter as to whether to score such a game as a loss (with implications for both players’ ratings) or a default.¹⁷

¹⁶ Compare Lehman 1981: 43, and Green 2004: 150.

¹⁷ I consulted with a professional chess arbiter about this and he suggested that in his opinion which of these verdicts would be most apt would depend on whether the person who made the illegal moves was drunk or otherwise incapacitated due to substance abuse, or whether instead he was suffering from some sort of mental illness.

The correct interpretation of the current official rules of chess is of course not essential here. The main point is that in any event there *could* be a game G such that a game of G can be played even though the players never engaged in *any* rule-following activity. If there were such a game then it would be enough that the players both turned up at the designated place at the designated time.

Finally, a competitive game is only played if the rules plus the “state of play” determine a result – win, loss or draw. Remember the distinction between whether someone has played G (for some game G) and whether some group of people have played *a game of G*. The latter requires that the game come to one of its conventionally intended outcomes. In the example just given a result (loss for the cheater) might follow from the fact that two illegal moves were played. What determines whether the game has reached such an outcome will, though, be fixed by the rules and the local ethos of the game.

Actually, the “state of play” is itself a delicate notion in this context too. Typically this is defined in terms of the moves that have so far been made, and the resulting score, legitimate moves available to each player, and so on. In fact, though, there is both a subjective and an objective way of understanding this notion, and the distinction matters to how we think about cases involving cheating. The two readings are nicely brought into sharp relief in David Lewis’s classic text, “Scorekeeping in a Language Game.” (Lewis 1979) Lewis suggests that we might understand the state of play (including but not limited to the score) as defined by the rules and the moves actually played up until that point. Alternatively, we might define the state of play in terms of what some canonical “scoreboard” says the state of play is. The canonical scoreboard might be the one on public display for the fans, or it might be an invisible one in the umpire’s head, as far as Lewis is concerned. The crucial point is that on the latter interpretation, however it is cashed out, it is impossible for the scoreboard to fail to give the correct score.

There is a sense in which the second sense of “state of play” is more subjective. In effect, the state of play in the second sense just is what the score in the first sense is represented as being by some privileged “scoreboard.” In some cases the idea that there is any such privileged scoreboard is tricky. In formal games with referees (or judges) it is perhaps not too hard to work out what should count as the relevant scoreboard. In informal games with no designated referee or judge, the scoreboard presumably must be understood in terms of what the players agree (perhaps implicitly) the score is at any given time. Obviously this means that in informal contexts there is more room for interminable dispute, of course, but that is borne out by practice! We should not demand more determinacy of our theory than is reflected in the practice.

To see how these two notions of a state of play can come apart, just think about a case in which the instant replay makes it clear that the referee’s verdict reflects a mistaken factual judgment. Suppose, for example, that an umpire in a baseball game calls someone as safe at home base when the instant replay makes it clear he was actually out. On the first notion, the player was actually out and the umpire’s judgment gets the score wrong. On the second notion of a state of play, this is incoherent – the umpire’s judgment defines the state of play.

Lewis himself does not take a stand on which of these two readings of ‘score’ (or ‘state of play’ to use my broader terminology) is correct. In fact, I am not sure there is any fact of the matter about which is correct, in general. In some cases, the practice strongly supports the more subjective conception. Even if it is very clear that the umpire’s judgment was mistaken, the practice will record the player as safe – and this may even determine the outcome of the game, for that matter. On the other hand, fans will with some cause complain about who “really” won such a game. I am

happy to allow that ‘state of play’ and ‘score’ are ambiguous as between these two different readings. All I am after here is *one* plausible definition of what it is to play a game such that playing a game in that sense is such that one can simultaneously play the game and cheat (and simultaneously play the game and trifle). The fact that there are other legitimate senses in which this is impossible is therefore no objection to my view, or my proposed resolution of the two paradoxes.

Here, then, is my proposed analysis of what it is for a group of people to play a competitive game over some specific span of time:

A set of people *S* (comprised of individuals *s*₁-*s*_{*n*}) have played a competitive game *G* in context *C* from *t*₁ to *t*_{*n*} if and only if:

- (1) Each member *m* of *S* has at some point prior to *t*₁ undertaken a commitment to the other members of *S* to follow rules *R* (where rules *R* are the constitutive rules of *G*) in context *C* from time *t*₁ until either (a) the state of play dictates that *m*’s participation in the game is over, or (b) the members of *S* agree to stop following rules *R* (that is, they decide to quit playing).
- (2) Each member *m* of *S* has at some point prior to *t*₁ undertaken a commitment to the other members of *S* to try to achieve what is in fact the lusory goal (in Suits’s sense) of *G*.
- (3) Each member of *s* shows up at the designated context of play *C* at the designated time *t*₁.
- (4) None of the members of *S* renounce their commitments (see 1 and 2 above) for the duration of *t*₁-*t*_{*n*}.
- (5) Each member of *S* engages in *enough* rule-following activity from *t*₁-*t*_{*n*}, where what counts as enough is context-dependent in the ways discussed in the text above.
- (6) At *t*_{*n*} the state of play is such that no further play is allowed, and a result is determined (e.g. someone won, someone lost or came in second (third, etc.), or there was a draw, or whatever).

This definition is already rather complex, with several moving parts, so I am going to overlook yet another complication here – that games can have adjournments. In the interest of not making the discussion here unmanageably complex, I am simply going to put the question of how to revise this definition to allow for the period during which a game is played to be broken up into discrete temporal chunks to one side. To be clear, though, I think this is a question of detail, and not one which threatens the core of the analysis.

Although this definition is couched in terms of groups of people it is easy enough to see how it fixes the truth conditions of claims about whether individuals have played the relevant game. It also provides a useful framework for the distinction drawn above between ‘playing a game of *G*’ and ‘playing *G*’. For it to be true that a group of people played *G* (for some game *G*) without necessarily finishing and thereby playing a game of *G*, the first five conditions must be met but the sixth (which requires that a result is fixed) need not be.

How does this analysis bear on our two paradoxes? Insofar as the paradox of trifling goes, the account in no way entails that one must try to win to count as playing the game. One must undertake a commitment to try to achieve the lusory goal (which constitutes winning), but nothing in this analysis entails that one follows through on that commitment. The account resolves this paradox simply by denying that one must try to win to count as playing the game.

What, though, about the intuitive data that I argued supported the thesis that one must try to win to count as playing a (competitive) game in the first place? Does that data not undermine the plausibility of my analysis, insofar as the analysis entails the rejection of this supposed necessary connection?

No. Recall how I motivated this necessary connection. I invited the reader to consider two people who knew the rules of chess and followed those rules until a checkmate position was on the board. However, they were not trying to deliver checkmate, but instead trying only to produce pretty patterns. Indeed, I suggested that they might not even have the idea of checkmate.

With the above analysis in hand, it should be clear that this thought experiment is insufficiently described in at least two important respects. First, do the players have the concept of checkmate? If they do not then they presumably did not have that concept before they started playing the game. In that case, though, they could not have undertaken a *commitment* to try to deliver checkmate. This, though, entails that condition 2 of my analysis is not met, and so they are not playing a game of chess.

Suppose they do have the concept of checkmate. We now must ask whether they undertook a commitment to try to deliver checkmate while following the rules of the game. If they did not then obviously they do not count as playing a game of chess on my account. What if they *did* undertake such a commitment? Even in that case, it does not follow that they have played a game of chess. Condition 4 of my analysis requires that the relevant commitments must not be renounced during the relevant time span. Insofar as they made it evident to one another that all they cared about was producing pretty patterns (while following the rules to do so), they will count as having renounced their commitments, and so again will not count as playing a game of chess.

Only in the scenario in which they *have* undertaken a commitment to follow the rules and deliver checkmate, and where they do not renounce these commitments (even implicitly) will it follow from my analysis that they are playing chess. In this very specific version of the case, though, it does seem eminently plausible that they are playing a game of chess – at least, it seems plausible that they are *in some sense* playing a game of chess – just very badly. As I will explain below, there is another sense in which they are not playing a game of chess; this reflects another way in which ‘play a game’ can have different contents in different contexts, or so I shall argue. My hope is that any residual intuition that the two people are not playing a game of chess in this version of my original example can be accommodated by this potential equivocation.

What about the paradox of cheating? Nothing in the proposed analysis entails that one must follow the rules to play a game. Granted, the analysis does require that the players together engage in *enough* rule-following activity, where what counts as enough is fixed by context. This, though, is a very minimal requirement. It will probably always be consistent with the inveterate cheat playing the game on the weak reading of ‘inveterate cheat’. It may even be consistent in some extreme cases with the inveterate cheat playing the game in the strong sense of that phrase.

Insofar as the analysis on offer is plausible, we have a principled basis for rejecting the thesis that one must follow the rules of a given game to count as playing that game. This, of course, is sufficient to disarm the paradox of cheating. Moreover, just as with the paradox of trifling we can now plausibly diagnose the intuitions that seemed to drive us to the thesis that one cannot count as playing a game. Recall that in my main example of the accidental chess “players,” the characters did not even know the rules but just moved the pieces in accord with the rules by sheer accident. Insofar as the players meet the analysis offered above, though, they must have undertaken a commitment to follow the rules *and* have engaged in “enough” rule-following behaviour. As I intend the first condition to be read, it requires that the players know the rules – at least, that they know the most important rules reasonably well. The intention to follow “rules R” must then be an intention to follow the relevant rules *de re*, and not a mere intention *de dicto* to follow the rules of chess, whatever those may be. Just how well one must know the rules in order to count as playing a game is not an easy question to know how to answer, and I suspect that our concept of playing a game is not terribly determinate on this point. The fact that in most cases the players must also engage in “enough” rule-following activity, where this will typically not be entirely vacuous, also should defeat the idea that it is some sort of sheer accident that their behaviour accords with the rules.¹⁸ In any event, it is not hard to see why the analysis on offer can exclude cases in which two people who know nothing about the rules and who move some chess pieces at random will not count as playing a game even if their moves do miraculously accord with the rules of chess.

However, this victory may well seem to have come too cheap for two related reasons. First, my account is not yet complete. I have so far only said what makes it true that some group of people have played a game over some span of time. I have not yet said what makes it true that some individual is playing the game at some specific moment within that span of time. Perhaps I can count as playing a game for a period of time even if I am not engaged in game playing behaviour throughout that period. Second, this incompleteness suggests a worry about my proposed resolution of the two paradoxes. Perhaps it is only when the players are actually following the rules and trying to win that their behaviour counts as game-playing behaviour. Indeed, it seems like there had better be *some* such distinction between that sort of behaviour and the behaviour of the cheat and the trifler. For there is a sense in which that sort of behaviour is game playing par excellence – the cheat and the trifler may count as playing the game in some sense, but at the same time they fall short of our intuitive paradigm of game players. If this is right, though, it may turn out that in the more central sense of ‘play a game’ that it is impossible to play a game and *at the very same time* be cheating or trifling. Yet in setting out the two paradoxes, I insisted that we needed to accommodate that very possibility.

The key to meeting this objection is to distinguish two ways of understanding talk of certain kinds of activities – activities which are such that to engage in one of them one must occupy a certain socially defined role which brings with it certain rights and obligations. In one sense, to count in engaging in such activities at a given time, one must simply occupy the relevant role at that time. In another sense, though, to count in engaging in such activities at a given time one must both occupy the relevant role *and* actively engage in rule-following behaviour where the relevant rules are the ones which define that role. Let me illustrate this distinction with a few examples before returning to the case of game playing.

Suppose someone is enrolled as a doctoral student in philosophy from 2009 until 2014. In one sense of ‘pursuing a PhD’ this is enough to make it true that from 2009 until 2014 they were pursuing a PhD. If, at the end of 2014, we asked how long such a person had been pursuing their PhD it would be natural simply to say, “five years.” In another sense, though, we might want to say

¹⁸ Compare Kreider 2011: 67-68.

that the person has not been pursuing their PhD for five years. Suppose that when I asked, I say that the student has been pursuing their PhD for five years, throughout the period from 2009 until 2014. Suppose that my interlocutor then sarcastically replied, “Oh really? Even when they were eating sleeping? Or when they went to the cinema, read a novel, visited their family, and when took that long holiday at the beach?”

Although the reply is pedantic and annoying, it makes a sound point all the same. In fact, the reply is annoying precisely because it makes a sound point. There is a more demanding sense of ‘pursuing a PhD’ in which the student counts as pursuing a PhD at a given point in time only if he is engaging in behaviour which is constitutive of trying to complete his degree – doing research, writing or revising a chapter, reviewing his notes, or whatever. At the same time, there is a sense in which he was pursuing his PhD throughout that period. In this much more minimal sense, all that must be true of him for him to count as pursuing his PhD is that he was enrolled in a doctoral program at that time. Because this sense of ‘pursuing a PhD’ is so minimal, it is intuitive to suppose that other sense (which requires actually doing something in service of completing one’s PhD) is somehow the more central or paradigmatic notion of pursuing a PhD.

Call these different senses the ‘normative role sense’ and the ‘rule-following sense’, respectively. Once this distinction is on the table, it should be clear that it cuts across a wide range of activity types. Indeed, plausibly any sort of activity which is defined in terms of occupying some norm-governed social role admits of some such distinction. We can make the same point with chairing a department meeting, teaching a class, raising a child, and indefinitely many other activities defined in terms of the occupation of some normatively defined role.

Most importantly for present purposes, the distinction applies to locutions of the form ‘play a game of G’ for any competitive game G. Given the analysis offered above, to be a game player at a certain point of time, it is enough that someone has undertaken the relevant commitments with like-minded others, shown up at the right time and that (between them) they have engaged in “enough” rule-following activity during the surrounding span of time. From the fact that someone is a game player (for some specific game G) at a given point in time, it follows *in the normative role* sense that he is playing G at that time. Crucially, this is consistent with its not being the case that he is playing the game at that point in time in the rule-following sense. Consider some examples which illustrate this point.

Suppose I am playing a game of tournament chess but I get up to go to the restroom. My clock is still ticking, the game is still in progress, so I am still playing a game of chess at that time even though I am not engaged in any rule-following behaviour of the relevant kind. I am therefore playing chess in the normative role sense but not in the rule-following sense. You might object that I will typically at least still be thinking about the position while I go to the restroom. Perhaps, but I might instead be thinking about where I will go for dinner that evening or any of a number of other topics that might pop into my head. Indeed, I might have gotten up from the board in part precisely to take my mind off the game altogether for a brief period of time so I could return to the board refreshed and better able to concentrate. Note, moreover, that I am still bound by the rules of the game even when I get up from the board in this way, and I do therefore still occupy a normatively defined role. If, for example, I covertly consulted a computer for advice on the best move in my game during my jaunt to the restroom then I would be guilty of cheating.

Indeed, one need not even get up from the board to fail to engage in any of the relevant sort of rule-following activity. I might sit at the board and simply be daydreaming. I might be focused on adjusting my chair, or enjoying a sip of my coffee at a given moment, with no thought whatsoever

of the game at that specific instant in time. Still, in the more minimal “normative role” sense I am still playing chess throughout the period during which the game takes place, even in these moments in which I am not engaged in any sort of rule-following activity of the relevant kind – even where this is construed broadly enough to include mental activity.

The point is easy to make in other contexts too. Anyone who has been to many Little League baseball games will have seen how easily distracted very young children are. The second baseman might be building a sand castle, while the left fielder might be counting blades of grass or simply daydreaming with no thought of the game whatsoever. Still, in the normative role sense they are playing baseball at this time. They are playing *poorly*, but they are still playing.

It should by now be easy to see how this distinction sheds instructive light on our two paradoxes. In each case, insofar as each of the theses which makes up the paradox is plausibly true there is an equivocation between the theses on the phrase ‘play a game’, thus avoiding the looming contradiction. It is true *in the rule-following sense* of ‘play a game’ that insofar as someone plays a game at time *t* that she is following the constitutive rules of the game at *t*. However, it is true *only in the normative role* sense of ‘playing a game’ that insofar as someone is cheating they are playing the game at that time. As this sense of ‘playing a game’ does not entail rule-following behaviour at that time, cheating does not entail both breaking and following the rules at one and the same time, thus our paradox is avoided.

The paradox of trifling is even easier to avoid on my account. For in neither of the senses I have so far articulated is it true that one must actually seek the lusory goal of the game at a given moment in time to count as playing the game at that time. I think this does capture one important sense of ‘play a game’, but someone might object that to play a game in the most full-blooded sense, one must not only follow the rules at that moment in time, one must do so in an effort to achieve the lusory goal. Perhaps this is right. Even if this is conceded, the paradox of trifling can be resolved. For it will only be in the robust sense of ‘play a game’ (which now requires both rule-following and pursuit of the lusory goal) that one must try to achieve the lusory goal of the game to count as playing the game at a given moment in time. The intuition that one can be a trifler but still count as playing the game can be accommodated because one is playing the game in the more minimal normative role sense of ‘play a game’.

One might at this stage worry that all of the work in my solution to the two paradoxes is being done by the appeal to this apparently much broader distinction between normative role and rule-following conceptions of activities. This might seem to make my positive account of what it is to play a (competitive) game unnecessary. One might also worry that the accounts I criticized in previous sections could also help themselves to this distinction to disarm the two paradoxes.¹⁹

The problem with this objection is that not all activity types plausibly admit of this distinction. While pursuing a PhD and chairing a meeting plausibly do admit of such a distinction, zipping up one’s trousers and running for one’s life do not, in my view. Nor does the distinction automatically apply whenever one has an activity which essentially involves rule-following. For example, to my ear at least, the distinction seems very strained in the case of an activity like following a recipe. While I am (in one sense) still playing a game of chess when I take a break and go to the toilet, I do not think I am still following a recipe when I take a break from my cooking to go to the toilet. My hunch is that this reflects the fact that there is a sense of ‘play a game’ which essentially involves undertaking an interpersonal commitment, and that this is how the normative role conception gets

¹⁹ Thanks to Sebastian Koehler for useful discussion of this objection.

a foothold. Since zipping one's trousers and following a recipe lack this essential interpersonal element, this may explain why the distinction does not carry over.

Insofar as this is right, then we need some account of what it is to play a game in each of the two senses on offer for at least two reasons. First, this is a useful strategy (perhaps not the only one, but a viable one) for establishing that the distinction really does apply in this case. Note that formalist theories like Suits' might make 'play a game' look very much like 'follow a recipe' in which case it would at the very least be unclear whether the needed distinction carries over. Second, spelling out the details of what it is to play a game in each of the two senses provides a deeper theoretical understanding of what it is to play a game, which is of independent interest.

My analysis has so far has been explicitly restricted to competitive games. A more satisfying account would be broader and allow us to see how talk of playing games should be understood more generally. In the next section I return to the idea that 'play a game' and cognates might be family resemblance words/phrases and argue that the analysis I have so far defended can plausibly be embedded in a broader family resemblance framework.

VI. VI. Family Resemblances Revisited

The analysis of 'play a game' and cognates developed above was restricted to competitive multi-player games (that is, games which admit of winners and losers). In this section, I want to provide a broader semantic framework which can handle both discourse about competitive multi-player games, but also non-competitive multi-player games as well as single-player games. In particular, I want to explore how Michael Pelczar's account of family resemblance words and phrases might usefully be deployed to provide a broader account of the meaning of 'play a game' and cognates. Before developing this idea I should briefly explain why the positive account developed in section V cannot be fully general.

On both the normative role and rule-following senses of 'play a game' laid out above, it is a necessary condition for someone to count as playing a game that they undertake a commitment, where this is understood as some sort of (often implicit) speech-act. For many non-competitive games, it is simply not plausible that one must undertake such a commitment to count as playing the game. In particular, this seems very strained for various single-person games. Of course, someone very enamoured of the theory laid out above could insist that when someone plays e.g. Solitaire that they thereby must have undertaken a commitment to themselves to follow the rules, and that something similar holds up for all other single-player games. To my mind, though, this is highly strained insofar so long as we continue to understand 'undertake a commitment' as roughly like making a promise. Perhaps people do sometimes make promises to themselves, though this is in some ways a problematic idea. Perhaps New Year's Resolutions are like this in some cases. However, the idea that one makes a promise to oneself to follow the rules *whenever* one plays a single-player game of some kind is simply not plausible.

However, as suggested in the previous section, we sometimes use 'undertake a commitment' in a somewhat different sense – one which need not involve speech-acts at all and which is not essentially interpersonal. Here I have in mind the sense invoked when we say someone has undertaken a commitment *simply by setting themselves some goal*. When I form an intention to learn to speak French fluently, or to get a paper published in a particular journal, there is a natural sense in which I have undertaken a commitment to do those things. Perhaps we should say that when one plays a single-player (I return to the broader category of non-competitive games below) game that this consists, in part, in undertaking a commitment in this sense.

A commitment 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 not sufficient to play the game, intuitively. One must also act on the commitment in pursuit of the relevant goal.

The attentive reader will by now have noticed that I have in effect reinvented Bernard Suits' classic definition of 'play a game' from *The Grasshopper*, or something very close to it. The only element of his definition I have not incorporated into the definition developed above is that the game player must accept the rules "just because" doing so makes the resulting sort of activity possible. I have not incorporated this particular element of Suits' theory here because on any natural interpretation this element of his theory is problematic. Admittedly, what Suits actually means by this clause does require a fairly unnatural reading of the plain language of the clause, so this may provide no objection to his view as he intends it to be understood. Unfortunately, a careful discussion of exactly how Suits thinks this clause of his definition should actually be understood, and whether it is defensible so understood, would require a level of detailed Suits exegesis which would require too much of a tangent. I will, though, briefly return to this issue below to try to explain how my proposed account can preserve what is plausible about this clause of Suits' theory as he understands the clause.

In fact, the similarity between the account emerging in the previous paragraphs and the core ideas of Suits' theory is no accident. In my view, the objections to Suits' insightful account associated with the paradox of cheating and the paradox of trifling are forceful only in the case of *competitive* games. In the case of a game like *Solitaire*, it is actually not implausible to say that someone who deliberately violates the rules has simply stopped playing the game. Similarly, if someone stops trying to produce four rows of cards of the relevant sort, it is not implausible to say that the person has thereby stopped playing *Solitaire*. In a game like *Solitaire* someone who stops following the rules or stops trying to achieve the relevant goal just is not playing the game any longer.

It is interesting to think about why we use the same word ('game') for activities which differ in these ways. Presumably it is because they are intuitively similar in certain important dimensions. Both involve constitutive rules and goals, and commitments to those rules and goals. They differ in the sense of 'commitment' in which they involve commitments to those rules and goals, and they also differ in the precise way in which they entail that playing a game is related to the relevant rules and goals. Still, the two notions are similar enough that one could see how they might reasonably be said to bear a "family resemblance" to one another, in some intuitive sense.

Nor is it hard to see why we might want a slightly different notion of 'play a game' in the case of non-competitive single-player games. Because they are single-player, the notion of an interpersonal commitment of the sort which figured in my positive account above simply does not arise. Because they are non-competitive, the idea of cheating also does not get a foothold. It may be that in the case of competitive games we need a way of holding people who break the rules accountable in a way that makes it natural to say that they are still playing the game when they cheat – which is why e.g. we can correctly say they lost the game rather than having defaulted in virtue of cheating, e.g. Whereas in the case of non-competitive single-player games, there would have been no pressure to make logical space for this characterization of the cheater. Because playing such games

does not raise the same interpersonal moral issues as competitive games, the notion of ‘play a game’ which has evolved to go along with such games is different in predictable ways.

I was therefore initially tempted to hold that ‘play a game’ is ambiguous as between the two senses I articulated above and something very close to the sense provided by Bernard Suits’ definition. In particular, my initial hypothesis was that the default interpretation of ‘play a game’ in the case of a competitive game would be the normative role sense laid out above (though the rule-following sense can also be invoked if the context is right). By contrast, I supposed that for non-competitive games like *Solitaire*, something very much like the sense of ‘play a game’ provided by Suits’s theory is the default interpretation.

The hypothesis holds up fairly well for a wide range of single-player non-competitive games. Not only *Solitaire* but most single player video games, for example, fit well with the Suits definition. The idea that I undertake a commitment, qua speech-act, to follow certain rules when playing (e.g.) a single-player game of *Pac Man* is highly implausible. I do, though, plausibly undertake a commitment qua intention to achieve the lusory goal (getting a high score, keeping my Pac Man alive, or whatever) only through the permitted means. If, for example, I opened up the Pac Man machine and reprogrammed it so it automatically gave me a massively high score and made my Pac Man invulnerable to the ghosts, then I would no longer be playing Pac Man. Similarly, if I did not try to keep my Pac Man alive or score points by eating dots and fruits, but instead tried to kill my Pac Men as quickly as possible (e.g.) then I would plausibly no longer be playing Pac Man (though I might be playing some other game, Kill the Pac Man, perhaps). So far the Suits account seems to work well, and indeed I think a similar account will work for many single-player video games.

However, the account unfortunately begins to break down when we turn our attention to games in which there seems to be no particular goal which is constitutive of playing the game. To be clear, I agree with Suits that we must take a broad view of the kinds of goals which can be constitutive of games. I entirely agree that we should allow for what he usefully calls “open games” – games which do not have a goal which is such that achieving the goal ends the game. Competitive games are closed games – they have a goal which is such that achieving it means one player either wins or loses, or it is a draw – and this of course ends the game. However, many non-competitive games do not have such a game-ending goal. Suits gives the nice example of what he calls “Endless Ping Pong.” Whereas ordinary Ping Pong is a competitive (and hence closed) game in which one tries to prevent one’s opponent from returning one’s shot, in Endless Ping Pong one has a partner rather than an opponent, and the constitutive aim is to keep the ball in play. In principle, this game can go on forever. Certainly achieving the goal, namely keeping the ball in play, does not end the game. Suits plausibly argues that the notion of an open game allows his account to classify games like “Cops and Robbers” and other such children’s make-believe games as games in his sense. In that case, the relevant goal is to keep the relevant sort of narrative going, where one does this without using certain means ruled out by the game (e.g. simply following a script from a crime film would not be playing Cops and Robbers).

Even taking this important point on board, though, there seem to be games which have no constitutive goal even if we are broad-minded enough to include the goals characteristic of open games like Endless Ping Pong. In particular, certain very popular computer games seem not to have any such goal. Here I have in mind games like *Minecraft* and *The Sims*. In these games, often referred to as “sandbox games,” there seems to be no particular goal one must seek to count as playing. Different people who play the game have different sorts of goals. Some people who play *The Sims* want their Sims to be very popular in the virtual world created, while others want

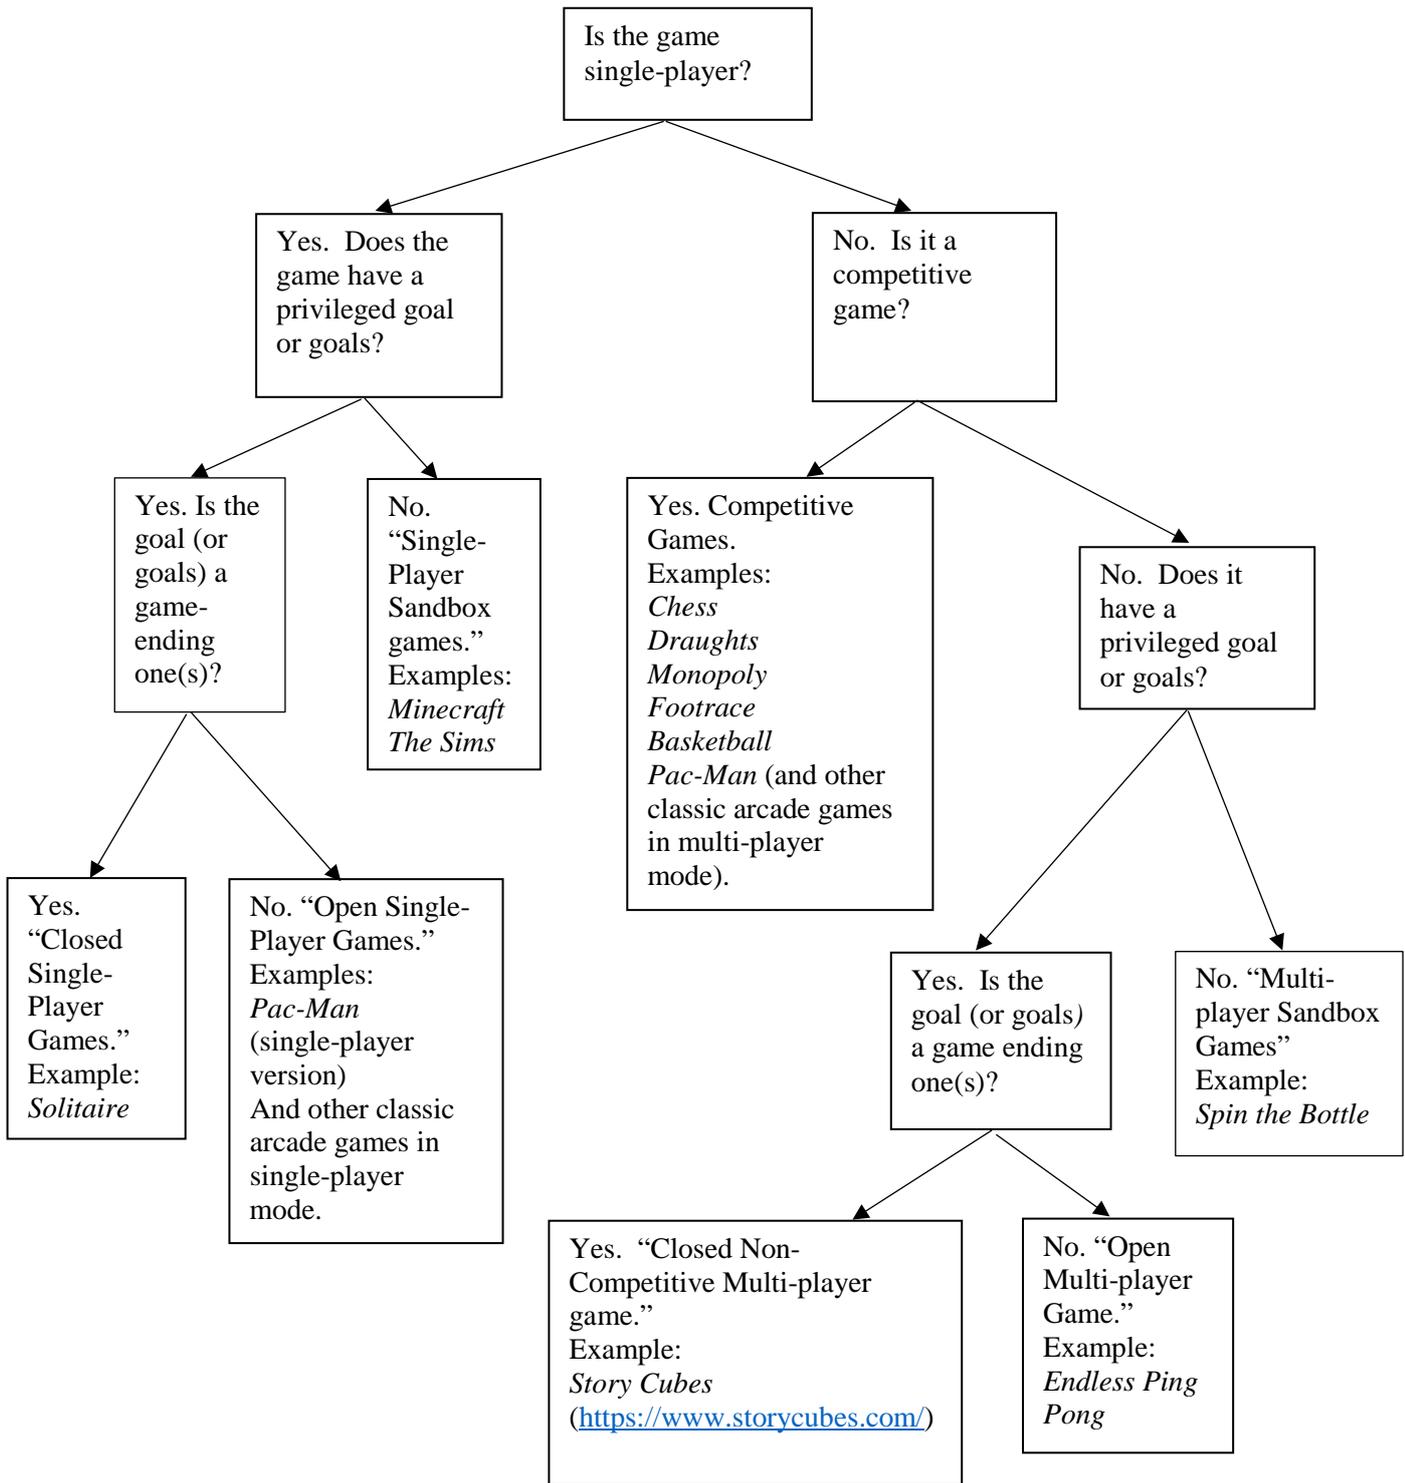
them to have extravagant homes and lifestyles, or to be rock stars, or whatever. Suits could insist that there is some more abstract or schematic goal which one must pursue to count as playing these games, such as “create an interesting life for your Sims” but to me this seems strained; the sort of thing one would only embrace if one were already deeply in the grip of a theory.

Although *Minecraft* and *The Sims* provide vivid illustrations of the point, we do not need to invoke such modern games to make the point. Much more old-fashioned activities which we do not hesitate to call ‘games’ seem to lack constitutive goals. Many so-called “party games” are like this. Consider, for example, *Spin the Bottle*. One might argue that the goal of *Spin the Bottle* is to make sure you kiss the person you find most sexually attractive, but one need not pursue this goal to play the game. Just think of someone who is married and playing with their spouse. Out of a sense of moral duty such a person might aim to kiss only his spouse even if there were someone else playing he found more attractive – or for that matter someone who aims to kiss someone else for the sake of “being naughty” even if he does find his spouse most attractive. In either case such a person could still count as playing the game.

One might then object that the (pre-lusory goal) is simply to kiss *someone*, and one then just restricts oneself to the means allowed by the rules of *Spin the Bottle* for achieving this goal. This does not seem constitutive of playing *Spin the Bottle* either, though. On some versions of the game, anyway, if the bottle is not unambiguously pointing towards any one player (as when it is equidistant between two of them) one need not kiss either person. Someone could count as playing *Spin the Bottle* even if he aimed to avoid kissing anyone; perhaps he is a hypochondriac who fears physical contact but who was coerced into playing, or perhaps he has undertaken an extreme vow of chastity which rules out even kissing.

Moreover, even though it is not competitive, *Spin the Bottle* is a multi-player game, and therefore is a social game. It is therefore not so implausible to suppose that for *this* sort of non-competitive game that undertaking a commitment qua speech-act to follow the rules of the game *is* constitutive of playing the game. Indeed, we can even imagine charging someone with cheating if they used a “rigged” bottle of some kind to ensure he always kissed whoever he most fancied. So while rules, rather than goals, do seem constitutive of this sort of non-competitive game they seem constitutive here at least in part in a way that is closer to the account developed in the previous sections for competitive games, rather than (simply) in the way that Suits’ theory suggests. In fact, one virtue of making commitments qua speech-acts constitutive of playing such non-competitive games is that it provides an elegant account of *with whom* one is playing a given game – namely those with whom the relevant commitments (qua speech-acts) are undertaken. On a more individualistic account like Suits’s, it is simply not obvious how this works.

At this point we have enough distinctions between relevantly different kinds of games that a flow chart might help provide an overview of the terrain:



With this taxonomy in place, I want to suggest that ‘play a game’ and cognates vary in content depending (at least in part) on what sort of game(s) is under discussion. If the topic is a multi-player game, whether competitive or not, then what we might usefully call the “social contract” conception of playing a game laid out in section V is the default content. Actually, of course, the social contract conception itself further divides as between what I above called the “normative role” and the “rule-following” conceptions. Which of these two more specific senses provides the right interpretation of a given utterance will depend on further more subtle contextual features. My working hypothesis, though one I will not be able to argue for here, is that the default interpretation of such phrases in the context of talk of competitive games is what I have called the normative role sense.

Actually, the fact that some multi-player games do not seem to privilege any particular goal requires a slight modification to the social contract account of game-playing from section V. Recall that on the account developed there to play a game one must undertake a commitment to pursue what Suits would call the lusory aim associated with the game. I have argued above that certain multi-player games, like *Spin the Bottle* simply do not have a lusory aim. To accommodate such games we could introduce yet another sense of ‘play a game’ but this seems unnecessary, and we should heed Occam’s Eraser insofar as we can.²⁰ A simpler solution would be to modify the account from section V so that it requires only that one undertake a commitment to pursue whatever aims are associated with the game in the way that lusory aims are associated with the games they define. For games which have no lusory aim this condition will simply be vacuously satisfied.

In contexts in which the game under discussion is a single-player game, my suggestion is that the default content for ‘play a game’ is something much closer to the account developed by Suits as I briefly suggested above. More specifically, in this sense of ‘play a game’, which we might call the individualist sense:

An agent A is playing a game G at time t if and only if at time t A is carrying out an intention to follow the rules of G, and if there is a lusory aim associated in the right way with G then A is also thereby trying to achieve that aim.

I have of course here made the obvious analogue of the move that I made in the case of multi-player games to deal with games (like *Minecraft*) which do not seem to have a lusory goal – namely just insisting that the player pursues (rather than merely commits to pursuing) whatever lusory aims there are, given the game under discussion.

Obviously a lot more could be said about this broader account, and I hope to develop it in more detail in future work. I certainly do not pretend to have argued for it at length here, much less to have defended it against the most important objections one might press against it. This broader account has a number of virtues. Because of the way it embeds the social contract conception developed above, it avoids the paradox of cheating and the paradox of trifling. It also accommodates what is insightful in the powerful theoretical framework developed by Suits but without making it vulnerable to objections from cheating or trifling by restricting its application to single-player contexts in which these objections are not forceful and where the account simply seems much more plausible. Finally, it helps make good sense of the intriguing idea that ‘play a game’ and cognates are family resemblance words and phrases, or so I shall now argue.

²⁰ “Do not multiply dictionary entries beyond necessity.” This nice maxim comes from Paul Ziff. See Ziff 1960: 44.

To begin with, note that the account on offer clearly satisfies two of the three conditions laid down by Pelczar for a word or phrase to count as a family resemblance word or phrase. In particular, on this account (a) the content of ‘play a game’ depends on context, and (b) the default way in which a content is selected is the topic of conversation – in particular, what sort of game(s) is (are) under discussion. An obvious complication is what to say when a speaker quantifies over several games, some of which are single-player and some of which are multi-player. This, though, is a complication facing *any* theory of family resemblance terms developed along the lines of Pelczar’s theory. In fact, this is a problem for everyone, really. After all, we need to make sense of statements like “Jones visited two banks today” uttered by someone who knows and means to convey that Jones visited both a financial institution and the side of a river on that day. Moreover, there are some obvious moves that can plausibly be made to deal with such cases. Perhaps a meta-linguistic account can be made to work, for example. That is, perhaps in such cases we should interpret the speaker as talking about several activities all of which can correctly be described as ‘playing a game’ even if the sense varies between activities. There are other moves that could be made here, but I put this complication to one side for now.

For ‘play a game’ to count as a family resemblance phrase in Pelczar’s sense, that phrase must also be *semantically open* in his sense. Pelczar uses analogies with the common law, improvisational jazz, and conversations to get the basic idea of an open practice across. In all of these cases, we find activities which are to some extent governed by norms, but which are also to a large extent open in the sense that those norms do not typically fix a single correct way of “going on.” Semantic openness reflects semantic practices which are open in this sense. When words or phrases are semantically open, it is “impossible to describe their literal use correctly through the precise specification of *any* content-determining rule.”²¹

With semantically open words and phrases, not only are whatever rules which fix the correct use of the term not fully determinate, the rules themselves are in a sense “up for grabs” and this is implicitly understood as part of their meaning. Pelczar discusses ways in which words which are open in his sense can acquire new senses. These include uses which begin as metaphorical but accrete over time to provide new literal senses, explicit stipulations, retroactive explications that make sense of usage, and speaker’s intentions. The accretion of precedents (here the analogy with the common law is important) can make it more plausible to allow that the word or phrase in question has acquired a new literal use.

I would add that it is an important part of the practice that participants can and do intelligibly argue about whether to extend the practice in one direction rather than another. Some, for example, might balk at the idea that *Minecraft* and *The Sims* are “really games” simply because they are so different from what one might reasonably take to be our paradigm of a game, namely a multi-player competitive game. Whether it makes sense to use the same discourse for both of these may depend on whether they have enough of the right sorts of things in common to make it useful for us to categorize them together in spite of their obvious and substantial differences. In particular, this might depend on what one takes to be *valuable* about game playing, since this might be common across single-player and multi-player games in spite of their obvious differences. No doubt such arguments have been had, and whether we “really” should classify such activities as instances of game playing was at one point very much up for grabs even though by now it seems hard to deny that there is a sense of ‘game’ which includes such “sandbox games.” A useful analogy might be with ongoing debates over whether chess, admittedly a game, should also count as a sport. It would also probably be instructive to juxtapose this account with Charles Stevenon’s insightful

²¹ Pelczar 2000: 498.

expressivist theory of what he calls “persuasive definitions,” but I lack the space to follow up on this idea here.²²

It might also be worth noting that the account developed here also fits well with Wittgenstein’s original idea that there are “family resemblances” between the various different things we call games.²³ The various senses of ‘play a game’ laid out here can be seen as all drawing on the same ingredients – rules, goals and commitments. Just what sort of commitment (speech-act or intention) varies between contexts, as do the precise role of the relevant rules and goals. Even so, the senses have enough in common that the metaphor of a “family resemblance” does not seem out of place.

I want to close by considering a final objection to my approach. My account might seem incapable of explaining why practices like taking part in a parliamentary debate or giving testimony at a criminal trial are not forms of game-playing. After all, in these contexts we do undertake commitments (quite explicitly in the legal case) to follow certain rules, and we may also undertake to pursue certain goals (justice, say). Clearly, though, these are not games.

Strictly speaking, these are not counter-examples to my account but for a reason that might suggest I have (ironically!) cheated. For in my canonical statement of the social contract conception (in section V) I formulated my definition in terms of rules R, “where rules R are the constitutive rules of G.” Since I have helped myself to the idea that some rules, but not others, are rules of a game, one might now worry that my account is viciously circular and therefore unhelpful. An analogous worry obviously applies to the individualist sense I have developed for single-player games, since here too I need to distinguish rules which are game rules from other sorts of rules one might intend to follow.

This is an important lacuna in my account as I have so far spelled it out, but I think it can plausibly be remedied. Although I cannot argue for this view in the detail it deserves here, I think we can understand what makes a set of rules count as the rules of a *game* is the *function* of that set of rules. Intuitively, 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 manifesting certain forms of human excellence (think of how we praise successful athletes and chess champions), and socializing in certain distinctive ways. Intuitively, these are what games are *for*. Insofar as we can distinguish game rules from other kinds of rules in this way, the worry about circularity can be met. The proposed account can therefore also use this conception of game rules to explain why taking part in a Parliamentary debate or giving testimony at a criminal trial are not games.

Obviously a lot more needs to be said about all of this. First, one would ideally like an account of what makes it the case that a set of rules has one function rather than another – and how this may vary depending on whether the rules in question are socially embodied or instead the private property (as it were) of some individual who comes up with a game he keeps to himself for his own private amusement or whatever.

Second, one would of course also ideally like a more precise account of what the relevant function or functions is (or are). I am tempted by the view that the most fundamental function of games is

²² See Stevenson 1938.

²³ Actually, as has been noted more than once, ‘family resemblance’ is probably not a great metaphor for Wittgenstein’s purposes anyway, since he seems to want to deny any single necessary commonalities – whereas for literal family resemblance of course common ancestry is a non-trivial necessary condition.

to enable people to play games in the sense of ‘play a game’ provided by Suits’s theory. That is, roughly, even though Suits has not provided a defensible and fully general definition of ‘play a game’ he has in effect identified the *function* of game rules – namely they enable us to make a “voluntary attempt to overcome unnecessary obstacles.” (Suits 2005: 55)

Actually, even more plausible is the view that the function of game rules is that they enable us to do what Tom Hurka calls “playing in a game.” Hurka plausibly argues (in the introduction to the third edition of Suits’s book) that Suits implicitly changes his conception of the “lusory attitude” when he discusses the role of games in Utopia. There it becomes clear that one must value the activities such rules make possible *for their own sakes* (Hurka 2005: 15). When one plays a game and values the relevant activity in this sense one is “playing in a game.” It seems very plausible to me that some such function captures an essential feature of game rules. The more intuitive functions mentioned in my initial gloss of the functional approach (having fun, manifesting certain sorts of excellences) are probably too broad, and would include rules we would not intuitively think of as the rules of a game. They are, however, still relevant since we value playing games in Suits’s sense because doing so provides opportunities for having fun, manifesting certain excellences, etc. Clearly, much more needs to be said about all of this, though. I intend to discuss these matters in much more detail in a separate article.

For now, though, it is worth noting that we can also intelligibly argue about which functions we *ought* to associate with ‘play a game’ and cognates. Moreover, we might do this by using forms of words like “that is not much of a game” and the like. This provides yet another way of understanding how the semantics for ‘play a game’ and cognates is open-textured in the way required by Pelczar’s conception of family resemblance words and phrases.

The appeal to these functions also strikes a nice balance between two extremes. Consider a world in which *nobody* followed the rules of chess in order to play in a game in the sense just identified (or have fun, or manifest certain distinctive virtues or whatever the most plausible candidate function is), but only did so out of a compulsion or at gunpoint, or whatever. In such a world it is plausible that they are not really playing chess because chess simply does not exist *as a game* in that world. In Suits’ useful terminology, the *institution* of chess may exist, but that institution is not deployed by people in that world in order to play a game of chess. My account can endorse this verdict, since the *function* of chess in that world (whatever it may be – to relieve compulsive anxieties, to avoid being shot) is not of the right sort for the rules of chess as socially embodied in that world to count as the rules of a game. At the same time, the account developed here can allow that in our world, where chess *is* a game because of the way in which it is socially embodied, a given individual can play chess out of a compulsion, or at gunpoint.²⁴ Playing chess for these kinds of reasons is possible in individual cases, on my account, but could not become the norm, since in that case the resulting socially embodied rules would no longer count as the rules of a game. The analogy with Kantian arguments against making a lying promise should be obvious here.

Yet another advantage of the proposed functional approach is that can explain why it is not a sheer *accident* that we use the same word (‘play’) when talking about playing a game and when talking about playing, full stop, as with a child playing with her food. I agree with Suits and others that

²⁴ Although I lack the space to argue for this here, I think this allows my account to capture what Suits was after with his ‘just because it makes such activity possible’ clause, but in a much more plausible way than Suits managed. The advantage of my approach is that by “going social” we do not need to impose any robust constraints on the psychology of individual game players. Or, rather, at least this will be so for games which are in fact socially embodied. Games which are the private property of a single person who plays it alone to amuse himself are different, but in that case something closer to what Suits has in mind also seems more plausible.

‘play a game’ does not conceptually *entail* ‘play’ (full stop). The jaded professional football player who is just in it for the money is still playing football but is not playing, full stop. At the same time, it seems unlikely to be a sheer accident that we use the same word in English in each case. My suggestion is that this reflects the fact that *one* of the functions of games is to provide a vehicle for having fun, and that a plausible definition of ‘play’ (full-stop) will advert in some way to doing something ‘for the fun of it’. These link between play and playing a game will in that case be functional, rather than semantic or conceptual. Again, this all requires more spelling out, and I shall return to these issues in separate work. Here I am simply taking a programmatic approach, and canvassing some of the key issues raised by the account developed here.

Conclusion.

This paper has had enough twists and turns that I will not try to summarize all of the main ideas or arguments here. Instead I shall simply close by indicating some further questions raised by the account I have developed here which will require further investigation. One question is whether my discussion of the way in which ‘play a game’ and cognates are family resemblance terms has exhausted the semantic contents associated with those phrases. For example, it is common to use the phrase ‘stop playing games with me’ as a kind of complaint which is apt when one’s interlocutor is manipulating or teasing you in some way. Similarly the phrase ‘gaming the system’ is a pejorative way of describing someone who is taking inappropriate advantage of loopholes and technicalities. Are these usages best construed as metaphorical, or are they by now sufficiently well entrenched in the right kinds of linguistic conventions to count as another sort of literal content? One might ask similar questions about uses like “That wasn’t a game, that was a massacre!” or “That wasn’t a game, that was batting practice for the Red Sox!”

Another set of questions is whether the semantic content of ‘play a game’ can shift depending on whether the game under discussion is a formal one, tightly defined by fairly strict rules, as with a FIFA football game, or whether it is much less formal, as with a pick-up game by a group of children after school. Although I have not explored this idea here, I think these kinds of contextual variations might well matter to semantic content too.

I have also here only touched very briefly on the important idea that what makes a set of rules qualify as the rules of a game is, at least in part, their function. Clearly, this requires a lot more spelling out, as do my all too brief remarks about the relationship between playing a game and playing (full-stop). Ideally, getting all of these further elements into view in a plausible framework would cast light on the appropriate role of both games and play (full stop) in a well-lived life. I hope to return to these important questions in future work.

Finally, in future work I should return to the analogy between the paradox of cheating and the Kantian “problem of bad action” to see whether the account defended here might offer some useful analogous strategies for the Kantian. My hunch is that the differences between morality and games will make this problematic, but it may still be instructive to see how these ideas fare when transposed to that context. More generally, Kantians often argue that a commitment to the moral law is constitutive of being a rational agent and these arguments sometimes put some weight on the analogy with the allegedly constitutive rules of games. The framework developed here may also cast some interesting light on these Kantian ideas, and on the very idea of constitutive rules more generally.

So there are many questions which still need investigation and in many ways the work done here sets the stage for a broader project. In the meantime, though, if the arguments here are sound

then I have established several important points. In particular, in this paper I hope to have done the following:

- (1) Laid out and motivated two neglected paradoxes for existing theories of what it is to play a game.
- (2) Plausibly argued that existing theories fare very poorly when viewed in light of these paradoxes.
- (3) Defended a theory of at least one semantic content for ‘play a game’ and cognates which elegantly dissolves these paradoxes.
- (4) Begun (however sketchily) to provide a broader framework for the meaning of ‘play a game’ and cognates which both incorporates the theory I use to resolve the two paradoxes as a proper part, and can at the same time vindicate the idea that this is in at least one useful sense a family resemblance phrase.

Even though there is obviously much more work to be done to answer the many questions raised by this account, it should not be without significant interest in its own right. Indeed, if I managed to achieve all four of those aims, I will have accomplished all I could reasonably hope to have accomplished in a single paper on this topic, as opposed to a monograph!

Works Cited

- Bambrough, R. 1960. "Universals and Family Resemblances." *Proceedings of the Aristotelian Society*.
- Chrisman, M. 2012. "The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions." *Journal of Philosophy*. 109 (10): 588-612.
- D'Agostino. 1981. "The Ethos of Games." *Journal of the Philosophy of Sport*.8: 1-18.
- Dellatre, E. 1976. "Some Reflections on Success and Failure in Competitive Athletics." 2: 133-139.
- Feezell, R. 1988. "On the Wrongness of Cheating and Why Cheaters Can't Play the Game." *Journal of the Philosophy of Sport*. XV: 57-68.
- Ganz, J. *Rules*. 1971. The Hague: Mouton.
- Gert, H. 1995. "Family Resemblance and Criteria." *Synthese*. 105: 177-190.
- Green, S. 2004. "Cheating." *Law and Philosophy*. 2: 137-185.
- Hurka, T. 2005. "Introduction." In *Suits 2005*: 7-20.
- Kenny, A. 1963. *Action, Emotion, and the Will*. New York; London: Routledge & K Paul.
- Korsgaard, C. 2009. *Self-Constitution*. Oxford: Oxford University Press.
- Kreider, A.J. 2011. "Game-Playing Without Rule-Following." *J. of the Philosophy of Sport*. 38: 55-73.
- Lance, M. 1998. "Some Reflections on the Sport of Language." *Philosophical Perspectives*. 12: 219-240.
- Lehman, C. 1981. "Can Cheaters Play the Game?" *Journal of the Philosophy of Sport*. VIII: 41-46.
- Lewis, D. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic*. 339-359.
- Pelczar, M. 2000. "Wittgensteinian Semantics." *Nous*. 34: 483-516.
- Ross, A. 1968. *Directives and Norms*. Lawbook Exchange Ltd.
- Sellars, W. 1963. "Some Reflections on Language Games," in *Science, Perception and Reality*. London: Routledge & Kegan Paul.
- Stevensen, C.L. 1938. "Persuasive Definitions". *Mind*. 47 (187): 331-350.
- Suits, B. 2005. *The Grasshopper: Games, Life and Utopia*. Broadview Press.
- Wittgenstein, L. 1958. *Philosophical Investigations*. Blackwell.
- Ziff, P. 1960. *Semantic Analysis*. Ithaca, NY: Cornell University Press.