

VINDICATING CHANCE: ON THE REDUCTIONISM/NON-REDUCTIONISM DEBATE

RAMIRO CASO
Instituto de Filosofía
Departamento de Filosofía
Facultad de Filosofía y Letras
Universidad de Buenos Aires
ramirocaso08@gmail.com

SUMMARY: The debate between reductionist and non-reductionist accounts of chance is presented, and the dialectical burdens acquired by each side of the debate are identified: the motivation problem and the explanation problem. It is argued that, while the motivation problem presents no challenge to non-reductionists, reductionists are unable successfully to answer it. Contrary to what has been suggested, both sides share the burden of the explanation problem. It is argued that non-reductionists can successfully answer it, whereas reductionists are unable to make the corresponding claim. Hence, the non-reductionist side has an important dialectical advantage.

KEY WORDS: Humean supervenience, Big Bad Bug, Principal Principle, credence-chance norms, open future

RESUMEN: Se presenta el debate entre reduccionismo y no reduccionismo respecto de la probabilidad objetiva y se identifican las cargas dialécticas adquiridas por cada posición: el problema de la motivación y el problema de la explicación. Se argumenta que, mientras que el problema de la motivación no presenta ningún desafío para los no reduccionistas, los reduccionistas no son capaces de responderlo exitosamente. Contrariamente a lo que se ha sugerido, ambos lados comparten el problema de la explicación. Se argumenta que los no reduccionistas pueden responderlo exitosamente, mientras que los reduccionistas no son capaces de hacerlo. Por lo tanto, el lado no reduccionista tiene una ventaja dialéctica importante.

PALABRAS CLAVE: superveniencia humeana, *big bad bug*, principio principal, normas de grados de creencia-probabilidad objetiva, futuro abierto

1. *Introduction*

It is customary to recognize two different concepts of probability: a subjective one, *credence* or *degree of belief*, which is the measure of reasonable partial belief, and an objective one, which may be called *objective chance* or *chance* (for short), and which may be thought of as the measure of the (objective) likelihood of an event's taking place (*cf.* Lewis 1980). An important question regarding these two concepts of probability concerns their relation, and in particular which restrictions are placed upon reasonable credence by objective chance: How should we let the objective probability of an event guide our

degree of belief in that event's taking place? A principle answering this question by spelling out exactly how we should let chance guide credence is called a *credence-chance norm*. By introducing chance as a guide to credence, credence-chance norms place further constraints on what reasonable partial belief amounts to, whereby they constitute substantive additions to classical Bayesian rationality.

In this paper, I'll address one of the main philosophical issues concerning the metaphysics of objective probability, the debate between reductionist and non-reductionist accounts of chance. As we'll see, credence-chance norms play an important role in this debate. Briefly, reductionist accounts of chance have it that facts about chance at a world are ultimately reducible to the global distribution of categorical facts throughout the entire history of that world, so that the world need not contain chances as primitive, non-reducible propensities or dispositions. Variation with respect to the nature of the supervenience base is possible, as long as no modal facts, no primitive dispositions or propensities are allowed among the fundamental properties of the universe. Non-reductionist accounts, on the other hand, claim that there are, at a fundamental level, non-categorical properties, primitive propensities or dispositions, responsible for the existence of a global chance distribution for a world at any given time. Are there any reasons to prefer one account of chance instead of the other?

Given that, as we'll see, one of the most intuitive ways of spelling out the connection between credence and chance —Lewis' 1980 *Principal Principle*— is inconsistent with reductionism about chance (a result that has come to be known as Lewis' *big bad bug*), credence-chance norms are central to answering this question and they play a key role in the debate between the two camps. As a consequence of the *big bad bug*, reductionists came to face the *motivation problem*, that is, the challenge of finding and motivating an alternative credence-chance principle compatible with their metaphysical assumptions, as well as the possible cost of being revisionary with respect to our ordinary way of epistemically relying on chance. Non-reductionists, on the other hand, face the *explanation problem*, that is, the challenge of explaining the rationality of endorsing any favored credence-chance norm, as well as the cost of positing it as a primitive principle of rationality if they cannot.

The formulation of conceptually well-motivated alternatives to the Principal Principle —Hall's 1994 *New Principle* and Ismael's 2008 recommendation— seemed to provide the reductionist with the upper

hand: the cost of moving to a new credence-chance principle was non-existent, or in any event negligible in view of the possibility of explaining the rationality of endorsing the favored credence-chance norm (thus apparently discharging or considerably diminishing the burden of the motivation problem). Non-reductionists, on the other hand, seemed still unable to meet their own challenge, thus providing a dialectical advantage to the reductionist side.

In what follows, I want to suggest that, contrary to what has been thought, it is the non-reductionist who actually has the dialectical advantage in this debate. I'll start by laying out some background in order to set the scene for the discussion to follow (section 2). Then I'll follow Briggs (2009) in identifying further problems for the reductionist, and I'll argue that, in view of these, the prospects for reductionists to be able to meet the motivation problem are slim at best (section 3). Then I'll turn to a fuller discussion of the explanation problem (section 4). As Hall (2004) argues, it turns out that the reductionist faces this problem too. In 4.1, I delineate a formal strategy, due to Pettigrew (2012), for justifying any favored credence-chance norm. In 4.2, I show how the non-reductionist is able to put in motion this formal strategy by providing a motivation for a key assumption upon which the strategy relies. In 4.3, I argue that the prospects for reductionists to be able to provide a similar justification are also slim. These considerations provide a dialectical advantage to the non-reductionist camp. I end with the discussion of a few objections that may be raised against the proposal (section 5) and with a summary of the discussion (section 6).

2. *Stage-Setting*

2.1. Humean Chances and the Big Bad Bug

Following Hall (2004), we lay down a few assumptions about chance. First, chances are probabilities defined over propositions. Second, chance is both time and world-dependent —thus, we speak of the chance of a proposition being true at a time t and at a world w . We represent the chance distribution at t at w by the probability function $P_{t,w}(\bullet)$. Third, chances are defined for arbitrary Boolean combinations of propositions, provided they are defined for the propositions so combined. Fourth, what's past is no longer chancy —that is, for any proposition A entirely about the past (relative to a time t), either $P_{t,w}(A) = 1$ or $P_{t,w}(A) = 0$, depending on whether A is true or false at t . Fifth, the chance distribution $P_{t',w}(\bullet)$ at a time

$t' > t$ may be obtained from the chance distribution at t by conditionalizing on the complete history of w from t up to and including t' , i.e., $P_{t',w}(\bullet) = P_{t,w}(\bullet|I_{t,t'})$, where $I_{t,t'}$ is a proposition completely describing the history of w between t and t' . In general, we can represent the chance distribution at a world at any given time in terms of an ur-chance function $urch_w$ and the complete history of w up to that time, so that, for any t , $P_{t,w}(\bullet) = urch_w(\bullet|H_{t,w})$, where $H_{t,w}$ is a proposition completely describing the history of w up to t . Finally, chances are law-governed, i.e., facts about chance at a world w at any given time are entailed by the laws of nature that hold at w and the complete history of w up to that time.

Lewis (1980) formulated a very plausible credence-chance principle, thus giving a precise expression to the idea that it is rational for us to let chance guide credence:

The Principal Principle (PP). Let C be any reasonable initial credence function, $ch_t(A) = x$ the proposition that the chance of A being true at time t equals x , and E a proposition compatible with $ch_t(A) = x$ that is admissible at t . Then:

$$C(A|ch_t(A) = x \ \& \ E) = x,$$

where E is the total evidence available to the agent at a given time. Lewis' assumptions concerning admissibility allowed him to derive a special case of (PP):

The Principal Principle (special case) (PP_S). Let C be any reasonable initial credence function, $P_{t,w}(\bullet)$ the chance distribution at w and t , $H_{t,w}$ the complete history of w up to t , and T_w the complete theory of chance for w . Then:

$$C(A|H_{t,w} \ \& \ T_w) = P_{t,w}(A),$$

where T_w is a conjunction of all the members of the maximally consistent set of history-to-chance conditionals that hold at w (that is, strong conditionals whose antecedents are propositions of the form $H_{t,w}$ (for some t in the history of w) and whose consequents specify facts about chance at w at t).¹ In this setting, chances are law-governed in the sense that, for any t , T_w , together with $H_{t,w}$, entail all the facts about chance at t at w .

¹ A proposition is admissible if and only if it is "the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes" (Lewis 1980, p. 272). Lewis gives three sufficient conditions for admissibility. First, a proposition is admissible at a time t (*modulo* the

Now, the Principal Principle conflicts with Lewis' own reductive account of chance, namely Humean Supervenience. According to this view, the fundamental features of reality are spatio-temporal relations and local qualities of spatio-temporal points or point-sized occupants of those points (cf. Lewis 1986b, p. ix). All else supervenes on those qualities and those relations. In particular, chance supervenes on the Humean base in an indirect way. As we saw, facts about chance at a time t depend both on facts about the history of the world w up to t and the laws of nature that hold at w (probabilistic or not). So, chances supervene upon the Humean base *via* the supervenience of both history and the laws of nature. There is no peculiar philosophical problem as to how history supervenes upon categorical facts. But there is a philosophical problem concerning how laws and, in particular, probabilistic laws, may depend on purely categorical properties. Lewis' answer is the *best system account of lawhood*: the laws at a world w are the theorems of the system that achieves the best balance between simplicity, strength, and fit. Even though they are not precise notions, we have enough grasp on the idea of simplicity and strength. However, the idea of fit is a tricky one. Lewis takes a two-step approach: first, discard the theories that make false statements about the history of w , and then look for the theory that assigns the highest probability to the actual course of events throughout the whole history of w —this is the one that best fits the whole history of w .

Here is where the big bad bug bites. For suppose that the complete theory of chance for a world w does indeed supervene on the global distribution of local, categorical properties throughout all of w 's history. Let G_w be one such global distribution for w . By the supervenience thesis, there is a complete theory of chance entailed by G_w —call it T_w . Now consider a time t well before G_w is complete. At t , G_w has a chance of obtaining, but also has a chance of not obtaining, so that $P_{t,w}(G_w) < 1$. Moreover, if t is adequately selected, there will be a chance of a very different global distribution G'_w obtaining,

possibility of foreknowledge, time travel, etc.) if it is entirely about the history of w up to t . Second, a proposition is admissible at t if it is a *history-to-chance conditional*. Finally, any Boolean combination of propositions admissible at t is itself admissible at t . This entails that the conjunction $H_{t,w} \& T_w$ is admissible at t . Then, we can let E be $H_{t,w} \& T_w$, so that $C(A \mid ch_t(A) = x \& H_{t,w} \& T_w) = x$. If $ch_t(A) = x$ holds at w , the strong conditional $H_{t,w} \square \rightarrow ch_t(A) = x$ also holds, hence $H_{t,w} \& T_w$ entails $ch_t(A) = x$; then, $ch_t(A) = x \& H_{t,w} \& T_w$ may be replaced without loss by $H_{t,w} \& T_w$. Moreover, if $P_{t,w}(\bullet)$ is defined on A , then $P_{t,w}(A) = x$, whence we get $C(A \mid H_{t,w} \& T_w) = P_{t,w}(A)$.

different enough from G_w so that T_w will no longer be the best fit for it. Rather, to it will correspond a different complete theory of chance T'_w incompatible with T_w . It follows from this incompatibility that $C(G'_w|H_{t,w} \& T_w) = 0$. However, for this particular t , we have that $P_{t,w}(G'_w) > 0$. The Principal Principle turns this into a full-fledged inconsistency:

Lewis' Result. (PP) and Humean supervenience are incompatible.

Proof. By T_w and $H_{t,w}$ we have that (1) $P_{t,w}(G'_w) \neq 0$. By the incompatibility between G'_w and T_w , we have that (2) $C(G'_w|H_{t,w} \& T_w) = 0$. By the admissibility of $H_{t,w} \& T_w$, $C(G'_w|ch_t(G'_w) = x \& H_{t,w} \& T_w) = x$ is a derivable instance of (PP). Since by hypothesis $ch_t(G'_w) = x$ holds at w for some x , $H_{t,w} \& T_w$ entails $ch_t(G'_w) = x$. Whence we get $C(G'_w|H_{t,w} \& T_w) = x$. Since $P_{t,w}(\bullet)$ is defined on G'_w , $P_{t,w}(G'_w) = x$, we get $C(G'_w|H_{t,w} \& T_w) = P_{t,w}(G'_w)$. Then, by (1), $C(G'_w|H_{t,w} \& T_w) \neq 0$. This contradicts (2).

This problem actually affects a broad class of reductionist accounts, for the culprit is the assumption that chance is reducible to the complete arrangement of categorical facts throughout the entire history of the world, regardless of the particulars of the reductive mechanism and the nature of the base for reduction. This has come to be known as Lewis' *big bad bug*: an analysis along reductionist lines is committed to the existence of undermining futures, that is, futures that (i) have a chance at t of occurring but (ii) whose occurrence would undermine the actual chances at t ; and the preceding argument shows that the existence of undermining futures is incompatible with an unrestricted acceptance of the Principal Principle. Obviously, something has got to give.

2.2. Alternative Credence-Chance Principles

The way out for the reductionist is to find a credence-chance principle compatible with the existence of undermining futures, and to provide a motivation for it, so that the move to the new principle doesn't involve an excessive cost.² The most promising alternatives

²The reductionist could also take issue with two further assumptions: that the chance function is defined for G_w and G'_w , and that $H_{t,w}$ and T_w both determine the chance distribution at t at w . This latter assumption is just the assumption that chances are law-governed, and we know of no good reason to reject it. The

are Hall's (1994) New Principle and Ismael's (2008) pair of principles, the Unconditional Principal Principle and the General Recipe for reasoning under uncertainty concerning chance.³

Hall (1994, 2004) proposed a replacement for (PP):

The New Principle (NP). Let C be any reasonable initial credence function, E the agent's total evidence, $P_{t,w}(\bullet)$ the chance distribution at w and t , $H_{t,w}$ the complete history of w up to t , and T_w the complete theory of chance for w . Then:

$$C(A|H_{t,w} \ \& \ T_w \ \& \ E) = P_{t,w}(A|H_{t,w} \ \& \ T_w \ \& \ E)^4$$

Hall derives (NP) from a deference principle that considers chance to be an *analyst-expert* (i.e., someone to whom we are willing to defer in virtue of considering her to be considerably better than us in evaluating the relevance of the information she's given), thus providing an independent conceptual motivation for (NP): it is (NP), rather than (PP), that adequately captures our understanding of chance, for it is grounded upon the epistemic role chance would

former assumption has been questioned, for example, by Ismael (2008). However, if a theory of chance fails to assign chances to undermining futures and hence to complete global patterns, the reductionist loses the grasp she has on the requirement of fit, and thus the grip she has on the best-system account of lawhood and chance. So, there is a dialectical concession in granting the assumption: without it, the most plausible versions of reductionism won't even get off the ground.

³The reductionist may also follow Thau's lead, alter the characterization of admissibility so as to make it relative to propositions as well as times (which seems to be the right thing to do anyway), and restrict (PPs) accordingly, so that $C(A|H_{t,w} \ \& \ T_w) = P_{t,w}(A)$ is not derivable when A describes an undermining future. However, this would cripple the applicability of (PP). Granted, it is not as crippling as declaring T_w inadmissible without qualification, for T_w will generally be admissible for futures too local or too restricted to be undermining. But it makes the application of (PP) quite limited nonetheless, for then it won't provide a guide on how to set our credence in propositions that express undermining futures. Briggs (2009, sect. 3.1) provides further reasons for thinking that restricting (PP) won't work, for even under suitable qualifications, it may give rise to contradiction.

⁴If E is true at w and entirely about the history of w up to t , then $H_{t,w} \ \& \ T_w$ entails E as well, so we can derive that $C(A|H_{t,w} \ \& \ T_w) = P_{t,w}(A|H_{t,w} \ \& \ T_w)$. Moreover, if $P_{t,w}(H_{t,w}) = 1$, this simplifies to $C(A|H_{t,w} \ \& \ T_w) = P_{t,w}(A|T_w)$, which is the special case of (NP) accepted by Lewis (1994), arrived at by considering our evidence to be entirely historical. If we further assume that $P_{t,w}(T_w) = 1$, then $P_{t,w}(A|T_w) = P_{t,w}(A)$, so that the special case of (NP) simplifies to $C(A|H_{t,w} \ \& \ T_w) = P_{t,w}(A)$. Hence, the special case of (NP) has the special case of (PP) as a special case, under the assumption that there are no undermining futures. However, if $P_{t,w}(T_w) < 1$, then, in general, $P_{t,w}(A|T_w) \neq P_{t,w}(A)$ —thus, if undermining is allowed, (NP) and (PP) come apart.

actually have for us. The problem with (PP) would be, then, that it treats chance as the *wrong kind of expert* (namely, as what Hall calls a *database-expert*, i.e., someone to which we are willing to defer in virtue of considering her to be substantially more informed than us).⁵

(NP) has some advantages over (PP). The most obvious one is that it is consistent with the existence of undermining futures, whereas (PP) is not, for even under the assumption that $P_{t,w}(G'_w) \neq 0$, by (NP) $C(G'_w|H_{t,w} \& T_w) = P_{t,w}(G'_w|H_{t,w} \& T_w) = 0$. Also, not only is it equivalent to (PP_S) under the assumption that there are no undermining futures, but it yields (PP_S) as a good approximation in garden-variety cases under the assumption that there are such futures, for in those circumstances $P_{t,w}(A|T_w) \approx P_{t,w}(A)$.⁶ So (NP) seems to be, in a sense, a deeper principle than (PP_S), for it corrects it, it yields (PP_S) as a good approximation in ordinary cases, and it explains why it fails in the extraordinary ones. Moreover, it doesn't depend, for its formulation, upon the notion of admissibility, which is not easy to characterize. And finally, Hall argues that (NP) correctly captures the purely epistemic role of chance.

But (NP) is not the only possible replacement. Ismael (2008) distinguishes between reasoning within a theory of chance and reasoning about which theory of chance is correct. When we reason in the first way, the rational thing to do is to equate our credence to the chances dictated by the theory. Thus, the correct deference principle for this case is a simplified version of the Principal Principle:

The Unconditional Principal Principle (UPP). Let C_t be the credence function of an agent at time t , $P_{t,w}(\bullet)$ the chance

⁵ By taking chance to be an analyst expert, Hall arrives at a particular version of a deference principle for chance, namely $C(A|CH^tP_{t,w} \& E) = P_{t,w}(A|CH^tP_{t,w} \& E)$, where E is the agent's total evidence and $CH^tP_{t,w}$ is the proposition that the chances at t are represented by the probability function $P_{t,w}$. This proposition says that chance is an expert in the sense that it states that we ought to set our credence in a proposition (conditional upon chance being the expert) to the "credence" chance would have in that same proposition, after being informed she is the expert, and after taking into account the rest of our evidence. The conditionalization upon $CH^tP_{t,w}$ reflects the idea that it is not part of what it is to be an expert to be certain that one is an expert. We arrive at (NP) by deriving from the previous principle that, for any F , $C(A|F \& CH^tP_{t,w} \& E) = C(A|F \& CH^tP_{t,w} \& E)$, by setting F to $H_{t,w} \& T_w$, and by simplifying $CH^tP_{t,w} \& H_{t,w} \& T_w$ to $H_{t,w} \& T_w$ due to entailment.

⁶ This also dispels, to some extent at least, the worry that (NP) is difficult to use in practice, for in ordinary cases all we need is the "rule of thumb" provided by (PP).

distribution at time t and world w , and $H_{t,w}$ the history of w up to t . Then:

$$C_t(A|H_{t,w}) = P_{t,w}(A)$$

We may tidy up this formulation a bit by dropping the time subscript for C and adopting talk of initial credence functions. Basically, what we get, restricting the evidence to $H_{t,w}$, is $C(A|H_{t,w}) = P_{t,w}(A)$, where C is any reasonable initial credence function and $H_{t,w}$ is the agent's total evidence. If we adopt the ur-chance notation, we may bring to the surface the parallelism between credence and chance, for (UPP) amounts to the claim that $C(A|H_{t,w}) = \text{urch}_w(A|H_{t,w})$. Importantly, Ismael's contention is that (UPP), rather than (NP) or (PP), is the principle that adequately captures the epistemic role of chance: if we are taking a theory of chance T_w as an epistemic policy rather than reasoning about its correctness, then our credence simply should follow chance. Now, given that we don't know which of the possible theories of chance is the correct one, our epistemic situation is not one of certainty concerning chance, and (UPP) is not much of a guide for us. Then, the deference principle we should use in attempting to adequate credence to chance, given our epistemic situation, is:

The General Recipe (GR). Let C be the credence function of an agent, $CH^{P_{t,w}}$ the proposition that the chances (at t at w) are given by $P_{t,w}$, and $P_{t,w}(\bullet)$ any possible chance distribution at t at w . Then:

$$C(A) = \sum_{P_{t,w}} C(CH^{P_{t,w}})P_{t,w}(A)^7$$

(UPP) and (GR) also present advantages with respect to (PP). Neither of them entails a contradiction under the assumption that there are undermining futures. The former entails that $C(G'_w|H_{t,w}) \neq 0$, given that $P_{t,w}(G'_w) \neq 0$, which is alright—for certainly $C(G'_w|H_{t,w})$ ought to be greater than zero. The latter entails that $C(G'_w) \neq 0$, which is also alright—for we are not sure which theory of chance is true, and as long as we give some positive credence to a theory of chance that assigns positive chance to G'_w (and we should, for G'_w is entailed by a consistent theory of chance given $H_{t,w}$), $C(G'_w)$ ought

⁷ Actually, Ismael formulates (GR) in terms of theories of chance only. Thus, for her, $C(A) = \sum_T C(T)P_T(A)$, where P_T is the chance distribution (at t at w) determined by $H_{t,w}$ and the theory T . However, this formulation doesn't account for the possibility of uncertainty concerning the history of the world. If we take into account this possibility, we arrive at the formulation in the body of the text.

to be greater than zero. (UPP) and (PP) are not, in general, equivalent. However, if we are certain with respect to the correct theory of chance, then (PP) yields (UPP), for $C(A|T_w \& H_{t,w})$ simplifies to $C(A|H_{t,w})$. Moreover, (PP) entails (GR), even though the converse entailment doesn't hold: (GR) is just (PP_S)'s recommendation for setting the unconditional degree of belief in A under uncertainty concerning chances (slightly re-written).

As it turns out, (NP) is incompatible with (UPP) and (GR), if we allow for undermining futures, for they entail different ways of determining unconditional credence. According to (NP), under certainty concerning the correct theory of chance, the credence in A should be $P_{t,w}(A|CH^t_{P_{t,w}})$, whereas both (UPP) and (GR) yield $P_{t,w}(A)$; and under uncertainty concerning which theory of chance is the correct one, the degree of unconditional credence in A should be $\sum_{P_{t,w}} C(CH^t_{P_{t,w}})P_{t,w}(A|CH^t_{P_{t,w}})$ according to (NP), whereas, according to (GR), once we introduce uncertainty concerning history, it should be $\sum_{P_{t,w}} C(CH^t_{P_{t,w}})P_{t,w}(A)$ —and, if we admit self-undermining chances, $P_{t,w}(A|CH^t_{P_{t,w}})$ and $P_{t,w}(A)$ yield, in general, different values. Hence, the reductionist has more than one recommendation to deal with, in facing the motivation problem. (If we do not allow undermining, then the recommendations are the same, for $P_{t,w}(A|CH^t_{P_{t,w}}) = P_{t,w}(A)$. So, the non-reductionist has only one recommendation to deal with.)

3. *Burdened Reductionism*

As already indicated, the challenge faced by reductionist conceptions of chance is to show that they are able to provide good motivation for (at least one of) the credence-chance principles proposed in the literature. It should already be obvious that one cannot accept (PP), for this will lead one directly into contradiction, and that restricting it is not a satisfactory solution either. What about (NP), (UPP) and (GR)?

Unfortunately, neither (NP) nor (GR) seem to be a good match for the reductionist. As for (NP), we may ask whether chance really functions as an analyst expert. As Briggs (2009, sect. 3.2) argues, the reductionist cannot accept that chance behaves that way. Rather, insofar as chance “makes statements” about chances at any given time by taking into account reliable information from the future (namely, that certain futures won't occur), chance behaves like a *database* expert who knows more than we could possibly know. So it seems that (NP), though consistent with reductionist assumptions, cannot be

given a good conceptual motivation within reductionist views. As for (GR), its adoption conflicts with Bayesianism, for, as Briggs (2009, sect. 3.4) remarks, it precludes certain uses of Bayes' Theorem. So it seems that the reductionist has a choice to make concerning (GR): either she preserves the applicability of Bayes' Theorem and forgoes this principle, or she preserves (GR) and forgoes the applicability of Bayes' Theorem, thus paying a pretty high cost.

Can the reductionist accept (UPP)? Indeed she can, but this is of little comfort. To see that she can, we just have to notice that, besides being consistent with the existence of undermining futures, (UPP) fits well with a conception of chance as a database expert. Indeed, all (UPP) says is, roughly, that, unless one knows better, one should adjust her credence to chance. Less roughly, (UPP)'s recommendation is that, unless one has reliable information *from* the future, one should guide her credence by the "credence" of chance. To see that this is of little comfort, recall that both (PP) and (NP) provide useful directions in conditional reasoning about chances. But (UPP) doesn't do this: its companion principle, (GR), has this burden. Moreover, with (UPP) alone, we are at a loss as to how to reason under uncertainty concerning the correct theory of chance. Settling just for (UPP) is quite crippling for the reductionist.⁸

It seems, then, that the reductionist is able to meet her challenge rather poorly: (PP) is out of her reach, (NP) is poorly motivated, and (GR) is in conflict with an important principle of Bayesian rationality. The only principle she seems to be able to motivate in line with her assumptions concerning chance is (UPP), a rather weak principle that cannot do all we require from a credence-chance norm. Her situation seems to be less than satisfactory. (The non-reductionist, on the other hand, can accept any of these principles. As for (PP), she can accept it without modification or restriction; as for (NP), she can conceive of chance as an analyst expert, since for her the chance distribution at any given time does not depend on the future, so it doesn't convey or rely upon information about it; finally, as for (UPP) and (GR), she can accept both of them —if she thinks chance works like a database

⁸This invites the question: May the reductionist accept both (UPP) and (GR), and restrict the applicability of Bayes' Theorem to those cases in which the evidence falls short of describing an undermining future? She may, but her situation will be less than satisfactory, for she won't be able to accommodate the kind of evidence that, were we to have it, would definitely discard a given theory of chance. Moreover, she won't be able to conduct conditional reasoning on the assumption that she's in possession of such evidence. It seems that this way out doesn't come without an important cost.

expert, she can embrace (UPP) without problem; and she can retain both (GR) and Bayes' Theorem, since for her there is no problem in setting $C(B|T)$ to be $P_T(B)$ in general, thus providing a well-defined quantity for the theorem to work with.)

4. *The Explanation Problem*

We've seen that reductionism doesn't cope well with its dialectical challenge. Let's turn now to the non-reductionist side. The non-reductionist's challenge is to explain why it is rational to adopt (PP), or any of the other credence-chance principles proposed, as a norm. And, were the non-reductionist unable properly to answer this challenge, the situation could provide the dialectical advantage to reductionism. Moreover, according to Lewis (1994, pp. 484–485), the prospects of her being able to do so seem to be quite slim, for it would require showing how knowledge of whatever feature of reality the non-reductionist identifies with chance could constrain rational credence.

However, as Hall (2004) points out, it is far from clear what the explanatory demand is supposed to be. For simply to say that we are in the presence of an analytic principle grounded upon the concept of chance won't do, for even though this answer may be acceptable by the non-reductionist's own lights, it is insufficient in this dialectical context. Maybe the way to go is to interpret the explanatory demand in terms of the derivability of the favored principle from what Hall calls the "categorical constraints" on reasonable initial credence functions, plus Bayesian assumptions concerning rationality. But then the situation ends in a dialectical stand-off. On the one hand, reductionist accounts do not actually *show* how the categorical facts may rationally constrain credence in the required way (Hall 2004, sect. V), so there is no real dialectical advantage here. On the other, while it is true that the non-reductionist cannot derive *any* credence-chance principle from the categorical constraints she recognizes, so that she has to posit them as primitive principles of rationality, it is not clear that the reductionist *can* provide the required derivation either. Moreover, if she cannot so derive her favored credence-chance principle, then she cannot even posit it as *primitive*: in a reductionist view, a credence-chance principle either is a categorical constraint placed upon credence by the feature of reality that underlies facts about chance, or it fails to be a constraint at all. So, any advantage the reductionist may claim on these grounds seems to be slim at best: the explanation problem is a problem for her also. The glimpse

of a potential advantage for reductionism now fades away: *vis-à-vis* the explanation problem, reductionism and non-reductionism are on equal ground. So, what now?

Taking a step back, we can restate the explanation problem. The aim is to offer an explanation along the following lines: *it is rational for an epistemic agent to set her credence to chance in the way recommended by such-and-such credence-chance principle because...*, where the dots are to be filled with some clause or other. The reductionist would then cite facts about how knowledge of the chance-makers she favors (frequencies, symmetries, best systems) may constrain credence in a rational way. The non-reductionist does not have such facts to cite. However, as I've just indicated, this doesn't provide much of an advantage to reductionist positions, for the reductionist hasn't really met this explanatory demand either: she has not shown *how* such a knowledge can indeed rationally constrain credence in the required way.

Now, let's bracket for the moment the explanation problem for the reductionist. Can the *non-reductionist* account for the rationality of abiding by the credence-chance principle she favors? If she can, and the reductionist cannot, this will provide an important dialectical advantage to non-reductionist views of chance. And I think she can, to a certain extent at least (and that reductionism cannot, but more on this later). The general answer will take the following form: it is rational for an epistemic agent to set her credence to chance in the way recommended by the favored credence-chance principle because *if she did otherwise, she would be in a worse epistemic situation than the one she would be in, if she abode by that principle*. How may the non-reductionist show this?

4.1. A Formal Justification of Credence-Chance Principles

Pettigrew (2012) seems to provide the beginning of a solution to this problem. The general strategy is to provide a formal justification of any of the credence-chance principles proposed, depending on the particular assumptions we make. The central idea is that, for each possible world $w \in \mathcal{W}$, there is a credence function v_w that is vindicated at w (i.e., that is perfect or exactly correct at w). It's possible to measure the distance of any credence function $b \in B$ (with B the set of all credence functions definable over a given algebra of propositions F) from being vindicated at w , for any w , in terms of its *inaccuracy measure*, that is, in terms of its distance at w with respect to v_w . Define V as $\{v_w : w \in \mathcal{W}\}$, and define the *convex hull of V* as

the smallest set V^+ which contains every vindicated credence function in V and, for any two credence functions which belong to V^+ , all of their mixtures (where the mixtures of two credence functions b and b' are all the credence functions $\lambda b + (1 - \lambda) b'$, for every $\lambda \in [0,1]$). Then, provided the inaccuracy measure satisfies certain constraints, it is possible to establish the following result:

Joyce's Theorem. Let B be the set of all probability functions definable over a (finite) algebra of propositions F , and $I : B \times W \rightarrow \mathbb{R}$ an inaccuracy measure. Then, for any $b \in B$:

1. $b \in B \ \& \ b \notin V^+ \rightarrow \exists b'[b' \in V^+ \ \& \ \forall w(I(b', w) < I(b, w))]$
2. $b \in V^+ \rightarrow \neg \exists b'[b' \in B \ \& \ \forall w(I(b', w) \leq I(b, w))]$

That is, an agent whose credence function is not in V^+ will do better (epistemically speaking), no matter how the world turns out to be, if she moves to a credence function in V^+ (i.e., her credence function will be closer to vindication in every world), and an agent whose credence function is in V^+ will do worse (epistemically speaking) in at least one world if she moves to another credence function (i.e., her credence function will be further from vindication in at least one world).

This theorem provides a formal justification for the following norm schema:

(GN) An agent ought to have a credence function in V^+ .

However, (GN) won't be of any use until the identity of V^+ is fixed, and this won't happen until the identity of v_w is fixed. There are a few natural candidates for being the vindicated credence function at a world. On the one hand, we may suppose that the vindicated credence function at a world w is the "truth" credence function for w , that is, the credence function that assigns 1 to all propositions true at w , and 0 to all propositions false at w . On the other hand, we may consider the credence function vindicated at w to be the ur-chance function for w , or something close to it. Then, it is possible to prove the following results:

1. If $v_w(\bullet) = \text{truth}_w(\bullet)$, V^+ is the set of all credence functions over F that satisfy (Prob)⁹

⁹Where (Prob) is the norm of Probabilism (that is, the norm that says that one ought to have a credence function that is also a probability function), and V^+ is the class of all credence functions that are also probability functions.

2. If $v_w(\bullet) = urch_w(\bullet)$, and there are no undermining chances, V^+ is the set of all probability functions over F that satisfy (PP)
3. If $v_w(\bullet) = urch_w(\bullet|CH_{urchw})$, and there are undermining chances, V^+ is the set of all probability functions over F that satisfy (NP)
4. If $v_w(\bullet) = urch_w(\bullet)$, and there are undermining chances, V^+ is the set of all probability functions over F that satisfy (GR)

These results yield four specifications of (GN):

(P-N) An agent ought to have a credence function satisfying (Prob)

(PP-N) An agent ought to have a credence function satisfying (PP)

(NP-N) An agent ought to have a credence function satisfying (NP)

(GR-N) An agent ought to have a credence function satisfying (GR)

(What about satisfying (Prob) in (PP-N) – (GR-N)? Since we've assumed that chance is a probability function, any credence function satisfying (PP), (NP) or (GR) will also satisfy (Prob).¹⁰ If the non-reductionist can show that $v_w(\bullet) = urch_w(\bullet)$, then she can show that it is rational to follow the credence-chance principle she favors. (Note that she need not decide between $urch_w(\bullet)$ and $urch_w(\bullet|CH_{urchw})$: since, for the non-reductionist, there need not be self-undermining chances, she may take it that $urch_w(\bullet) = urch_w(\bullet|CH_{urchw})$.) More precisely, if she assumes (as she most certainly will) that there are no self-undermining chances and she identifies $v_w(\bullet)$ with $urch_w(\bullet)$, then we have a formal justification of (PP): it is rational for an agent to follow (PP) because, if she does otherwise, she will risk being in a worse epistemic situation. Moreover, since under the assumption that there are no undermining futures, (PP), (NP) and (GR) are equivalent, this amounts to a formal justification of all the credence-chance principles so far proposed. Note that if we allow self-undermining chances, we lose the possibility of justifying (PP) (as it should be), and which principle we end up justifying depends on whether we


¹⁰ I owe this observation to Rachael Briggs.

consider $urch_w(\bullet)$ or $urch_w(\bullet|CH_{urchw})$ to be the credence function vindicated at w . So, this strategy is open to the reductionist also (we'll come back to this in 4.3).

4.2. Vindicating Chance

To solve the explanation problem, the non-reductionist needs to show that $urch_w(\bullet)$ is the credence function vindicated at w . So, the question for her is: Why not truth? That is, why not $v_w(\bullet) = truth_w(\bullet)$? Let's start by noticing that we have two different conceptions of indeterminism. One is epistemic, and consists in the impossibility of knowing all the truths about the future (relative to a time t) of a world w from the complete knowledge of the laws of nature that hold at w and the complete knowledge about the entire history of w up to t . The other is metaphysical, and considers all possible futures (relative to a time t) as "real" or "objective" alternatives (Belnap *et al.* 2001, chap. 6), rather than as epistemic possibilities arising from the unknowability of future history. There seems to be little point, if any, in discussing which of these conceptions reflects "true" indeterminism, or which of them better captures our ordinary conception of an indeterministic world. What's important for our purposes is the constraints each of these conceptions place on the structure of history.

As Thomason (1970), Belnap *et al.* (2001), and MacFarlane (2003), among others, have variously argued, a metaphysically indeterministic conception of history calls for a forward-branching time structure without a designated history. The picture is well known, but it is worth reviewing. Time in a metaphysically indeterministic world may be represented as a tree:



where the graph is oriented from left to right. In the graph, t_0 represents the present; the absence of branching to the left of t_0 encodes the assumption that what's past is no longer chancy; h_1 , h_2 and h_3 are possible histories with respect to t_0 , that is, possible ways

the history after t_0 might go; the absence of a designated history represents the view that no history is “ontologically privileged”, that none is “actually going to take place”.

The non-reductionist may endorse a metaphysical understanding of indeterminacy, and I think she should, for two reasons. First, for the possibility of justifying the assumption that $v_w(\bullet) = urch_w(\bullet)$, as we’ll shortly see. Second, because it is more congenial with primitivism about chance. This will become clearer when we move to considering how reductionism may address the explanation problem. For the moment, just notice that the non-reductionist doesn’t rely on the existence of a metaphysically privileged future to make sense of her conception of chance, so she’s not constrained to posit such a thing. For her, the global pattern of outcomes and frequencies need not be given before chances can be given —her direction of explanation is the other way around: first she recognizes chances as primitive features of the world, and then she considers any pattern of outcomes and frequencies that may eventuate to be a consequence of a real, objective openness of future history.

From this non-reductionist point of view, truth may strike us as the vindicated credence function for a world only if we adopt an external perspective on that world: we imagine the world as if it were in front of us, in its entirety, so to say, and from that vantage point we ask which credence function is the correct one for that world. And we adopt a “wait and see” attitude: the correct credence function is the one that assigns 1 to any proposition that turns out true at w , and 0 to any proposition that turns out false at w . But things look importantly different when we adopt the internal perspective. From this perspective, the metaphysically indeterministic nature of future history gives rise to semantically defective discourse: unless a given aspect of the future is settled at t , that is, unless it is determined at t that *that* aspect will come about, any proposition describing that aspect of the future lacks truth value at t . Consequently, at that time there is no fact of the matter as to which credence function “gets it right” in terms of truth. True: we can always, in making retrospective assessments, claim that one credence function “got it right”, as it were. We can even say that the credence function that corresponds to truth at w is the one that gets everything right throughout the whole history of w . But this doesn’t mean that it was the credence function that was right at that moment. It is this possibility of making retrospective assessments what strikes me as generating the impression that truth is the vindicated credence function in a metaphysically indeterministic world. But once we look closer at what happens with

truth values in a metaphysically indeterministic world w , it turns out that the chance function for w is better suited to play the role of the credence function vindicated at w .

How does retrospective assessment work? Suppose S_1 asserts at t_0 a certain proposition A (say, that there will be a sea-battle within the day, that the toss of the coin will come up heads), whereas S_2 asserts, also at t_0 , the negation of A . Basically, we ascertain the correctness or the incorrectness of the assertions by waiting and seeing how history actually goes: if history continues in a way that makes A true (say, h_1), then we say that S_1 was right after all, that what she said was true. We may even say things such as that S_1 was correct, or spoke truly, because she said of history that it was going to continue the way it actually did. If history continues in a way that makes A false (say, h_2 or h_3), then we say that S_1 was mistaken, or that what she said was false. Again, we may even say that S_1 spoke falsely because she said of history that it was going to continue in a way it actually didn't. (The same goes, *mutatis mutandis*, for S_2 's assertion.) That is, in order to assess an assertion about the indeterministic future, we stand from a time after the assertion is made, a time at which the events that yield a definite truth value for the asserted proposition either definitely occurred or definitely did not, and then, from that vantage point, we determine whether the assertion was true or false.

The potential *non sequitur* comes when we slide from these formulations to saying that S_1 was right, or spoke truly, because she said of history that it was going to continue *in the way it was actually going to continue*. Note here the change in the verb—for there is a natural yet inadmissible move from talk of the (possible) future that actually came to pass, to talk of the (possible) future that, at the time of the assertion, was actually going to take place. It is this last formulation that is inadmissible: if the world is metaphysically indeterministic, then it is not settled, at t_0 , which of all the possible continuations of history is going to take place. Hence, at t_0 , there is no unique way in which the history after t_0 is *actually* going to develop. So, even though we may say, from the vantage point of the way history actually developed, that S_1 was correct or incorrect, this doesn't imply that there was, at t_0 , a fact of the matter as to which prediction was the correct one.

Now, to drive the point home, consider all super-asserters for a world w at a given time t . A super-asserter at t for w asserts two very complex propositions: the proposition $H_{t,w}$, and a proposition $F_{t,w}$ that completely describes a possible continuation of w 's history

after t , that is, a proposition completely describing some possible future of w . (We leave aside assertions of propositions not about the history of w , even though they are part of what is true or false at a world, for the complication adds nothing of importance.) Let's suppose, for simplicity, that, for each complete proposition about the future of w after t , there is exactly one super-asserter that asserts it, and let's stipulate that, if two propositions $F_{t,w}$ are necessarily equivalent, they have the same super-asserter. Then, super-asserters and possible futures are in one-to-one correspondence. Obviously, super-asserters represent all the possible truth-value assignments to propositions about the entire history of w at a given time. For entirely historical propositions, there is only one such assignment. For propositions about the future, each possible distribution of truth values is represented by one super-asserter.

We may see the unfolding of w 's history as a way of discarding super-asserters. Each time a new event takes place, one or more super-asserters are discarded, namely the ones that made assertions incompatible with the way history actually continued. The super-asserters are reduced to one only when the history of w gets closed, that is, only when there are no more chancy events in the history of w . This may happen because the history of w comes to its end, or because it reaches a time t_c at which all the events after t_c are already necessitated by the state of the world at t_c plus the laws that hold at w . It doesn't matter for our purposes which one it is. All that matters is that we can say which super-asserter "got everything right" only from the vantage point of this "end of history", not a moment before. Crucially then, we cannot say, at any given time before the history of w gets closed, that there is a super-asserter that got it right already, but we do not and cannot know which: this would be endorsing some kind of "thin red line" view on the future that is incompatible with a metaphysically indeterministic understanding of the world. At any given time before history gets closed, there is no actual way history is going to develop. Consequently, there is no super-asserter that actually got it right.

Obviously, super-asserters are just literary devices for representing all the candidate "truth" credence functions for w at any given time t . They are useful devices because what happens to super-asserters happens to the candidate credence functions for truth in a metaphysically indeterministic world: it is not settled which credence function is the "truth" credence function for w until the history of w gets closed, that is, until there are no more chancy events. This is so because there is no actual way history is going to develop at

any time prior to the time at which history gets closed. It's not that propositions about the indeterministic future have a determinate truth value that is inaccessible to us; it's not that they have a truth value, but it is indeterminate which one it is: rather, they have no truth value at all —their truth or falsity is not determined yet.¹¹ Hence, there is no fact of the matter as to which candidate for being the “truth” credence function for w is the correct one. In a sense, then, no one is: reality lacks the determinateness that is required for it to select one of them as the credence function that reflects truth for w . But this is so, crucially, because what is truth for w itself is underdetermined by reality.

This suggests that the vindicated credence function at w cannot be the credence function that corresponds to truth at w , for there is no such unique credence function —rather, we have quite a good deal of credence functions that compete for being the truth credence function for w , none of which enjoys any metaphysical or epistemological advantage over the others. The chance function for w at t (for any t), on the other hand, is perfectly well defined at t itself (provided the chance distribution at t is well defined). This makes it the perfect candidate for being the credence function vindicated at w —for, at any given time, there is a perfectly determinate sense in which the chance function “gets it right”: it is exactly correct in what it says about chances at w .

4.3. Unvindicating Reductionism

So, non-reductionists are able to meet their share of the explanation problem. What about the reductionist? Suppose she successfully answers the problems identified by Briggs (2009) and summarized in section 3, or suppose she accepts them as a cost she has to live with. Can she successfully meet her share of the explanation problem? The strategy outlined in 4.1 is open to her, provided she can show that either $urch_w(\bullet)$ or $urch_w(\bullet|CH_{urchw})$ is the vindicated credence function at w , rather than $truth_w(\bullet)$. So, if she can show that $truth_w(\bullet)$ is not the vindicated credence function at w , she'll have a formal justification of the credence-chance norm she favors (provided she takes the extra step of adjudicating vindication between $urch_w(\bullet)$ and $urch_w(\bullet|CH_{urchw})$). I want to suggest that she won't fare well at this task.

¹¹ Determinate yet inaccessible truth values are just another way of treating indeterminacy as an epistemic phenomenon. Indeterminately having a truth value, on the other hand, is something of which not much sense can be made in this framework.

Reductionists seem to be committed to an epistemic understanding of indeterminism. After all, the way reductionists explain chance requires that there be an ontologically privileged history, a unique way history is actually going to develop, for there to be a unique global pattern of outcomes and frequencies. The direction of explanation is the inverse with respect to non-reductionism: first the pattern, then the facts about chance. That is, the pattern has to be already given, so that chance may supervene upon it. So it seems that the epistemic construal of indeterminacy is the view she has to endorse.

Let's start by noticing that epistemic indeterminism is compatible with two different time structures: a linearly ordered history and a forward-branching time structure with a designated history (informally, the history that will actually take place).¹² These structures allow for a perfectly well-defined truth distribution at any given point of history. In the case of a linearly-ordered time structure, sentences (propositions) are evaluated with respect to the only possible continuation of history, thus being true or false at any time within that history. In the case of a forward-branching time structure with a designated history, matters are a bit more complex but essentially the same: at any given time t , sentences (propositions) not involving operators such as "It is settled that" and "It is possible that" will be evaluated with respect to the designated history, while sentences (propositions) involving those operators will be evaluated with respect to every history passing through t , to prevent collapse of operators.¹³ In any event, the alternative, non-designated histories in this framework have the status of epistemic possibilities, which are not ontologically equivalent to the designated one. So, in either case, facts about chance are settled by the actual history of the world, be it because there is only one possible history, be it because there is an ontologically privileged one.

Hence, at any given time t within the history of w there is a fact of the matter as to which, of all the truth distributions compatible with the history of w up to t , is the truth distribution that matches the complete history of w . Now, if the complete truth distribution for w is well-defined, then it seems that $truth_w(\bullet)$, and not $urch_w(\bullet)$ or $urch_w(\bullet|CH_{urchw})$, is the credence function vindicated at w . For $truth_w(\bullet)$ contains all the information chance has, and more besides

¹² Lewis endorses a non-branching history in 1986a, but the other way of understanding time is also open to the reductionist.

¹³ In a linearly-ordered time structure, the collapse may be handled by giving these operators a straightforwardly epistemic reading.

—with respect to the past, relative to any time t , it agrees entirely with chance; with respect to facts about chance at t , it states all of them by assigning a truth value to statements reflecting those facts; and with respect to outcomes concerning the future relative to t , it gives *more information* than chance does. So, if one is exactly correct with respect to the complete history of w , it seems to be truth.

This invites the question: Can the reductionist adopt a forward-branching time structure without a designated history as her conception of time, so that truth be undefined? If at all compatible with her understanding of indeterminacy and chance, the cost of doing so would be excessive. For one thing, absent any privileged history, there is no unique, privileged global distribution of categorical properties for chance to supervene on. Also, as we saw in 4.2, such a structure gives rise to semantically defective discourse: certain propositions will lack truth value at certain points in history. The problem is that not only statements about outcomes will be semantically defective in this sense, but also statements about the chances of outcomes will turn out to be so. This is, at least, one lesson to draw from the existence of undermining futures in a forward-branching framework without a designated history. For statements about chances will receive their truth values in the same way other statements do: by supervaluating over possible histories. If, following the reductionist, we posit that chances are relative to complete histories, and we allow histories that assign different chances to the same events, then, relative to different histories, statements of the general form $ch_t(A) = x$ will have different truth values. This means that these statements will, in general, turn out to be undefined also.¹⁴

There are a couple of further moves the reductionist may attempt. First, she might remark that, under uncertainty concerning which outcome will come about, it is only reasonable to guide oneself by chance, not by truth. Second, she may reinforce this by saying that, since truth is in principle inaccessible to us (for complete knowledge of the laws of nature and complete knowledge of the entire history up to any given time do not yield complete knowledge about the future with respect to that time), it cannot be the credence function vindicated at w —hence, it must be chance (or something close to it). She should resist the temptation of giving “reasonable” a prudential

¹⁴ The reductionist may choose a third way: she may claim that there are various possible histories, but that each of them determines the same chances. This way, semantically defective sentences are safely bounded so as to exclude statements about chance. This won't be much of a refuge, though, for chances will be subject to the implausible kinds of fluctuation described by Arntzenius and Hall (2003).

reading: it should not be claimed that it is reasonable in a practical sense to do so, for the reasonableness or rationality involved in the explanation problem is epistemic. Her best hope seems to be to argue that, due to our less-than-ideal epistemic situation, agreement with chance is all we get, and that this is enough for the vindication of chance. However, such considerations seem somewhat extraneous to the concept of vindication, and we are owed a fuller justification. So, until such a justification is offered, the dialectical advantage seems to be on the non-reductionist's side.

5. *Some Loose Ends*

In the previous section, I argued that, contrary to what's commonly held, it is the non-reductionist who has the dialectical advantage. In this section, I would like to address a few loose ends that might give rise to objections to the way in which I mounted the case for non-reductionism pertaining chance.

The first problem that should be addressed concerns the status of the Law of Excluded Middle (LEM): Are sentences of the form $\varphi \vee \neg\varphi$ true when φ expresses a chancy proposition? Since at a time t , a proposition is true just in case it is true at every history that passes through t , such sentences must be true, for every history passing through a time t will be a history at which $\varphi \vee \neg\varphi$ holds (for either it will be a history at which φ holds, or a history at which $\neg\varphi$ holds) —hence, t itself will be a time at which $\varphi \vee \neg\varphi$ holds (this is the result of supervaluating over histories in order to characterize truth).

But then, it seems that counterintuitive consequences follow, for we have to reject that the T-sentences hold for this truth predicate —otherwise, bivalence will hold for chancy propositions. Indeed, assume that $\varphi \vee \neg\varphi$ holds. Also assume that we have the corresponding T-sentences: $True(\langle\varphi\rangle) \leftrightarrow \varphi$ and $True(\langle\neg\varphi\rangle) \leftrightarrow \neg\varphi$. Since classically valid arguments are supervaluationally valid, we can move from $\varphi \vee \neg\varphi$ to $True(\langle\varphi\rangle) \vee True(\langle\neg\varphi\rangle)$. If we understand being false as having a true negation, this conclusion amounts to $True(\langle\varphi\rangle) \vee False(\langle\varphi\rangle)$ —and this a statement to the effect that bivalence holds for φ . So, given LEM and supervaluational validity, we have to reject the T-sentences (at least) for chancy propositions.

And more counterintuitive consequences will also follow. For assume that A expresses a chancy proposition at a time t , so that $P_{t,w}(A) = r$, for some $r \in (0,1)$. Then, the rational thing to do, e.g. if we take (PP_S) as the favored credence-chance norm, will be to set

$C(A|H_{t,w} \& T_w)$ to r . However, since A is neither true nor false at t , $True(\langle A \rangle)$ will be false at t , so that $P_{t,w}(True(\langle A \rangle)) = 0$. Hence, $C(True(\langle A \rangle)|H_{t,w} \& T_w)$ should be 0, not r . And all of this seems counterintuitive, for this amounts to saying that, even though the chance of the coin landing heads is r (and even though we should believe to degree r that the coin will land heads), the chance of the proposition that the coin will land heads being true is 0 (hence, we should believe to degree 0 that it is true that the coin will land heads).

Now, this might strike one as unpalatable, for it seems that the T-sentences are constitutive of our understanding of the truth predicate, so that rejecting them is a costly philosophical move. And related to this point, it seems that, for any sentence A , A and “It is true that A ” are in some sense equivalent, so that making different assignments of chance and credence to A and $True(\langle A \rangle)$ would amount to assigning different probabilities to equivalent sentences, something *definitely* odd. In any event, it seems, accepting LEM for chancy propositions is a costly philosophical move.¹⁵

But this need not be so, for the truth predicate characterized by supervaluating over histories is equivalent to the settledness operator in the sense that whenever “It is settled that A ” is true at a time t , so will be “It is true that A ” and vice versa. So the claim that a proposition p is true at a time t is really equivalent to the claim that it is settled that p at that same time. So, $True$ is a strong truth predicate, in the following sense: $True(\langle A \rangle)$ is true (at a time t) just in case it is settled that A at t , and it is false (at a time t) if either it is settled that $\neg A$ at t or it is indeterminate at t whether A or $\neg A$. Yet there is no problem if we introduce a weak truth predicate Tr for which the T-sentences hold.¹⁶ For this truth predicate, it will hold that A and $Tr(\langle A \rangle)$ are equivalent, so it will hold that $P_{t,w}(A) = P_{t,w}(Tr(\langle A \rangle))$.

¹⁵ Rejecting LEM for chancy propositions would also be costly, for where A expresses a chancy proposition, $A \vee \neg A$ won't be logically true —hence, $P_{t,w}(A \vee \neg A) < 1$, thus requiring a deviant probability theory.

¹⁶ See Fine 1975 and McGee and McLaughlin 1994 for details. They introduce a disquotational truth predicate for a language with supervaluational semantics for vagueness, but the case is parallel to that of supervaluational semantics for future contingents. It is noteworthy that Thomason (1970), in introducing truth into the object language, introduces a truth operator semantically equivalent to double negation. Hence, all T-sentences hold for this truth operator, which thus expresses a weak, “disquotational” concept of truth. Then, strong truth could be made to coincide with the inevitability operator, for which the T-sentences do not hold.

The philosophical interpretation of this situation is up for grabs: we may claim, as in Field (1994), that we have two ways of using our ordinary truth predicate, a weak (disquotational) and a strong (non-disquotational) one, that our ordinary truth predicate has two senses, that we have two different concepts of truth (as Fine (1975) seems to think), or that our truth predicate is guided by two different principles (a Correspondence Principle and a Disquotational Principle) that become incompatible when truth-value gaps are introduced, as in McGee and McLaughlin (1994).¹⁷ Be that as it may, we may distinguish (regiment) two different truth predicates, a weak, disquotational truth predicate Tr , and a strong, non-disquotational one, $True$. The T-sentences will hold for Tr , but not for $True$. Now, given that $True$ is a strong truth predicate, it is not odd that $P_{t,w}(A) \neq P_{t,w}(True(\langle A \rangle))$. It would be odd for the weak truth predicate but, since the T-sentences hold for Tr , $P_{t,w}(A) = P_{t,w}(Tr(\langle A \rangle))$, as we've already remarked.

So, it seems that the acceptance of LEM for chancy propositions is not really a costly move, for the seemingly unpalatable consequences are actually palatable when we are dealing with the strong truth predicate $True$, and they do not arise when we are dealing with the weak truth predicate Tr .¹⁸

The second problem that should be addressed concerns the criticism of the truth norm that was developed in section 4, for there we considered the truth norm and the chance norm(s) as if they were in competition. However, it is far from clear that they are, for all the truth norm establishes is that credence functions should be probability functions, while the chance norms establish that some probabilistic credence functions (those that obey a further constraint related to chance) are better than others. Do we have to reject the truth norm as vindicated, in defending the chance norm(s)? Moreover, in rejecting the truth norm, haven't we relied on an excessively metaphysical notion of truth, one that is not required by the truth norm at all? We've stated the truth norm as the injunction to have a credence function such that $C(A) = 1$ just in case A is true, and $C(A) = 0$ just in case A is false. But the disquotational notion of truth is enough to describe this credence function, we don't need a strong

¹⁷ For yet another take on the thin concept of truth, see Lewis 2001.

¹⁸ "But wait, if the T-schema holds for Tr , and if LEM holds, then for any sentence S whatsoever, $Tr(S) \vee Tr(\neg S)$ will also hold. But this entails that bivalence holds for any sentence." Not quite, for once we distinguish between a strong and a weak truth predicate, the weak truth predicate is not suitable to express bivalence. See McGee and McLaughlin 1994 and De Vidi and Solomon 1999 for details.

notion of truth to do that. We don't even need the notion of truth at all in order to specify what the credence function is supposed to be: it is simply the credence function such that $C(A) = 1$ just in case A , and $C(A) = 0$ just in case $\neg A$. Doesn't this show that our rejection of the truth norm was unduly metaphysical?

It seems to me that the answers are “yes” and “no”, respectively: in defending the chance norm(s), we must reject the truth norm, and our rejection of the truth norm is not unduly metaphysical (it *is* metaphysical, in a sense, but not because we've understood the truth norm in unduly metaphysical terms).

Let's turn to the first question. What Pettigrew (2012) shows is that Joyce's theorem can be used to vindicate different norms, according to which credence function we deem vindicated at any given world: if we deem that the vindicated credence function is the “truth” credence function, i.e., the function C such that $C(A) = 1$ just in case A , and $C(A) = 0$ just in case $\neg A$, then we have a vindication of the norm of Probabilism, which says that we ought to have a credence function that is a probability function (i.e., which obeys no constraints beyond the Bayesian ones); if we deem that the vindicated credence function is the one that tracks chance in a certain way (and we've seen that there are two such functions, $urch_w(\bullet)$ and $urch_w(\bullet|CH_{urchw})$), then we have a vindication of the norm of chance, i.e., the norm that says that we ought to have a credence function that tracks objective chance (*modulo* the correct story concerning how we epistemically rely on chance). Given that we've assumed that chance is a probability function, the norm of chance says that we ought to have a credence function that is a probability function *that tracks chance in a certain way*. Thus, the norm of chance requires us to have a credence function that obeys a substantive constraint *beyond* the constraints of the probability calculus.

Now, the norm of probabilism and the norm of chance are in competition in the following sense: while it is true that, in complying with the norm of chance, we also comply with the norm of probabilism, in complying with the norm of probabilism we may *violate* the norm of chance. So, if we want to vindicate the chance norm, then we must reject the norm of probabilism as a *complete* characterization of what a reasonable credence function is. The following is a striking example (due to Hall (2004, pp. 106–107)). Let A express a proposition to which one would have to be insane in order to assign a low probability, like the proposition that one's own hand is not going to turn into a rubber chicken within the next five minutes. Let E be

a proposition representing a highly rich course of experience (such as your total experience, or mine). Now assume that E is your total evidence up to now. The norm of probabilism makes it permissible for you to have a credence function C such that $C(A|E)$ has arbitrarily low value, since such a credence function need not violate any axiom of the probability calculus. That is, the constraints placed by the truth norm on rational credence functions make it permissible (hence rational) for you to be almost certain that your hand *will* turn into a rubber chicken within the next five minutes. But this is not a permissible credence function according to the norm of chance. So we may comply with the norm of probabilism while violating the norm of chance in really odd ways. If we deem $truth_w(\bullet)$ to be the credence function vindicated at w , then such violations of the norm of chance are permissible. Hence, if we want to vindicate the norm of chance, we have to reject the truth credence function as the credence function vindicated at a world. Of course, if we comply with the norm of chance, we also comply with the norm of probabilism (and in this sense they are not in competition), but mere compliance with the norm of probabilism is not enough to secure compliance with the norm of chance (and in this sense, they are).

As for the second question, indeed we can avoid all truth talk in specifying the truth credence function, for the truth predicate was being employed in a purely expressive sense. Thus, we may indeed describe the truth credence function $truth_w(\bullet)$ as the function such that $C(A) = 1(0)$ just in case $A(\neg A)$. However, this doesn't really help in defending the truth norm, for we can restate the criticism using this way of specifying $truth_w(\bullet)$: at any time t before the history of w gets closed, there will be no fact of the matter as to which putative "truth" credence function for w is *the* truth credence function for w —i.e., the credence function that assigns credence 1(0) to A just in case $A(\neg A)$ — simply because relative to t , for any proposition A that is chancy at t , there will be no fact of the matter as to how the world will go A -wise (if we assume metaphysical indeterminism, as we are assuming the non-reductionist will).

Now, this may still strike us as unduly metaphysical, for we've assumed a heavy-weight, metaphysical notion of truth in order to carry on the criticism. As we've pointed out in this section, failures of bivalence make sense only in the presence of a strong, non-disquotational notion of truth. Since our criticism to the truth norm relied on such a notion of truth, it could be seen as unduly metaphysical. However, this doesn't strike me as correct. The fact that we may state the truth norm without the notion of truth (or by deploying only a thin,

disquotational notion of truth) doesn't mean that we cannot make sense of failures of bivalence for chancy propositions in terms of a stronger, non-disquotational notion of truth, under a metaphysical understanding of indeterminism about the future. And it is failure of bivalence what makes it possible to reject $truth_w(\bullet)$ as the vindicated credence function, not an unduly metaphysical understanding of the truth norm itself. Thus, the objection to the truth norm is indeed metaphysical, but there is no illicit inflationary or metaphysically-laden understanding of the truth norm. What would have to be argued in order to block the objection to the truth norm is that the non-reductionist about chance cannot make sense of failures of bivalence, so that she cannot reject the truth norm on those grounds. But it seems that she *can* make sense of such failures, for a strong, non-disquotational notion of truth is within her reach.

6. Summary

In the preceding sections, I essayed a defense of non-reductionism concerning chance by arguing that, contrary to what has usually been assumed, it is the non-reductionist who has a dialectical advantage over the reductionist. The strategy was two-fold: I identified the dialectical challenges faced by each camp and I argued that non-reductionists fare better in answering those challenges.

The challenges were the *motivation problem* and the *explanation problem*. The motivation problem belongs exclusively to the reductionist side of the debate: after realizing that one of the most natural ways of spelling out the epistemic role of chance, Lewis' Principal Principle, is incompatible with the metaphysical assumptions of reductionism pertaining chance, the reductionist side came to face the problem of finding and motivating a credence-chance principle that is both consistent with its metaphysical assumptions and that is conceptually well-motivated. Following Briggs (2009), I argued that the reductionists do not fare well in answering this problem.

The explanation problem, on the other hand, has been thought by some (e.g., Lewis 1994) to be a problem for the non-reductionist exclusively, for it is the problem of explaining the rationality of letting chance guide credence in the way spelled out by the favored credence-chance norm —and it has been assumed that knowledge of the chance-makers posited by reductionist accounts of chance can somehow explain the rationality of letting chance guide credence. However, as Hall (2004) argues, it is a problem for the reductionist

as well, for it is not clear how the chance-makers that the reductionist posits can actually constrain credence in the required way.

After concluding, with Hall, that both reductionists and non-reductionists face the explanation problem, I sought for a possible way out by considering how to fill the dots in an explanation of the following form: *it is rational for an epistemic agent to set her credence to chance in the way recommended by such-and-such credence-chance principle because...* The key ingredient was provided by a generalization of Joyce's (1998) nonpragmatic vindication of probabilism due to Pettigrew (2012). Informally, Pettigrew's strategy allows for the following completion: *if the agent did otherwise, she would be in a worse epistemic situation than the one she would be in, if she abode by that principle*. Less informally, Pettigrew shows that, if we identify the credence function vindicated at a world w with the chance function for that world (or something close to it), we have a formal justification of what we could call the norm of chance, i.e., the norm that says that we ought to have a credence function that tracks chance in the way specified by the favored credence-chance norm.

Now, the key assumption in Pettigrew's strategy is that we can make the identification of the credence function vindicated at a world with the chance function for that world. However, another plausible candidate is the truth credence function for that same world. How may we justify the assumption that sets the formal justification of credence-chance norms in motion?

Time structure came to the rescue of the non-reductionist: if the non-reductionist endorses an indeterministic view of time (which she can do, and is congenial to her overall view of the openness of the future, chancy events), then she can substantiate the claim that the credence function vindicated at a world is the credence function that tracks chance in the way required by the credence-chance norms she favors (provided she doesn't accept undermining futures, all credence-chance norms boil down to the same epistemic recommendation). More precisely, I argued that the truth credence function for an indeterministic world w cannot be the credence function vindicated at w for the simple reason that there is no fact of the matter as to which of the candidates for being the truth credence function is the right one —reality is indeterminate with respect to truth, so chance is the only thing we get.

Time structure, on the other hand, gets in the way of the reductionist: given that the reductionist has to posit a deterministic time structure in order to provide a supervenience base for chance, truth is

well defined for the reductionist. So, it seems that the reductionist has to stay with truth as the vindicated credence function. That means that all she can formally justify is the norm of probabilism (i.e., the injunction to have a credence function that is a probability function), but not the norm of chance (i.e., the injunction to have a credence function that tracks chance in the epistemically proper way). So, unless she can provide a different, properly epistemic consideration for vindicating chance instead of truth, the dialectical advantage seems to reside with the non-reductionist side.¹⁹

REFERENCES

- Arntzenius, F., and N. Hall, 2003, “On What We Know about Chance”, *British Journal for the Philosophy of Science*, vol. 54, no. 2, pp. 171–179.
- Belnap, N., M. Perloff and M. Xu, 2001, *Facing the Future*, Oxford University Press, Oxford.
- Briggs, R., 2009, “The Anatomy of the Big Bad Bug”, *Noûs*, vol. 43, no. 3, pp. 428–449.
- De Vidi, D., and G. Solomon, 1999, “On Confusions about Bivalence and Excluded Middle”, *Dialogue*, vol. 38, no. 4, pp. 785–799.
- Field, H., 1994, “Disquotational Truth and Factually Defective Discourse”, *The Philosophical Review*, vol. 103, no. 3, pp. 405–452.
- Fine, K., 1975, “Vagueness, Truth and Logic”, *Synthese*, vol. 30, no. 3–4, pp. 265–300.
- Hall, N., 2004, “Two Mistakes about Credence and Chance”, *Australasian Journal of Philosophy*, vol. 82, no. 1, pp. 93–111.
- , 1994, “Correcting the Guide to Objective Chance”, *Mind*, vol. 103, no. 412, pp. 505–518.
- Ismael, J., 2008, “**Raid!** Dissolving the Big Bad Bug”, *Noûs*, vol. 42, no. 2, pp. 292–307.
- Joyce, J., 1998, “A Nonpragmatic Vindication of Probabilism”, *Philosophy of Science*, vol. 54, no. 4, pp. 575–603.
- Lewis, D., 2001, “Forget about the ‘Correspondence Theory of Truth’”, *Analysis*, vol. 61, no. 4, pp. 280–289.

¹⁹Special thanks are due to Rachael Briggs for her insightful and useful comments. This paper was written as a result of my stay at the NYU Philosophy Department during the 2011 Spring Semester, made possible by the Balzan Fellowship for Graduate Students I was awarded. I gratefully acknowledge the support of the Fondazione Internazionale Premio Balzan, and of the winner of the Balzan Prize 2008, Thomas Nagel. This paper was completed with the support of the CONICET, in the form of two Doctoral Research Scholarships granted for the period 2009–2012 (Beca Interna de Postgrado Tipo I) and for the period 2012–2014 (Beca Interna de Postgrado Tipo II).

- Lewis, D., 1994, "Humean Supervenience Debugged", *Mind*, vol. 103, no. 412, pp. 473–489.
- , 1986a, *On the Plurality of Worlds*, Blackwell, Oxford,
- , 1986b, *Philosophical Papers*, vol. II, Oxford University Press, Oxford.
- , 1980, "A Subjectivist's Guide to Objective Chance", in R. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, vol. II, University of California Press, Berkeley, pp. 263–293.
- MacFarlane, J., 2003, "Future Contingents and Relative Truth", *The Philosophical Quarterly*, vol. 53, no. 212, pp. 321–336.
- McGee, V., and B. McLaughlin, 1994, "Distinctions without a Difference", *The Southern Journal of Philosophy*, vol. 33, Supplement, pp. 203–251.
- Pettigrew, R., 2012, "Accuracy, Chance, and the Principal Principle", *The Philosophical Review*, vol. 121, no. 2, pp. 241–275.
- Thau, M., 1994, "Undermining and Admissibility", *Mind*, vol. 103, no. 412, pp. 491–503.
- Thomason, R., 1970, "Indeterministic Time and Truth-Value Gaps", *Theoria*, vol. 36, no. 33, pp. 264–281.

Received: December 18, 2014; revised: July 8, 2015; accepted: September 10, 2015.