Chances, Credences and Counterfactuals

Richard Bradley
LSE

April 24, 2016

Abstract

1 Introduction

This paper examines the relation between three concepts: rational degrees of belief (or credences), counterfactuals and chances. All three notions are hotly debated in philosophy and I will necessarily have to take lot for granted if any progress is to be made on the question of how they are related. In particular I will assume that the degrees of belief of a rational agent can be represented by a probability measure on a Boolean algebra of propositions, sufficiently rich as to contain both counterfactual propositions and propositions about chances. Then my task becomes the somewhat more modest one of saying how degrees of belief in counterfactual propositions and in chance propositions are related to each other and to degrees of belief under suppositions. My hope is that doing so will shed light on all three concepts by establishing the constraints that any interpretation of them should satisfy.

The focus of this investigation will be a particular claim about how counterfactuals and chances are related that I will call Skyrms' Thesis in honour of its original proponent, Brian Skyrms (in Skyrms (1980) and Skyrms (1981)). Skyrms' Thesis says, very roughly, that rational degrees of belief in a counterfactuals go by the expected conditional chances of their consequents, given the truth of their antecedents. For instance, suppose that I have in my hand a coin that might be fair or two-headed or two-tailed and that you believe each possibility to be equally likely. Then, according to Skyrms' Thesis, the degree to which you should believe that if I were to toss the coin it would land heads is given by the expected conditional chance of it landing heads given that it is tossed. Its conditional chance of landing heads equals 0.5 or 1 or 0, according to whether it is fair or two-headed or two-tailed. So its expected chance of landing heads is $(\frac{1}{2} \cdot \frac{1}{3}) + (1 \cdot \frac{1}{3}) + (0 \cdot \frac{1}{3}) = 0.5$. Hence it is probable to degree one-half that if the coin were tossed it would land heads.

Skyrms' Thesis is a special case of what has come to be called the Ramsey Test hypothesis, a principle that relates credence in conditionals (counterfactual or otherwise) to credence under a supposition. I will discuss the Ramsey Test hypothesis and the notion of a supposition in more detail in section 4. A second claim—that I will term the Principal Suppositional Principle—relating credence under suppositions to conditional chances, is required to derived Skyrms' Thesis. It is examined in section 5 along with one further claim about the independence of credences in chances from suppositions of a particular kind. To get to the Principal Suppositional Principle some background is required. To this end, in section 3, I recall Lewis' treatment of the relation between chances and credences, suggesting a formulation of the relationship between them that allows for elimination of the problematic notion of admissibility, so central to Lewis' theory. This proposal depends however on understanding chances in a particular kind of way, and motivating this interpretation will be my first task.

Some terminology that will be used throughout. I will take the degrees of belief of a rational agent to be given by a probability function P defined on a Boolean algebra of propositions. Propositions will be denoted by italicised capitals and the conjunction of any two propositions X and Y by either their concatenation XY or by the pair (X, Y). The operations of negation and disjunction will be denoted by $\neg$ and $\vee$ respectively and the logical contradiction and tautology by $\perp$ and $\top$ respectively. The counterfactual 'if A were the case then X would be' will be denoted by $A \rightarrow X$. Similarly the proposition that the chance that X is true (at certain point in time t) is equal to x will be denoted by $Ch_t(X) = x$ and the proposition that the conditional chance that X, given that A, at t, is equal to x will be denoted by $Ch_t(X|A) = x$ .For simplicity the time index will often be dropped.

2 Chance as Ideal Probability

We face uncertainty about a good many things. In a sense all uncertainty stems from lack of information. But there is clearly a big difference between the uncertainty I might have about the time that the bus from Tel-Aviv to Jerusalem departs in the morning, which derives from a simple failure to consult the timetable, and uncertainty that is structural or irreducible in some way. In indeterministic systems such as those described by quantum mechanics, uncertainty is deeply structural because there is simply no information to be obtained that will settle the question of where particles are located (prior to their measurement). In other cases such information may exist in principle but in practice is impossible to obtain. Our uncertainty about the rainfall in, say, Kinshasa on the 1st of January 2050, is hardly less severe for all our knowledge of the deterministic meteorological system governing the weather in Zaire, for this system is chaotic and accurate predictions about distant events in such systems is impossible. So too even for homely, deterministic and non-chaotic systems like those governing coin tosses or the development of cancers or political insurrections. The uncertainty we face regarding such events deserve the label 'structural' because they reflect not the idiosyncratic state of knowledge of a particular individual but physical constraints on all of us on the accessibility of certain kinds of information. With regard to such uncertainty we are all in the same boat.

These observations give strong support to Lewis' widely-shared view that there are two sorts of probabilities—the subjective degrees of belief of a Bayesian agent and the objective chances of events—and that an understanding of uncertainty and how to manage it, rests on the relationship between them. There is much about this view that I think is correct. In particular, I will defend a version of what Lewis called the Principal Principle relating objective and subjective probability: informally, that a rational agent should set her degree of belief in any proposition to what she expects its objective chance to be. But contrary to the mainstream interpretation of Lewis' view, I will argue that the objectivity of chances does not stem from the fact that chances are physical properties of the world. Chances are probabilistic judgements whose objectivity resides in their being expert or 'best-possible' judgements given the physical facts. They are thus neither simply frequencies, nor the propensities that purportedly explain them. Rather they are the judgements of an ideal reasoner who is fully informed of all the propensity and/or frequency facts, but not of the truth of the events that are the bearers of chances.

The idea of an expert probability for an agent goes back to Haim Gaifman (1988), who characterised it as a probability assignment the agent was committed to tracking in the sense of taking as a constraint on her attitudes, the principle that her degrees of belief in some proposition X should equal the expert's probability for X. In a sense, Truth is an expert of this kind, requiring an agent who aims at the truth to obey the principle that:

$$P(X|X \text{ is true}) = 1$$

Similarly, as Hall (2004) and Joyce (2007) suggest, the Principal Principle can be read as saying that Chance is an expert probability, requiring those who seek objectivity to align their credences with the objective chances in the sense of satisfying for all events X:

$$P(X|Ch(X) = x) = x$$

when Ch is a probability measure of the true chances of events.

This proposal is, on the face of it neutral about what chances are, and indeed different authors have filled it in different ways. Hacking (1965), for instance, suggested that relative frequencies were expert probabilities and hence that rational degree of belief in any repeatable event X should equal its relative frequency in the appropriate reference class of events. More often though, Lewis's principle has been read through the prism of a propensity interpretation of chances. Neither, it seems to me, offer the possibility of a sufficiently general interpretation because of some well-known limitations of each.

The main limitation of frequentism is that it does not allow for single case chances. Yet there are many non-repeatable events for which talk of objective probability of its occurrence seems perfectly sensible. In debate about the effects of climate change, for instance, there is much discussion of the chances of human extinction and other catastrophic events. Is all such talk purely subjective?

The propensity interpretation does not suffer from this limitation; indeed several variants of it are explicitly designed to deal with single-case chances. The main problem here is that propensities are not probabilities at all in the strict sense (as first observed by Paul Humphreys (1985)). A propensity is a disposition of a set-up to produce certain kinds of outcomes: of spins of roulettes wheels to cause to ball to land on even numbers, of weather systems to produce snow, of levels of sugar consumption to result in diabetes. Such talk is causal in nature and, indeed, the propensity interpretation provides a natural home for accounts of probabilistic causation. But causation by its nature is typically one-directional, while probabilities are always two-directional. The chance of a window shattering if a stone is thrown at it might sensibly be viewed as a physical propensity of a set-up involving flying stones and windows, but the chance of the stone being thrown, given that the window shattered, cannot. But the two chances are equally well-defined.

To elaborate, consider a simple example in which a fair coin will be tossed five times. What is meant by fair depends, of course, on the interpretation that is given to chances. According to the propensity theorist, the coin will have a certain disposition to land heads whose magnitude will depend on the set-up: the properties of the coin, the manner in which it is tossed and various environmental factors. A coin is fair therefore when the set-up is such as to ensure that the coin is equally disposed to land heads as to land tails. The frequency theorist on the other hand will say that the coin is fair because on half the tosses in the relevant reference class of tosses of this coin it lands heads, and on half it lands tails.

Let's consult our intuitions on some basic cases. What is the chance of the coin landing heads on the first toss? One-half is the only reasonable answer in view of the fairness of the coin. And the chance of it landing heads on the last (fifth) toss, given that it has landed heads on the first four tosses? Again the answer is one-half, absent any grounds for thinking that the tossing of the coin has undermined the conditions for its fairness. Finally what is the chance of it landing heads on the last toss given that it has landed heads on the first four tosses and that it will land heads in only four out of the five tosses? The answer it seems to me, must be zero. For we cannot accommodate this information about the proportion of heads landings without drawing this conclusion.

Let H_i be event of the ith coin toss landing heads and Ch be a chance function on the Boolean algebra based on the events $\{H_1, ..., H_5\}$. Let E be the event of four out of five tosses landing heads, i.e. $E = H_1H_2H_3H_4\neg H_5 \vee ... \vee \neg H_1H_2H_3H_4H_5$. Now the description of the set-up plus our answers to the three questions constrain Ch as follows:

1. $Ch(H_1) = 0.5$
2. $Ch(H_1H_2H_3H_4) = (\frac{1}{2})^4 = 0.0625$
3. $Ch(H_5|H_1H_2H_3H_4) = 0.5$

4. $Ch(H_5|H_1H_2H_3H_4 \wedge E) = 0$

Now the propensity view cannot give an interpretation of the (fragment of) a chance function I have just implicitly defined on the algebra of coin landing events. For the fact that the coin has landed heads on the first four tosses and that it will land heads in only four out of the five tosses does not give us much grounds for thinking that the coin is not fair. One should expect that frequencies in small classes of events will diverge from propensities and so such divergences are slim evidence for a change in the dispositional facts. So on the propensity view, the chance of a fair coin landing heads on the fifth toss is still one half, no matter what the frequency facts are. The root problem here, it seems to me, is that the conditional probability of the last toss landing heads, conditional on four out of five tosses landing heads is not really a propensity at all. It is a judgement that is made in the light of knowledge of the relative frequencies, knowledge that in this case overrides anything that we know about the physical propensities.

Finite frequentism also cannot give an interpretation of these chances, though for somewhat different reasons. From the frequentist's point of view, it makes no sense to talk of the chance of a fair coin landing heads given that the frequency of heads landing is greater than one-half. If the coin has landed heads four out of five times then it is not a fair coin (by the frequentist definition of fair). So on the interpretation of chances as frequencies, constraint 3 is meaningless. It does not matter whether the relevant frequencies come from a finite reference class or an infinite one. Whatever the relevant reference class, the frequency of heads must be one-half if the coin is fair. But in the reference class picked out by the condition that $H_1H_2H_3H_4$ the frequency of heads is not one-half.

The deeper problem here is that the fact that frequentism doesn't attribute chances to events such as a particular coin landing heads independently of a class to which that event belongs. Strictly there is no chance of a particular coin landing heads, just a chance of that coin, qua member of a particular reference class, landing heads. To put it somewhat differently, the relevant empirical facts concern classes of events, not particular ones. This of course means that frequentists cannot allow for single-case chances, as we noted before. But even on their home ground, where events are repeatable, it would be less confusing to say that the chances of particular events are *inferred from* the frequencies than to say that they *are* frequencies. That coins like to one to be tossed land heads in the relevant reference class with a frequency of one-half grounds the judgement that its chance of landing heads is one-half. But then chances are the judgements mandated by the frequency facts, they are not themselves frequencies.

In summary, I do not contest that both frequencies and propensities constitute physical facts which in many circumstances determine the chances. Nor that they indirectly constrain the others. But neither directly determines all the chances that we are interested in, so neither can deliver a comprehensive interpretation of chance. In contrast, by taking chances to be expert probabilities we can make sense both of the role that these physical facts play in determining chances and account for the chances that are determined by inferences from these facts.

3 Chance-Credence Principles

It is widely agreed that chances and credences should be related by some version of what Lewis (1980) calls the Principal Principle. Roughly the principle says that a rational agent should set her degrees of belief in accordance with their expected chances or, equivalently, that the probability of any proposition conditional on the truth of some chance hypothesis is just whatever chances that hypothesis accords it. More formally, let Ω be a Boolean algebra of propositions and Z be a subset of it consisting of the propositions governed by chance. Then:

Principal Principle: Let $X \in Z$ and $x \in [0,1]$. Let P be any reasonable credence function, t be any time and E be any proposition consistent with X that it admissible at t. Then:

$$P(X|Ch_t(X) = x, E) = x$$

There are two notions in Lewis' principle that need filling out: that of a reasonable credence function and that of an admissible proposition. A reasonable credence function, says Lewis, is a regular probability function (i.e. assigning 0 only to logically false propositions) such that taking it as one's initial degrees of belief and learning by conditionalisation on one's total evidence leads only to beliefs that are reasonable given one's evidence. This is not a very helpful characterisation at all since he says nothing about what beliefs are reasonable to acquire from experience. Others have suggested something rather different: that P should not encode any inadmissible information about the chances.

An admissible proposition is one that doesn't include information that pertains to the truth of X except through its chance. As Lewis (1980, p. 92) puts it:

> 'Admissible propositions are the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes.'

The usefulness of the Principal Principle therefore depends on how much information is admissible. It is standardly assumed that, at least, historical information prior to t and information about how possible histories and possible laws bear on chances is admissible. But information relevant to the truth of X that is not encoded in the chances is not. An obvious example is the information that X itself is true, for $P(X|Ch_t(X) = x, X) = 1$ if $P(Ch_t(X) = x, X) \neq 0$. So either the truth of X implies that its chance of truth is one at any time t, or rational credence does not, in this case, track the chances. More generally, information about future chances are also inadmissible, as is any information about the truth value of X not screened off by the chances. As

Nissan-Rozen (2013) has shown this may include many types of information yielded by experience that do not intuitively seem inadmissible by the lights of Lewis' quoted explication of this notion.

It turned out that on Lewis own best-systems account of chances, propositions about the chances themselves are inadmissible, which led him and others to suggest a somewhat different relationship between chances and credences. We will not pursue this debate, however, as a slightly different principle, first proposed by Ned Hall Hall (2004), better captures that relationship we are after and shares with the New Principle the advantage of dispensing with the need to talk of admissibility. It is the following:

Belief-Chance Principle: Let $X, A \in Z$, $x \in [0,1]$ and t be any time. If $P(A, Ch_t(X|A) = x) \neq 0$, then:

$$P(X|A, Ch_t(X|A) = x) = x$$

The Belief-Chance Principle says that the degree of belief a rational agent should have in a proposition X, given that A is true and given some hypothesis as to the conditional chance of X given that A, should just equal whatever the hypothesis says that the conditional chance is. Once again the idea is that Chance functions as a probability expert. This being so, you should set your conditional degrees of belief for X, given that A, in accordance with what Chance says they should be.

Let us see why the notion of admissibility can now be dropped. Consider our previous example. The coin has an equal propensity to lands heads and tails, a fact which grounds an expert judgement of probability one-half for the coin landing heads on the final toss, absent any further information. But, given the information that the coin has landed heads in four out of five tosses, the expert must judge that the coin has landed tails on the final toss, even though its propensity to land tails is no greater than its propensity to land heads. What makes the information inadmissible on Lewis' account is that it determines judgement without going via the physical chances. But this matters not at all for the Belief-Chance Principle which relates subjective degrees of belief to expert judgement.

3.1 The Semantics of Chance

We have been speaking rather informally both about the chances of truth of propositions and about credences in such chances. It is time to make such talk more precise and in particular to specify the content of propositions about chances. In possible worlds semantics, which I take as my starting point, propositions are modelled as sets of possible worlds; intuitively the set of worlds at which the proposition is true. In similar fashion I propose to model chance propositions as sets of probability functions; intuitively the set of probability functions that make the chance proposition true. For instance the proposition that the chance of rain tomorrow is greater than 0.4 is just the set of probability functions assigning probability greater than 0.4 to rain.

To capture this more formally, let $\mathcal{Z} = \{Z, \models\}$ be a Boolean algebra of factual propositions; intuitively Z contains those propositions to which it is meaningful to ascribe chances. Let $\Pi = \{ch\}$ be the set of all probability functions on $\mathcal{Z}$ and let $\Delta = \wp(\Pi)$ be the set of all subsets of Π. The elements of Δ serve here as the chance propositions. In particular, for any $X \in Z$, and $x \in [0, 1]$, $Ch(X) = x$ is the proposition defined by $\{ch \in \Delta : ch(X) = x\}$. A maximally specific chance hypothesis is simply the conjunction of a consistent and exhaustive set of propositions regarding the chances (at some point in time) of the factual prospects. In particular, Ch will denote the hypothesis that says that the chances are as given by the probability function ch, and Ch_A the hypothesis that the conditional chances, given that A, are as given by ch. Note that while Ch is the singleton set $\{ch\}$, the proposition Ch_A is not, since different chance functions can agree on the conditional chances, given A.

The focus of our interest is the product set $Z \times \Delta$ whose elements are combinations of factual and chance propositions. For instance the proposition $(Y, Ch(X) = x)$ is the element of this set that is true when it is both the case that Y and that the chance of X is x. Hereafter, for simplicity, I will write Y for (Y, Δ) and Ch for (Z, Ch). The question we have been pursuing is: what attitudes are agents rationally permitted to hold with respect to these prospects? Lewis' Principal Principle is one such an answer to this question, correct subject to specification of the notion of admissibility. Another answer, the one I am defending, takes the form of the Belief-Chance Principle, which we can now state more exactly as:

Belief-Chance Principle: For all $X, A \in Z$ and any $Ch_A \in \Delta$, if $P(X) \in (0, 1)$ and $ch \in Ch_A$ then:

$$P(X|A, Ch_A) = ch(X|A)$$

So formulated, the Belief-Chance Principle says that the degree of belief a rational agent should have in a factual prospect X, conditional on any factual prospect A and corresponding maximally specific conditional chance hypothesis Ch_A, equals the conditional chance of X, given that A, according to that hypothesis.

4 Conditionals and Suppositions

The Ramsey Test hypothesis is the name given to a thesis that has figured prominently in contemporary debate in both the semantics and pragmatics of conditionals, much of it fuelled by widespread dissatisfaction with the material conditional as a rendition of the semantic content of ordinary language conditionals.[1] In its probabilistic version it asserts that the probability that if α

[1] The literature on both probabilistic and non-probabilistic version of the Ramsey Test hypothesis is now very large. See for instance Gärdenfors Gärdenfors (1988) and Edgington Edgington (1995).

then β equals the probability that β on the supposition that α is true. More formally, let * be a function on pairs of probability functions and propositions, that maps any probability P and proposition α to a suppositional probability P^*_α, where a suppositional probability is real-valued function on a Boolean algebra of propositions satisfying, for any proposition α:

P^*1 (*Probability*) P^*_α is a probability function

P^*2 (*Certainty*) $P^*_\alpha(\alpha) = P^*_\alpha(\top)$

P^*3 (*Anchoring*) $P^*_\top(\alpha) = P(\alpha)$

Then the hypothesis can be rendered more formally by:

(*Probabilistic Ramsey Test*) Rational degrees of belief in conditionals equal degrees of belief in their consequent on the supposition of their antecedent; i.e.:

$$P(\alpha \to \beta) = P^*_\alpha(\beta)$$

In general, when we suppose that α is true, we form a set of (suppositional) beliefs that includes the belief that α and diverges as little as possible from our actual beliefs. But there are many different standards for minimal divergence. We might suppose that as a matter of fact α is true, such as when I suppose, to help with my financial planning, that I won't have enough money at the end of the month to pay the rent. Suppositions of this kind should respect to as great a degree as possible current unconditional beliefs: I should not, for instance, adopt the belief that I will secure a large inheritance to cover the rent. Things are quite different when we suppose or imagine that, contrary to the facts, α is true. A supposition of this kind may well be best accommodated by giving up some of one's beliefs not contradicted by α, to allow retention of well-entrenched ideas about the way that the world works. For example, when supposing that it rained yesterday, in order to think about what I would have done had this been the case, I might have to give up my belief that I went for a walk in the mountains that day, even if I did in fact do so (and have sore feet to prove it).

A closely related distinction—between evidential and interventional suppositions—has played an important role in the development of causal decision theory. When you make a supposition as part of evidential reasoning, you reason as if you have received evidence that implies the truth of the supposition. In contrast when you suppose something interventionally you imagine that there has been some intervention in the course of events which makes the supposition true. In this latter case, unlike the former, you do not revise your degrees of belief in any of the causes of the condition supposed true because you do not treat your supposition as positive evidence for them.

Consider, for instance, the situation modelled by the causal graph in Figure 1, in which the arrows represent relations of causal influence between variables.[2]

[2] The use of graph to model causal relations is now well entrenched in many accounts of causal inference. See, for instance, Spirtes *et al.* (2000) and Pearl (2009).

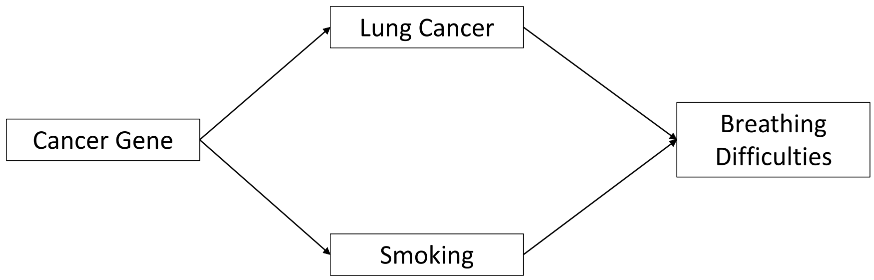

Figure 1: Causal Graph for Breathing Difficulties

According to the graph the presence or otherwise of a certain gene is a probabilistic cause of both lung cancer and of smoking, while both of these are causes of breathing difficulties. These causal relations will induce probabilistic correlations between the variables relevant to evidential reasoning. In particular, evidence that the agent will in fact smoke makes its more probable that they have the gene in question, which in turn makes it more likely that they will get lung cancer. On the other hand, if we suppose that there is an intervention from outside of the causal system represented by the graph (for instance, a freely made choice) which makes it true that the agent will smoke, then their smoking no longer provides evidence for the presence of the gene. So interventional supposition of the agent smoking should not lead to revision of the degree to which we believe that the agent will get cancer.

How are these two sets of distinctions—factual versus counterfactual and evidential versus interventional—related? Evidential supposition and matter-of-fact supposition are just the same thing I think. But the relation between counterfactual and interventional supposition is less straightforward. What counterfactual and interventional supposition have in common is that they lead to revision of beliefs about the facts in a way which retains entrenched beliefs about causal connections. They do so in slightly different ways however. When we engage in interventional supposition we do not presuppose that the hypothesised prospect is as a matter of fact false. When I think about what would happen if I were to cancel my classes and take the day off, for instance, I need not be sure that I won't do it. On the contrary, the very reason why I engage in the supposition is to help me to decide whether I should do it. In doing so I suspend belief about whether the supposition is true, rather than presume that it is false. So interventional supposition is not a form of counterfactual supposition.

Equally not all cases of counterfactual supposition seem to involve interventional reasoning (though many clearly do). When we make an interventional supposition we don't revise our probabilities for the causes of the things supposed true. In entertaining the supposition that I smoke (by choice), I must not, as we saw, revise my probability for having the gene. On other hand, when I suppose, say contrary to the facts, that I didn't get lung cancer, then I don't need to imagine that someone intervened to prevent it. Consequently, I may revise my degrees of belief for the possible causes of the lung cancer in a way that would be inappropriate for interventional supposition. For example, I might infer that I must not have had the gene, even though having it is a causal antecedent, not consequent, of having lung cancer.

The differences between counterfactual and interventional supposition can be subtle and will not matter to our subsequent discussion. So I will focus on the distinction between evidential, matter-of-fact supposition and interventional (potentially contrary-to-fact) supposition. Now one of the advantages of the Ramsey Test hypothesis is that it allows us to link the fact that there are different kinds of suppositions or ways of supposing something true to the widely accepted distinction between indicative and counterfactual conditionals. Indeed the Ramsey Test is best viewed as a test schema with different types of belief revision associated with different modes of supposition being suitable for testing the credibility of different kinds of conditionals: evidential supposition for indicative conditionals and interventional and/or contrary-to-fact supposition for counterfactual conditionals.

Consider, for example, the following two conditionals concerning Jim, a canny investor who very rarely loses money, and an investment in ostrich farming futures, a currently fashionable financial instrument, the market for which the wise consider to be a bubble.

1. If Jim invests in ostrich futures, then he will make a packet.
2. If Jim were to invest in ostrich futures, he would make a loss.

Although this pair of conditionals make 'opposite' claims, assertion of both is quite reasonable from the point of view of the Ramsey Test schema. The first conditional, being indicative, is evaluated by supposing that Jim will as a matter of fact make the investment. Since he rarely makes a mistake, it is reasonable to conclude that the investment will be profitable. The second conditional, being an interventional counterfactual, is evaluated by supposing, (potentially) contrary to the facts, that Jim makes the investment. Since we expect the bubble to pop, it is reasonable to conclude that he would lose money if he did so.

Most of the literature on the Ramsey Test hypothesis has focused on a version appropriate to indicative conditionals, known as Adams' Thesis and which asserts that the probability of an indicative conditional is the conditional probability of its consequent given its antecedent. Some have argued that Adams' Thesis holds for counterfactual conditionals as well, with appropriate allowance for temporal considerations. But this is clearly refuted in our example. The conditional probability of Jim making money, conditional on him investing in ostrich futures, is high because in almost every credible world in which Jim makes an investment, the investment is profitable. So while Adams' Thesis correctly predicts the high probability of the first conditional, it also incorrectly predicts the low probability of the second conditional.

In contrast, the aforementioned proposal of Brian Skyrms correctly predicts the high credibility of the second conditional. Recall that Skyrms' Thesis says that the probability of a counterfactual conditional equals the expected conditional chance of its antecedent given the truth of its consequent. Formally:

Skyrms' Thesis: $P(A \rightarrow X) = \sum ch(X|A) \cdot P(Ch_A)$

Hence to ascertain how credible it is that Jim would lose money if he were to invest in ostrich futures, we should, according to this thesis, consider various possible hypotheses regarding the chances of losing money given such an investment, and then weight them by how likely we think each such hypothesis is to be true. In this case the relevant chance hypotheses will reflect combinations of causal conditions that determine the success of the ostrich farming industry (demand for ostrich products, costs of animal feed, capital costs and so on). In bubbles of the kind postulated, supply of a good rapidly outstrips demand, leading to losses. So, in expectation, the conditional chance of Jim making money, given an investment, is low.

5 Suppositions and Chances

Skyrm's Thesis seems to accord well with our intuitive judgements as to credibility of counterfactual conditionals. What I want to do now is show how it can be derived from our basic postulates about chances and conditionals (the Chance-Credence principle and the Ramsey Test hypothesis). To do so we must first say something about how chances behave under suppositions. Consider a variant of the previous example in which smoking (S) and lung cancer (L) have a common genetic cause (G), but in which smoking is also a probabilistic cause of lung cancer. Our interest is in the effect of taking up smoking on our prospects for lung cancer and for this reason we need to separate the contributions of genetic factors to lung cancer from those of the smoking itself. Consider a set of mutually exclusive and exhaustive hypotheses $\{CH_S\}$ concerning the conditional chances of relevant prospects given smoking and let *ch* be a representative chance function picked out by hypothesis CH_S. For definiteness let smoking mean 20 a day for 10 years.

Because smoking and lung cancer have a common genetic cause, $P^*_S(L)$, the probability of lung cancer on the interventional supposition of smoking, will not equal $P(L|S)$, the conditional probability of lung cancer given smoking (i.e. the probability of lung cancer on the evidential supposition of smoking). Hence:

$$P^*_S(L|S) = \frac{P^*_S(LS)}{P^*_S(S)} = P^*_S(L) \neq P(L|S)$$

So the conditional probability of lung cancer given smoking in not invariant under the supposition of smoking.

In contrast, I claim, the conditional probability of lung cancer given smoking *and* any hypothesis about the conditional chances given smoking, *is* invariant under the supposition of smoking. This is obviously true when we suppose that we will, as a matter of fact, smoke, since the fact of smoking is already being conditioned on. But it is equally true if the mode of supposition is interventional. For entertaining the possibility of smoking does not alter the conditional chance, given smoking, that the chance hypothesis under consideration confers on lung cancer. On the contrary, the probabilistic causal effect of smoking is already built into this hypothesis. And the latter uniquely determine what degrees of belief for lung cancer it is rational to adopt.

We can put the claim slightly differently. Suppose that your conditional chance of lung cancer, given smoking, is 60%. This fact concerns a relationship that holds between you smoking and you developing lung cancer that holds irrespective of whether you smoke or not. In taking up smoking I have a causal effect on lung cancer (i.e. I affect my chances of lung cancer) *via, but not on,* the conditional chances of lung cancer given smoking. This is precisely what makes these conditional chances a structural feature of our environment, rather than an epistemic one.

The point is a general one. The conditional probability attaching to any proposition, given some condition, that derives from the possible truth of a hypothesis concerning its conditional chances given the condition in question is not affected by the supposition that the condition is true. More formally:

Rigidity of Chances: For all $X, A \in Z$ and any $Ch_A \in \Delta$:

$$P^*_A(X|A, Ch_A) = P(X|A, Ch_A)$$

I take Rigidity of Chances to be a fundamental condition on chances in the following sense. If we are ask ourselves what propositions play the role of chances in an agent's deliberations, and especially deliberations about what to do, then only propositions satisfying this condition can be considered candidates. Other propositions can satisfy this condition—for instance, propositions of the form 'God ordains that X'—but any proposition not satisfying it cannot express a structural constraint on the agent's interventions.

Now from Rigidity of Chances and the Belief-Chance Principle it follows that:

$$\begin{aligned} P^*_A(X|Ch_A) &= P^*_A(X|A, Ch_A) \\ &= P(X|A, Ch_A) \\ &= ch(X|A) \end{aligned}$$

This yields a natural generalisation of the Belief-Chance Principle:

Principal Suppositional Principle: For all $X, A \in Z$ and any $Ch_A \in \Delta$, if $P(X|A) \in (0,1)$ and $ch \in Ch_A$ then:

$$P^*_A(X|Ch_A) = ch(X|A)$$

The Principal Suppositional Principle (PSP) says that the conditional probability of X, given the truth of both A and a hypothesis as to the conditional chances of X, under the supposition that A is or were true, equals the conditional chances of X, given that A, according to the hypothesis. For instance, suppose that the conditional chance of lung cancer, given smoking, equals 60%. Then, given this, your degree of belief in developing lung cancer, on the supposition of smoking should just be 60%.

Now from PSP and the law of total probability, it follows that:

$$\begin{aligned} P^*_A(X) &= \sum P^*_A(X|Ch_A) \cdot P^*_A(Ch_A) \\ &= \sum ch(X|A) \cdot P^*_A(Ch_A) \end{aligned} \tag{1}$$

In other words the subjective probability of X under the interventional supposition that A equals the expected conditional chance of X, given that A, calculated relative to the probabilities of the conditional chance hypotheses induced by the supposition that A. When supposition is evidential in form this just reduces to:

$$P(X|A) = \sum ch(X|A) \cdot P(Ch_A|A) \tag{2}$$

i.e. to the requirement that one's conditional degrees of belief in X given that A should equal the expected conditional chance of X given that A. (For this to hold it is sufficient that $P^*_A(X|A) = P(X|A)$. See Bradley (2012)).

In contrast, to get a fix on the relationship between conditional chance and credence under interventional supposition one further assumption is required. It is that the maximally specific hypotheses regarding the conditional chances given A, are probabilistically independent of A, i.e. that:

Chance Independence: $P^*_A(Ch_A) = P(Ch_A)$

When Chance Independence hold it follows from equation (1) that:

$$P^*_A(X) = \sum ch(X|A) \cdot P(Ch_A) \tag{3}$$

In other words, the probability of X under the interventional supposition that A is just the expected conditional chance of X, given the truth of A.

Finally, we can return to the rational credence in conditionals. Recall that by the Ramsey Test Hypothesis, $P(A \to X) = P^*_A(X)$. Together with equation 2, this gives us a version of Adams' Thesis. Together with equation 3 it gives us Skyrms' Thesis, that the probabilities of counterfactuals equal the expected conditional chances of their consequents given the truth of their antecedents.

The status of Skyrms' Thesis thus depends on that of Chance Independence. The latter is an attractive principle, but I don't think that it holds with complete generality. Failure of the principle would imply that the probabilities of some chance hypotheses are not independent of the supposition that certain kinds of interventions will be performed. This might happen if an intervention is probabilistically correlated with factors that determine the conditional chances given the intervention. Suppose for example that the causal influence on the development of lung cancer due to genetic factors is correlated with smoking, even though smoking is not a cause of these factors. Then the conditional chances of lunch cancer given smoking may not be invariant under the supposition that one smokes. And Skyrms' Thesis will fail.

References

Bradley, Richard. 2012. Multidimensional Possible-world Semantics for Conditionals. *Philosophical Review*, **121**(4), 539–571.

Edgington, Dorothy. 1995. On Conditionals. *Mind*, **104**(414), 113–128.

Gaifman, Haim. 1988. A theory of higher order probabilities. *Pages 191–219 of: Causation, chance and credence*. Springer.

Gärdenfors, P. 1988. *Knowledge in flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge.

Hacking, Ian. 1965. *Logic of statistical inference*. CUP Archive.

Hall, Ned. 2004. Two mistakes about credence and chance. *Australasian Journal of Philosophy*, **82**(1), 93–111.

Humphreys, Paul. 1985. Why propensities cannot be probabilities. *The philosophical review*, 557–570.

Joyce, James M. 2007. Epistemic deference: The case of chance. *Pages 187–206 of: Proceedings of the Aristotelian Society*. JSTOR.

Lewis, David. 1980. A Subjectivist's Guide to Objective Chance. *In:* Jeffrey, Richard C. (ed), *Studies in Inductive Logic and Probability*. University of California Press.

Nissan-Rozen, Ittay. 2013. Jeffrey Conditionalization, the Principal Principle, the Desire as Belief Thesis, and AdamsŠs Thesis. *The British Journal for the Philosophy of Science*, **64**(4), 837–850.

Pearl, Judea. 2009. *Causality*. Cambridge university press.

Skyrms, B. 1980. *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. Yale University Press.

Skyrms, Brian. 1981. The prior propensity account of subjunctive conditionals. *Pages 259–265 of: Ifs*. Springer.

Spirtes, Peter, Glymour, Clark N, & Scheines, Richard. 2000. *Causation, prediction, and search*. Vol. 81. MIT press.