28

Temporal Indexicals

JOHN PERRY

Take any event – the death of Queen Anne, for example – and consider what change can take place in its characteristics. That it is a death, that it is the death of Anne Stuart, that it has such causes, that it has such effects – every characteristic of this sort never changes. . . . in every respect but one it is equally devoid of change. But in one respect it does change. It began by being a future event. It became every moment an event in the nearer future. At last it was present. Then it became past, and will always remain so, though every moment it becomes further and further past.

(McTaggart 1908)

1. Introduction

The expressions "now," "today," "tomorrow," "yesterday," "last month," "a year ago," "past," "future," "present," and others like them are *temporal indexicals*.[1] Such expressions, in their standard use, refer to periods of time – moments, minutes, hours, days, months, eras, and so on. Which period a given utterance of an indexical refers to depends in part on the meaning of the expression and in part on time of the utterance. (I include among utterances not only spoken episodes of language use, but also those involving writing, typing, signing, and other alternatives to speaking.)

Temporal indexicals contrast with what I shall call "dates." This will include not only dates ordinarily so called, like "July 4, 1786," "April, 1999," or "2008," but also such phrases as "before 2007" or "after 2043." There is a useful ambiguity between dates as expressions and dates as the days or other periods the expressions stand for; I will try to use the term only for expressions. The period of time a date refers to does not depend on the time of the utterance – with certain qualifications mentioned below.

Dates are akin to descriptions, in that they incorporate conditions that pick out a period of time. They are quite unlike ordinary descriptions, like "the day Queen Anne

A Companion to the Philosophy of Time, First Edition. Edited by Heather Dyke and Adrian Bardon.

died," because of their systematic nature. Temporal indexicals and dates are both quite different from *names*. There is no reason a particular day couldn't be named, and we do have names for various eras.

July 5, 1766, is a particular day that happened long ago. But what about July 5 – the "day" that comes once a year? In America it is usually devoted to cleaning up debris from fireworks, recycling beer cans, and putting away flags. What sort of thing is it? An abstract day? A property of days? We have names like "Independence Day" that pick out such days. One can refer to this sort of day with a temporal indexical. One could say, come 4th of July, "Today we always celebrate American Independence by reading the Declaration of Independence, shooting off fireworks, and drinking; then tomorrow we clean up the mess, recycle the beer cans, and sober up; by the 6th, we have forgotten all about it." Interesting – but for this essay I shall put the question of what such days are to one side.

Another important topic I will not discuss is tense. Tense is usually reckoned as an indexical phenomenon. The past-tense marker indicates that the state or activity in question occurred before the time of utterance, and so forth. I will assume, without a lot of confidence, that the discussion of temporal indexicals will illuminate at least most of the philosophically significant issues to which tense gives rise. The literature on tense, especially in linguistics and logic, is enormous. Antony Galton's "Temporal Logic" (Galton 2008) is a good place to start your investigation.

The difference between dates and temporal indexicals is connected with importantly different ways we acquire, store, and use information about things – not only about periods of time, but objects of all sorts. Temporal indexicals share a feature with our most primitive and indispensable ways of thinking: they are *role-based*. I will explore this in Section 3.

Temporal indexicals often play an important part in philosophical arguments about time. In such contexts they are often used in a sort of one-off way; their indexical sense contributes to, but does not quite capture, what the author intends to say. An example is this claim of McTaggart's in his famous essay about the unreality of time:

> But in one respect it [the death of Queen Anne] does change. It began by being a future event. It became every moment an event in the nearer future. At last it was present. Then it became past, and will always remain so, though every moment it becomes further and further past.
>
> (McTaggart 1908)

Suppose that McTaggart wrote these words late on New Year's Eve, 1907. A straightforward utterance of "future" at that time would stand for to the period of time that began on January 1, 1908 and still continues. Queen Anne died on August 1, 1714. Whatever McTaggart meant by saying, in 1907, that the death of Queen Anne "began as a future event," it surely was not meant to imply that her death ever occurred later than 1907. His words manage to say something that makes sense to us, and even seems plausible; but what? I will try to figure out what this thing is, and how the temporal indexicals "past," "present," and "future" contribute to conveying it, in Section 4. I go on to briefly consider McTaggart's whole argument about the unreality of time in Section 5.

First, however, I will develop a framework for thinking about temporal indexicals and how they differ from dates.

2. Temporal Indexicals: Basic Ideas

(2.1) provides a list of simple rules that seem to capture the standard meaning of the word 'today' as well as some other common indexicals[2]:

(2.1) An utterance *u* of "today" refers to the day on which *u* occurs.
An utterance *u* of "yesterday" refers to the day before the day on which *u* occurs.
An utterance *u* of "tomorrow" refers to the day after the day on which *u* occurs.
An utterance *u* of "present" or "now" refers to the time at which *u* occurs.
An utterance *u* of "past" [future] refers to the stretch of time up until *u* [subsequent to *u*].
An utterance *u* of "here" refers to the place at which *u* occurs.
An utterance *u* of "I" refers to the speaker of *u*.

Consider a simple sentence involving "here" and "today":

(2.2) It is sunny here today.

(2.3) gives the truth-conditions of utterances of this sentence:

(2.3) An utterance *u* of "It is sunny here today" is true if it is sunny in the place in which *u* occurs on the day on which *u* occurs.

Let *u* be an utterance of (2.2) made by Rip Van Winkle on July 3, 1766 in the Catskills.[3] We can assign two different, but consistent, truth-conditions to *u* based on our rules.

(i) (Given that *u* is an utterance of English in which the words have their ordinary meanings), *u* is true if:
(A) It is sunny in the place in which *u* occurs on the day on which *u* occurs.[4]
(ii) (Given all of that, *plus* the fact that *u* occurred in the Catskills on July 3, 1766) *u* is true if:
(B) It is sunny in the Catskills on July 3, 1766.

(A) gives us what I will call the "reflexive" or "utterance-bound" truth-conditions of *u*. It identifies the truth-conditions of *u* by identifying the time and place referred to in terms of *u itself*. The truth-conditions (A) provides are conditions *on* the utterance *u*. (B) on the other hand provides truth-conditions *of u* that are not conditions *on u*, but on the time and place referred to in *u*. I call these the *referential* truth-conditions of *u*.

The utterance-bound and referential truth-conditions of *u* are completely consistent. (B) tells us *what else* the world must be like, given the truth-conditions specified by (A), *plus* the additional facts about when and where *u* occurred. Given those additional facts, the truth-conditions of (A) will be met if and only if those of (B) are met.

Still, the two kinds of truth-conditions get at significantly different properties of the utterance *u*. To explore this point, it is convenient to introduce propositions[5] that encode the truth-conditions, which we can call the utterance-bound and referential *contents* of *u*.

(2.3u-b) That it be sunny in the place in which ***u*** occurs on the day on which ***u*** occurs.
(2.3ref) That it be sunny in the **Catskills** on **July 3, 1766**.

I have used boldface to indicate what these propositions are about. (2.3u-b) is about *u* itself; (2.3ref) is about the Catskills and the day, July 3, 1766.

With dates, we can specify which period of time an utterance refers to in a way that does not depend on when the utterance occurs. For example:

> An utterance *u* of "July 3, 1766" refers to the third day of the seventh month of the sixty sixth year of the eighteenth century.

It does not matter when the utterance occurs, as long as the date is being used in the standard way.[6] For the sake of this essay, we can also assume that each utterance of "the Catskills" can also be given a reference, independently of any additional facts about the utterance, namely, the mountains that begin west of the Hudson River and a bit north of Kingston, New York, and extend west as far as Delaware County. Thus, given the meanings of the words and phrases in (2.4):

(2.4) It is sunny in the Catskills on July 3, 1766.

It seems we can give referential content, without needing any additional information about the time of utterance:

(2.4ref) [= (2.3ref)] That it is sunny in the Catskills on July 3, 1766.

An utterance of (2.4) will have the same referential content as Rip's utterance *u*, but not the same utterance-bound content.

Now consider another utterance *u*' of (2.2) a day later, also by Rip, still hiking in the Catskills. The two utterances *u* and *u*' will have different referential contents: that it was sunny on July 3 in the Catskills, and that it was also sunny on July 4. If the weather changed, the second might be false. The two utterances will also have different utterance-bound contents: that it was sunny at the time of *u* in the location of *u*, and that it was sunny at the time of *u*' in the location of *u*'. But of course there is something quite similar about these propositions; they require *different* utterances to meet the *same* condition:

> being such that it was sunny on the day *u* was made, in the place *u* was made.

We need a name for what they have in common. I will say that the utterances share the same "truth-requirement." The truth-requirement on an utterance is a property that it may or may not have, which is determined by its syntax and semantics. The utterance-bound truth-condition is a proposition determined by: (a) the

truth-requirement; and (b) the utterance itself. Rip's utterances *u* and *u*' have the same truth-requirement, but not the same utterance-bound content, and not the same referential content.

We now have two equivalence classes for Rip's utterance *u*. There are those utterances that have the same truth-requirement, and those that share referential truth-conditions. The first set includes *u*', the second does not. The second contains an utterance of (2.4), the first does not. These two classes help us think about importantly different properties of the utterance, that project onto other utterances and events connected with them in quite different ways.

Utterances that share truth-requirements will share various epistemic, cognitive, and pragmatic properties. The truth or falsity of any utterance of (2.2) can be determined by looking around where it takes place on the day it takes place. But such utterances may differ in truth-value. Utterances that share referential truth-conditions will always have the same truth-value, but may be associated with quite different sorts of evidence and ways of thinking and acting.

Frege notes that to say the same thing tomorrow, and one says today using "today," one will need to use "yesterday" (Frege 1918). If Rip were to say, on July 4:

(2.5) It was sunny here yesterday.

In the circumstances we are imagining, while he is still hiking in the Catskills, then his utterance would have the same referential content as the utterance of (2.2) made the day before: that it was sunny in the Catskills on July 3, 1766. Intuitively, we would say that he said the same thing, in two different ways. If I say today, "It was sunny in the Catskills on July 3, 1766" then, at least for many purposes, we would reckon that I said just what Rip said by saying "It is sunny here today" on July 3, 1766, and by saying "It was sunny here yesterday" on July 4, 1766. But my utterance would be based (at least if I were more scrupulous than I am) on consulting historical records, not by looking around or remembering (as Rip's were).

For many purposes, it seems that utterances with the same referential contents count as cases of saying the same thing, in spite of differences in the sort of evidence, cognition, and action with which they are associated. That is, for many purposes, we can equate "what is said" by an utterance, or perhaps more carefully, by the speaker of the utterance, with the referential content of the utterance. In such cases we may focus on the "counterfactual possibilities." Consider the counterfactual circumstance in which Rip does not utter *u*. In that circumstance, the utterance-bound truth-conditions of *u* would clearly not have been met, for there would be no utterance to meet them. But in such counterfactual circumstances, it seems that *what Rip actually said* would still have been true; Rip's silence wouldn't have had an effect on the weather. This test for what is said argues against taking the utterance-bound content as what is said; at least in many cases, referential content works better.

On the other hand, suppose I ask you when a committee of which we are both members next meets. This conversation occurs on May 5, 2011; the next meeting is May 6, 2011. You might truthfully say either, "It meets tomorrow," or "It meets May 6, 2011." If you say the latter, and I miss the meeting, not realizing the date, I might feel you did not tell me something important. The two utterances have the same refer-

ential content. But, at least for certain purposes, it does not seem correct that you told me the same thing you would have told me, if you had said "It meets tomorrow." You left out an important piece of information – that it would meet *tomorrow* – and included a piece that, knowing my disorganized way, you might have guessed was irrelevant – that it would meet *on May 6*. This is a version of Frege's classic problem about identity (Frege 1892). "Tomorrow" and "May 6, 2011" provide different "modes of presentation" of the same day. The difference has cognitive significance, and seems relevant, in this case, to what we take you to have said.

"It is sunny here today" is the sort of thing one is inclined to say on the basis of a certain sort of evidence: namely, looking around one and seeing that it is sunny. When one opens one's eyes and looks around, one finds out about the weather on a certain day and in a certain place: namely, the day on which one does the looking around, and the place where one does it. When one says, "It is sunny today here," one says something about a certain day and a certain place: namely, the day on which one makes the utterance and the place where one makes it. Now if one looks around to discover what the weather is like, and then with reasonable promptness announces what he has discovered, there will be a certain phenomenon I call *role-linking*. The day and place play two roles in the agent's life: being the day on which he looks around and the place where he looks around. These are linked to a second pair of roles: being the day he makes his utterance and being the place where he makes it. In virtue of the meanings of "'today'" and "'here'," these pairs are linked to a third: being the day referred to and being the place referred to. Given this role-linkage, what he says will do a good job of expressing what he discovers.

If you had said, "The next meeting is tomorrow," you would have linked the occupant of two roles: being the day after the conversation took place, and being the day to which you were referring, and saying it had the property of being when the committee next met. When you said, "The next meeting is May 6, 2011," you left that role-linking up to me. This is the reason I feel you left something out, even though the unhelpful remark you made and the helpful one you could have made have the same referential content.

As we shall see, role-linking is the key to understanding why temporal indexicals (and other sorts of indexicals) are so useful, and the sorts of thoughts they express so important.

3. Two Ways of Thinking about Things

Our system of temporal indexicals and our system of dates correspond to two different but interconnecting ways we think about things, which I call *role-based* and *descriptive*. Both contrast with a third way, which I call *via detached notions*.

3.1. Roles

Roles are relations, represented in a certain way. Where R is a binary relation, we distinguish between its domain and its range. A relation is functional if for each object in the domain there is one and only one object in the range that stands in that relation to

it. It is a partial function if there is at most one object that stands in that relation to it. Partial functions, indexed to a particular "subject" or member of the domain, are roles.

By "indexed" I mean that the subject is not supplied by the usual grammatical techniques (names, descriptions, quantifiers), but in some extra-grammatical way. Some examples of indexing:

1. We have a two-place relation symbol $R(,)$ and a domain of objects a, b, c, ... We form one place predicates $R_a(\ )$, . . . ,$R_b(\ )$ . . . . "$R_a(z)$" is true just in case $R(a,z)$. Here the subscripted "a" is the index, R is the role, and z *plays* the role (relative to a) or is the occupant of the role (relative to a).
2. We have a fairly large card or piece of paper that stands for an object or event – say a party we are throwing. On the surface we write down facts about the object or event (Figure 28.1). We do not have to write down expressions that refer to the event. We simply identify the objects that play various *roles* relative to it:
3. At a junction on a path, we might find an arrow-shaped sign (Figure 28.2) stuck in the ground.

The sign somehow conveys to us that we can find a mess hall in the direction the sign points. Presumably there are mess halls in almost any direction from any location on earth, if one goes far enough, so to have a function we have to suppose that the role is being the direction of the nearest mess hall, or the mess hall for campers in the campsite where the sign is. The indexed subject is the location into which the sign is stuck, or

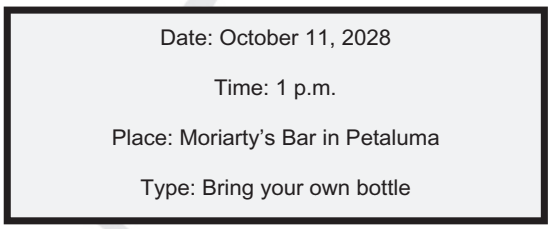

Figure 28.1

Figure 28.2

the location of a camper who sees the sign. The role is being the direction of the appropriate mess hall. The occupant of the role is the direction in which the sign points.

All the rules of (2.1) fall into this pattern. The utterance itself is the indexed subject. The roles are utterance-based roles: being the speaker of, being the time of, and so forth. The occupants of the roles are the objects referred to by the utterances of the expressions.

3.2. *Role-based Cognition*

We also have role-based *cognition*. Indeed, the impressions and ideas that are involved in the most primitive kinds of thinking are clearly role-based representations. Imagine an animal, a hen say, seeing a kernel of corn on the ground in front of her. She has a certain kind of impression. Her impressions differ, we may suppose, depending on the *kind* of object that she sees: whether it is something like a kernel of corn, or something quite different, like a golf ball. But the impression not only carries information about the kind of the object the hen sees, but also about its distance and direction from the eyes of the hen, and so from the hen itself. This information is crucial if the hen is to get herself into a position to peck at the kernel and eat it.

Perception gives us information about the kinds of objects that play various roles in our lives: the things we are touching, seeing, chewing on, listening to, and so forth. Consider the sort of perception our hen had of the kernel of corn. Any perception of this sort will be veridical if there is a kernel of corn at a certain distance and direction from *the hen that has it*. So here, as I look at it, the hen is the indexed subject for the representation. Certain aspects of the visual impression the hen has represent the distance and direction from *the hen whose impression it is*. Other things about the impression represent additional or *incremental* information about the object that is in that direction at the distance: that it is a kernel of corn, or something small, peckable, and most likely edible, not a golf ball, for example. The impression is a role-representation, indexed to its possessor.

The hen's brain will connect this sort of impression with certain kinds of actions, ones that will bring the hen that performs them into a position to peck at an object that distance and direction from it. In this system there is an important bit of *architectural role-linking*; *the hen* that has the perception, *the hen* that the perception motivates to move, *the hen* that ends up in a position to peck the corn, and *the hen* that gets nourished if all goes well, will all be the same hen. The burden of keeping these roles linked does not fall to the hen; it is simply a matter of nature's architecture for chickens.

As the hen walks towards the kernel of corn, there is another kind of role-linking involved. The objects seen by the hen at successive moments, as she approaches the kernel of corn, will be the same. This I will call *environmental role-linking*. Given the way things work in the hen's natural environment these perceptions will be of the same object. The demands the whole transaction puts on hens' cognitive capacities is not too great, thanks to architectural and environmental role-linking.

Another kind of architectural role-linking is also important. The hen's sensory states contain information about the distance and direction of the kernel of corn from her eyes. Given the ways normal hens work in normal environments, the effect of moving her legs in the way she does will get her a certain distance in a certain direction. A

normal chicken will move her legs so that she advances in the same direction as the kernel of corn is relative to her eyes. The chicken's architecture, its anatomy and nervous system, link the role of *being the direction the perceived edible is relative to her eyes* and *being the direction in which the movements of leg and claw normally put in motion by such perceptions will carry her*. Some clever animal scientist could, no doubt, screw this up. The linking is contingent, but architectural. As long as the hen stays out of laboratories and philosophy thought-experiments, she need not worry about it.

The similar impressions of different hens, having different perceptions, will have the same "perception-bound truth-requirement," but different truth-conditions. This is the hallmark of role-based cognitions: the same state, at different times or in different agents, gives information about different objects – different kernels of corn, in the case we are imagining. But the role these different kernels play is the same, and will be architecturally connected with other roles they play. This sort of cognition is all that is involved in our most primitive "information-games," in which we pick up information through our senses about the world around us, and then – the same agent at pretty much the same time – behave in ways that make sense given that information.

3.3. Detached Cognition

But humans also have a number of different representational systems, which put rather complex cognitive burdens on them, and which provide what I call *detached* ways of thinking about things. By "detached," I mean: not attached to a current perception or short-term memory.

I meet you at a party. I see you and talk to you. I learn about you as the occupant of the roles: person looked at, person talked to, and so forth. I do not need to keep track of the fact that the same person is seeing you and talking to you, any more than the hen had to keep track of who her perceptions were going to affect. I have a perceptual-conversational-role-based *cognitive fix* on you. Then you leave. I remember you. I have a cognitive fix on you that is *detached*, rather than attached to any particular perception or short-term memory of you.

In this case, my detached idea, or *notion* – the term I use for ideas of things – is *of* you because of how it was formed. During our conversation I have noted several important things about you. I attach ideas of these things to the notion of you, to create sort of inner files. Some things, like your profession, name, hometown, and position on, say, compatibilism in free-will, will be useful next time we meet; I can build on what I already know about you, in deciding what to call you, which topics to ask you about, and perhaps which topics to avoid. For this file to be of much use, it will have to include things that will allow me to recognize you, perhaps an image, or get in touch with you, perhaps a phone number. What makes this notion *of* you, however, is neither the accuracy of the information *in* the file associated with the notion, nor its providing a unique characterization of you, which it may not do, but rather the fact that you were the origin of the notion, and the source of most of the information in it. If things work out as they should, you will also be the person the information is applied to on subsequent occasions. The system is of value because it works the majority of the time, although misidentification and lack of recognition are common phenomena.

3.4. *Thinking about Periods of Time*

These distinctions apply to our perception of and thoughts about events and periods of time. Consider Rip van Winkle again. He wakes up on July 5, 1786, twenty years and two days after he fell asleep. He remembers July 3, 1766 "as if it were yesterday"; so much so, that he thinks he can refer to it that way. He says, "I fell asleep yesterday," "Yesterday was a nice day in the Catskills," and so forth. What he says is false; he did not fall asleep on July 4, 1786, and, as it happens, it was not a very nice day in the Catskills. Still, he has a memory of how things were on the day he fell asleep, July 3, 1766. He has a detached file of that date, with some misinformation in it; although he is thinking of that day, he cannot refer to it in the way that he thinks he can.

Dates are the most common way we have for referring to and thinking about days more than a few days in the past, or more than a couple of days into the future. Periods of time, unlike people, are related to one another in stable and orderly ways that are incorporated into our system of dates. Because of the orderly way periods of time are related to one another, and dates are related to the periods they stand for, we can deduce a lot about the relation among days by considering relations among the dates that identify them. Two people with adjacent social security numbers may be thousands of miles apart at any given time. But July 4, 1766 happened one day after July 3, 1766 and has reliably remained one day later since then.

Dates identify days in virtue of their properties, and so are of the same general type as descriptions. But the fact that they are systematic, and that there are no missing days, makes them a very special type of description. If we can frame a coherent date, we have a way of referring to a period of time, and saying things that are true or false, or will turn out to be true or false, based on what happened, or will happen, during that period of time.

Given human practices and institutions, dates are a powerful way of learning about days. The information about the past in libraries, newspapers, and other sorts of chronicles is organized by date. What we know or think we know about the future is organized into almanacs, calendars, day-books, schedules, agendas, conference programs, and the like, all organized by date.

All this information is of no use to us, however, unless we can fit it into our system of role-based ideas, and temporal indexicals, which is to say, it is of no use to us unless we can fit it into our lives. We do not have a view from nowhere or from nowhen, much less the ability to do things from nowhere or nowhen. What we directly perceive about the world is how things are *here* and *now*, and we directly act on the world by moving bodies and limbs and mouths here and now. When we represent things in ways that abstract from the role those things play in our lives, our representations are communicable, portable, and potentially quite useful – but only when they are re-attached to someone's *here* and *now*, *this* and *that*. A traveler's guide to London can be read and enjoyed, and even corrected and expanded, by people all over the world. But to use it, one needs to be in London and connect the things named and described in the book with things one sees around one. Then the guidebook tells us *what else* these things are like, other information about them. *Given* that *that* building is The Royal Albert Hall, *it* was opened by Queen Victoria in 1871.

Suppose it is October 7, 2011. I wake up. I can easily know a lot about that day, under the role it plays on that day in my life, *being today*. I look out the window and see that it is sunny *today*, at least in the morning. I notice that I am stuffed up, and seem to have a cold *today*. I know that I am awake, that I am alive, and I am out of bed and functioning, however inadequately – all this is happening *today*. I stumble down to the kitchen, to get the coffee brewing. I know if I can brew it *now*, I can drink it *today*, indeed, *this morning*.

There in the kitchen, on the refrigerator we will suppose, is a huge calendar, covering at least the month of October, 2011. It has to be huge, for the purposes of this essay, since I also use it as a diary, adding information to the cells about what happened on a given day. It contains lots of information about today, under the date "October 7, 2011." Alistair and Johanna are in town and she is giving a talk. I have a tentative lunch date with Alistair. But the calendar on the refrigerator does not contain one important item of information about October 7 – that it is *today*. I need to figure that out, before I can incorporate the information on the calendar into the plans and expectations I have for the rest of the day. Until I do that, having the calendar is like having a guidebook which tells me that Queen Victoria opened The Royal Albert Hall in 1871, while I am still lost in the streets of London, having no idea which building is which.

In either case, what is needed is role-linking. *That building*, the one I see before me, is the one the book calls "Prince Albert Hall" and provides further information about. *Today*, the day I perceive facts about when I look around, the day on which actions I now undertake will happen, is the one my calendar calls "October 7," and so the day that cell on my calendar provides further information about.

A lot of the roles today plays in my life are architecturally linked. I do not need to worry about whether the day whose weather I now perceive is the day whose weather will determine how wet I get when I go outdoors. But this is not the case with *being today* and *being the day the cell of my calendar I am looking at contains information about*. I have to take steps to get those roles linked.

Assuming I do not just remember what yesterday's date was, and so infer today's date, various things help me with this role-linking task. Perhaps I keep only the current month on the refrigerator; that narrows things down a bit. I remember working yesterday, so it is not Monday. I may find entries in some cells that I am sure I have not done yet, and in other cells things I am sure I have done. I may go out on the front porch and pick up the paper, and look at the date. One way or another, I can usually figure out which cell corresponds to *today*.

Of course, these methods are a bit old-fashioned. Nowadays I can look at the news on the TV, and the date will probably be displayed across the bottom. Or I can just look on my cell-phone. And then there is the method I am often tempted to use, but usually do not: wait until I have missed my first appointment, when someone will call me and tell me what day it is.

Once I have figured out which day it is, there is a transfer of information, from date-based ways of thinking to role-based ones, and vice versa. As I am imagining my practices, for the purposes of this essay, I also look around and jot down on the calendar what I see: "Nice day; no wind; sunny sky; neighbor's dog barking, etc." And I plan my day based on the appointments and reminders I find already in the cell, and write down everything I decide to do.

3.5. *Dynamic Representations*

Now let us change things a bit. Since the calendar is on the metal door of my refrigerator, I can use a little magnet to mark what day it is. Alternatively I can write "today" in the right cell, once I figure it out. Each night at midnight, as I go to bed, I simply move the magnet to the next cell, or erase "today" and rewrite it in the next cell. The next morning the calendar provides *more* information than it would without the magnet or the notation. As long as my midnight practice can be counted on, I can look at the calendar and get an answer to the question, "What is today's date?" This is just the sort of thing that computers now do for us effortlessly. But I will just think about my calendar, refrigerator, and magnet.

We can think of the magnet being stuck to a cell as an utterance that lasts twenty-four hours, from midnight to midnight. The indexed subject is the day of the utterance, that is, the day it is stuck to the cell for. This will be the same day on which it is perceived, at least if the perceiver is standing in front of my refrigerator. The role is the same as with "today": being the day of the utterance. The presence of the magnet on the cell amounts to an indexed identity statement.

The calendar, without the magnet, is basically what we might call, following McTaggart, a "B-representation." Suppose I have another huge chart on the wall. In the middle of the chart there are three columns, labeled "Yesterday," "Today," and "Tomorrow." There are additional columns to the left of these three, labeled "the day before yesterday," and "the day before the day before yesterday," and so on, and similarly on the right with "the day after tomorrow," and so on. All the same entries that are on the October page of my calendar are found in the appropriate columns on this chart. This chart is an "A-representation." The B-representation seems like it could represent everything that happens. After all, whatever happens, happens on some day or another. And it seems, if the entries are correct, they never need to change. But it seems to leave something important out: which day is today. The A-representation tells me that, so it does not leave something out. But unlike the B-representation, it must change every night at midnight to stay accurate.

When I use my magnet on my refrigerator B-representation, I turn it into sort of a hybrid between A- and B-representations. This works, so it must be coherent. But is it not puzzling? If what the magnet's being there represents is true, it must be because of something that happens on October 7. But if it happened on October 7, then that is when it happened. So why do I have to move the magnet at midnight, to keep it from containing a falsehood?

The right way to look at the calendar, without the magnet, is as systematic but detached information. It is partial. It is like the guidebook; the entry on Prince Albert Hall tells us what *some* building is like, and the rest of the book tells us what other streets and building related in other ways to it are like. But it just tells us what *else* some building is like; we have to fit what it tells us on to the right building in our experience, on to perceptual and pragmatic roles that it occupies, before we have complete information. The additional information is the role the things described *in* the book have to something outside, the use of the book.

The magnet – as long as I keep it up do date – links the role of *being the day the information in a certain cell is about*, and *being the day on which the calendar is being perceived.*

What changes is not what happens on October 7, but what role that day plays in my life – or the life of anyone who observes the calendar on that day, with the magnet in that cell. The day that occupies the role, *being the day on which the calendar is used*, changes every twenty-four hours. In the same way, but less regularly and systematically, which building plays the role, the *building I am looking at*, changes as I walk around London, and I need to keep this role linked with the *building I am reading about* to find the guidebook useful.

The idea of incorporating the information that the magnet gives, the role-linking information, into the B-representation in some permanent way makes no sense. The magnet links the day a cell is about to the occupant of the role of being the day the calendar is viewed. And that changes.

The magnet's being on the date in question is a representation, an utterance of sorts, one designed to last for twenty-four hours; call it *u*. It has utterance-bound and referential content. The referential content is the same as an utterance of "October 11, 2011 is October 11, 2011" would have. That is not very helpful. The utterance-bound content is more promising: that *u* occurs on October 11, 2011. Since I am seeing *u*, that is I am seeing the magnet in that calendar cell, I know that my perception is occurring on October 11, 2011.

The calendar, plus the magnet, plus my practice of moving the magnet at midnight, constitutes a "dynamic representation." Of course, computers now provide us with dynamic representations – automatically updated calendars. And watches and clocks are dynamic representations, although the information they offer about the time at which they are examined is rather sparse – simply the "o'clock" properties of the time.

4. Moving from Future to Present to Past

The second part of the remark by McTaggart at the beginning of this essay seems right. We know what he is getting at when he says the death of Queen Anne was, for a long time, a future event, and then it became a present event, and since then it has been a past event. This seems sort of right. But it is not altogether easy to see *why* it seems sort of right. The most natural interpretation of "future," "present," and "past" treats them as indexicals. An utterance of "future" in 1907, when McTaggart was finishing his article, would have referred to the era that began then and still continues. The death of Queen Anne was *never* part of that era.

Suppose you lived in England in 1714. Queen Anne had a stroke on July 30, and died on August 1. The heir apparent was a German, the elector of Hanover, who became George I, but there were Catholic claimants, including a half-brother of Anne's, so you might well have worried about political turmoil. But things went relatively smoothly, at least for a while. By 1716 you might have said, perfectly naturally, with no hidden philosophical agenda:

> When Queen Anne's death was imminent, clearly in the near future, I worried if the nation would stay calm. But now that it is in the past, I have no such worries.

So consider:

(4.1) On July 30, 1714 the death of Queen Anne was in the future.
(4.2) On August 2, 1714 the death of Queen Anne was in the past.

The intuitive truth-conditions of (4.1) are not mysterious. (4.1) is true because Queen Anne died after July, 1714. Had she died on or before July 30, 1714 it would be false. (4.2) is true because Queen Anne died before August 2, 1714. Had she died on August 2, 1714, or later, it would have been false. What is a bit puzzling is how these statements come to have those truth-conditions, since it is not what our theory of temporal indexicals, as developed so far, predicts.

What is at issue are complex phrases such as:

(4.3) was future
was in the future
was past
will be present
was future on March 30, 1714
was past by August 2, 1714
and so forth

In each case we have a temporal indexical ("past," "present," "future") combined with a tensed auxiliary verb ("is," "was," "will be"). What is the tensed auxiliary verb doing there?

I call this phenomenon "de-indexing." We have an expression, like "tomorrow," that has two features: (a) it is associated with a relation between days; and (b) it is indexed to the day of utterance. When we "de-index" the expression, we preserve (a) but give up (b). We still use it for a relation between days, that rather than the subject-day being supplied by indexing, it is supplied with some more ordinary grammatical device.

If I say:

> Tomorrow I go to the dentist.

I say, about the day after the day on which I make the remark, that I will go to the dentist then. If I say:

> Tomorrow is always a lousy day to go to the dentist; yesterday is much better.

I am saying that on any given day *d*, it is lousy on *d* to have to go to the dentist on the next day; it is better to have gone the previous day. "Tomorrow" is de-indexed from the day of utterance, and the domain object is supplied by the quantifier "always."

If you and I are having a debate, with you saying:

> One should never put off until tomorrow what you can do today.

While I insist:

> One should never do today what can be done tomorrow.

It is most likely that we are debating the merits of the general policies:

> For any given day *d*, do not put off until the day after *d* what you can do on *d*.
> For any given day *d*, do not do on *d* what can be put off until the day after *d*.

Rather than some peculiar feature of the day of our conversation and the one following.

So too it seems that "past" has a non-indexical use that derives from its indexical use. In its indexical use, an utterance of "past" refers to the period of time before that utterance. In its non-indexical use, "past" refers to a property an event has at a time, if it occurred before that time.

With "today," "tomorrow," and "past" the most natural use of the words is as indexicals, so that they need to be "de-indexed" before they can be used with non-indexically supplied subjects. With other words, things go the other way around.

Consider the expression "father." *Being the father of* is a relation. But "father" has a use as a role-word, where it refers to the father of the speaker of the utterance. If Mitt Romney says, "Father saved American Motors," his use of "father" refers to his father, George Romney. But if he supplies another person to serve as the son in question, "father" functions as a normal relation word. If Romney says, "Newt's father saved green stamps," he will have referred to Gingrich's father, not his own.

Other expressions seem neutral between their use as indexicals or simply as relation words. Consider "local." This has an indexical use, as when I say, "Let's try the local bar." "Local" provides the relation of *being in the same neighborhood as*. The bar I am suggesting we try is the one located in the neighborhood in which my utterance occurs. But I can also say, "When we get to Amy's house, let's try the local bar." Here the bar has to be local relative to Amy's house, not to the place of my utterance. In the first use, the subject neighborhood is supplied indexically; in the latter, it is supplied, still somewhat indirectly, by the reference to Amy's house.

When we de-index words like "tomorrow" and use them as relation words rather than for roles, there is a psychological element that is quite important. As we saw, utterance-relative roles are closely associated with epistemic and pragmatic roles, and also with the states of perception, memory, expectation, and intention that are associated with picking up information via those roles, and acting in accordance with such information. These connections are retained in de-indexing. The implication of the advice I urged above something like:

> When you find yourself thinking, "I can do it tomorrow," resist, and try to think "I will do it today" instead.

To make sense of McTaggart's remark we need to do two things. First, we have to realize that "past," "present," and "future" are de-indexed; they are not functioning as indexicals, but as relation words with cognitive overtones due to their standard indexical use.

Second, we have to realize that McTaggart is unapologetically doing what Bergson called "spatializing time." He is treating the future and the past as places. Whatever McTaggart thought, this is only a metaphor. The future is not a sort of place where events that have not happened yet wait for their time on stage, and the past is not a sort of place where events go after they leave the stage. To say an event is *in* the future is to say that such an event has not yet happened, but will. To say that an event is *in* the past is to say that it has happened. To say that an event was in the future, but is now in the past, is to say that while there was a time before such an event happened, such an event now has happened.

For convenience here is what he says about Queen Anne's death:

> But in one respect it does change. It began by being a future event. It became every moment an event in the nearer future. At last it was present. Then it became past, and will always remain so, though every moment it becomes further and further past (Taggart 1908, 460).

If we suppose that the phrase "when it began" amounts to something like "when it began to be an object of concern," the following seems a pretty good paraphrase of the thought that the quote conveys:

> The first time anyone worried about the death of Queen Anne, it was a future event. It became every moment an event in the nearer future. At last it was present. Then it became past, and will always remain so, though every moment it becomes further and further past.

If we get rid of the metaphor, this amounts to the following, and seems true:

> There was a period of time, starting with the first moment the possibility of Queen Anne dying occurred to anyone, and lasting until the time of her death (August 1, 1714), relative to which the death of Queen Anne was in the future, if that simply means it happened after that period. Then there was a period, (or perhaps just an instant, depending on the metaphysics of death), relative to which it was present, that is, it happened during that period. And for some time before now, and from now on (whether we take this to have been written in 1907, when McTaggart wrote, or in 2012, when I do), it is in the past, that is, it happened before every part of that period of time.

If when one reads McTaggart's words, they seem true, this is the truth they are conveying.

So, similarly, it makes perfect sense to say, as I stand before my calendar at the stroke of midnight on October 7:

> October 7 is no longer today; now it is yesterday, and October 8 is today.

Perhaps I feel a bit sad, as it seems these changes happen more quickly as I get older.

"No longer" is an indexical, meaning roughly "before now, but not now." I am saying "October 7 was today, but it is now not today," using a de-indexed "today." The "was" in the first clause, and the "now" in the second, provide us with the members of the domain that the de-indexed "today" requires. Relative to October 7, October 7 played the *today* role; relative to October 8, it does not play that role. But it still plays that role

relative to October 7, and it never did play that role relative to October 8. So this does not mean that there has been some change in the characteristics of the events of October 7 or October 8. The sense in which, as the minute hand moves past twelve, October 7 is no longer today, is the de-indexed sense of "today." It is no longer today not because it has changed, but because the relevant domain object has changed, from October 7 to October 8.

5. Temporal Indexicals and the Passage of Time

McTaggart's argument for the unreality of time has two sides. One we have just been looking at: he says we think that events do change, going from future to present to past, and that this is essential to time as ordinarily conceived. Queen Anne's death was once future, then present, and now it is in the past, using "past" indexically, but "present," and "future" in a de-indexed manner. This does not mean that something happened on a certain day, and then as time passed, no longer happened on that day. Nor does it mean that some event occurred during some period of time, and then subsequently did not occur during that period of time. Events do not have to move about in time; events do occur, but whatever period of time they occur in, there they stay.

The other side of his argument is that it is equally essential to time, as we conceive it, that events stand in temporal relationships of happening before, being simultaneous with, and happening after, which *never* change. According to McTaggart, these two essential properties of time cannot be combined; our concept of time is incoherent; there is not and could not be anything that fits the conception; time is unreal.

Leaving Queen Anne aside, let us consider the elimination of the San Francisco Forty-Niners' chance of going to the SuperBowl in 2012. The elimination occurred about 6:30 PST January 22, 2012, when the New York Giants kicked a field goal to win in overtime. I was observing this sad event as it occurred (on television, but I will ignore that complication).

I might say, McTaggart style, that I worried about the elimination all day up until it occurred; all that time it was a future event. Then, in an instant, it became a present event. Since then it has been a past event, no longer feared, merely regretted. So what happened when I saw it go from being future, to being present, to being past? Did the event of the Niners' elimination change in some way that our concept of time both requires and precludes?

One view that is sometimes entertained is that the event already existed; the idea that I see it coming into existence is a bit of an illusion; what changed was that I perceived it or came to know it. But then we seem to have the same problem for my perception, as we did for the Niners' elimination. My perception of the Niners' elimination, like the Niners' elimination, seemed to come into existence. This whole idea seems quite hopeless.

A second view is that what I see is not an illusion, but not quite the event coming to exist either. There is some property of presentness, which events acquire and then lose; they do not come into existence, but come to have the property of being present. This is perhaps part of the conception of time that McTaggart thought we had; that events start in the future and *move into* the present, and that is what it is for them to become

present, and what it is for them to happen. He was right that this view is incoherent, or at any rate, quite wrong.

The third view is that things are just as they seem. When the ball flew between the uprights, I saw the elimination of the Niners *occur* or *happen*. The event did not exist before then, because for an event to exist is for it to happen, and it had not happened yet. I saw it come into being. I did not see it *move* from one place, the future, to another, the present.

To avoid McTaggart's dilemma, on this view, we need merely to deny that, until it happened, the event of the Niners' elimination stood in any relation to anything. When the field goal was kicked, the Niners' elimination happened, that is, came into existence, and it was only then that *it* stood in any temporal relation to any other event. What has changed about the event since then are not its characteristics. It has come to stand in new relations to new events, as they have happened after it did.

If this is the right way to look at it, then McTaggart's second claim is not right. As new things happen, all the events that have already happened come to stand in new relations to these new events. Queen Anne's death was not subsequent to the beheading of Charles I, until she died. It was not prior to the death of Queen Victoria, until Queen Victoria died. If I ever finish this paper, then the event of my agreeing to do it will pleasantly enter into a new fact, the fact that the agreement is followed by its fulfillment. What McTaggart had right is that, on our ordinary conception of time, once these facts come into being, they do not change. But he was wrong to deny that new events, and hence new facts, and new facts about old events come into existence. They do not move from the future into the present; they just happen.

But if we take this line about this side of McTaggart's argument, does that not undermine what we said about Queen Anne's death? We agreed that, before August 1, 1714, it was future. Now we are saying that it did not exist until it happened, and so did not stand in any relations to anything else until it happened. If so, it was not future relative to, say, July 1, 1714 on July 1, 1714, because it did not exist yet.

Suppose some member of the court said, on July 1, 1714, "Queen Anne's death will occur in August, 1714." From our current point of view, this can seem pretty baffling. For one thing, what is the phrase "August, 1714" doing? It seems to be referring to a period of time. But does that period of time already exist, in July? Are we to suppose that none of the events that occur in a future month exist, until they happen, but that the month is somehow already there waiting for them?

What to say about this is complicated, not because there is nothing plausible to say, but because there are a lot of alternatives. One can hardly be confident about one's choice. Nevertheless, here is the way I look at it. Call our courtier's utterance on July 1, 1714 *u*. On July 1, *u* had reflexive truth conditions:

> *u* is true if there will be a time *t*, and a period of time *T*, both subsequent to the time of *u*, such that Queen Anne dies at *t*, *t* occurs during *T*, and *T* is August, 1714.

Given that *u* occurs on July 1, 1714, we get what we can call the time-bound truth-conditions:

> *u* is true if there is/was/will be a time *t*, and a period of time *T*, both subsequent to July 1, 1714, such that Queen Anne dies at *t*, *t* occurs during *T*, and *T* is August, 1714.

Once we get past August, 1714, and that the relevant period of time exists, *u* has referential truth conditions:

> *u* is/was/will be true if there is/was/will be a time *t*, such that Queen Anne dies at *t*, and *t* occurs during August, 1714.

I am inclined to think that August 1714 did not exist as of July 1, 1714, but rather came into existence as the things continued to happen. It was nevertheless possible to refer to it ahead of time. As an analogy, suppose on July 1, 1714 a hunter shot a very powerful gun. The bullet flew for days, until a week later, on July 8, it hit and killed a fawn who had been born the day before. One thing we might say is that the hunter shot a fawn that did not exist at the time he fired. To kill a fawn, one might suppose, is to initiate a course of events that leads to the death of the fawn, and the hunter did this. The hunter shot a particular fawn, even though that fawn did not exist when he did the shooting. When, in early 2012, I said, "The next Presidential election will occur in November, 2012" I referred to November, 2012. I did everything I could do to complete the conditions necessary for my utterance to be such a reference, and I was confident the rest of what was necessary would be supplied in due course. I referred to a month that did not yet exist at the time I referred to it. I was unlike the hunter in one important respect. Whereas it seems unlikely that he could have known ahead of time that his present movement would, eight days hence, come to have effects that made them a case of shooting a fawn, I was supremely confident that as time unfolded the conditions necessary for my utterance to be a reference to November, 2012 would develop. That is, I was quite certain that November would eventually happen.

Now what of the use, on July 1, 1714, of "Queen Anne's death"? We could say something analogous. Our member of the court referred to a death, and a period of time, and said that the one would occur during the other, and was quite right as it turned out. However, I do not think this is the right thing to say in this case.

Suppose our courtier was a member of the anti-German cabal. He said, calling the future George I "the elector of Hanover":

> The coronation of the elector of Hanover as King of England will never happen.

Here it does not seem that we want to say that he referred to the event of the coronation of the elector, since if what he said had been true, there would have been no such event. It seems that he wrongly predicted that there would be no such event, no coronation of the elector as King. In doing so, he unintentionally used an event-description that was eventually satisfied by something. Analogously, in the earlier case, we should say that he rightly predicted that there would be an event, which was a death of Queen Anne. He did not refer to the event, although in this case he intended to use an event-description that would eventually be satisfied.

If, on July 1, 1714, someone were to say, "The death of Queen Anne will occur in August," they clearly would have said something true. It seems that what makes their remark true is that there had as yet been no event that meets the description, "The death of Queen Anne," when they spoke, but subsequently, the following August, such an event occurred.

There is no conflict between McTaggart's insight about how events do not change, as amended, and his insight about how they do change, as we understood it. So McTaggart gives us no reason to doubt the reality of time.

This reasoning pretty much follows C.D. Broad, in *Scientific Thought* (1923). There are objections to the view of time it expresses, stemming from the special theory of relativity. If the view is right, the difference between past and future is metaphysically rather important, amounting to the temporal edge of a growing universe. On the theory of relativity, however, the difference is relative to the place and time of an observer. And, at least according to certain philosophers, the relativity is very far-reaching; any event that is future relative to me now, and so does not yet exist, will be past relative to some event that is past relative to me, and so does exist, given a clever enough choice of observation points. Exactly how the leading edge of the universe could be so relative to points of observation is not easy to see.

I am somewhat inclined to a Moorean resolution. I am now observing things happen. Which is to say I am observing events come to exist, that did not exist up until now. Any argument to the contrary must contain a mistake, even if I cannot say what it is. If Einstein said this makes no sense, he must have been wrong; perhaps he misinterpreted his own theory. That seems a bit of a cheeky thing for a philosopher to say about a great scientist. However, we say such things about great philosophers all the time. Be that as it may, Moorean arguments are surely most effective when the possibility to be swept aside is something philosophers have come up with in the throes of skepticism, not something that has become a part of accepted science.

But there is hope for time. Credible philosophers, Michael Tooley and Steven Savitt, among others, think there is some hope for reconciliation between something like Broad's view and the special theory (Tooley 2000; Savitt 2008). I hope this can be worked out. I think there is plenty of time.

Notes

1 Quite a while ago Thomas Hofweber and I worked on a paper tentatively entitled "The View from Nowhen." The present essay owes a great deal to our discussions, and have borrowed examples and ideas from the joint paper. If the joint paper ever comes to be, we will know how much of what I say here Hofweber agrees with.

2 These rules correspond to what David Kaplan calls "character" in his seminal works on indexicality, especially *Demonstratives* (1989). For a comparison of Kaplan's framework and the one developed here, see chapter 2 of *Critical Pragmatics* (Korta and Perry 2011).

3 From Washington Irving's "Rip Van Winkle" (1819) I have added details to the story about specific dates and utterances. On themes from Van Winkle, see also Perry 1997.

4 The "is" in (a), (b), and other similar contexts should be read as "is, was, or will be."

5 These propositions may be sets of worlds, or functions from worlds to truth-values, or persistent sets of situations, or structures of relations and objects, or virtually any other approach the reader prefers, as long as they are abstract objects that encode truth-conditions.

6 This is not quite right. From 1582 until the last century dates referred to different days in countries that had not yet accepted the Gregorian calendar than they did in countries that had.

References

Broad, C.D. (1923). *Scientific Thought*. London: Routledge and Kegan Paul.

Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik* 100:25–50. Trans. by M. Black as On Sense and Reference. In *Translations from the Philosophical Writings of Gottlob Frege*, P. Geach and M. Black (eds. and trans.). Oxford: Blackwell. (Third edition 1980.)

Frege, G. (1918). Der Gedanke. Eine Logische Untersuchung. *Beiträge zur Philosophie des deutschen Idealismus* I (1918–1919):58–77. Translated by A.M. Quinton and M. Quinton (1956), as The Thought: A Logical Inquiry. *Mind* 65, 289–311.

Galton, A. (2008). Temporal Logic. In E.N. Zalta (ed.). *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/fall2008/entries/logic-temporal (accessed October 1, 2012).

Irving, W. (1819). Rip Van Winkle. Many editions. First appeared in *The Sketch Book of Geoffrey Crayon*.

Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (eds.). *Themes From Kaplan* (481–563). New York: Oxford University Press.

Korta, K. and Perry, J. (2011). *Critical Pragmatics*. Cambridge: Cambridge University Press.

McTaggart, J.E. (1908). The Unreality of Time. *Mind* 17, 456–473.

Perry, J. (1997). Rip Van Winkle and Other Characters. *The European Review of Analytical Philosophy* 2, 13–39. Reprinted in Perry (2000), enlarged edition.

Tooley, M. (2000). *Time, Tense and Causation*. Oxford: Oxford University Press.

Savitt, S. (2008). Being and Becoming in Modern Physics. In E.N. Zalta (ed.). *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/win2008/entries/spacetime-bebecome (accessed October 1, 2012).

Further Reading

Perry, J. (2000). *The Problem of the Essential Indexical and Other Essays*. Stanford: CSLI.

Perry, J. (2006). How Real are Future Events? In F. Stadler and M. Stöltzner (eds.). *Time and History* (13–30). Frankfurt: Ontos Verlag.