Coordinating through Precedent Without Common Inductive Standards

Abstract

It has been argued that predicting others' behavior in the context of a coordination game relying on precedent as a standard of inference presupposes so-called 'common inductive standards'; i.e. a piece of (common) knowledge on the part of each player that one's co-participants employ a similar standard of inference. My concern in this paper is to show that this claim is implausible. Precedent-based predictions are incompatible with `common inductive standards'; the main reason being that common inductive standards are characterized in terms of higher-order beliefs (or knowledge); i.e. beliefs about what others believe yet others will do. In coordination games, however, justifications based on higher-order beliefs defeat justifications based on precedent. For this reason, I suggest that, when coordinating through reasoning from precedent, the common inductive standards requirement should be cast negatively as the common *absence* of higher-order beliefs. This analysis extends a nascent line of research according to which common knowledge can be harmful to certain cooperative activities.

1. Introduction

Consider the following case:

> **Fast Food.** You and I want to meet for lunch. We have two options: McDonald's or Wendy's. We don't care where we'll have lunch as long as we'll have lunch together.

In situations such as this one, each of us is trying to predict what the other is going to do in an attempt to match our choices. These predictions can be justified in various ways. Maybe we have antecedently agreed or promised to each other that we would go to McDonald's. But suppose that we have no explicit agreement to go on, but that, instead, we have always gone to McDonald's in the past. In this case, you may rely on precedent and simply reason as follows: "Since that's where we've always gone, that's where she'll go again". For ease of presentation, let's call the proposition "We've always gone to McDonald's" proposition A, and the proposition "You (or I respectively) will go to McDonald's" proposition $x_{y/I}$. We can, then, represent the evidential relation between A and $x_{y/I}$ as follows:

(1) $A \rightsquigarrow x_{y/I}$

The squiggly arrow indicates that the inference is defeasible; i.e. valid only by default; the fact that we've always gone to McDonald's does, of course, not *necessitate* that this is what we shall do.[1] In predicting your behavior, I may use what's stated in **(1)**. Let $K_{I/y}$ denote "I/you know that ...". We may, then, represent my reasoning process as follows:

(2) $K_I(x_y)$ because $K_I(A)$ and $K_I(A \rightsquigarrow x_y)$

I know that you'll go to McDonald's, because I know that this is what we've always done, and because I know that people tend to continue to act as they have in the past.

Although, these claims are couched in terms of knowledge (instead of, say, certainty, belief, degrees of belief, or reason to believe) this should be taken more as a convention rather than a substantive commitment. As Harvey Lederman (forthcoming) stresses, in debates about coordination the term ``common knowledge" is often used as a catch-all term for "common

[1] For this reason Cubitt and Sugden (2003) say that A gives a "reason to believe" that x is the case. Alternatively, using Lewis' preferred language, we can say that A "indicates" that x is true.

knowledge, common belief, and common certainty". In this paper, knowledge will often be important only insofar as it entails belief; and many of the examples that follow will be couched in terms of (common) belief.

Now, it has been said that, at least in the context of coordination games, the line of reasoning presented in **(3)** is incomplete. Famously, David Lewis has argued that predicting another agent's behavior using precedent as a source of evidence depends on the "mutual ascription of some common inductive standards […]." (Lewis, 1969, 56f). Similarly, Sugden and Cubitt, analyzing Lewis on convention, require that the players "have reason to believe that, in particular relevant respects, they have common background information and common inductive standards" (Cubitt & Sugden, 2003, 185). A common inductive standard is a piece of knowledge (belief, or reason to belief) that a particular salient feature (e.g. precedent) is a "projectible regularity" (Cubitt & Sugden, 2003, 198) such that "each person must have *reason to believe that each other person shares his own standards* about what can be inferred inductively from A" (my italics) (Cubitt & Sugden, 2003, 198). Extending this thought, Christina Bicchieri has suggested that, to select an equilibrium in a coordination game, "we must introduce some salience criterion of choice, *and common knowledge thereof* […]. Salience may be provided by precedent [...]" (my italics) (Bicchieri, 2005, 36). Similarly, Cyril Hédoin has argued that the players need "*common knowledge* […] that they share the same reasoning modes" (Hédoin, 2014, 380). In summary, the thought is that simply *having* the same standards of inference is not enough for successful coordination; rather, the players need additional beliefs or knowledge that the respective others are disposed to reason similarly.

To see the putative importance of a common inductive standard we should examine a case in which it is missing: Reconsider **Fast Food** and suppose for a minute that I am initially inclined to predict your behavior using precedent; suppose also that, this time around, I think that you falsely believe that I intend to *defy* precedent predicting that I have decided to go to Wendy's instead. Of course, in this case, it would seem quite silly to stick to my guns and use precedent as my inductive standard. I know that you are going to try to match what you believe I will do and, thus, go to Wendy's. If my goal is to coordinate with you, then I should likewise go to Wendy's. Hence, in predicting your behavior, I ought not only to know what we've done in the past, but, additionally, that you are also inclined to predict that I will go to McDonald's using precedent as a standard of inference.

Adding a common inductive standard to formulation **(2)** we obtain:

> **(3)** $K_I(x_y)$ because $K_I(A)$ and $K_I(A \leadsto x_y)$ and, as one precondition, $K_I(K_y(x_I)$ because $Ky(A)$ and $K_y(A \leadsto x_y))$.[2]

This statement reads as follows: I know that you will go to McDonald's in the future, (a) because I know that we've gone to McDonald's in the past, (b) because I know that people tend to act in the future as they have in the past, and (c) because I know that you think that I will go to McDonald's for the same reasons (i.e. because you know that we've done so in the past, and because you know that people tend to act in the future as they have in the past).

In this paper, I will argue that the common inductive standard requirement expressed in **(3)** is implausible for two reasons. First, it implies a higher-order belief about what the other player thinks oneself will do. However, predicting another player's behavior based on precedent is incompatible with the presence of such higher-order beliefs. The following argument, which I will explicate in this paper, shows this:

---

[2] Vanderschraaf and Sillari (2013) formulate this idea in terms of "symmetric reasoning". Symmetric reasoners know that an inference that the agent herself can draw can also be drawn by another agent. Formally, this idea is expressed as follows: $[\mathbf{K}_i(A) \Rightarrow \mathbf{K}_i(E)$ and $\mathbf{K}_i(A) \Rightarrow \mathbf{K}_i\mathbf{K}_j(A')$, then $\mathbf{K}_i(A')] \Rightarrow \mathbf{K}_i\mathbf{K}_j(E)$; i.e. if an agent can infer E from A, and the agent can infer that the other agent knows A, then she is also in a position to know that the other agent knows E. The definiens says that for each agent i, if i can infer from A' that E is the case and that everyone knows that A' is the case, then i can also infer that everyone knows that E is the case.

P1 – Interdependence. In a two-player pure coordination game between player *A*[3] and player *B*, *A*'s first-order prediction about what she takes *B* to choose provides sufficient reason for her own rational strategy choice.

P2 – Double Justification. *A*'s belief about what she takes *B* to choose can be justified by appeal to (a.) precedent, or (b.) higher-order beliefs suitably characterized (e.g. what *A* believes *B* believes that *A* will choose).

P3 – Higher-Order Defeat. In pure coordination games, justifications based on higher-order beliefs always defeat justifications based on precedent.

—

C – Absence Precondition. Precedent can justify *A*'s belief about *B*'s strategy choice only if *A* has no second-order belief about what *B* thinks *A* will choose (e.g. beliefs about what *B* believes *A* will choose).

The argument makes explicit an evidential relation between first-order behavioral predictions, precedent, and second-order beliefs. As the paper unfolds, we will see that the argument generalizes for second, third, and, ultimately, *n*th level behavioral predictions.

The second reason rendering **(3)** implausible is that, sometimes, we want to *explain* common knowledge that a certain inductive standard is used by appeal to the independent predictive power of precedent-based inferences. If these inferences were to presuppose such common knowledge, then these explanations would turn out to be circular.

[3] The argument is stated from player *A*'s perspective. This is just to avoid clutter. The reader should fill in the exact same argument as given from *B*'s perspective.

Before continuing further, I should add a caveat. The "common inductive standards" requirement depends on various idealizing assumptions. For instance, the idea that reasoning from precedent depends on *infinitely many* nested higher-order beliefs about the inductive standard may[4] depend on the assumption that the players are unbounded reasoners. Furthermore, the idea that a player may only base her predictions of the other's behavior on precedent provided that the other does so as well, depends, at least on the face of it, on the idea that the player is rational, and, for further iterations, on the idea that their rationality is common knowledge.

Many have found it implausible that successful and reliable coordination should be premised on infinitely complex higher-order beliefs. The remedy has usually been to relax at least one of the various idealizing assumptions. Harvey Lederman (2017) has argued that coordination can be facilitated by letting go of the assumption that the players' rationality is common knowledge. Players can be rational, but they need not be smug; i.e. they might not know that they are rational. Others (e.g. Kneeland, 2012) have argued that bounded reasoners can coordinate without common knowledge. Yet others (e.g. Skyrms, 2004) have explored coordination in entirely non-strategic contexts.

In this paper, I won't follow these approaches and, thus, keep all rationality assumptions. These are strong assumptions and the decision to keep them needs a bit of justification. First, it is simply worth investigating whether fully rational agents would need common inductive standards to coordinate. Second, in the philosophical literature, such common inductive standards are often added because rational agents would need them to coordinate their actions. Hence, a discussion of these standards under these idealized circumstances marks a natural

[4] Harvey Lederman (2017) has denied this claim.

extension of the extant literature. Third, the reason these assumptions are often dropped is to lend empirical validity to a particular model. Actual agents, it is sometimes argued, simply aren't unbounded reasoners and don't commonly know that they are rational. In this paper, my primary concern is not empirical validity, but, rather, a principled investigation into the tension between reasoning from precedent and higher-order expectations. Lastly, any successful theory of coordination should hold up under idealized circumstances, as it would be quite surprising if successful coordination were to *require* cognitive limitations.

In the early days of research on coordination (e.g. Rubinstein, 1989) and cognate cooperative activities such as joint action (e.g. Bratman, 1989), and conventional behavior (e.g. Lewis, 1969), common knowledge that each participant will choose a particular strategy was seen as a requirement. This sentiment finds its most rigorous expression in Rubinstein's *Coordinated Attack* and *Electronic Mail* games. Subsequently, many had noticed that such common knowledge requirements depend on various idealizing assumptions (see above) and that relaxing these assumptions may render these requirements unnecessary. In each case, it is nevertheless argued that common knowledge is, although not necessary for coordination, always compatible with it. This idea has recently been challenged. Common knowledge requirements can, at times, be *harmful* to coordination (see Lederman, 2017; Schönherr (forthcoming)). The present paper extends this nascent line of research, arguing that solving a coordination game relying on precedent is in tension with common knowledge that this standard is used.

In the first part of this paper (section 2 – 4), I will show why the common inductive standard requirement expressed in **(3)** is implausible. In the second part (section 5), I will provide an alternative. More concretely, in the next section, I shall detail **P1** and **P2**. In section three, I will explicate **P3**, which will put us in a position to see why **C** is true, and, thus, why **(3)** is false. In section four, I will show how common inductive standards obscure the explanatory relation between precedent as a plausible rule of inference, and common knowledge that this inference rule is used by the players. Lastly, in section 5, I will sketch a positive picture describing how coordinating agents should think of one another in the relevant situations. Put coarsely, predicting other players' actions using precedent presupposes a form of mutual belief *suspension* about the inductive standard used by one's co-participants.

2. Interdependence and Double Justification

The games I will be talking about are two-player, conflict-free, pure coordination games; i.e. games with multiple strict Nash equilibria[5] in which one player's gain does not require the other player's sacrifice. Such games can be represented by the following matrix[6]:

		Player 1	
		X	Y
Player 2	X	1,1	0,0
	Y	0,0	1,1

Figure: 1

In this game, players have to solve the *equilibrium selection problem*. There are two relevant pure[7] equilibria, {X,X} and {Y,Y}, and the players have to figure out a way to settle on one of them. Ultimately, each player is trying to match what she takes the other player to choose, which is

[5] A Nash equilibrium is a set of strategies such that no individual has an incentive to change her choice given the choices of the others.
[6] This matrix should be read as follows: The labels 'Player I', and 'Player II' represent the players. The labels "X", and "Y" represent the players strategies. The numbers represent the players' utilities. These utilities are a function of the players' strategies. In this paper, I will be making use of players' *pure* strategies; i.e. a player's decision to play a strategy with probability 1. The expected utility of choosing a particular strategy is the utility associated with this choice given the expected pure strategy of the other player(s).
[7] Mixed equilibria won't matter for our purposes.

why each player's choice depends only on estimates (beliefs, credences, or knowledge) about the other player's choice. This is all I shall say in defense and illustration of the **Interdependence** premise.

Let's move on to the second premise. Beliefs about the other player's choice can be justified in several ways. In this paper, I will focus on only two sources of evidence: Higher-order beliefs and precedent (and, concomitantly, the combination of both). Let's start with precedent as a source of evidence.

Predicting behavior using precedent means inferring future behavior based on a past behavioral regularity. The validity of such reasoning can perhaps be explained by the fact that "we may tend to repeat the action that succeeded before if we have no strong reason to do otherwise" (Lewis, 1969, 36). To see how precedent might do this, consider first an example from David Lewis: "I know very well that I have often seen cars driven in the United States, and almost always they were on the right. [...] Given a regularity in past cases, we may reasonably extrapolate it into the (near) future" (Lewis, 1969, 41). Many (Sillari, 2008; Lewis, 1969; Sugden, 2015; Bicchieri, 2005) acknowledge that precedent can provide coordinating agents with evidence despite the fact that there is no theory with regard to what it is that makes a particular feature salient. Past and future actions are never alike in all, but, rather, merely in some respects. Reasoning from precedent is therefore dependent on certain salient features of actions. This, however, does not detract from the fact that precedent is real; in any case, in this paper I will assume that it is. Lastly, precedent-based reasoning is defeasible; it is a mere "last resort [for the players], when they [the players] have no stronger ground for choice" (Lewis, 1969, 35). This idea was represented above by the squiggly arrow. Higher-order beliefs, as I will argue this paper, turn out to be a "stronger ground for choice".

Let's move to higher-order beliefs as a source of evidence. To start seeing how higher-order beliefs can guide predictions of the other's behavior, consider the following vignette:

Fast Food 2. You and I want to meet for lunch. We have two options: McDonald's or Wendy's. We don't care where we'll have lunch as long as we'll have lunch together.

I learn that a source, who you believe to be infallible, told you that I will go to Wendy's this time around.

For ease of understanding, we can depict my epistemic situation as follows:

I believe that

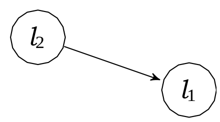

you believe that I will be at Wendy's.

Figure 2

The indexed letter l_n indicates the depth of reasoning (e.g. l_2 represents a second-order belief). Note also that the arrow simply indicates that direction of the nesting of beliefs; it does not indicate a rule of inference of any sort. I reason as follows: Since you believe that I will go to Wendy's, you will intend to match my decision. Thus, I expect you to go to Wendy's, which is why the rational decision on my part is to go to Wendy's. My second-order belief about you justifies my first-order belief about where you will go. My first-order belief, in turn, settles my decision.

Such second-order beliefs can, in turn, be justified by a third-order belief as the following figure illustrates:

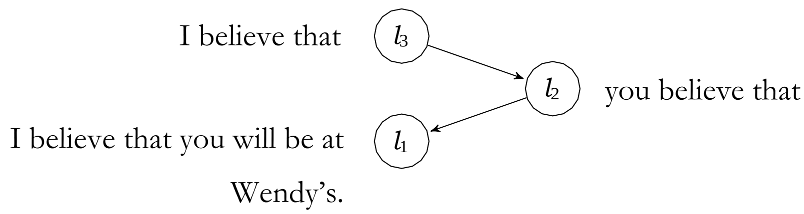

Figure 3

Of course, third-order beliefs can be justified by fourth-order beliefs and so on *ad inf.* Thus, the following general picture emerges: Any level-*n* belief about a player's choice can be justified, presumably, in various ways. One particular way to justify a level-*n* belief is in terms of a level *n* + 1 belief. If each belief is indeed justified in this way, the result will be an infinite hierarchy of justifying beliefs. Such an infinite hierarchy of beliefs marks the crucial characterization (not a definition[8]) of common belief relevant for our purposes. Furthermore, if these beliefs are true, then this amounts to a characterization of common *knowledge*.

Although discussions of coordination games make frequent reference to common knowledge simpliciter that each will choose their part of a particular equilibrium, we should instead rely on the slightly amended notion of common *reciprocal* knowledge (and belief) which is gleaned from Robert Sugden (2015). The problem with good old common knowledge is its reflexivity; i.e. if *p* is common knowledge among individuals in a population *N,* then each individual in *N* knows that *p*, knows that she *herself* knows, etc. In coordination games such as **Fast Food**, however,

[8] Although common knowledge (belief) can be *characterized* by an infinite hierarchy of (actual, potential, or dispositional) nested higher-order true beliefs, it should be noted that this is really just a *characterization*; not a definition. Definitions of common knowledge have, for instance, been given in terms of public events, or inference patterns between symmetric reasoners (for an overview consult Vanderschraaf (2014)). These definitions, however, need not concern us, be- cause, although common knowledge is not defined in terms of iterated beliefs, it nevertheless entails these iterations. Hence, by contraposition, a failure of such iterated knowledge likewise entails a breakdown of common knowledge.

each player is just concerned with what she believes *the others* are going to do, what others take yet others do and so on *ad inf.* As a reminder, this type of reasoning is illustrated in *Figure 2* and *3*. The Sugden-inspired notion of 'common reciprocal knowledge' captures this idea by taking reflexivity out of the definition. A group of players have common reciprocal knowledge (belief) that p is true, if, and only if, for all players i and j in N, where $i \neq j$, i knows (believes) that p holds for j; all individuals $i, j,$ and k in N, where $i \neq j$ and $j \neq k$, i knows (believes) that j knows (believes) that p holds for k, and so on *ad inf.*.

Lastly, we should note that both types of justification can be combined; e.g. third-order predictions can justify second-order beliefs that can, in turn, justify precedent. Here is a figure illustrating this thought:

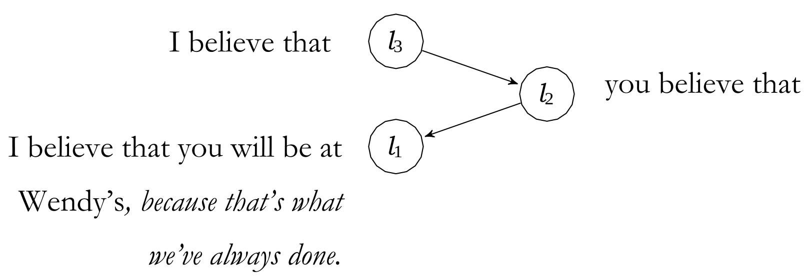

Figure 4

Similarly, we should say that each level n belief can, in principle, be justified with regard to a level $n+1$ belief, or, alternatively, precedent.

3. Higher-Order Defeat

Precedent can justify predictions about the other player's actions only in the absence of any reciprocal belief that would likewise justify or undermine this prediction; this is because, as premise three of the above argument states, justifications based on higher-order beliefs defeat precedent-based justifications. Let's add some color and precision to this idea.

Let RB_n stand for any such *n*th level reciprocal belief. Let *Precedent* stand for a piece of precedent-based evidence bearing on this belief. Lastly, let RB_{n+1} stand for the reciprocal belief justifying RB_n. To illustrate, consider the following propositions:

> RB_1 **–** I believe that you will go to Wendy's.
>
> *Precedent* **–** In that past, we have always gone to McDonald's.
>
> RB_2 **–** I believe that you believe that I will go to Wendy's.

The idea, then, is that RB_1 can neither be justified nor undermined by *Precedent* given the presence of RB_2. The reason is simple: I expect you to think that I, trying to match your choice, will go to Wendy's. Hence, I expect you to go to Wendy's, which is why I myself ought to go to Wendy's. Given this distribution of beliefs, it shouldn't matter to me that we've always gone to McDonald's in the past. This doesn't change even if we make explicit the fact that your prediction of what I'm going to do has been based on precedent. Consider the following proposition:

> *RB_2* – I believe that you believe that I will go to Wendy's, because I believe that: you know that that's what we've always done, and you believe that people tend to act in the future as they have in the past.

Once I know that precedent-based reasoning has led you to a conclusion about where I'll go, I should not use precedent in predicting your behavior. Once again, my second-order prediction has defeated precedent as a reason for justifying my first-order prediction. The fact that I also happen to know *how* you came to settle on your belief does not change that.

Defeat relations among reasons can be grounded in various ways. For instance, in standard cases, more specific information defeats less specific information (e.g. Horty, 2012, 216). To provide just one example, consider the fact that Tweaty is a bird. This is a reason for believing that Tweaty can fly. Suppose further that Tweaty is also a Penguin which acts as a defeater for the aforementioned inference. Once a reasoner knows that Tweaty is a Penguin she is not justified in concluding that Tweaty flies even though Tweaty is a bird. One might, thus, wonder whether the defeat relation between precedent and higher-order beliefs can be explained in similar ways. I think this is not so. Rather, the defeat relation in our case is simply grounded in basic assumptions about the structure of the game. In a coordination game, as I've explained above, each player is trying to match the other's choice; i.e. each player will act on what she believes the other is going to choose. If the structure of the game, as well as the players' rationality, are common knowledge, then each player knows that the same holds true for the other player. Each player knows that the other will act on her belief about what she thinks the other is going to do, which is why precedent has, at this point, been defeated.

Thus, we can state the intended evidential relations as follows using standard Bayesian notation:

(4) $Pr(RB_n \mid Precedent \wedge RB_{n+1}) = Pr(RB_n \mid RB_{n+1})$

Condition **(4)** captures the idea that precedent cannot raise the probability of a prediction given the presence of any higher-order belief that would likewise bear on this prediction. The example just stated illustrates that **(4)** is true. But let's expand on this idea a bit more. Consider a case in which I merely believe that you have *some* pertinent reciprocal belief but that I don't know which one. In this case, **(4)** will still be true as the following vignette will make clear:

Fast Food 4. You and I want to meet for lunch. We have two options: McDonald's or

Wendy's. We don't care where we'll have lunch as long as we'll have lunch together.

In the past, we've always gone to McDonald's. I learn that a source, who you (perhaps falsely) believe to be infallible, flipped a fair coin, and either told you that I would be at Wendy's this time (if it came up heads), or she told you that I would be at McDonald's (if it came up tails).

In this case, my rational response is debatable. However, one thing is clear; I should not rely on precedent in making my decision. You have a belief about where I'm going to be, and you will try to match my decision based on this belief. This is conclusive for you and, thus, precedent should not be invoked. This is different from the case in which I don't know whether you have any belief about what I'm going to choose. Consider the following case an illustration of this thought:

Fast Food 5. You and I want to meet for lunch. We have two options: McDonald's or Wendy's. We don't care where we'll have lunch as long as we'll have lunch together.

In the past, we've always gone to McDonald's. I learn that a source, who you (perhaps falsely) believe to be infallible, flipped a fair coin, and either told you that I would be at Wendy's this time (if it came up heads), or *she didn't tell you anything at all* (if it came up tails).

In **Fast Food 5**, the intuition that precedent may permissibly be invoked is strong. The guiding intuition, I think, is that in the case in which she didn't tell you anything at all, you have no belief about what I'm going to do which is why precedent may permissibly be invoked.

These examples show that first-order predictions based on precedent are defeated by the presence of a second-order reciprocal belief. Similar examples can easily be constructed for each level; i.e. *n*th-order predictions based on precedent are defeated by *n+1*th order reciprocal beliefs. The common inductive standard requirement (i.e. expressed in statement **(3)**[9] from above), however, introduces such a higher-order expectation about what the other believes oneself will do, which is why it is implausible.

Now, statement **(3)** frames the common inductive standard in terms of second-order knowledge. However, nothing in the analysis changes if we substitute second-order for *common* knowledge (as Bicchieri has suggested). After all, if the agents have common knowledge about the inductive standard they use, then each level-*n* belief about a player's choice is accompanied by a defeating level *n+1* belief, which is why no level-*n* belief can be justified via precedent.

The cases discussed in this section present the reader with situations in which both types of evidence *conflict*; cases in which precedent favors one response and higher-order beliefs favor a different response. In these cases, precedent is defeated. I should emphasize, however, that conflict is not required for defeat. To see why reconsider **Fast Food 4**; once I learn that you have a belief about where I will be, I simply don't care about precedent anymore; precedent becomes a non-issue. This is independent of whether precedent aligns or opposes the conclusions I drew based on higher-order reasoning. For the reader who prefers to reserve "defeat" language for conflict cases, it may be more accurate – although less common – to say that evidence provided by higher-order beliefs "positively screens off" (e.g. Tal & Comesana, 2016) precedent-based evidence.

4. Explanatory direction

There is a second objection suggesting that common inductive standards cannot be a precondition for reasoning from precedent in coordination games. Suppose for a minute that

[9] $K_I(x_y)$ because $K_I(A)$ and $K_I(A \rightsquigarrow x_y)$ and, as one precondition, $K_I(K_y(x_i)$ *because* $Ky(A)$ *and* $K_y(A \rightsquigarrow x_y))$.

two coordinating agents have common inductive standards; i.e. they commonly know that they use precedent as a standard of inference. In this case, we may wonder where this piece of common knowledge came from. We may ask "Why is it that the players have common knowledge that they use *this particular* standard, and not some other, perhaps more outlandish, one?". One compelling candidate answer is that precedent is simply a good standard to use; i.e. it works as a predictor of people's actions. In short, precedent is not a good behavioral predictor because people think that they use it as an inductive standard, but, rather, precedent is a common inductive standard, because it is a good predictor.

Now, if precedent-based inferences were to *presuppose* such common knowledge, then the fact that precedent is simply a good predictor could, on pains of circularity, hardly be recruited for such explanatory purposes. In other words, if precedent were a good standard of inference only if it were already presupposed that the agents commonly know that this is the standard is invoked, then the fact that that precedent is simply a good predictor could hardly be invoked in explaining why the agents commonly know that this is the standard they invoke. For this reason, predicting behavior using precedent should be allowed without presupposing that the agents commonly know that it is invoked as a standard of inference.

A critic might want to argue that the independent appeal of precedent as a plausible predictor derives from contexts other than coordination games. In coordination games, a player's action, this critic might say, depends on predictions of the other agent's behavior. In other contexts, however, this is not so. For instance, in predicting what wine I'm going to have tonight, you might invoke precedent and predict that I will choose red; and since precedent is a good standard of inference to use in these non-strategic contexts, it is, one might conjecture, also a good standard in the context of coordination games.

I don't think that this thought is plausible though. After all, rational agents should not be inclined to use a standard of inference in a context in which it is simply not suitable. If it were true that reasoning from precedent requires common inductive standards when playing a coordination game, it would simply be a mistake to suppose that such reasoning has independent appeal (i.e. is a plausible standard of inference even without presupposing common inductive standards).

5. Reasoning from precedent presupposes common belief suspension

In this paper, I started with the thought that predicting others' behavior in a coordination game using precedent as an inductive standard seems to be quite straightforward. This simple inference was captured in statement **(2)**[10]. In the philosophical literature, however, it has been argued that this simple inference ought to be enriched by a common inductive standard; i.e. higher-order expectations about the other players' expectations. The guiding thought throughout this paper was that this addition is implausible: predicting an agent's behavior in a coordination game using precedent as a standard of inference is incompatible with such higher-order expectations, which is why we should say that precedent-based predictions are permissible only in the absence of higher-order beliefs about the players' choices.

To see how this would go, let's go back to **Fast Food**. When thinking about whether you believe that I will go to McDonald's, I may lack both, the belief that you think that I will, and the belief that I won't go to McDonald's. In short, I might suspend belief about where you think I will go. Belief suspension about a proposition *P* entails, almost everybody agrees, neither believing or disbelieving that *P* is the case (e.g. Bergmann, 2005: 420, Wedgwood, 2002). Only Jane Friedman has doubts (see Friedman 2017). The connection between suspension and not believing, she contends, is ***normative,*** not descriptive. An agent who is suspended about *P* ***ought***

[10] Here a reminder for statement **(2):** $K_I(x_y)$ because $K_I(A)$ and $K_I(A \leadsto x_y)$

not believe nor disbelieve it; but since we're operating under the assumption of perfect rationality, we can sidestep these subtleties. Of course, *simply* not believing and disbelieving is not sufficient for belief suspension. After all, a person who has never even considered a certain proposition is not suspended about it; she simply doesn't entertain this proposition. Belief suspension requires some form of cognitive contact with the pertinent proposition. It is, as Scott Sturgeon puts it, a state of "committed neutrality" (Sturgeon 2009, 90). The exact form of cognitive contact is, of course, controversial. It has been said that suspending requires "refraining" (Moore, 1979), "withholding", or "resisting" (e.g. Oconnor, 2010) believing. This is obviously not the place to adjudicate between these issues; however, I think it is important to keep in mind that an agent may consider a proposition and yet neither believe nor disbelieve it.

Is there a candidate situation in which coordinating agents suspend belief about the other's beliefs concerning one's own choice, suspend belief about what she beliefs oneself beliefs she will do, and so on *ad inf.*.? – a situation in which precedent-based predictions go undefeated? I think there is; namely when both agents *commonly know that they have just started deliberating about what to do.*

To start seeing this, we should note that, at the start of their deliberation, the agents haven't considered any evidence yet, which is why they should suspend belief about what the other will do, why she will do it, what the other thinks oneself will do and why she thinks oneself will do it, and so on *ad inf.* In short, at the start of one's deliberation, one has simply not formed any beliefs yet, and one should resist forming these beliefs until one has suitably considered the pertinent evidence. There are two reasons in support of this thought. First, it seems that if attitude suspension is ever appropriate, then this should be when an agent has not considered any evidence. Friedman emphatically states that "it is hard to think of evidential circumstances more appropriate for suspension" than situations in which an agent has no evidence whatsoever. Second, it is reasonable to think that belief suspension is appropriate in deliberative contexts.

This is the position recently defended by Friedman (2017) who states:

> "[W]e can say that there is nothing more to "opening a question in thought" than simply suspending judgment on that question. In suspending about *Q* we make *Q* an object of inquiry. From there we can wonder or be curious or deliberate (and so on) about *Q*. Suspending about a question puts that question on our research agenda." (Friedman 2017, 26)

Deliberating, or "inquiring", about whether *Q* is true is most appropriate when we haven't settled on either believing or disbelieving it. Deliberation is, then, the kind of activity that aims at resolving this neutral state.

Suppose that, in deliberating about how the respective other will act, we're initially suspended, because we haven't considered any evidence, and, as a corollary, have not formed any higher-order belief about the other's choice. Suppose next, that we (commonly) know that we've just started our deliberative process, and, thus, (commonly) know that we are so-suspended. In this case, we don't have any higher-order beliefs about where the other thinks oneself will go. I don't have any belief about what you think I will do, and I also know that you don't have any such belief about what I think you will do. In this initial state, there are no higher-order beliefs. All potentially defeating higher-order beliefs are absent and we may, at least as far as the relevant defeaters go, permissibly predict the respective other's behavior using precedent as a standard of inference.

The following picture emerges: It is epistemically permissible for agents to predict each other's behavior based on precedent only in the absence of higher-order reciprocal beliefs about their actions. The latter condition is (for instance) satisfied when agents commonly know that they've just started deliberating about how to act and are thus suspended about what the other thinks oneself will do, what she thinks oneself thinks the other will do etc.

6. Conclusion

David Lewis wrote that salience (e.g. grounded in precedent) can support coordination by providing reasons for choosing a strategy when there is "no stronger ground for choice" (Lewis 1969, 35). Higher-order predictions about what the other player thinks oneself will choose present, I have argued in this paper, such a "stronger ground for choice". For this reason, precedent-based predictions are legitimate only in the *absence of such higher-order behavioral predictions.* More concretely, I have argued that this absence requirement is satisfied when the agents commonly know that they both suspend belief about what the respective other is going to do and why she's going to do it. This claim is directed against a philosophical doctrine according to which precedent-based predictions require a "common inductive standard"; e.g. higher-order predictions about what the other player thinks oneself will choose.

The idea that common knowledge requirements should, in the context of coordination games, by couched in terms of belief *absences* has rarely been noticed. In fact, I only know of two authors who have recognized such absences to be relevant in spelling out common knowledge in coordination games. First, in a discussion of "mutual" knowledge in communicative contexts, Martin Davies (1987, 717) suggests that "the philosophical work which was to be done by the notion of mutual knowledge should instead be assigned to a negatively characterized notion: mutual absence of doubt." The second reference is from Richard Moore's (2013, 492) discussion of common knowledge in conventional behavior. He notes that "the extent to which common knowledge is necessary for conventional activity will be determined by its coordinative role. Such a role might consist in *protecting participants in a convention from higher-order doubts* about the conformity of others" (my italics). These somewhat cursory remarks merely hint at the structural importance belief-absences have for solving coordination games. In this paper, I've tried to elaborate on this idea. Importantly, the present analysis showed that common knowledge was not simply an unnecessarily baroque theoretical element, but, rather it's presence was said to act as a defeater for reasoning from precedent in the context of solving coordination games.

References

Bicchieri, C. (2005). The grammar of society: The nature and dynamics of social norms. Cambridge University Press.

Blomberg, O. (2016). Common knowledge and reductionism about shared agency. *Australasian Journal of Philosophy*, 94 (2), 315–326.

Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information.

Cubitt, R. P., & Sugden, R. (2003). Common knowledge, salience and convention: A reconstruction of David Lewis' game theory. *Economics and Philosophy*, 19 (02), 175–210.

Davies, Martin 1987. Relevance and Mutual Knowledge, Behavioral and Brain Sciences 10/4: 716_17.

Friedman, J. (2013). Suspended judgment. *Philosophical Studies,* 162 (2), 165–181.

Friedman, Jane (2013b). Rational Agnosticism and Degrees of Belief. *Oxford Studies in Epistemology* 4:57.

Friedman, J. (2017). Why suspend judging? *Nou ̂s*, 51 (2), 302–326.

H ́ajek, A. (2016). Deliberation welcomes prediction. *Episteme*, 13 (4), 507–528.

Hájek, A. (1998). Agnosticism Meets Bayesianism. *Analysis*, *58*(3), 199-206.

Horty, J. F. (2012). *Reasons as defaults*. Oxford University Press.

Kneeland, T. (2012). Coordination under limited depth of reasoning. University of British

Columbia Working Paper.

Lederman, H. (2017). Uncommon knowledge. *Mind*, fzw072. doi: 10.1093/ mind/fzw072

Lewis, D. (1969). Convention: A philosophical study. Harvard University Press.

Moore, R. (2013). Imitation and conventional communication. *Biology and Philosophy*, 28 (3), 481–500.

O'Hear, Anthony (ed.) (2009). *Epistemology*. Cambridge University Press.

Rubinstein, A. (1989). The electronic mail game: Strategic behavior under "almost common knowledge". The American Economic Review, 385–391.

Schönherr, J. (forthcoming). Lucky Joint Action. Philosophical Psychology.

Sillari, G. (2008). Common knowledge and convention. *Topoi*, 27 (1-2), 29–39.

Sugden, R. (2015). Team reasoning and intentional cooperation for mutual benefit. *Journal of Social Ontology*, 1 (1), 143–166.

Tal, E., & Comesana, J. (2016). Is evidence of evidence evidence? *Nouˆs*, 50 (4).

Vanderschraaf, P., & Sillari, G. (2014). Common knowledge. In E. N. Zalta (Ed.), The stanford encyclopedia of philosophy (Spring 2014 ed.). Metaphysics Research Lab, Stanford University.

Ralph Wedgwood, R. (2012). The Aim of Belief. *Philosophical Perspectives*, 36 (s16), 267 – 297.