A Dichotomic Analysis of the Surprise Examination Paradox

Paul FRANCESCHI
University of Corsica

revised april 2002

p.franceschi@univ-corse.fr
http://www.univ-corse.fr/~franceschi

The surprise examination paradox finds its origin in an actual fact. In 1943-1944, the Swedish authorities plan to implement a civil defence exercise. The radio broadcasted then an announcement according to which a civil defence exercise would take place during the following week and nobody would be able to know in advance the date of the exercise. The mathematician Lennart Ekbom noticed the problem arising from this announcement and exposed it to his students. A wide diffusion of the paradox then ensued.

Currently, the paradox is most frequently described as a professor's announcement of a surprise examination. A professor announces to her students that a surprise examination will take place on the next week. The intelligent student reason as follows. Obviously, a surprise examination cannot take place on the last day of the term, say Saturday. But since Saturday is ruled out, the clever student will expect it on Friday. Thus, Friday is also ruled out. By a similar reasoning, the students rule out successively Thursday, Wednesday, Tuesday and Monday. Finally, all days are ruled out by the same reasoning. However, this does not prevent the examination from finally occurring surprisingly, say on Wednesday. Thus, the student's reasoning proved to be fallacious. However, such a line of reasoning appears intuitively valid. The paradox lies here in the fact that the student's reasoning seems valid, whereas it is in contradiction with the facts, namely that the examination can occur truly surprisingly, in accordance with the professor's announcement.

In the literature, several solutions to the *surprise examination paradox* (thereafter SEP) have been proposed. There does not exist however, at present time, one consensual solution. I shall present in this article what appears to me as an original solution to SEP. In section 1, I analyse the surprise notion in detail. I introduce then in section 2, the distinction between a monist and dichotomic analysis of the paradox. I also present there a dichotomy leading to distinguish two basically and structurally different versions of the paradox, respectively based on a conjoint and a disjoint definition of the surprise. In section 3, I describe the solution to SEP corresponding to the conjoint definition. Lastly, I expose in section 4, the solution to SEP based on the disjoint definition.

1. The surprise notion

Let us begin with the *surprise* notion.[1] What constitutes an instance of surprise? How can the surprise be defined in the context of SEP? On the one hand, it can be observed that the surprise obviously emerges when the student does not predict that the examination will occur on day k and that the examination occurs on day k. Such a situation can take two different forms: either (i) the student predicts that the examination will take place on day p (with $1 \leq p \leq n$)[2] such that $p \neq k$ and the examination occurs on day k; or (ii) the student considers that the examination will not take place at all (the associated prediction begin then denoted by $p = 0$) and the examination occurs on day k. These two situations can then be taken into account in a unified way, by considering in general that the surprise arises without ambiguity when the student predicts that the examination will occur on day p such as $p \neq k$ (with $0 \leq p \leq n$) and that the examination occurs on day k. Intuitively, in this last situation, the surprise emerges from to the fact that the student has made an erroneous prediction. Thus, such a situation can be termed *surprise by error*.

Consider, on the other hand, another type of situation. The student decides to implement the following strategy. On day 1, I will predict that the examination will occur on day 1; on day 2, I will predict that the examination will occur on day 2; ...; on day n, I will predict that the examination will occur on day n. In this manner, the student thinks, I am sure that my prediction will be always true. Let us call *incremental* such a strategy. It appears indeed that this last strategy leads to the fact that the corresponding prediction will prove to be systematically true.[3] Under these conditions, one observes first that the surprise *by error* cannot appear any more. Nevertheless, isn't the surprise possible all the same when the incremental strategy is implemented? The corresponding situation can be put in sharpest relief by considering for example 1000-SEP[4]. In this case, whatever the final date of the examination, the student's prediction will prove to be true, because of the incremental strategy used by the latter. However, one can consider that the examination will nevertheless possibly occur surprisingly in such circumstances, for example on day 127. Admittedly, the student's prediction will prove true. But this last prediction will not be justified, so that the corresponding prediction will finally appear unjustified. Consequently, one can consider that the surprise can also occur in the context where the incremental strategy is implemented. But this last type of surprise, unlike the surprise by error, is not due to an erroneous prediction. The surprise which appears in the context of the incremental strategy is due to the lack of justification of the corresponding prediction. Consequently, this type of surprise can be termed *surprise by non-justification*.

At this stage, we are in a position to give a definition of the surprise notion. Such a definition includes the two cases of surprise which has just been defined: surprise by *error* and by *non-justification*. It is thus a definition whose structure is disjunctive: the *surprise* corresponds either to an erroneous prediction or to an exact but unjustified prediction. Conversely, the *non-surprise* corresponds to an exact and justified prediction.

In what follows, the sequence associating the prediction p performed on day d will be denoted by $d\alpha p\beta$ (with $1 \leq d \leq n$, $0 \leq p \leq n$ and $\alpha, \beta \in \{0, 1\}$), by posing $\alpha = 1$ (respectively, $\alpha = 0$) if the examination takes place (respectively, does not take place) on day d and $\beta = 1$ (respectively $\beta = 0$) if the student's forecast is justified (respectively unjustified). For commodity, one will use {°, *} to denote respectively the values {0, 1} respectively taken by α and β. Thus, the surprise by error will be denoted by k*p° (with $1 \leq k \leq n$, $0 \leq p \leq n$ and $k \neq p$), the surprise by non-justification by k*k°, and the non-surprise by k*k*. More generally, the surprise - when it is not necessary to distinguish between surprise by error or by non-justification - will be denoted by k*p° (for $0 < p \leq n$). Let us give a few examples. With this last notation, 6*6* denotes the fact that the examination has occurred on day 6 and the student has justifiably predicted on day 6 that the examination would occur on day 6; 6*6° denotes the fact that the examination has occurred on day 6 and the student has made on day 6 an exact but unjustified prediction that the examination would occur on day 6; 6*7° denotes the fact that the examination has occurred on day 6 and the student has predicted erroneously on day 6 that the examination would occur on day 7; 6*0° denotes the fact that the examination has occurred on day 6 and the student has predicted erroneously on day 6 that the examination would not take place on any day of the week.

Lastly, it should be noticed that the preceding definition of the surprise is based on *positive* instances, i.e. cases where the examination occurs truly. Thus, k*0° and k*k° denote a case where the examination occurs truly on day k ($\alpha = 1$). In effect, the sequences corresponding to a day where the examination does not occur ($\alpha = 0$) are not taken into account. Such sequences which correspond to a prediction performed on a day k where the examination does not occur, have the structure k°$p\beta$. But these last cases correspond to *negative* instances (the examination does not take place on the given day) and their case can consequently be legitimately ignored.

2. Monist or dichotomic analysis of the paradox

The majority of the classical analyses intended to solve SEP are based on a general solution, which globally applies to the situation which is that of SEP. In this type of analysis, a single solution is presented, which is supposed to apply to all variations of SEP. Such a type of solution has a unitary nature and appears based on what one can be called a *monist* theory of SEP. The majority of the solutions to SEP proposed in the literature are monist analyses. Characteristic examples of this type of analysis of SEP are the solutions suggested by Quine (1953) or Binkley (1968). Robert Binkley notably exposes a reduction of SEP to *Moore's paradox*. He points out that on the last day, SEP reduces to a variation of the proposition 'P and I do not know that P' which constitutes Moore's paradox. Binkley extends then his analysis concerning the last day to the other days of the week. In a similar way, the solution under consideration by Dietl (1973) which is based on a reduction of SEP to the *sorites paradox* also constitutes a monist solution to SEP.

Conversely, a dichotomic analysis of SEP is based on a distinction between two different scenarios of SEP and on the formulation of an *independent* solution to each of the two scenarios. In the literature, the only analysis which has a dichotomic nature, this seems to me, is that of Wright and Sudbury (1977). In effect, the analysis developed by Crispin Wright and Aidan Sudbury[5] leads to distinguish two cases: on the one hand, on the last day, the student is in a situation which is that which results from Moore's paradox; on the other hand, on the first day, the student is in a basically different situation where she can validly believe the professor's announcement. Thus, the fact of emphasising of these two types of situations leads to the rejection of the *temporal retention principle*, according to which what is known at time T_0 is also known at time T_1 (with $T_0 < T_1$). In what follows, I will present a dichotomic solution to SEP. This solution is based on the distinction of two scenarios of SEP, associated with different structures of the definitions of the *surprise* and the *non-surprise*.

At this step, it proves useful to introduce the matricial notation. If one considers for example 7-SEP, the various cases of surprise and non-surprise can be thus rendered with the help of the following table S[k, s], where k denotes the day where the examination takes place and s denotes whether the situation is the non-surprise ($s = 0$) or the surprise ($s = 1$):

(D1)

	S[k, 0]	S[k, 1]
S[7,s]	7*7*	7*p°

S[6,s]	6*6*	6*p°
S[5,s]	5*5*	5*p°
S[4,s]	4*4*	4*p°
S[3,s]	3*3*	3*p°
S[2,s]	2*2*	2*p°
S[1,s]	1*1*	1*p°

The dichotomy on which the present solution is based results directly from the analysis of the structure which is that of (D1). It appears indeed that such a structure corresponds to a maximal definition, where all the cases of surprise and non-surprise are made possible. One can however conceive of variations of SEP associated with more restrictive definitions, where certain cases of surprise and non-surprise are not authorised by the statement of SEP. In order to highlight the various situations that can be met, it is worth beforehand setting up a notation making it possible to describe the *structure* of the corresponding definitions. The latter can thus be modelled[6] with the help of a matrix S[k, s] where k denotes the day where the examination takes place and S[k, s] denotes whether the corresponding case of non-surprise (s = 0) or of surprise (s = 1) is made possible (S[k, s] = 1) or not (S[k, s] = 0) by the conditions of the statement. With this notation, the structure of (D1) can be rendered as follows:

(D2)

	S[k, 0]	S[k, 1]
S[7,s]	1	1
S[6,s]	1	1
S[5,s]	1	1
S[4,s]	1	1
S[3,s]	1	1
S[2,s]	1	1
S[1,s]	1	1

It appears here that (D2) is such that all cases of non-surprise or surprise are made possible by the statement, i. e. formally $\forall\ k\ (1 \leq k \leq n)$ S[k, 0] = S[k, 1] = 1. The associated matrix can be defined as a *rectangular* matrix. However, one can conceive of variations of SEP associated with a more restrictive definition. It is worth considering now the structure of these more restrictive definitions. These latter are such that there exists at least one impossible case of non-surprise or of surprise, i. e. formally $\exists\ k\ (1 \leq k \leq n)$ S[k, 0] = 0 $\vee$ S[k, 1] = 0. Such a condition leaves room for a certain number of variations, the characteristics of which it is worth studying.

Preliminarily, one can notice that certain extreme definitions can be discarded since the beginning. It is allowed to think indeed that any definition associated with a restriction of (D2) is not appropriate. Thus, there exist minimal conditions for the emergence of SEP. In this sense, a first condition is that the *base step* being present. This base step is such that S[n, 0] = 1. It presents the general form n*n* and corresponds to 7*7* at the level of (D1). In the absence of this base step, one does not have the paradoxical effect of SEP. Consequently a definition such as S[n, 0] = 0 can be discarded.

One second condition is that the examination can finally occur by surprise. This makes it possible for the professor's announcement to be finally vindicated (*vindication step*). Such a condition is classically mentioned as a condition for the emergence of the paradox. Thus, a definition which would be such that all cases of surprise are made impossible, i. e. formally $\forall\ k\ (1 \leq k \leq n)$ S[k, 1] = 0 would also not be appropriate.

Given the above elements, one is now in a position to describe accurately the minimal conditions which are those of SEP:

(C3) S[n, 0] = 1 (*base step*)
(C4) $\exists\ k\ (1 \leq k \leq n)$ such as S[k, 1] = 1 (*vindication step*)

At this stage, it is worth considering the structure of the versions of SEP based on the definitions satisfying the minimal conditions for the emergence of the paradox, which have just been detailed. It appears here that the structure of the definition associated with SEP can present two forms of a basically different nature. A first form of SEP is associated with a structure of definition of the surprise and the non-surprise such that there exists during the n-period at least one day where the surprise and the non-surprise are made simultaneously possible, i. e. formally $\exists\ k\ (1 \leq k \leq n)$ S[k, 0] = S[k, 1] = 1. Such a definition can be called *conjoint*. A second form of SEP whose structure is basically different is such that for each day of the n-period, it is impossible to have at the same time the surprise and the non-surprise, i. e. formally $\forall\ k\ (1 \leq k \leq n)$ S[k, 0] + S[k, 1] = 1 [7]. A definition of this nature can be termed *disjoint*. Consequently, one will be led in what follows to distinguish two structurally different versions of SEP respectively based on: (i) a *conjoint* definition of the proper instances of surprise and non-surprise; (ii) a *disjoint* definition of the same instances. The need to make such a dichotomy finds its legitimacy in the fact that in the original version of SEP, it is not specified whether one must take into account a *disjoint* or *conjoint* definition of the cases of surprise and non-surprise. Does the professor take into account a disjoint or conjoint definition? With regard to this particular point, the statement of SEP appears ambiguous. In the original presentation of SEP, an accurate answer to this question is lacking. In what follows, I will thus consider successively two versions of SEP, respectively based on a conjoint or disjoint definition of the instances of surprise and non-surprise.

3. A conjoint definition of the surprise

Let us begin with the version of SEP based on a conjoint definition of the surprise and the non-surprise. Call SEP(I) such a version. Intuitively, such a variation corresponds to a situation where there exists at least one day during the n-period where both surprise and non-surprise are made possible by the statement. The corresponding structure is such that $\exists\, k\ (1 \leq k \leq n)\ S[k, 0] = S[k, 1] = 1$. Several types of definitions are susceptible to meet this criterion. I shall examine them in turn.

3.1 The definition associated with the rectangular matrix and Quine's solution

To begin with, it is worth considering the definitions which are such that all cases of non-surprise and surprise are possible, i.e. formally $\forall\, k\ (1 \leq k \leq n)\ S[k, 0] = S[k, 1] = 1$. The corresponding matrix is a *rectangular* matrix. Let SEP(I□) be such a version. The definition associated with such a structure is maximal because all the cases of non-surprise and surprise are authorised. (D2) itself corresponds to this general structure. But there is also a version of SEP based on $n = 1$ which satisfies this definition. The structure associated with 1-SEP(I□) is as follows:

(D5)	S[1, 0]	S[1, 1]
S[1,*s*]	1	1

Thus, 1-SEP(I□) is the minimal version of SEP which satisfies not only the above condition, but also (C3) and (C4). For this reason, (D5) can be regarded as a canonical form of SEP(I□).

At this step, it is worth proceeding to provide a solution to SEP(I□). For this purpose, let us recall first Quine's solution. The solution to SEP proposed by Quine (1953) is well-known. Let us distinguish first two parts in the structure of the professor's announcement:

(A6) the examination will take place during the week
(A7) the day of the examination will constitute a surprise[8]

Quine highlights the fact that the student eliminates successively the days n, n-1, ..., 1, by a reasoning based on backward induction (*backward induction argument*, BIA) and concludes then that (A6) is false. The student reasons as follows. On day n, she will predict that the examination will take place on day n, and consequently the examination cannot take place on day n; on day n-1, she will predict that the examination will take place on day n-1, and consequently the examination cannot take place on day n-1; ...; on day 1, she will predict that the examination will take place on day 1, and consequently the examination cannot take place on day 1. Finally, the student concludes that the examination cannot take place on any day of the week. Consequently, on day 1, she predicts that the examination will not take place; on day 2, she predicts that the examination will not take place; ...; on day n, she predicts that the examination will not take place. The sequence of the successive predictions of the student is thus: $1\alpha0\beta$, $2\alpha0\beta$, $3\alpha0\beta$, ..., $(n\text{-}1)\alpha0\beta$, $n\alpha0\beta$. She predicts thus that the examination will not take place on any day of the week, and thus that $\neg$(A6). But this last conclusion makes it finally possible for the examination to occur surprisingly, including on day n. According to Quine, the flaw in the student's reasoning lies precisely in the fact of not having taken into account this possibility since the beginning. Because on this last assumption, the student would have considered the two principal assumptions which are n*n* and n*0°, instead of only n*n*, thus preventing the fallacious reasoning.[9]

On the other hand, Quine directly applies his analysis to 1-SEP. The student's error resides, according to Quine, in the fact of having considered only one hypothesis, namely 1*1*. In fact, the student should have considered 4 cases (i) 1*1*; (ii) $1°1\beta$; (iii) $1°0\beta$; (iv) 1*0°. And the fact of considering the hypothesis (i) but also the hypothesis (iv) which is compatible with the professor's announcement would have prevented the student from concluding that the examination would not finally take place.[10] Consequently, it is the fact of having only taken into account the hypothesis (i) which can be identified as the cause of the fallacious reasoning. Thus, the student did only take partially into account the set of hypotheses resulting from the professor's announcement. If she had apprehended the totality of the relevant hypotheses compatible with the professor's announcement, she would not have concluded fallaciously that ¬(A6).

At this stage, it appears clearly that Quine's solution applies adequately to the version of SEP(I□) based on the surprise *by error* as it has just been defined. Such a notion corresponds to the general definition S[*k*, 0] = {*k***k**} and S[*k*, 1] = {*k**0°}. On this assumption, the corresponding definition of the non-surprise and the surprise is as follows:

(D8)

	S[*k*, 0]	S[*k*, 1]
S[7,*s*]	7*7*	7*0°
S[6,*s*]	6*6*	6*0°
S[5,*s*]	5*5*	5*0°
S[4,*s*]	4*4*	4*0°
S[3,*s*]	3*3*	3*0°
S[2,*s*]	2*2*	2*0°
S[1,*s*]	1*1*	1*0°

However, it appears that these last conditions are restrictive with regard to the general definition (D1), since the surprise by non-justification is not taken into account here. It is thus allowed to wonder whether Quine's solution also applies to the version of SEP(I□) based on the surprise by *non-justification*. This corresponds to the general definition: S[*k*, 0] = {*k***k**} and S[*k*, 1] = {*k***k*°}. On this assumption, the corresponding definition is as follows:

(D9)

	S[*k*, 0]	S[*k*, 1]
S[7,*s*]	7*7*	7*7°
S[6,*s*]	6*6*	6*6°
S[5,*s*]	5*5*	5*5°
S[4,*s*]	4*4*	4*4°
S[3,*s*]	3*3*	3*3°
S[2,*s*]	2*2*	2*2°
S[1,*s*]	1*1*	1*1°

This situation takes place within the framework of a version of SEP where the surprise by error is not possible any more. Such a version corresponds for example to a situation where the student implements the *incremental* strategy described above. In this case, the student proceeds as follows: on day 1, she predicts that the examination will take place on day 1; on day 2, she predicts that the examination will take place on day 2; ...; on day *n*, she predicts that the examination will take place on day *n*. Thus, the successive predictions made by the student are 1α1β, 2α2β, 3α3β, ..., *n*α*n*β. With the help of this strategy, the student is sure that her prediction will finally be true. By proceeding this way, the student prevents the surprise by error from occurring. However, as observed above, this does not prevent the surprise from finally appearing, but under the form this time of the surprise by non-justification. It also proves that this situation can be rendered with the help of a version of SEP which takes place within the framework of a game where the student is obliged to perform on each day *k* a prediction *k.* In this situation also, the emergence of the surprise by error is made impossible. But this does not prevent the examination from occurring surprisingly, under the form this time of the surprise by non-justification. For the prediction made on each day by the student can appear justified or unjustified. And it is thus the parameter β that constitutes the criterion of the surprise or the non-surprise.

Thus, in the two types of situations which have been just described, one has indeed a version of SEP where the surprise by non-justification can finally occur. But in this new context, can Quine's solution from now on apply? Consider the hypothesis that SEP is based on a game where the rule obliges to perform on each day *k* a prediction *k* and where the criterion that determines the state of surprise or non-surprise is the fact that the forecast is justified or not. The surprise occurs then if the forecast is unjustified (β = 0) and the non-surprise if the forecast is justified (β = 1). It appears then that the student can reason as follows: on day *n*, I will consider justifiably that the examination will take place on day *n*, therefore the examination cannot occur by surprise on day *n*; on day *n*-1, I will consider justifiably that the examination will take place on day *n*-1, therefore the examination cannot occur by surprise on day *n*-1; ...; on day 1, I will consider justifiably that the examination will take place on day 1, therefore the examination cannot occur by surprise on day 1. Consequently, the examination cannot occur by surprise on any day of the *n*-period. Thus, on day 1, I will predict justifiably that the examination will take place on day 1; on day 2, I will predict justifiably that the examination will take place on day 2; ...; on day *n*, I will predict justifiably that the examination will take place on day *n*. In this last version, it appears that the student performs the prediction *k* on day *k* because she is obliged to do so by the game rule, but her forecast is not justified because she is convinced that the examination will not occur by surprise. Here, the student cannot conclude as in the original quinean solution that ¬(A6) because the initial data prevent her to do so, but she concludes this time that ¬(A7). Except for this last difference, the student's reasoning is completely identical to the one observed in the original quinean version. Thus, it appears finally that Quine's solution also finds to apply in this type of situation. Just as in a traditional quinean situation, one thus notes that Quine's solution also applies to a version of SEP based on (D9).

Thus, it proves finally that Quine's solution applies not only to a definition based on the surprise by error, but also to a definition based on the surprise by non-justification. Hence, Quine's solution generally applies to the disjunctive definition of the surprise presented above and to the situation associated with a rectangular matrix corresponding to SEP(I□).

3.2 The definition associated with the triangular matrix and Hall's reductio

The preceding analysis of Quine's solution shows that it is not finally necessary to distinguish between surprise by error and surprise by non-justification. One can thus validly consider a notion of unified surprise. Consequently, one can limit the study of what refers to Quine's solution to the matrix associated with the corresponding definitions. It is worth considering now matrices the general structure of which is not rectangular, but which meet however the conditions (C3) and (C4) mentioned above. Such matrices have a structure that can be described as *triangular*. Let thus SEP(IΔ) be the corresponding version.

Consider first 7- SEP(IΔ), which corresponds to a definition whose structure is as follows:

(D10)		S[*k*, 0]	S[*k*, 1]
	S[7,*s*]	1	0
	S[6,*s*]	1	1
	S[5,*s*]	1	1
	S[4,*s*]	1	1
	S[3,*s*]	1	1
	S[2,*s*]	1	1
	S[1,*s*]	1	1

It is apparent here that 7-SEP(IΔ) corresponds to a situation where the surprise cannot occur on day *n*. The general structure corresponding to this type of definition is:

(D11)		S[*k*, 0]	S[*k*, 1]
	S[*n*,*s*]	1	0
	S[*n*-1,*s*]	1	1
	....	....	....

And one can in the same way consider the canonical structure:

(D12)		S[*k*, 0]	S[*k*, 1]
	S[2,*s*]	1	0
	S[1,*s*]	1	1

It is worth now presenting a solution to the versions of SEP associated with the definitions having the structure of (D12). Such a solution is based on a reduction recently presented by Ned Hall, the context of which it is worth beforehand recalling.

Quine's solution has led to an objection, presented in particular by A. J. Ayer (1973) and Christopher Janaway (1989). Ayer imagines thus a version of SEP where a person is told that when a pack of cards is dealt out, she will not know beforehand when the Ace of spades will arrive. Nevertheless, she is authorised to verify the presence of the Ace of spades before the pack of cards is shuffled. These objections aim at highlighting a situation where the paradox is indeed present but where Quine's solution finds no more to apply, because the student knows indubitably, given the data initial of the problem, that the examination will indeed take place.

In the version of SEP under consideration by Quine (1953), it appears clearly that the fact that the student doubts of (A6), at a certain stage of the reasoning, is authorised. Quine thus deliberately places himself in a situation where the student has the faculty to doubt of (A6). The versions described by Ayer (1973), Janaway (1989) but also Scriven (1951) convey the intention to prevent this particular step in the student's reasoning. Such scenarios correspond, in spirit, to SEP(IΔ). Also related is the variation of the *Designated Student Paradox* described by Sorensen (1982), where five stars - one gold star and four silver stars - are put on the back of five student, given that it is indubitable that the gold star is placed on the back of the designated student.[11]

However, Ned Hall (1999) has recently exposed a reduction, which tends to refute the objections classically opposed to Quine's solution. The argumentation developed by Hall is as follows:

> We should pause, briefly, to dispense with a bad - though oft-cited - reason for rejecting Quine's diagnosis. (See for example Ayer 1973 and Janaway 1989). Begin with the perfectly sound observation that the story can be told in such a way that the student *is* justified in believing that, come Friday, he will justifiably that an exam is scheduled for the week. Just add a second *Iron Law of the School*: that there must be at least one exam each week. (...) Then the first step of the student's argument goes through just fine. So Quine's diagnosis is, evidently, inapplicable.
>
> Perhaps - but in letter only, not in spirit. With the second Iron Law in place, the last disjunct of the professor's announcement - that E_5 & $\neg J(E_5)$ - is, from the student's perspective, a *contradiction*. So, from his perspective, the *content* of her announcement is given not by SE_5 but by SE_4: $(E_1 \ \& \ \neg J_1(E_1)) \vee ... \vee (E_4 \ \& \ \neg J_4(E_4))$. And now Quine's diagnosis applies straightforwardly: he should simply insist that the student is not justified in believing the announcement and so, come Thursday morning, not justified in believing that crucial part of it which asserts that if the exam is on Friday the it will come as a surprise - which, from the student's perspective, is tantamount to asserting that the exam is scheduled for one of Monday through Thursday. That is, Quine should insist that the crucial premise that $J_4(E_1 \vee E_2 \vee E_3 \vee E_4)$ is *false* - which is exactly the diagnosis he gives to an ordinary 4-day surprise exam scenario. Oddly, it seems to have gone entirely unnoticed by those who press this variant of the story against Quine that its only real effect is to convert an *n*-day scenario into an *n*-1 day scenario[12].

Hall thus puts in parallel two types of situations. The first corresponds to the situation, based on the surprise by error, in which Quine's analysis classically takes place. The second corresponds to the type of situation under consideration by the opponents to Quine's solution and in particular Ayer (1973) and Janaway (1989). On this last hypothesis, a stronger version of SEP is taken into account where one second *Iron Law of the School* is considered, according to which it is admitted that the examination will necessarily take place during the week. The argumentation developed by Hall leads to the *reduction* of a second type version of *n*-SEP to a quinean type version of (*n*-1)-SEP. This equivalence has the effect of annihilating the objections of the opponents to Quine's solution.[13] For the effect of this reduction is to make finally possible for Quine's solution to apply in the situations described by Ayer and Janaway.

In order to apprehend the scope of Hall's reduction, it is worth reformulating Ayer and Janaway's objection in order that it applies at the same time to a version of SEP(I) based on the surprise by error and non-justification. For this purpose, one will consider that, in spirit, the scenario under consideration by Ayer and Janaway corresponds to a situation where the surprise by error or non-justification is not possible on day n (i.e. $S[n, 1] = 0$). This indeed has the effect of neutralising Quine's solution based on *n*-SEP(I□). But Hall's reduction finds then to apply. And its effect is to reduce a scenario corresponding to (D11) to a situation based on (D5). Consequently, Hall's reduction entails that *n*-SEP(IΔ) reduces to (*n*-1)-SEP(I□). It involves: *n*-SEP(IΔ) ≡ (*n*-1)-SEP(I□) for $n > 1$. Thus, Hall's reduction makes it finally possible for Quine's solution to apply to SEP(IΔ).

At this stage, it is worth considering other variations, and notably the one corresponding to the following structure:

(D13)

	$S[k, 0]$	$S[k, 1]$
$S[7,s]$	1	0
$S[6,s]$	1	0
$S[5,s]$	1	1
$S[4,s]$	1	1
$S[3,s]$	1	1
$S[2,s]$	1	1
$S[1,s]$	1	1

whose general structure is :

(D14)

	$S[k, 0]$	$S[k, 1]$
$S[n,s]$	1	0
$S[n\text{-}1,s]$	1	0
$S[n\text{-}2,s]$	1	1
....	....	....

and to which one can associate the following canonical structure:

(D15)		$S[k, 0]$	$S[k, 1]$
	$S[3,s]$	1	0
	$S[2,s]$	1	0
	$S[1,s]$	1	1

It appears here that this last structure corresponds to a variation of SEP to which Hall's reduction also applies straightforwardly. The corresponding scenario is the one where the surprise can neither occur on day n nor on day n-1. This corresponds for example to a variation of the situation described by Ayer and Janaway where the professor announces to the students that on the next week will not take place only one examination, but rather *two* examinations, and that these examinations will occur surprisingly. The corresponding version of SEP presents a paradoxical nature, in the same way as the traditional version based on only one examination. Such a situation appears even more restrictive than the one described by Ayer and Janaway, because the student knows that the days n and n-1 necessarily correspond to a situation of non-surprise. Thus, the structure of (D14) itself shows that Hall's reduction can be extended to a situation corresponding to a definition even more restrictive, such that the surprise can neither occur on day n nor on day n-1, i. e. formally not only $S[n, 1] = 0$ but also $S[n-1, 1] = 0$. In this case, Hall's reduction entails: n-SEP(IΔ) ≡ (n-2)-SEP(I□).

More generally, one can conceive of a version of n-SEP(IΔ) associated with a matrix such that (i) $\exists\ k\ (1 \leq k \leq n)\ S[k, 0] = S[k, 1] = 1$; (ii) $\forall\ p > k\ S[p, 0] = 1$ and $S[p, 1] = 0$; (iii) $\forall\ q < k\ S[q, 0] = S[q, 1] = 1$. In this new situation, an *extended Hall's reduction* applies to the corresponding version of SEP. Such a version corresponds for example to a variation of SEP where the professor announces to the students that n-k examinations will take place on the next week. In this case, Hall's extended reduction leads to: n-SEP(IΔ) ≡ (n-k)-SEP(I□).

4. A disjoint definition of the surprise

Let us turn now to the situation where SEP is based on a *disjoint* definition of the cases of surprise and non-surprise. Call SEP(II) the corresponding version. Intuitively, such a version corresponds to a situation where for each day of the n-period, it is not possible to have at the same time the surprise and the non-surprise. On each day one has thus exclusively either the surprise or the non-surprise. The structure of the associated matrix is such that $\forall\ k\ (1 \leq k \leq n)\ S[k, 0] + S[k, 1] = 1$. The minimal version which corresponds to this definition is 1-SEP(II), with which the following matrix is associated:

(D16)		$S[1, 0]$	$S[1, 1]$
	$S[1,s]$	1	0

However, it can be noticed that such a matrix does not meet the condition (C4). However (C3) and (C4) must be satisfied in order that one has truly the emergence of the paradox. And in the situation associated with (D16), it appears indeed that the student can reason as follows: on day 1, I will justifiably know that the examination will take place on day 1. Consequently, I predict that the examination will take place on day 1. Thus, (A6) and (A7) are finally falsified. Consequently, in the case of 1-SEP(II),[14] one does not observe the emergence of the paradox, since the professor's announcement is finally falsified. In effect, the condition (C4) being not satisfied in 1-SEP(II), the examination cannot finally occur by surprise. Thus, the situation associated with 1-SEP(II) is not truly paradoxical, but only constitutes a situation where the professor's announcement appears falsified.

Nevertheless, it appears that 2-SEP(II) corresponds to a disjoint definition and meets at the same time the conditions (C3) and (C4). In fact, intuitively, one has indeed the emergence of the paradox for $n > 1$. It suffices to consider for example (i) $n = 7$ and 3*3° or (ii) $n = 10$ and 5*5° or (iii) $n = 1000$ and 127*127°. In all these cases, one notes that the professor's announcement is lastly vindicated: an examination finally takes indeed place by surprise. Moreover, one interesting characteristic of SEP(II) is that the paradox emerges intuitively in a clearer way when large values of n are taken into account. Several authors also mention this characteristic.[15]

At this stage, the following question can be posed: can't Quine's solution apply to SEP(II)? However, the preceding analysis of SEP(I) shows that a necessary condition in order to Quine's solution finds to apply is such that there exists during the n-period at least one day where both surprise and non-surprise are made possible, i. e. formally $\exists\ k\ S[k, 0] + S[k, 1] = 2$. However such a property is that of a *conjoint* structure and corresponds to the situation which is that of SEP(I). But in the context of a *disjoint* structure, the associated matrix satisfies conversely $\forall\ k\ S[k, 0] + S[k, 1] = 1$. Consequently, this prohibits applying Quine's solution to SEP(II).

In the same way, one could pose the question whether Hall's reduction cannot apply equally to SEP(II). Hence, does one not have thus: n-SEP(II) ≡ (n-1)-SEP(I)? It also appears not. In effect, as observed above, Quine's solution cannot apply to SEP(II). However the effect of Hall's reduction is to reduce a given scenario to a situation where Quine's solution finally finds to apply. But, since Quine's solution cannot apply in the context of SEP(II), Hall's reduction is also in the impossibility of producing its effect.

4.1 A precise disjoint definition

The structure of the disjoint definition of the cases of surprise and non-surprise leads in fact to distinguish whether the concept of surprise is *precise* or *vague*. Let us envisage first the hypothesis where this definition is *precise*. Let SEP(II\) be such a version.

A characteristic instance of this type of version is the one based on the following definition (for $n = 7$):

(D17)

	S[k, 0]	S[k, 1]
S[7,s]	1	0
S[6,s]	1	0
S[5,s]	1	0
S[4,s]	0	1
S[3,s]	0	1
S[2,s]	0	1
S[1,s]	0	1

It clearly appears that in this last case, the BIA leads to a false conclusion. The particular step in the reasoning which can be identified as fallacious is the one where the student passes from 5*5* to 4*4*. The flaw in the reasoning resides precisely here in not taking into account the fact that the surprise notion is precise. If the student had reasoned correctly, she would have considered that the surprise is here a precise notion and that consequently, a precise cut-off exists between the concepts of surprise and non-surprise. Thus, at a certain level, the passage of k*k* to (k-1)*(k-1)* is not valid because k*k* is followed by (k-1)*(k-1)°, *by definition*.

It should also be noted that the canonical matrix associated with SEP(II\) which meets the conditions (C3)-(C4) is the following one (with $m < n$):

(D18)

	S[k, 0]	S[k, 1]
S[n,s]	1	0
....	....	....
S[m,s]	0	1

Such a precise definition of the surprise notion can also be put in sharper relief by choosing n large and by considering a definition based on an approach by degrees. In this case, the associated matrix comprises m columns, thus allowing a progressive passage of the non-surprise to the surprise, with m-2 intermediate values. One can thus choose intermediate[16] values equal to 0.5, 0.25, 0.75, etc. and more generally intermediate values s_i such that $0 < s_i < 1$. For m rather large, one has then $\forall\, k$ S[k, 0] + S[k, s_1] + ... + S[k, s_{m-2}] + S[k, 1] = 1.

Lastly, it should be observed that the definitions corresponding to SEP(II\) which have been just described, present a *linearity* property, i. e. formally, are such that $\forall\, k, i$ ($1 < k \leq n$, $1 \leq i < m$) if S[k, s_i] = 1 then S[k-1, s_i] = 1 or[17] S[k-1, s_{i+1}] = 1. It appears indeed that a structure of definition which would not present such a property, would not capture the intuition corresponding to the surprise notion. For this reason, it appears sufficient to limit the present study to the structures of definitions satisfying the property of linearity.

4.2 A vague disjoint definition

As mentioned above, the disjoint definition of the cases of surprise and non-surprise can also present a *vague* nature. Let SEP(II≈) be the corresponding version. In this case, a proper instance of surprise and non-surprise are, just as previously, at least present. But unlike SEP(II\), there also exist in SEP(II≈) one or more *borderline* cases. In this case, the definition has then the following structure[18] (for $n = 7$ and $0 < \iota < 1$):

(D19)

	S[k, 0]	S[k, ι]	S[k, 1]
S[7,s]	1	0	0
S[6,s]	1	0	0
S[5,s]	1	0	0
S[4,s]	0	1	0
S[3,s]	0	0	1
S[2,s]	0	0	1
S[1,s]	0	0	1

It appears here that 7*7* ($\beta = 1$) constitutes a proper instance of non-surprise, that 1*1° ($\beta = 0$) is a proper instance of surprise and that 4*4ι ($\beta = \iota$) constitutes a borderline case. Here also, some variations can be take into account, by choosing in particular n large and a much higher number of borderline cases, in order to make emerge more clearly the phenomenon of *higher-order vagueness*, characteristic of vague concepts. However, one can consider that a proper instance of surprise and non-surprise as well as a borderline case, are sufficient to characterise SEP(II≈). The associated canonical matrix[19] has then the following structure ($m < v < n$):

(D20)	S[k, 0]	S[k, ι]	S[k, 1]
S[n,s]	1	0	0
....	....	....	....
S[v,s]	0	1	0
....	....	....	....
S[m,s]	0	0	1

It should be observed that such a structure corresponds to the classical characterisation of a vague predicate. In effect there is on the one hand the extension of the surprise (S[k, 1]), its anti-extension (S[k, 0]) which are such that the extension and the anti-extension of the surprise are mutually *exclusive*. Moreover, the extension and the anti-extension of the surprise are not sufficient to capture exhaustively the definition of the surprise notion. Thus, this definition is completely in conformity with the classical definition of a vague predicate, characterised by a mutually exclusive and non-exhaustive extension and anti-extension.[20]

At this stage, one can verify that a version of the *sorites paradox* (SP) can be built with the elements of SEP(II≈). The corresponding instance of SP presents the following form (for n = 1000):

(21) 1000*1000* is a case of non-surprise
(22) if k*kβ is a case of non-surprise, then (k-1)*(k-1)β is a case of non-surprise
(23) ∴ 127*127* is a case of non-surprise

This version conforms completely with the classical version of SP, leading to a conclusion which contradicts the intuitive fact according to which 127*127° is a proper instance of surprise. Thus, it appears that 1000*1000* with non-surprise and 127*127° with surprise, are in the same type of relation as 10000 (grains of sand) with *heap*, and 0 (grains of sand) with *non-heap*. In this particular version of SEP, the BIA plays the role of the induction step in SP. Taking into account these elements, it appears that one has a direct reduction of SEP(II≈) to SP.

One can have a confirmation of this analysis by highlighting a version of SP built from the converse predicate of that of *non-surprise* which is used in (21)-(23). It is indeed known that vague predicates P which lead to SP are such that one can build a converse version of SP with the help of the predicate ~P. One has thus, straightforwardly, a version of SP based on the *surprise* notion, which constitutes the converse version of (21)-(23). This version is thus the following one:

(24) 127*127* is a case of surprise
(25) if k*kβ is a case of surprise, then (k+1)*(k+1)β is a case of surprise
(26) ∴ 1000*1000* is a case of surprise

It proves now necessary to take into account a certain number of objections, which can be raised against the present analysis. Roy Sorensen (1988, p. 324-7) enumerates in particular certain objections against the assimilation of SEP with SP, which can also concern, in the present analysis, the reduction of SEP(II≈) to SP. In the literature concerning SEP, the assimilation of SEP with SP results from the solutions suggested by Dietl (1973) and Smith (1984). Concerning these last solutions,[21] the first objection mentioned by Sorensen is that the phenomenon of vagueness does only arise for large values of n. In the present analysis, the version of SP corresponding to SEP(II≈) moves too fast, could it be objected.[22]

However, it appears that the existence of a proper instance, a borderline case and a proper counter-instance of the concept of surprise constitute sufficient criteria for the existence of a vague notion and for the emergence of the conditions of SP. Consequently, it does not appear necessary to take into account large values of n.[23]

Consider, on the other hand, what Sorensen regards as the principal objection to the assimilation to SP. The assimilation of SEP with SP is based on the fact that n*n* constitutes a proper instance of non-surprise. However, Sorensen continues, one does not have veritably the certainty that the examination will not occur by surprise on the last day. Thus, the objection goes, the premise (21) could be false, entailing the fact that the base premises of the two paradoxes are not logically equivalent.[24]

However, such an objection does not find to apply in the present context, given that one places oneself exclusively under the conditions of a version of SEP based on a disjoint definition. In these circumstances, one has indeed S[n, 0] = 1. But the fact that it is a version of SEP(II≈) involves S[n, 1] = 0. The consequence is that the premise n*n* is necessarily true, since the surprise cannot occur on day n. Hence, it appears that Sorensen's objection only concerns the monist solutions proposed by Dietl and Smith, and more generally, those solutions that do not distinguish between the versions of SEP based on a conjoint or disjoint definition of the surprise and the non-surprise. Given that they apply indifferently to SEP(I) and SEP(II), these solutions are vulnerable to Sorensen's objection. But such an objection doesn't apply to the present analysis, because the assimilation to SP is only effective on the more restrictive assumption of a version of SEP based on a disjoint and vague definition of the surprise, where the base premise reveal necessarily true.

It is worth finally considering another objection to the present analysis. If SEP(II≈) reduces to SP, according to this objection, one must expect that the fact of making the surprise notion precise dissolves the paradox, in the same way as making the vague concept entirely precise in SP leads to the disappearance of the paradox.[25] Does one veritably observes any such situation?

One can answer to that that in effect, the fact of making the surprise notion precise involves the dissolution of SEP(II≈). Indeed, but the fact of removing the borderline case(s) has the unique effect of transforming an instance of SEP(II≈) into an instance of SEP(II\).

5. Conclusion

I should mention lastly that the solution which has been just proposed applies, as far as I can see, to the variations of SEP mentioned by Sorensen (1982). Indeed, the structure of the canonical forms of SEP(I□), SEP(IΔ), SEP(II\) or SEP(II≈) shows that whatever the version taken into account, the corresponding solution does not require to appeal to a principle of temporal retention. The above solution is also independent of the order of elimination and can finally apply when the duration of the *n*-period is unknown at the time of the professor's announcement.

Lastly, it should be noted that the strategy implemented in the present study appears structurally similar to the one implemented in my analysis of *Hempel's problem* (1999): first, establish a dichotomy which makes it possible to divide the given problem into two distinct classes; second, show that each of the resulting versions admits a specific resolution[26]. Similarly, in the present analysis of SEP, a dichotomy is made and the two resulting categories of problems lead to an independent solution. This suggests that the fact that two structurally independent versions are inextricably entangled in the philosophical paradoxes could be a more widespread feature than one could think at first glance and could also explain partly their intrinsic difficulty.[27]

References

AYER, A. J. 1973. On a Supposed Antinomy. *Mind* 82: 125-6.
BINKLEY, R. 1968. The Surprise Examination in Modal Logic. *Journal of Philosophy* 65: 127-36.
CHALMERS, D. 2002. The St. Petersburg two-envelope paradox. *Analysis* 62: 155-7 .
CHOW, T. Y. 1998. The Surprise Examination or Unexpected Hanging Paradox. *The American Mathematical Monthly* 105: 41-51.
DIETL, P. 1973. The Surprise Examination. *Educational Theory* 23: 153-8.
FRANCESCHI, P. 1999. Comment l'Urne de Carter et Leslie se Déverse dans celle de Hempel. *Canadian Journal of Philosophy* 29: 139-56. English transaltion at http://www.univ-corse.fr/~franceschi
HALL, N. 1999. How to Set a Surprise Exam. *Mind* 108: 647-703.
JANAWAY, C. 1989. Knowing About Surprises: A Supposed Antinomy Revisited. *Mind* 98: 391-410.
MONTAGUE, R. & KAPLAN, D. 1960. A Paradox Regained. *Notre Dame Journal of Formal Logic* 3: 79-90.
O' CONNOR, D. 1948. Pragmatic paradoxes. *Mind* 57: 358-9.
QUINE, W. 1953. On a So-called Paradox. *Mind* 62: 65-6.
SCRIVEN, M. 1951. Paradoxical announcements. *Mind* 60: 403-7.
SHAW, R. 1958. The Paradox of the Unexpected Examination. *Mind* 67: 382-4.
SMITH, J. W. 1984. The surprise examination on the paradox of the heap. *Philosophical Papers* 13: 43-56.
SOAMES, S. 1999. *Understanding Truth*. New York, Oxford: Oxford University Press.
SORENSEN, R. A. 1982. Recalcitrant versions of the prediction paradox. *Australasian Journal of Philosophy* 69: 355-62.
SORENSEN, R. A. 1988. *Blindspots*. Oxford: Clarendon Press.
WRIGHT, C. & SUDBURY, A. 1977. The Paradox of the Unexpected Examination. *Australasian Journal of Philosophy* 55: 41-58.

[1] The object of the present section is not the study of the surprise notion *per se*. It rather aims at disentangling the surprise notion from its psychological underpinnings, in order to make finally more apparent the structure of the paradox.
[2] In what follows, *n* designates the last day of the period.
[3] Given that only positive instances (i. e. a prediction associated with a day where the examination truly occurs) are taken into account, and that negative instances (i. e. a prediction associated with a day where the examination does not occur) are simply ignored. This point is analysed in more detail at the end of the present section.
[4] Let 1-SEP, 2-SEP..., *n*-SEP denote the problem for 1 day, 2 days, ..., and more generally *n* days.
[5] Though I oversimplify here.
[6] For simplification, one uses here the same matricial notation as previously.
[7] Cases where S[*k*, 0] + S[*k*, 1] = 0 can be purely and simply ignored.
[8] As mentioned above (see fn. 3), emphasis on the surprise notion is only used here by way of simplification. One could consider alternatively:

(A7') you will not be in a position to perform an exact an justified prediction of the examination day
[9] See p. 65: 'It is notable that *K* acquiesces in the conclusion (wrong, according to the fable of the Thursday hanging) that the decree will not be fulfilled. If this is a conclusion which he is prepared to accept (though wrongly) in the end as a certainty, it is an alternative which he should have been prepared to take into consideration from the beginning as a possibility.'
[10] See p. 66: 'If *K* had reasoned correctly, Sunday afternoon, he would have reasoned as follows: "We must distinguish four cases: first, that I shall be hanged tomorrow noon and I know it now (but I do not); second, that I shall be unhanged tomorrow noon and do not know it now (but I do not); third, that I shall be unhanged tomorrow noon and know it now; and fourth, that I shall be hanged tomorrow noon and do not know it now. The latter two alternatives are the open possibilities, and the last of all would fulfill the decree. Rather than charging the judge with self-contradiction, let me suspend judgment and hope for the best."'
[11] 'The students are then shown four silver stars and one gold star. One star is put on the back of each student.' (1982, p. 357).
[12] Hall (1999, pp. 659-60).
[13] Hall refutes otherwise, but on other grounds, the solution suggested by Quine.
[14] In full rigour, this last formulation is inappropriate, since this particular version is not paradoxical.
[15] See notably Hall (1999, p. 661).
[16] It should be noted here that such *intermediate* values are precise, unlike the *borderline* cases considered below (see §5.2), which have a vague nature.
[17] It consists of an *exclusive or*.
[18] This structure of definition of the surprise as a vague concept proves sufficient to be used as a support to the present discussion. It leads however to difficulties when a thorough study appears necessary, for example in the context of the study of the sorites paradox. Such difficulties are mentioned by Scott Soames (1999, p. 207): "There is, however, more at work in Sorites paradoxes, and also more to vagueness, than just partiality. If partiality were the whole story and o were an object in the undefined range of a predicate F, then we would have no choice regarding how to characterize o. We would have to reject the claim expressed by 'it is F' as well as the claim expressed by 'it is not F'. But if we think of actual vague predicates, like *green*, *looks green*, *bald*, *rich*, and so on, we do not feel this way. (...) we can imagine circumstances in which it is perfectly correct to characterize the predicate as applying to the object or as failing to apply to it.".
[19] For the reasons mentioned above, the present study will be limited to the matrices presenting the linearity property.
[20] This definition of a vague predicate is borrowed from Soames (1999, p. 210). Considering the extension and the anti-extension of a vague predicate, Soames points out: "These two classes are mutually exclusive, though not jointly exhaustive".
[21] In the literature relating to SEP, the solutions suggested by Dietl and Smith are based on the assimilation of SEP with SP. However, these are monist analyses, which do not lead, unlike the present solution, to two independent solutions based on two structurally different versions of SEP.

In addition, with regard to the analyses suggested by Dietl and Smith, it does not appear clearly whether each step of SEP is completely assimilable to the corresponding step of SP. This characteristic has motivated the following remark from Sorensen:

> Indeed, no one has simply asserted that the following is just another instance of the sorites.
>
> i. Base step: The audience can know that the exercise will not occur on the last day.

ii. Induction step: If the audience can know that the exercise will not occur on day n, then they can also know that the exercise will not occur on day $n - 1$

iii. The audience can know that there is no day on which the exercise will occur.

> Why not blame the whole puzzle on the vagueness of 'can know'? (…) Despite its attractiveness, I have not found any clear examples of this strategy. (1988, p. 292-3).

Conversely, in the context of SEP(II≈), this step-by-step assimilation is fully endorsed.

[22] See 1988, p. 324: 'One immediate qualm about assimilating the prediction paradox to the sorites is that the prediction paradox would be a very 'fast' sorites. (...) Yet standard sorites arguments involve a great many borderline cases.'

[23] See Sorensen (1988, p. 324): 'A (...) reply would be to simply concede that the prediction paradox only contains one borderline case and then deny that a sorites paradox must have more that one borderline case. After all, some might insist that a single borderline case is sufficient to produce the familiar sorites pattern of clear cases being buffered by unclear cases.'

[24] See 1988, p. 324-5: 'A more serious objection emanates from the epistemological approach to the prediction paradox. Those who assimilate the prediction paradox to the sorites paradox accept the base step of both paradoxes. In the surprise test version, they agree with the clever student's claim that the test cannot be a surprise if given on the last day. This accords well with most people's initial reaction to the puzzle. But as proponents of the epistemic approach have shown, there are grounds for doubt about this premiss. Neither Dietl nor Smith provides any replies to these objections. Yet it is clear that the analogy with the sorites is considerably weakened if the base step of the prediction paradox is uncompelling.'

[25] This objection is due to Timothy Chow.

[26] A characteristic example of this type of analysis is also the solution to the *two-envelope paradox* presented by David Chalmers (2002, p. 157): 'The upshot is a disjunctive diagnosis of the two-envelope paradox. The expected value of the amount in the envelopes is either finite or infinite. If it is finite, then (1) and (2) are false (...). If it is infinite, then the step from (2) to (3) is invalid (...)'.

[27] I am grateful to Ned Hall, Timothy Chow and Professor Claude Panaccio, for very useful comments on earlier drafts of this paper.