

Divergent Level-Relative Chances and the General Principle

Patryk Dziurosz-Serafinowicz

Abstract

This paper shows how to apply a particular chance-credence principle, called the General Principle, to the case of two divergent level-relative chances of a given proposition in order to single out level-relative chance that should underpin one's credence in that proposition. It is shown that the key move hinges on the use of admissibility clause in the GP. This principle is, then, tested against the case of two divergent viability fitnesses understood as level-relative survival chances. The case is taken to show that there are situations that fall outside the scope of the admissibility clause which is essentially *qualitative*. As a remedy, I suggest that the GP should be endowed with a *quantitative* notion of admissibility. My suggestion exploits a line of thought employed by Brian Skyrms in his work on the notion of the resiliency of chance.

1 Introduction

It is sometimes claimed that a proposition A does not have one unconditional chance to come out true, but instead many conditional chances that could disagree with each other. One way this can happen is that the properties of a chance set-up at its various levels of description confer different chances on A 's coming out true. So, an organism's chance of surviving given its individual phenotype may differ from its chance of surviving given that it is a member of a group with a certain advantageous property. Call such conditional chances *level-relative chances*. Given a hierarchy of levels of properties, there will be higher- and lower-level chances. Whether a chance is higher- or lower-level is a relative matter. To give an example, an organism's chance of surviving given its genotypic property is higher-level with respect to the chance of surviving given its genic property, but it is a lower-level chance with respect to the chance of surviving given its phenotypic property. For simplicity's sake, let us assume that there are only two levels of properties, i.e., the i -level and the j -level. Now, suppose that an epistemic agent knows A 's i - and j -level chances that disagree with each other. The problem arises: which one of these divergent chances should underpin her credence about A ?

Many philosophers have recognized the significance of this problem.¹ Recently, Alan Hájek (2007, p. 579) has argued that the problem looms because 'you can't serve all your masters at once, so you have to play favorites. But who trumps whom, which trumps which?'. In the light of this recognition, to resolve this problem, we must come with a principled way to single out the 'right' condition (the 'right' level-relative property of the chance setup) for a conditional chance of A that should underpin a credence about A . Here, 'right' does not mean the ontologically privileged one, for all the level-relative chances are ontologically equal (there is no sense to treat some of them as less real than other). Rather, the idea is that by singling out the right level-relative property, we determine a level-relative chance of A that is a better guide for one's credence than the other competing level-relative chance of A .

It seems natural to think that some reasonable principle relating one's credence and one's evidence about chance—the so-called *chance-credence* principle—will provide an answer to our

¹See, for example, Price (1984), Hájek (2007), Hoefer (2007), Glynn (2010).

concern. However, most of the well-known chance-credence principles tell us how to deal with unconditional chances. That is, by providing a link between one's credence about chance and one's credence about the behaviour of an experimental set-up, these principles tell you that you should set your credence equal to the unconditional chance, all things being equal. This is true of Miller's Principle (Miller 1966), and the Principal Principle (Lewis 1986). Could there be a principle that tells us how to deal with conditional chances, and particularly with level-relative chances?

One possibility is the so-called General Principle (GP) which, as proved in Vranas (2004), entails Lewis's Principal Principle, Lewis's (1994) and Hall's (1994) New Principle, and van Fraassen's Conditional Principle (van Fraassen 1989, p. 202).² Roughly, it says that an agent with evidence E ought to set her credence about A equal to the conditional chance of A given B , if she knows B and her evidence E is admissible with respect to the proposition that the conditional chance of A is $ch(A|B)$. And E is admissible with respect to that proposition if, roughly speaking, it does not provide information 'overriding' A 's conditional chance (or, it tells us nothing 'over and above' what is told by this chance). Interestingly, we can use the admissibility clause in the GP to decide which one of A 's divergent conditional chances should constrain the agent's credence in A . The key move is as follows. Given knowledge of two divergent level-relative chances, $ch(A|B_i)$ and $ch(A|C_j)$, and of propositions about level-relative properties of a chance set-up, B_i and C_j , the GP is applicable to $ch(A|B_i)$ if C_j is admissible to the proposition about $ch(A|B_i)$. If this is so, then, on Lewis's construal of admissibility which will be explained precisely in this paper, it is precluded that the GP is applicable to $ch(A|C_j)$. Once C_j is admissible relative to the proposition about $ch(A|B_i)$, it cannot be true that B_i is admissible with respect to the proposition about $ch(A|C_j)$. Consequently, we get A 's conditional chance that 'trumps' the other one and so deserves to be called a better guide to one's epistemic life. This solution is possible because the GP is supplemented with the admissibility clause, a specific *ceteris paribus* clause. The thought is that each of the two divergent level-relative chances can underpin one's credence, all else being equal. But all the other things cannot be equal for both these chances. This raises an interesting possibility: we might single out a level-relative chance for which the other things are equal and declare that this chance should underpin one's credence. It seems, then, that by using the admissibility clause, we can in principle decide upon which one of the divergent level-relative chances should guide one's credence.³

My main task in this paper is to examine the extent to which the use of admissibility clause succeeds in picking out a level-relative chance to which the GP is applicable. In particular, I first identify cases of divergent level-relative chances in which such application of the GP appears straightforward. I then go on to argue that scientific reasoning deals with cases that cannot be adequately handled by appealing to the admissibility clause. I illustrate this claim by means of a case of level-relative chances in evolutionary theory. This case involves two different viability fitnesses understood as survival chances of an organism taken from a population consisting of groups, where each group is a mix of altruists and selfish types. That is, in each group an organism has an individual-level fitness in virtue of its having an individual phenotype; and also, *qua* member of a group, it possesses a group-level fitness, which is its fitness in virtue of being in that group. I then provide a reason for why the GP armed with the admissibility clause cannot provide a satisfactory answer to the question of which of these two divergent level-relative chances should underpin an agent's credence about the proposition that some

²In fact, the GP as presented in Vranas (2004) corresponds to Strevens's (1995) chance-credence principle called CP.

³Ideally, one could try to argue that the GP armed with the admissibility clause promises to resolve what Hájek (2007) calls the *epistemological reference class problem*, to wit, the problem of which of many conditional probabilities should guide our credences. This paper, however, will not try to tackle this issue.

organism, randomly selected from the population, will survive. As a diagnosis, I claim that the construal of admissibility clause in the GP is essentially *qualitative* and as such is too narrow to handle cases of this type. As a remedy, I suggest that the admissibility clause in the GP should be understood *quantitatively* rather than qualitatively. My suggestion exploits a line of thought employed by Brian Skyrms in his work on resiliency. A secondary goal in this paper is to motivate the claim that, upon the quantitative reading of admissibility, highly resilient chances can underpin credences.

The structure of this paper is as follows. In section 2, I will show how the GP applies to any pair of level-relative chance functions. In section 3, I will provide, by drawing on the idea of screening off, a precise account of the key notion of admissibility and then will identify cases in which the use of admissibility clause so understood seems to be perfectly adequate. In section 4, I will test the GP and its admissibility clause against a case taken from evolutionary theory and will argue that the solution based on the admissibility clause pales. In section 5, I will discuss two reactions to this problem, and will argue that the GP's inability to deal with this case stems from the qualitative nature of admissibility clause. In section 7, I will sketch, by drawing on Skyrms's idea of the resiliency of chance, a quantitative approach to admissibility. Finally, in section 8, I will provide some reasons for why a highly resilient chance might underpin one's credence despite the lack of maximal admissibility.

2 The GP and Level-Relative Chances

Let \mathcal{A} be a finite algebra of propositions, subsets of some set W . For all propositions $A \in \mathcal{A}$, let $cr(-)$ be an agent's credence function over \mathcal{A} . It assigns to each proposition A a credence, a number in $[0, 1]$, that measures the agent's degree of belief in A . Let $ch(-|B)$ be the conditional chance function over \mathcal{A} , for some fixed conditioning proposition B .⁴ This function assigns to each proposition A a conditional chance, a number in $[0, 1]$. (Nothing substantial in my presentation hinges on whether this conditional chance measures A 's limiting relative frequency to come out true given B , or a chance set-up's propensity to display A given B .) Suppose that, for some fixed $A \in \mathcal{A}$, $\langle ch(A|B) = x \rangle$ is the proposition that A 's chance given B equals x . Further suppose that the agent knows $\langle ch(A|B) = x \rangle$, B and has evidence E . Assume that $cr(E \wedge B \wedge \langle ch(A|B) = x \rangle) > 0$. Then:

(GP) An agent ought to have a credence in A such that

$$cr(A|E \wedge B \wedge \langle ch(A|B) = x \rangle) = x,$$

if E is admissible with respect to $\langle ch(A|B) = x \rangle$.⁵

For example, the agent is about to form a credence about the proposition that a given coin will land heads (A). If one knows that (B) the coin is biased in favour of heads, that the chance of A given B equals x , and that (E) the coin landed heads yesterday, then one's credence in A

⁴If $B \in \mathcal{A}$ and $ch(B) > 0$, the function $ch(-|B)$ may be defined by the ratio of unconditional chances, that is, for any $A \in \mathcal{A}$, $ch(A|B) = \frac{ch(A \wedge B)}{ch(B)}$. If one regards the ratio analysis of conditional chance as flawed, one may treat $ch(-|B)$ as a primitive notion which is not reducible to unconditional chances. For example, one may define $ch(-|B)$ over $\mathcal{A} \times \mathcal{B}$, where \mathcal{B} is a non-empty subset of \mathcal{A} , as satisfying Popper's axioms for conditional probability.

⁵Strictly speaking, we should require the conjunction $E \wedge B$ to be admissible with respect to $\langle ch(A|B) = x \rangle$. However, we can omit the reference to B in the admissibility clause because B is always admissible with respect to $\langle ch(A|B) = x \rangle$. That is, since the conditional chance $ch(A|B)$ must reflect all the information conveyed by B , B cannot tell us anything about A over and above what is told by $ch(A|B)$.

ought to be x , provided that E is admissible with respect to the proposition that A 's chance given B is x .

Now, let us apply the GP to the case in which there are two level-relative chance functions over \mathcal{A} . Let the propositions B_i and C_j stand for, respectively, some i - and j -level property of a chance set-up. Assume that B_i and C_j belong to \mathcal{A} . We denote the i -level chance function over \mathcal{A} given B_i as $ch(-|B_i)$. Similarly, we take $ch(-|C_j)$ to stand for the j -level chance function over \mathcal{A} given C_j . We take it that these two chance functions disagree on the chance assignment over \mathcal{A} . Suppose further that, for some fixed $A \in \mathcal{A}$, $\langle ch(A|B_i) = x \rangle$ is the proposition that A 's chance given B_i is x , $\langle ch(A|C_j) = y \rangle$ is the proposition that A 's chance given C_j is y , and $x \neq y$. Assume that the agent knows propositions about the chance set-up's properties and propositions about the requisite level-relative chances. What she does not know is A 's chance conditional on the conjunction $B_i \wedge C_j$, $ch(A|B_i \wedge C_j)$, though it might be perfectly defined. The question arises: should she defer to $ch(A|B_i)$, or to $ch(A|C_j)$?

The GP, applied to $ch(A|B_i)$, requires the following:

(GP _{$ch(A|B_i)$}) An agent ought to have a credence in A such that

$$cr(A|\langle ch(A|C_j) = y \rangle \wedge C_j \wedge B_i \wedge \langle ch(A|B_i) = x \rangle) = x,$$

if C_j is admissible with respect to $\langle ch(A|B_i) = x \rangle$.⁶

Thus the agent who knows two divergent level-relative chances of A , $ch(A|B_i)$ and $ch(A|C_j)$, ought to set her credence in A equal to the chance $ch(A|B_i)$, if C_j is admissible with respect to the proposition that A 's chance given B_i is x .

When applied to $ch(A|C_j)$, the GP requires the following:

(GP _{$ch(A|C_j)$}) An agent ought to have a credence in A such that

$$cr(A|\langle ch(A|C_j) = y \rangle \wedge B_i \wedge C_j \wedge \langle ch(A|B_i) = x \rangle) = y,$$

if B_i is admissible with respect to $\langle ch(A|C_j) = y \rangle$.

Thus, the agent faced with the same situation ought to set her credence in A equal to y , if B_i is admissible with respect to the proposition that A 's chance given C_j is y .

The crucial point about this case is that once the requisite admissibility relations are determined, these requirements cannot both be satisfied. That is, if the chances $ch(A|B_i)$ and $ch(A|C_j)$ disagree, and C_j is admissible with respect to the proposition about $ch(A|B_i)$, then the GP is applicable to $ch(A|B_i)$ and inapplicable to $ch(A|C_j)$. Likewise, if B_i is admissible with respect to the proposition about $ch(A|C_j)$, then the GP is applicable to $ch(A|C_j)$ but inapplicable to $ch(A|B_i)$. In either case, we get A 's level-relative chance that 'trumps' the other one. And it is this trumping chance that should constrain the agent's credence in A . This holds because the two admissibility relations cannot both be true. So, if C_j is admissible with respect to $\langle ch(A|B_i) = x \rangle$, it is not true that B_i is admissible with respect to $\langle ch(A|C_j) = y \rangle$. In the next section, I give a more precise explanation of why this is so by clarifying Lewis's idea of admissibility.

⁶For simplicity's sake, we speak of the admissibility of C_j instead of the admissibility of the conjunction $\langle ch(A|C_j) = y \rangle \wedge C_j \wedge B_i$. This move is justified as follows. B_i is always admissible to $\langle ch(A|B_i) = x \rangle$, since it cannot tell us anything over and above what is told by $\langle ch(A|B_i) = x \rangle$. Further, it seems reasonable to require that if C_j is admissible, then so must be $\langle ch(A|C_j) = y \rangle$. After all, if C_j tells us nothing over and above what is told by $\langle ch(A|B_i) = x \rangle$, then a fortiori $\langle ch(A|C_j) = y \rangle$ must do the same. For example, if the proposition that Usain Bolt is 100m world record holder is admissible to the proposition about the chance of a coin-flip landing heads, then so is the proposition about the chance of this coin-flip landing heads given that Usain Bolt is 100m world record holder. Finally, if all three propositions are admissible, then their conjunction is also admissible. This corresponds to Lewis's (1986, p. 96) idea that any Boolean combination of admissible propositions is also admissible.

3 Admissibility and Screening-Off

What makes proposition E admissible? David Lewis characterized the notion of admissibility as follows:

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. Once the chances are given outright, conditionally or unconditionally, evidence bearing on them no longer matters. (Once it is settled that the suspect fired the gun, the discovery of his fingerprint on the trigger adds nothing to the case against him.)(Lewis 1986, p. 92)

Lewis's characterization of admissibility does not amount to a definition, but it can be regarded as a vague approximation of it. What is the leading idea behind it? Lewis suggested that two kinds of information seem to fit his characterization: historical information and information about chance itself. Historical information concerns facts about the past history of events up to the point where the experiment in question is about to be performed. Intuitively, information about initial conditions of a coin toss, which is essentially historical, should not override information about the chance for heads. But, as Strevens (2006) points out, in a deterministic world this information would be inadmissible, since it entails an outcome-specifying proposition. If so, then historical information is not always admissible. Also, it seems that information about chance itself is not always admissible. If you endorse the view that a chance distribution at a given time supervenes on the past, present and future distribution of outcomes, then it appears that information about chance itself is inadmissible: it says what the outcome will be before the experiment runs.⁷ It seems then that the intuition behind Lewis's account of admissibility cannot be fully captured by these two sorts of information. Could we come with a better understanding of this intuition? And, more importantly, given such an understanding, could we tell when information about a chance set-up's property is admissible to information about its level-relative chance?

Before we suggest an account of Lewis's idea of admissibility, we have to make some qualifications concerning the notion of admissibility. Firstly, a proposition E is admissible always relative to other propositions (Thau 1994). But to which other propositions? Although it is typically taken that E should be regarded as admissible with respect to outcome-specifying propositions A , I take it, heeding the suggestions in Vranas (2002), that E should be judged admissible relative to propositions about A 's chances. The case of divergent level-relative chances seems to provide a reason for why this should be so. For E may be admissible to $\langle ch(A|B_i) = x \rangle$, but inadmissible to $\langle ch(A|D_j) = y \rangle$. So, by judging E 's admissibility only relative to A one does not take into account possible admissibility relations in this case.

Secondly, E 's admissibility must be relativized to time (Lewis 1986). A proposition that says about a result of a coin toss is inadmissible at any time before the result of this toss is settled. But it is admissible at any time after this toss.⁸ For ease of exposition, whenever I speak about a proposition's admissibility, I mean its admissibility at time before an outcome-specifying proposition is known to be true or false.

The account of Lewis's idea of admissibility, I propose, goes as follows. A proposition E is admissible with respect to the proposition about A 's chance if A 's chance gives a *complete* probabilistic prediction about A given E .⁹ Once A 's chance is determined, the proposition E

⁷For a recent discussion of this issue, see Strevens (1995), Vranas (2004).

⁸One has to be cautious not to generalize this observation: it does not mean that any evidence about the future is inadmissible at the time before the result of a chancy experiment is settled. Suppose that E is about an event which will occur after the result is settled, yet it is not in any way correlated with that result. For this reason, we would not call it inadmissible, though it is about the future.

⁹I should emphasize that this account is not a definition of admissibility. Rather, it provides a necessary condition for Lewis's notion of admissibility. But it might not be sufficient. For we might additionally require

provides no help in predicting probabilistically the truth value of A . This is because it does not yield information that goes beyond what A 's chance says about the truth of A . On this account, if E entails an outcome-specifying proposition A , and A 's chance is non-trivial, then E adds information beyond what A 's chance says. Thus, A 's chance does not yield a complete probabilistic prediction about A given E . Thus, the specific cases of historical information and information about chance discussed above would count as inadmissible. Further, it seems that both relevant and irrelevant information may be admissible in the sense just given: the chance for heads gives complete probabilistic prediction given that Usain Bolt is 100m world record holder (irrelevant information) as well as given that the coin is made of bronze (relevant information). Importantly, on this account of admissibility, it is not required that A 's chance must be predictively complete *simpliciter*. For we might not be certain about all the factors that could possibly affect A . We should keep in mind that chance is a guide to our epistemic life given that we think of ourselves as finite and temporally conditioned beings. If we were certain of all the relevant factors that could influence a chance set-up's outcomes, we probably would not need chances to guide our credences about these outcomes. Consequently, we should judge whether chance is predictively complete relative to those factors that are epistemically possible for us. Moreover, A 's chance may be predictively complete relative to E , but incomplete relative to a different proposition F , where both E and F are epistemically accessible.

In the case of unconditional chance, the suggestion to understand admissibility *qua* chance's predictive completeness can be made precise by using the notion of *statistical irrelevance*.¹⁰ If $\langle ch(A) = x \rangle$ is the proposition that A 's chance is x and E is any evidence, then:

(Admissibility_{uncond}) E is admissible with respect to $\langle ch(A) = x \rangle$ if

$$ch(A|E) = ch(A),$$

provided that the conditional chance is defined. Thus, E is admissible with respect to $\langle ch(A) = x \rangle$ if conditionalizing on E does not change A 's chance. For example, given a die which is about to be rolled, the proposition that the die is yellow is stochastically irrelevant to an outcome A (and so $ch(A)$ should remain unchanged upon conditionalizing on E), but the proposition that A is even seems to be stochastically relevant and so conditioning on this proposition could make a difference to the chance of A .

Let's apply the account of admissibility just given to the case of two divergent level-relative chances. Whether C_j is admissible to $\langle ch(A|B_i) = x \rangle$ depends on whether $ch(A|B_i)$ gives a complete probabilistic prediction about A given C_j . Like in the case of unconditional chance, we can make this more precise by using the notion of *screening-off*, or *conditional independence*. Let us assume that B_i and C_j are stochastically relevant to the truth of A , i.e., $ch(A|B_i) \neq ch(A|C_i)$. Then:

(Admissibility_{cond}) C_j is admissible with respect to $\langle ch(A|B_i) = x \rangle$ if

$$ch(A|B_i \wedge C_j) = ch(A|B_i),$$

provided that the conditional chances are defined. In other words, C_j is admissible with respect to the proposition about A 's chance given B_i if A 's chance conditional on B_i is the same as its chance conditional on B_i and C_j . That is to say, the conditioning proposition C_j is *stochastically irrelevant* to A given the conditioning proposition B_i , or the information it conveys does not

that the agent must be *certain* or *justified* in believing that chance gives a *complete* probabilistic prediction given E .

¹⁰In doing so, I follow Levi (1980, chap. 12) and Strevens (1999).

help to probabilistically predict whether A once B_i has been taken into account.¹¹ We say then that B_i *screens off* C_j from A 's chance.¹² To use Lewis's phrasing, we may say that information conveyed by C_j comes by way of A 's chance given B_i . And this holds because B_i renders A and C_j conditionally stochastically independent.

Now, given the notion of screening-off, we can account for the fact that if C_j is admissible with respect to $\langle ch(A|B_i) = x \rangle$, it is not true that B_i is admissible with respect to $\langle ch(A|C_j) = y \rangle$. This can be explained by the fact that the screening-off relation, which underlies admissibility relations between these propositions, is construed as an asymmetric relation: if B_i renders C_j stochastically irrelevant to A , then it does not hold true that C_j does the same with respect to B_i . For if B_i renders C_j stochastically irrelevant to A , and if C_j renders B_i stochastically irrelevant too, then we have

$$ch(A|B_i) = ch(A|C_j),$$

which could mean the following: to determine A 's conditional chance, it does not matter whether one uses the conditioning proposition B_i or C_j . In any case, one ends up with the same conditional chance of A . But this contradicts our assumption that C_j and B_i are stochastically relevant to the truth of A , i.e., $ch(A|B_i) \neq ch(A|C_j)$.

The asymmetry of screening-off is crucial to adjudicating between the divergent level-relative chances. For if the admissibility relation in $(GP_{ch(A|B_i)})$ is true, it follows that the admissibility relation in $(GP_{ch(A|C_j)})$ must be false. Whether the winner of our competition is chance $ch(A|B_i)$ or chance $ch(A|C_j)$ depends on which one of these two admissibility relations holds true. The crucial question, however, is: Could we use the admissibility clause so understood in any case in which we have to decide which one of the divergent level-relative chances trumps which?

There seems to be cases in which by using the admissibility clause we can pick out a level-relative chance to which the GP is applicable. To identify them, we first prove the truth of the following proposition (proof in the appendix):

Proposition 1. *If $B_i \subseteq C_j$, then $ch(A|B_i \wedge C_j) = ch(A|B_i)$.*

That is, if B_i entails C_j , then it follows that B_i renders A and C_j conditionally stochastically independent. Now given the asymmetry of screening-off, C_j cannot render A and B_i conditionally stochastically independent. If this is the case, then we get, from a pair of divergent level-relative chances, the level-relative chance, $ch(A|B_i)$, that should underpin by the GP one's credence in A . Two examples of this type of case come into mind. First, B_i may describe some maximally fine-grained property displayed by a chance set-up, while C_j describes some more coarse-grained property of this chance setup. Admissibility clause, then, prescribes that in such a situation it is A 's chance conditional on the maximally fine-grained property that should guide one's credence in A . Second, B_i may stand for the chance set-up's property that asymmetrically necessitates its other property described by C_j . To put this in other words, the latter property supervenes on the former property, or whenever the chance set-up displays the latter, it also must display the former but not vice versa. By using the admissibility clause, we have that A 's chance conditional on the subvenient property will trump A 's chance conditional of the supervenient property. The question whether the admissibility clause provides a correct solution in these situations, though philosophically important, cannot be undertaken

¹¹Notice that this is not to say that C_j is stochastically irrelevant to A *simpliciter*. We assume that both C_j and B_i make difference to A 's holding true. The condition only says that C_j is stochastically irrelevant to A given B_i .

¹²*Cognoscendi* will recognize this as a special case of the notion of screening-off proposed by Reichenbach (1956) and defended by Salmon (1971). On their account, given three propositions A , B , and C , A is said to screen off B from C if conditionalizing on A renders B_i and C_j stochastically independent.

here. Rather, the point I am promoting is that in these situations we can use the admissibility clause to pick out a level-relative chance to which the GP is applicable.

So far so good. But often there are cases in which we know divergent level-relative chances and yet it seems that the admissibility clause cannot be used in this way. In the next section, I introduce such a case. By doing so, I will argue that a successful use of the admissibility clause is essentially limited.

4 The GP and Viability Fitnesses in Evolutionary Theory

This section introduces a case from evolutionary theory—a case of two divergent viability fitnesses of an organism—and its bearing on the use of admissibility clause in the GP. Firstly, to set the stage, I briefly describe the basics of the idea of natural selection acting at different levels of biological organization and the concept of fitness relativized to such levels. Secondly, I introduce the case. Thirdly, I show what is required for the admissibility clause in the GP to bring a solution in this case. Finally, I argue that this solution cannot succeed.

Natural Selection at Multiple Levels: Fitnesses as Level-Relative Chances. Let us assume that natural selection is the only force influencing a trait’s or property’s evolution in some population. For natural selection to act it is required that the objects of selection must vary with respect to heritable fitness—the objective chance of surviving.¹³ Fitness differences among objects are causally affected by differences among the objects’ properties that are subject to selection. For example, when we say that, in a given population, butterflies have different chances for surviving, we mean that they vary with respect to a certain property (e.g., being camouflaged), and that this variation has a causal effect upon their fitnesses. This property and the fitness covary because differences in the former affect causally differences in the latter. This causal relation ensures that if the objects vary with respect to having some property, then they also vary in their fitness values.¹⁴

The simple picture just presented may be made more complex by acknowledging that the biological world exhibits a nested or hierarchical organization (genes are properly included in organisms that in turn are properly included in groups of organisms). If so, natural selection may act for or against properties at more than one level, e.g., for genic, genotypic, phenotypic or group properties. And there cannot be selection at a given level unless there is variation in fitness at that level. That is, there is group selection for or against some group property if there is variation in fitness among objects depending on whether or not they have this property, there is genic selection for or against some genic property if there is variation in fitness among objects depending on whether or not they have this property, and so on for other levels.

But what are the objects that vary with respect to fitness? This issue is important, for even if we agree upon the fact that selection can act for genic, phenotypic or group properties, it is left open what the objects are to which these properties and thus the fitnesses are attached. For example, we may agree that selection favours being fast over being slow, yet we may disagree whether being fast is advantageous (increases the level of fitness) for genes, organisms or groups.

¹³Typically, quantitative models in evolutionary theory deal with the concept of fitness as having two components, namely the *fertility* which captures the expected number of offspring, and the *viability* which expresses the chance of surviving. Nothing essential hinges on the fact that my focus is only on the viability component of fitness.

¹⁴For a similar view, see especially Okasha (2006, chap. 3). A somewhat different view is defended in Rosenberg (1978) and Sober (2000, chap. 3). They hold that the relation between fitness and an organism’s properties amounts to *simultaneous determination* or *supervenience*, that is to say, two organisms that are identical in their properties must have the same fitness, but the fact that they have the same fitness does not entail that they must be identical with respect to their properties.

This worry is known as the problem of the ‘benchmarks’ of selection (Sober 1984, chap. 8) or the ‘focal’ units of selection (Okasha 2006, chap. 2). For my purposes, I take it that it is organisms that are the objects to which we attach properties for or against which selection may act, and it is organisms that vary with respect to fitness. In particular, I assume that a given organism has a group-level property by the fact that it belongs to that group.

Given that differences in fitness among organisms are causally influenced by differences in properties they possess and that selection may act for or against properties at various levels, organisms may have different *level-relative* fitnesses in virtue of its having different level-relative properties. For example, it is perfectly possible that an organism may suffer in virtue of belonging to a group with a certain disadvantageous group property, but benefit in virtue of its genotypic property. Intuitively, a zebra may benefit in virtue of its being fast, but suffer in virtue of its belonging to a group with a high proportion of slow zebras that are less able to avoid predators. These facts about zebra’s level-relative properties give rise to its different fitness values: one of them is the fitness conditional on its having individual phenotypic property of being fast, the other is the fitness conditional on its belonging to a group with slow zebras. I suggest to think of these various fitnesses as of an organism’s conditional chances of survival, with the conditioning propositions describing the organism’s level-relative properties that are subject to selection force. I will then argue that once this quite intuitive way of thinking about fitnesses is granted, the use of admissibility clause in the GP faces a serious challenge.

Case. Consider a population divided into groups each consisting of n organisms. Suppose that groups vary in the percentage of altruistic (a) and selfish (s) individuals they contain. I take it that a and s are two measurable phenotypic properties at the individual level. If the propositions A and S stand for, respectively, the altruistic and selfish property of an organism at its individual level, then we can determine the first pair of fitness functions, namely $ch(-|A)$ and $ch(-|S)$. If the proposition O stands for “an organism o survives until next month”, these functions give O ’s conditional chances. I call them *individual-level* fitnesses, where:

- $ch(O|A)$ is o ’s fitness conditional on its altruistic phenotype a .
- $ch(O|S)$ is o ’s fitness conditional on its selfish phenotype s .

Also we may define o ’s fitness conditional on o ’s belonging to a group with a certain group property. It reflects an intuitive idea that o ’s viability depends on the group to which it belongs. For the sake of simplicity, I take it that groups are individuated by a certain characteristic concerning the average distribution of phenotypes a and s in them. If a_i denotes the altruistic phenotype of the i th individual in a group of n individuals, then the average altruism is given by $\bar{a} = \frac{1}{n} \sum_1^n a_i$. By analogy, the average selfishness is given by $\bar{s} = \frac{1}{n} \sum_1^n s_i$, where s_i denotes the selfish phenotype of the i th individual in a group of n individuals. Further, I assume that groups vary with respect to \bar{a} and \bar{s} . For our purposes, let the group with a higher average altruism than selfishness be denoted by g_a and the group with a higher average selfishness by g_s . If the propositions G_a and G_s encode for, respectively, the property of belonging to group g_a and the property of belonging to group g_s , we can determine the second pair of fitness functions, $ch(-|G_a)$ and $ch(-|G_s)$. They give two other conditional chances, which I call *group-level* fitnesses, that is:

- $ch(O|G_a)$ which is o ’s fitness conditional on its belonging to g_a .
- $ch(O|G_s)$ which is o ’s fitness conditional on its belonging to g_s .

Evolutionary biology teaches us that, under certain conditions, the following two statements hold true (e.g., Sober 1988; Sober 2000, chap. 4):

(Within-Group Fitness) For each organism o , regardless of what group in the population it is in,

$$ch(O|A) < ch(O|S).$$

(Between-Group Fitness) For each organism o in the population, regardless of its individual phenotype,

$$ch(O|G_a) > ch(O|G_s).$$

The first statement is a comparative claim about individual-level fitnesses within each group of our population. It says that for an organism, regardless of what group it is in, it is better to be selfish than to be altruistic: selfish organisms are fitter than altruists within each group. This is a simple consequence of the evolutionary definition of altruism and selfishness. That is to say, a trait is altruistic when it increases the fitness of others in a group and decreases the fitness of an organism possessing it, while the selfish trait always increases the fitness of an organism possessing it because it receives benefits from all altruists but never reciprocates. The second statement says that, regardless of its individual phenotype, an organism in a group with a higher proportion of altruists is fitter than an organism in a group with a higher proportion of selfish types. An obvious corollary of the second statement is the following comparative claim: altruistic organisms that belong to groups with a higher proportion of altruistic types are fitter than selfish organisms that belong to groups with a higher proportion of selfish types. This corollary when combined with the first statement shows that (i) within each group in the population at hand an altruist is *less fit* when compared to a selfish type and (ii) in the population at hand an altruist that belongs to a group with a higher proportion of altruistic types is *fitter* than a selfish organism that belongs to a group with a higher proportion of selfish types. We see, then, that the same altruistic organism in the ensemble of groups under consideration may have two different fitnesses, i.e., the individual and the group-level one.

Table 1 illustrates this case. There are two groups, denoted by G_a and G_s , each with a different proportion of altruists and selfish organisms. In each group $ch(O|A) < ch(O|S)$ ($0.4 < 0.7$ and $0.2 < 0.3$). But groups also vary in fitness, i.e., the group denoted by G_a is fitter than that denoted by G_s ($0.43 > 0.29$), which has effects on o 's fitness conditional on its belonging to a certain group, i.e., on its group-level fitness. Here the group-level fitness is given by the *average* individual fitness in the group, weighted by the frequency of altruistic and selfish types in that group. So, when we say that an organism o has a group-level fitness we mean by that its having such average individual fitness within that group. If the average altruism in a group, denoted by G_k , of n organisms is \bar{a} and the average selfishness is \bar{s} , then the group-level fitness is given by

$$ch(O|G_k) = (\bar{a}) ch(O|A) + (\bar{s}) ch(O|S) \tag{1}$$

where $\bar{a} + \bar{s} = 1$. Our case can then be explained by the fact that the higher the proportion \bar{a} of altruistic types in a group with higher viability, the greater the average individual-level fitness of that group; and the higher the proportion \bar{s} of selfish types in the group with lower viability, the lower the average individual-level fitness of that group. This holds despite the fact that altruists are less fit than selfish types in each group.

Next, I will argue that this case brings troubles in an epistemic context, namely when we ask which one of these level-relative chances should constrain an agent who has credence about O 's holding true.

The GP and our Case. Let $\langle ch(O|A) = x \rangle$ be the proposition that O 's chance given A is x , and $\langle ch(O|G_a) = y \rangle$ be the proposition that O 's chance given G_a is y . Suppose that an experimenter wants to conduct a trial on our population: she wants to know whether a given organism selected from this population will survive until next month. She has some

Group	\bar{a}	\bar{s}	$ch(O A)$	$ch(O S)$	$ch(O G_k)$
G_a	0.9	0.1	0.4	0.7	0.43
G_s	0.1	0.9	0.2	0.3	0.29

Table 1: An example of our case.

expectation about this and her expectation is her credence assigned to O . Further, she knows that the organism have the properties denoted by A and G_a , knows that $ch(O|A) = x$ and $ch(O|G_a) = y$, and, in addition, she knows that A and G_a are stochastically relevant to the truth of O , i.e.,

$$ch(O|A) \neq ch(O|\neg A)$$

and

$$ch(O|G_a) \neq ch(O|\neg G_a).$$

More specifically, she knows that (i) A is negatively correlated with O , i.e., $ch(O|A) < ch(O|\neg A)$, and (ii) G_a is positively correlated with O , $ch(O|G_a) > ch(O|\neg G_a)$. She does not know, however, the chance of O conditional on a logically stronger proposition $A \wedge G_a$. The question arises: which one of the two level-relative chances for O should constrain her credence about O ?

It follows from our case that these chances do not inform your credence in the same way. Her level of expectation in O when adjusted to $ch(O|G_a)$ is higher than her level of expectation in O when adjusted to $ch(O|A)$, as shown in Table 1. Speaking in a comparative way, deferring to the chance $ch(O|A)$ would set her expectation at a lower level when compared to one's credence adjusted to the chance $ch(O|S)$. But deferring to the chance $ch(O|G_a)$ would set your credence at a higher level when compared to one's credence adjusted to the chance $ch(O|G_s)$.

To answer our question by appealing to the admissibility clause in the GP, it is required that either A renders O and G_a conditionally stochastically independent, or G_a renders O and A conditionally stochastically independent. But since neither A entails G_a nor G_a entails A , this cannot be settled a priori. However, it seems that this issue can be decided upon evidence that the experimenter could acquire.

Evidence Against Admissibility. One of our initial assumptions was that when an organism's level-relative fitnesses covary with its level-relative properties, it is precisely because its fitnesses are causally influenced by these properties. That is, if the covariance between its fitness and one of its level-relative properties is non-zero, this is because the fitness and that property are *causally correlated*. To recall, in our peculiar case the agent knows that there is a (negative) causal correlation between A and O 's chance, and there is a (positive) causal correlation between G_a and O 's chance. Hence, she knows that both level-relative properties covary with the two corresponding level-relative fitnesses. To recall, the role of admissibility clause is to tell us that in the presence of one of these level-relative properties, causal correlation between O 's chance and the other property disappears, having the covariance set equal to zero. That is, its role is to tell which level-relative property is *causally sufficient* to determine A 's chance and thus which leaves the other property causally irrelevant. This in turn tells which property suffices to predict O 's chance.¹⁵ Once the property's causal sufficiency has been settled, the fitness conditional on that property is *predictively complete*.

Although this issue cannot be settled a priori on our case, the experimenter can gain evidence that bears on it. Specifically, she may manipulate or intervene on the experimental set-up to

¹⁵The converse is not necessarily true. For example, information about barometer reading may suffice to predict the storm, yet it is not even causally relevant to its occurrence.

		Individual Phenotype	
		A	S
Group	G_a	r	s
	G_s	t	w

Figure 1: Possible values of four conditional survival chances, i.e., $r = ch(O|A \wedge G_a)$, $s = ch(O|S \wedge G_a)$, $t = ch(O|A \wedge G_s)$ and $w = ch(O|S \wedge G_s)$.

obtain evidence about the set-up prior to manipulation, evidence about how the setup's chances to yield an outcome are causally correlated with its level-relative properties. That is, the goal of this manipulation test is to show, say, that if G_a is admissible with respect to the proposition $\langle ch(O|A) = x \rangle$, then manipulating the truth-value of G_a , while holding fixed O , will not change the value of the conditional chance function $ch(O|A \wedge G_a)$, that is:

$$ch(O|A \wedge G_a) = ch(O|A \wedge \neg G_a). \quad (2)$$

This test is counterfactual, i.e., it provides what-if-things-had-been-different information, and is aimed at providing *evidence*, not proof, about the causal correlations between O 's chance and the set-up's level-relative properties denoted by A and G_a . We ask: what would happen to O 's chance if we wiggled the truth-value of G_a while holding fixed A ? Here the operative idea is that, by manipulating G_a , we ask whether conditionalizing on G_a makes a difference to O 's chance conditional on A . And if such a manipulation leaves this chance unaltered, we have evidence that G_a makes no difference to O 's chance given A .¹⁶ In other words, we are entitled to say that A screens off G_a from O 's chance.

Note, however, that in the case under consideration, the experimenter does not know the values of conditional chances $ch(O|A \wedge G_a)$ and $ch(O|A \wedge \neg G_a)$. Nevertheless, for the purpose of manipulation, she might estimate these values from the observed frequencies. That is, we can assume that she has run a series of similar experiments that resulted in some relative frequency with which organisms of type A and G_a survive until next month and in some relative frequency with which organisms of type A and $\neg G_a$ survive until next month.

With these ideas in mind, we may formulate our manipulation test as follows:

(MAN1) Would o 's fitness change if o belonged to a group with a different group property, but had the same individual altruistic phenotype?

(MAN2) Would o 's fitness change if o had a different individual phenotype, but belonged to a group with the same group property?

To answer these questions, one may consider groups with different group properties by looking at the possible shifts in the proportion of altruists and selfish types in each group, and we may consider different individual phenotypes by looking at the available phenotypes in the population under consideration. We do not need to consider all possible scenarios to answer (MAN1) and (MAN2). I consider one scenario in the light of which the admissibility clause pales, that is, a scenario in which an alternative to A is S and an alternative to G_a is G_s . Figure 1 gives fitnesses that an organism o would have in each of the four possible individual phenotype/group combinations, where their values are estimated from the observed frequencies. To answer our

¹⁶For the idea that causal relations are revealed by manipulation and intervention, see especially Woodward (2003).

two counterfactuals, it suffices to focus on the values t and s . These are accessible alternatives for our altruist whose actual fitness is given by r . If both being an altruist A and being in the group G_a causally contribute to o 's survival, then the manipulations we are considering would have to show that $r \neq t$ and $r \neq s$. Such a scenario is not far to seek. Not only is it conceptually possible but also it is implied by empirically well-confirmed theory of multi-level natural selection advocated by Sober and Wilson (e.g., Sober and Wilson 1998). According to this theory, natural selection can *really* operate *simultaneously* at more than one level of biological organization. If this is so, then the upshot of natural selection—a change in trait frequency in a given population—is due to natural selection operating at each of those levels. And if natural selection operates at each of those levels, then there must be fitness variation at those levels simultaneously. So, in our case, for natural selection to be at work at the individual level it must be true that altruistic and selfish organisms vary within each of the groups, i.e.:

$$ch(O|A \wedge G_a) \neq ch(O|S \wedge G_a). \quad (3)$$

$$ch(O|A \wedge G_s) \neq ch(O|S \wedge G_s). \quad (4)$$

If simultaneously natural selection operates at the level of groups, then there also must be fitness variation at that level, i.e.:

$$ch(O|A \wedge G_a) \neq ch(O|A \wedge G_s). \quad (5)$$

$$ch(O|S \wedge G_a) \neq ch(O|S \wedge G_s). \quad (6)$$

Importantly, this model provided by the theory of multi-level selection is an answer to our manipulation test. In particular, (5) is an answer to (MAN1) and (3) is an answer to (MAN2). Both answers are in the affirmative. That is to say, (5) tells us that if o belonged to a group with a different group property described by G_s , but had the same individual altruistic phenotype, then its fitness would be different. Likewise, (3) tells us that if o had a different individual phenotype, a selfish one, but belonged to a group with the same group property described by G_a , then its fitness would change. This is a simple consequence of the fact that, in the population at hand, there might be both natural selection for individual phenotypes and natural selection for group-level properties.

The manipulation test just given provides evidence that neither the altruist's individual phenotypic property nor her group-level property is causally sufficient with respect to her chance of survival. Thus, the thought runs, we are entitled to claim that neither one screens off the other from O 's chance. It follows then that neither G_a is admissible with respect to $\langle ch(O|A) = x \rangle$ nor A is admissible with respect to $\langle ch(O|G_a) = y \rangle$. This has the following epistemic payoff. Someone who adjusts her credence to only one of these level-relative chances is not epistemically rational. She does not take into account all the evidence about level-relative chances that has bearing on the proposition she has credence about. This is an outright violation of the requirement of total evidence, which tells us that, as a matter of epistemic rationality, a rational agent's credence should reflect all the available evidence. It permits the agent to hold a credence that reflects only some subset of all her evidence if the totality of her additional evidence is irrelevant (Carnap 1962, p. 211). Of course the admissibility clause in the GP is intended to prevent us from violating the requirement of total evidence. Yet it fails in our case and so the GP is inapplicable. We are left with a hardly acceptable verdict. We are told that since the admissibility relations do not hold, no level-relative chance under consideration can constrain your credence. But intuitively we are far from saying that these level-relative chances have no bearing at all on our credences. After all, they are given by the well-confirmed evolutionary theory.

5 Two Reactions

There are at least two possible reactions to our verdict. First, one might argue that the case just given does not threaten the plausibility of admissibility clause in the GP. That is, one may argue, following Lewis's account of chance, that the level-relative fitnesses under consideration are not chances at all. And they are not chances precisely because they fail to satisfy Lewis's platitude about chance which says 'don't call any alleged feature of reality "chance" unless you've already shown that you have something, knowledge of which could constrain rational credence' (Lewis 1994, p. 484). On this view, the level-relative fitnesses are chances not because they figure in stochastic models of evolutionary change but because they satisfy the GP which is an intuitive generalization of Lewis's Principal Principle. Because they fail to satisfy the GP in our case, we should not treat them as genuine chances and thus perhaps we should not take information they convey seriously. There is nothing wrong with the admissibility clause and so, our intuition notwithstanding, the verdict should be sustained.

I take this Lewis-style reaction to be questionable. It seems that evolutionary theorists do not investigate rational credences, but biological phenomena. Hence, they assign level-relative chances not because they justify or underpin rational credences, but because they explain these phenomena. If so, chances they assign remain genuine chances even though they fail to justify rational credences in certain situations. Moreover, it seems that Lewis's platitude is not the only one that can be applied to level-relative chances. For example, one might well require that a candidate for chance should explain singular occurrences or observed frequencies.¹⁷ Now, could we say that a given level-relative chance that satisfies this platitude, but does not satisfy the platitude captured by the GP, is not a genuine chance? There seems to be no definitive answer to this question.

The second reaction, which I concur with, is that the level-relative fitnesses under consideration are genuine chances even though they might fail to satisfy Lewis's platitude. One may hold that indeed a level-relative chance should constrain rational credence, yet claim that the use of the GP is ill-founded. The key is to recognize that it hinges on the admissibility clause whose plausibility may be called into question. How can it be questioned? One way this can be done is to acknowledge that the admissibility clause works for a proposition's level-relative chances whose conditioning propositions cite properties that can, in principle, compete for causal sufficiency with respect to the determination of that proposition's chance. That is, it presupposes that, for any pair of level-relative chances, we can say, in principle, which one of them constrains our credence by performing a kind of bracketing operation on the other competing level-relative chance. The bracketing operation hinges on the screening-off condition. By performing it, we carve off the irrelevant information conveyed by one conditioning proposition and so expect to have the relevant one left by another. This, however, might be a false assumption. As we have seen, Sober and Wilson's theory of multi-level selection provides a model in which neither the organism's individual phenotypic property nor her property of being in a group is, in principle, *causally sufficient* to determine its fitness. In this light, the way in which the screening off relation, which underlies admissibility, apportions causal responsibility between these level-relative properties is not tenable. For it purports to single out a property which is *causally sufficient* to determine the organism's survival chance, while leaving the other property *causally irrelevant*, in a situation in which it is in principle impossible. On this view, then, we are entitled to appeal against the verdict, keeping alive the possibility that evidence about the two level-relative fitnesses have bearing on our credences about other propositions.

But how can our case be ruled on after the appeal? I suggest that the GP should be

¹⁷Other such platitudes have been defined in the case of unconditional chance in Schaffer 2007, and some of them seem to be equally applicable to level-relative chances.

endowed with a different admissibility clause. The admissibility clause based on screening-off is essentially *qualitative* and as such is too restrictive. That is, by using it, we ask whether B_i screens off C_j , or whether C_j screens off B_i from A 's chance. It thus excludes the possibility that these screening-off relations might not hold absolutely, i.e., $B_i (C_j)$ may screen off $C_j (B_i)$ to a certain degree. But the gist of our case is that such a possibility should not be excluded, for we ask: is the organism fitness due *more* to its phenotypic property or to its being in a group? What we need, then, is a *quantitative* admissibility clause which would allow us to determine a degree to which a given proposition is admissible with respect to a proposition about chance.

6 Quantitative Admissibility: Explanation *via* Resiliency

The admissibility relation between propositions may not be an all-or-nothing matter, but it may come in degrees. It seems that we may speak about some propositions being more or less admissible than other. That is, instead of asking whether E is admissible or not, we may ask to what extent E is admissible to a proposition about chance. Consider the following example. You are about to throw a die today and you know that the chance that it lands up '2' is one sixth. Additionally you know the following propositions:

1. The die landed up '3' yesterday.
2. The outcome of the die throw is even.
3. A reliable crystal ball predicts that '2' will come up.

It seems that (1) is more admissible than (2), which in turn is more admissible than (3). (1) conveys historical information about the past behaviour of the die which should have no impact on its future behaviour. (2) gives us information over and above the information contained in the chance of one sixth, but unlike (3) it does not tell us what the outcome of the die throw will be. (3) seems to be maximally inadmissible with respect to the chance of one sixth, since it reveals what the outcome will be. In other words, whereas (1) reveals *nothing* about the future behaviour of the die, (2) reveals *something* and (3) reveals *everything* by telling us now what that future behaviour will be.

In what follows, I propose an account of admissibility that, I believe, can capture this gradual aspect.¹⁸ In general, it goes along Brian Skyrms's notion of the resiliency of chance. I show that such quantitative conception of admissibility is a natural generalization of the qualitative account of admissibility.

Skyrms (1977; 1978; 1980) has taken the resiliency of chance, both unconditional and conditional, to be a mark of its stability, and has gone on to argue that well-confirmed statistical theories should posit resilient chances.¹⁹ According to Skyrms, resiliency of chance means its approximate invariance under variation of experimental factors. The most important features of Skyrms's notion of resiliency are the following:

- resiliency of chance comes in degrees, and the degree of resiliency reflects how much chance wiggles under conditionalization on experimental factors,
- chances should approximate the highest degree of resiliency; in other words, the higher the resiliency, the better,

¹⁸Interestingly, Thau (1994), and after him Lewis (1994), recognized the quantitative nature of admissibility.

¹⁹Skyrms (1984, chap. 3) has also applied the idea of resiliency to a Bayesian conception of chance, it wit, the view that chance is one's degree of belief which is a 'mixture' of physical probabilities. Skyrms has argued that to count as chance such epistemic probability must be resilient.

- resiliency is always relative to a set of experimental factors that constitute the *scope* of resiliency.
- the idea of resiliency applies both to unconditional and conditional chances.

Let us present Skyrms's idea of resiliency more precisely, with the focus on conditional chances. Let \mathcal{F} be a set of propositions F_1, \dots, F_n , each describing an experimental factor. Then, the degree of resiliency of the conditional chance $ch(A|B)$ over \mathcal{F} , denoted by $R(ch(A|B), ch_{F_i}(A|B))$, is given by

$$R(ch(A|B), ch_{F_i}(A|B)) = 1 - \max_{F_i \subseteq \mathcal{F}} |ch(A|B) - ch_{F_i}(A|B)|,$$

where F_i ranges over the experimental factors and $ch_{F_i}(A|B)$ comes from $ch(A|B)$ by conditionalizing on F_i .²⁰ That is, the degree of resiliency $ch(A|B)$ is one minus the maximal possible difference between the chance $ch(A|B)$ and the chance $ch_{F_i}(A|B)$. Maximal resiliency of the conditional chance over \mathcal{F} equals 1. According to Skyrms (1977), this happens when B *screens off* F_i from A 's chance. This is an ideal which conditional chances should approximate.

Let us now introduce a link between the quantitative admissibility and the resiliency of conditional chance. For simplicity's sake, suppose that an agent's evidence E consists of one experimental factor denoted by F . To recall, given the account of qualitative admissibility presented in section 3, F is admissible with respect to $\langle ch(A|B) = x \rangle$ if B *screens off* F from A 's chance. This in turn means that F is admissible with respect to $\langle ch(A|B) = x \rangle$ if the resiliency of $ch(A|B)$ over F equals 1. However, if this ideal is not attainable and the resiliency is less than 1, it seems that we can legitimately speak about the degree of F 's admissibility with respect to $\langle ch(A|B) = x \rangle$. The degree of F 's admissibility, then, will be the extent to which the conditional chance $ch(A|B)$ is invariant upon conditionalizing on F , or simply, its degree of resiliency over F . More precisely:

(Quantitative Admissibility) F is admissible with respect to $\langle ch(A|B) = x \rangle$ to degree r if

$$R(ch(A|B), ch_F(A|B)) = r.$$

where $ch_F(A|B)$ comes from $ch(A|B)$ by conditionalizing on F . Here, the idea is that the higher the value of resiliency r over F , the more admissible F is with respect to $\langle ch(A|B) = x \rangle$. The idea of quantitative admissibility, hence, is a natural generalization of the idea of qualitative admissibility based on the notion of screening off.

Where does this leave us vis-à-vis the question of which one of the divergent level-relative chances should underpin one's credence? If we emancipate the admissibility clause from the restrictive screening-off condition, it seems that we might escape the stricture that chance underpins one's credence if one's evidence is *maximally admissible*. In the next section, I will make some headway towards defending this claim.

7 Near-Admissibility and Highly Resilient Chances

One might be tempted to say that the quantitative approach to admissibility should be preferable to the qualitative one in cases involving divergent level-relative chances. It appears that both these approaches lead to the same result in cases where either B_i entails C_j , or C_j entails B_i . Moreover, whereas the use of qualitative admissibility fails to give a solution in cases like

²⁰More generally, we may define the degree of resiliency of the conditional chance $ch(A|B)$ over Boolean combinations of propositions F_1, \dots, F_n . Also we may want A and B to be in \mathcal{F} . If so, we need to require any $F_i \in \mathcal{F}$ to be consistent with $A \wedge B$ and $A \wedge \neg B$.

the one discussed in section 4, the use of quantitative admissibility leaves open the possibility of giving such solution. Even if C_j is not maximally admissible with respect to $\langle ch(A|B_i) = x \rangle$, it might be admissible with respect to that proposition to a certain degree r . And this degree of admissibility reflects the extent to which the chance $ch(A|B_i)$ is resilient over C_j . Likewise, B_i , though not maximally admissible, may still be admissible to $\langle ch(A|C_j) = y \rangle$ to a certain degree r^* , where r^* reflects the extent to which the chance $ch(A|C_j)$ is resilient over B_i . More generally, by using the quantitative admissibility clause, we can rank divergent level-relative chances by the criterion of resiliency. Now given that the degrees of resiliency are different, could we reasonably claim that one of these level-relative chances underpins one's credence? This is a highly nontrivial question. In the remainder of this section, I promote a positive answer to this question. Regrettably, I have no conclusive argument to offer in this regard. But I do want to give reasons for the claim that chance, whether unconditional or level-relative, whose degree of resiliency is sufficiently high can underpin one's credence. Again, my point is not to argue that this is *the* correct answer, but merely to show that, unlike the use of qualitative admissibility, the use of quantitative admissibility can succeed in picking out a level-relative chance that underpins the agent's credence in such situations.

Interestingly, Lewis (1994, p. 486) claimed that even if E is not perfectly admissible to a proposition about chance, but is *nearly* admissible, the conclusion that chance underpins credence should hold. And, E is nearly admissible just in case it tells us so little 'over and above' what is told by chance so that its impact of the agent's credence is negligible. I take Lewis's claim to have a certain plausibility. Of course, his striking statement can hardly be called an argument. In terms of resiliency, Lewis's statement, arguably, amounts to saying that chance can underpin credence if its degree of resiliency over E is *sufficiently high*. But what value should r have to count as sufficiently high? I suspect that not definitive answer can be given. Whatever this value should be, it seems that a more pressing question to answer is: how could we motivate Lewis's claim? That is, why can highly resilient chances, albeit not maximally resilient, underpin one's credence? After all, Lewis's claim stands in tension with the stricture that a rational agent should proportion her credence to all the available evidence.

Firstly, we should recognize the fact that evidence about more resilient chance tends to make one's credence more *stable* in the face of one's additional evidence. That is, an agent who sets her credence equal to a highly resilient chance is more reluctant to change her credence upon conditionalizing on her additional evidence. Consequently, her credence in a given proposition would tend to concentrate on a certain value given by the chance of that proposition when she conditionalizes on the rest of her evidence. Ideally, we could require that evidence about chance makes one's credence maximally stable in this sense so that no additional evidence will alter its value. As it has recently been emphasized by Lyon (2010), one of the intuitions we seem to have about chance is that evidence about chance makes one's credence in other propositions stable in some very strong sense. Lewis expressed this intuition as follows:

To the extent that uncertainty about outcomes is based on certainty about their chances, it is a stable, resilient sort of uncertainty—new evidence won't get rid of it. (Lewis 1986, p. 85)

For example, if one knows that the chance of a die landing up '2' is one sixth, one's credence for '2' should be one sixth and it is a stable credence in the sense that it does not change when one acquires additional evidence. I think that this intuition is only partially correct. For it seems equally plausible to claim that chance is not a perfect expert for one's credence. To put this in other words, chance is not like truth, a function that assigns 1 to all actual truths and 0 to all actual falsehoods. As Joyce (2007) has pointed out, truth advises us rightly no matter what evidence we have. But often this is not so in cases when we adjust credences to chances. For one might possess evidence concerning experimental factors over which chance is not maximally resilient. For example, if one knew all the micro-details of the die, the resiliency of the chance

of a die landing up ‘2’ would drop significantly. But still there might be chances whose degree of resiliency may be sufficiently high so that one’s credence informed by such chance, upon conditionalizing on one’s additional evidence, would tend to concentrate on the value given by that chance. As a matter of having stable credences, such chances should guide our credences. Of course, it is not clean-cut how the stability of credences relates to the requirement of total evidence. Both these requirements appear to be important. In the case we have discussed in section 4, no level-relative chance can be maximally resilient. Both individual- and group-level chance encode information that is relevant to the outcome. The requirement of total evidence tells us that both chances should regulate our credences. But it is equally possible that one of these chances would be highly resilient over experimental factors. There might be only *some* selection operating at the individual level while the *main* selective force is on the group level. If so, it is the group-level chance that would have a profound impact on one’s credence while the impact of the individual-level chance would be negligible. In this regard, stability tells us that it is the group-level chance that should regulate our credences. Which recommendation has the upper hand? The best, I think, one can offer is the following conditional answer: if we gave more weight to stability, the latter recommendation would have the upper hand.

Secondly, and more generally, we should not assume a priori that statistical theories posit maximally resilient chances. Skyrms (1980, p. 17) has postulated that, at best, well-confirmed statistical laws should posit highly resilient chances. He has motivated this claim by the following observation:

We find no absolute resiliency in physics, nice as it would be if we did. And philosophers have no business trying to lay down, a priori, standards for the scope of resiliency appropriate for physical theories. Standards for resiliency evolve along with physical theory in a big virtuous circle, and in our dealings with nature we take what we can get. (Skyrms 1980, p. 19)

In the light of Skyrms’s remarks, no standard for the scope and the value of resiliency appropriate for statistical laws seems to be a priori justifiable. Much depends on the nature of things we deal with. Sometimes high resiliency is all we can get. But this does not necessarily mean that high resiliency is a poor standard in scientific practice. Often chances that are highly, albeit not maximally, resilient over all the experimental factors provide satisfactory explanations. Take, for example, chances in classical statistical mechanics. Consider an explanation of why a cup of coffee cools down. Classical statistical mechanics tells us that this is because given a chance distribution over initial conditions it is overwhelmingly probable that its micro-state is one that lies on a trajectory that deterministically takes it into the macro-state “cooled down”. Explanations based on such chances are perfectly satisfactory, even though these chances are not maximally resilient over the micro-details of a given chance set-up. What makes them satisfactory is, among other things, that these chances are highly resilient: we are told that the overwhelming majority of micro-states compatible with a given macro-state would evolve to a higher- entropy macro-state. So, even if we knew these micro-states, this would have a negligible impact on our credences. The point is that, for explanatory purposes, there is no sense to move to a “truer” chance by conditionalizing on a further specification of these micro-details.

8 Concluding Remarks

I have argued that an application of the GP to a pair of divergent level-relative chances for a given proposition is not so simple a matter as one might think. In so doing, I have explained the admissibility clause in the GP by drawing on Lewis’s idea of admissibility. I have shown that the use of admissibility clause is essentially limited. This is due to the restrictive screening-off condition that underlies admissibility. I have illustrated this point by discussing a case of level-relative chances, i.e., viability fitnesses, in evolutionary theory.

As a remedy, I have, firstly, argued that we can revise the GP by providing a quantitative rather than a qualitative understanding of the admissibility clause. To bolster this appeal, I have developed a quantitative account of admissibility by drawing on Skyrms's idea of the resiliency of chance, and have shown that this account is a natural generalization of the qualitative approach. Secondly, I have suggested and motivated the claim that highly, albeit not maximally, resilient chances can underpin credences. Maximal admissibility may be too stringent a condition for chance, whether unconditional or conditional, to guide our credences.

Appendix

Proof of Proposition 1. Suppose that $B_i \subseteq C_j$ and suppose that ch is a probability function. By probability theory, we have that $ch(C_j|B_i) = 1$. Then

$$\begin{aligned}
ch(A|B_i) &= \frac{ch(A \wedge B_i)}{ch(B_i)} \\
&= \frac{ch((A \wedge B_i \wedge C_j) \vee (A \wedge B_i \wedge \neg C_j))}{ch(B_i)} \\
&= \frac{ch(A \wedge B_i \wedge C_j) + ch(A \wedge B_i \wedge \neg C_j)}{ch(B_i)} \\
&= \frac{ch(A|B_i \wedge C_j)ch(B_i \wedge C_j) + ch(A \wedge \neg C_j|B_i)ch(B_i)}{ch(B_i)} \\
&= \frac{ch(A|B_i \wedge C_j)ch(B_i \wedge C_j)}{ch(B_i)} + \frac{ch(A \wedge \neg C_j|B_i)ch(B_i)}{ch(B_i)} \\
&= ch(A|B_i \wedge C_j)ch(C_j|B_i) + ch(A \wedge \neg C_j|B_i) \\
&= ch(A|B_i \wedge C_j) + \frac{ch(A \wedge \neg C_j \wedge B_i)}{ch(B_i)} \\
&= ch(A|B_i \wedge C_j)
\end{aligned}$$

as required. □

References

- Carnap, R. (1962). *Logical Foundations of Probability*. Chicago]University of Chicago Press.
- Glynn, L. (2010). Deterministic Chance. *British Journal for the Philosophy of Science*, 61(1):51–80.
- Hájek, A. (2007). The Reference Class Problem is Your Problem Too. *Synthese*, 156(3):563–585.
- Hall, N. (1994). Correcting the Guide to Objective Chance. *Mind*, 103(412):505–518.
- Hoefer, C. (2007). The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance. *Mind*, 116(463):549–596.
- Joyce, J. M. (2007). Epistemic Deference: The Case of Chance. *Proceedings of the Aristotelian Society*, 107(2):187–206.
- Levi, I. (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. The Mit Press.

- Lewis, D. (1986). A Subjectivist's Guide to Objective Chance. In *Philosophical Papers*, volume 2, pages 83–132. Oxford University Press.
- Lewis, D. (1994). Humean Supervenience Debugged. *Mind*, 103(412):473–490.
- Lyon, A. (2010). Deterministic probability: Neither chance nor credence. *Synthese*, 182(3):413–432.
- Miller, D. (1966). A Paradox of Information. *British Journal for the Philosophy of Science*, 17(1):59–61.
- Okasha, S. (2006). *Evolution and the Levels of Selection*. Oxford: Oxford University Press.
- Price, H. (1984). Mellor, Chance and the Single Case. *British Journal for the Philosophy of Science*, 35(1):11–23.
- Reichenbach, H. (1956). *The Direction of Time*. Dover.
- Rosenberg, A. (1978). The Supervenience of Biological Concepts. *Philosophy of Science*, 45(3):368–386.
- Salmon, W. C. (1971). *Statistical Explanation And Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.
- Schaffer, J. (2007). Deterministic Chance? *British Journal for the Philosophy of Science*, 58(2):113–140.
- Skyrms, B. (1977). Resiliency, Propensities, and Causal Necessity. *The Journal of Philosophy*, 74(11):704–713.
- Skyrms, B. (1978). Statistical Laws and Personal Propensities. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1978:551–562.
- Skyrms, B. (1980). *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. New Haven: Yale University Press.
- Skyrms, B. (1984). *Pragmatism and Empiricism*. Yale University Press.
- Sober, E. (1984). *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Chicago: University of Chicago Press.
- Sober, E. (1988). What is evolutionary altruism? In *New Essays on Philosophy and Biology (Canadian Journal of Philosophy Supp. Vol. 14)*. University of Calgary Press.
- Sober, E. (2000). *Philosophy of Biology*. Westview Press.
- Sober, E. and Wilson, D. S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Strevens, M. (1995). A Closer Look at the 'New' Principle. *British Journal for the Philosophy of Science*, 46(4):545–561.
- Strevens, M. (1999). Objective Probability as a Guide to the World. *Philosophical Studies*, 95(3):243–275.
- Strevens, M. (2006). Probability and Chance. In Borchert, D. M., editor, *Encyclopedia of Philosophy, second edition*. Macmillan.

- Thau, M. (1994). Undermining and Admissibility. *Mind*, 103(412):491–504.
- van Fraassen, B. (1989). *Laws and Symmetry*. Oxford University Press.
- Vranas, P. B. M. (2002). Who's Afraid of Undermining? *Erkenntnis*, 57(2):151–174.
- Vranas, P. B. M. (2004). Have Your Cake and Eat It Too: The Old Principal Principle Reconciled With the New. *Philosophy and Phenomenological Research*, 69(2):368–382.
- Woodward, J. (2003). *Making Things Happen*. Oxford: Oxford University Press.