The New Tweety Puzzle: Arguments against Monistic Bayesian Approaches in Epistemology and Cognitive Science

Matthias Unterhuber & Gerhard Schurz
University of Duesseldorf, Germany

20th May 2012

Abstract. In this paper we discuss the new Tweety puzzle. The original Tweety puzzle was addressed by approaches in non-monotonic logic, which aim to adequately represent the Tweety case, namely that Tweety is a penguin and, thus, an exceptional bird, which cannot fly, although in general birds can fly. The new Tweety puzzle is intended as a challenge for probabilistic theories of epistemic states. In the first part of the paper we argue against monistic Bayesians, who assume that epistemic states can at any given time be adequately described by a single subjective probability function. We show that monistic Bayesians cannot provide an adequate solution to the new Tweety puzzle, because this requires one to refer to a frequency-based probability function. We conclude that monistic Bayesianism cannot be a fully adequate theory of epistemic states. In the second part we describe an empirical study, which provides support for the thesis that monistic Bayesianism is also inadequate as a descriptive theory of cognitive states. In the final part of the paper we criticize Bayesian approaches in cognitive science, insofar as their monistic tendency cannot adequately address the new Tweety puzzle. We, further, argue against monistic Bayesianism in cognitive science by means of a case study. In this case study we show that Oaksford and Chater's (2007, 2008) model of conditional inference – contrary to the authors' theoretical position – has to refer also to a frequency-based probability function.

Keywords: New Tweety puzzle, probability, frequency, probabilism, monistic Bayesianism, objective Bayesianism, Bayesian rationality, Oaksford and Chater, conditional inference, MP-MT asymmetry, cognitive science

1. Introduction

Uncertainty seems to be an inherent feature of all human affairs, both in every-day life and scientific inquiry. It is, hence, not surprising that probabilistic approaches have found widespread application as a normative framework in epistemology (cf. Earman, 1992; Howson & Urbach, 2006) and in cognitive science (Chater, Tenenbaum, & Yuille, 2006a; Oaksford & Chater, 2007; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010). Prominent approaches in epistemology and cognitive science are subjective Bayesian accounts. These approaches share that they employ subjective interpretations of probabilities (Talbott, 2008; Bovens & Hartmann, 2003). A subjective interpretation renders probabilities as agent-relative degrees of belief in propositions, such as the degree of belief of the next coin toss ('a') landing heads ('Ha'), abbreviated as '$P(Ha)$'. This interpretation of probabilities contrasts with objective frequency-based interpretations, such as the limit of the (relative) frequency of a dice landing heads $p(Hx)$ in a sequence of tossings (where 'Hx' stands for 'x lands heads' and 'x' is an individual *variable*). Thus, we use 'P' for a degree of belief function and 'p' for a frequency-based probability function.

With 'Bayesianism' we refer in the following always to subjective Bayesianism. Monistic Bayesianism is a brand of subjective Bayesianism, which takes subjective probabilities as the only fundamental probabilistic notion. Monistic Bayesianism has been defended, for example, by de Finetti (1973; Gillies, 2000, 70ff). Hájek (2012, in §3.3.4) calls this position 'orthodox Bayesianism'. Besides subjective Bayesianism there exists also the less well-known position of objective Bayesianism, which defends the notion of intersubjectively rational (and in this sense "objective") degrees of belief (cf. Williamson, 2010). Since all objective Bayesians accept – besides intersubjective probabilities – subjective (agent-relative) degrees of belief, objective Bayesians are not monistic but dualistic Bayesians.

In this paper we will criticize monistic Bayesianism. We do not criticize the use of subjective probabilities in general. We accept that any fully adequate probabilistic approach in epistemology and cognitive science has to employ subjective probabilities. We, hence, do not argue in favor of an alternative monistic account to Bayesianism, such as either frequentistic approaches (which take limiting frequencies as fundamental) or a purely qualitative approach, for example Kripke models or ranking functions. We rather argue that *monistic* Bayesianism is inadequate, inasmuch as both subjective probabilities and relative frequencies have to be taken as primitive notions for any descriptively and normatively adequate theory of epistemic states.

1.1. The Tweety Puzzle and Theories of Epistemic States

At the heart of our argument lies the new Tweety puzzle. The basic idea behind the puzzle goes back to a well-known exception structure from the non-monotonic reasoning literature (Brewka, 1991, p. 2f):

Generic Form:		Singular Form:	
(Ag)	Birds can fly.	(As)	If Tweety is a bird, it can fly.
(Bg)	Penguins cannot fly.	(Bs)	If Tweety is a penguin, it cannot fly.
(C)	Tweety is a bird and a penguin.		

(Ag), (Bg), and (C) represent the generic formulation of the Tweety case and (As), (Bs), and (C) its singular formulation. If we – in what follows – refer to (A)-(C) without qualification in terms of the generic vs. the singular version, we mean that our claim in this context applies to both versions of the new Tweety puzzle.

Ordinary language conditionals often express uncertain conditionals that admit exceptions. This is the case in the Tweety example: It is an exception structure inasmuch as Tweety is an exceptional bird, which cannot fly (since it is a penguin), although birds can (in general) fly. Since penguins are birds (i.e., "penguin" is more specific than "bird"), the unambiguous conclusion is that Tweety cannot fly.[1] We call it the "new" Tweety puzzle, since it is not – unlike its original formulation – meant as a challenge to default reasoning, but as a challenge to probabilistic theories of epistemic states.

To explain the new Tweety puzzle, we will, first describe, the standard account of (A)-(C) in terms of classical (propositional or predicate) logic. The problem of classical logic is that it renders (A)-(C) jointly inconsistent, despite the fact that, intuitively, (A)-(C) are conjointly consistent. Classical logic does so for both the generic and the singular version of (A)-(C), as can be seen from the following formalization:

Generic Form:		Singular Form:	
(A1g)	$\forall x(Bx \rightarrow Fx)$	(A1s)	$Bt \rightarrow Ft$
(B1g)	$\forall x(Px \rightarrow \neg Fx)$	(B1s)	$Pt \rightarrow \neg Ft$
(C1)		$Bt \& Pt$	

Here 'Bx', 'Px', 'Fx' and 't' abbreviate 'x is a bird', 'x is a penguin', 'x can fly', and 'Tweety', respectively. '$\forall$', '$\rightarrow$', '&' and '$\neg$' are the universal quantifier, the material implication, the conjunction and the negation operator, respectively. The contradiction arises in the generic case, since (A1g) and (C1) imply Ft, while (B1g) and (C1) give us $\neg Ft$. Analogously, in the singular case (A1s) and (C1) on the one hand and (B1s) and (C1) on the other hand imply Ft and $\neg Ft$, respectively.

Bayesians (whether monistic or not) as well as proponents of non-monotonic logics agree that classical logic is not an appropriate means to represent (A)-(C). At the core of the classical logic representation of the Tweety case lies the following problem: The classical logic formalization presupposes that (A) and (B) do not allow exceptions, while (A) and (B), as applied in everyday and scientific discourse, seem to contradict this assumption.

Probabilistic representations of (A)-(C) are more flexible, because they allow for exceptions, and inasmuch do not necessarily render (A)-(C) inconsistent. In fact, one motivation for probabilistic approaches to modeling epistemic states – besides making it possible to model degrees of beliefs rather than categorical beliefs – is the fact that probability theory allows us in principle to represent exceptions (cf. Pearl, 1988, p. 2). This is also seen as

[1] We do not have to add the specificity information 'penguins are birds' as a further premise in order to make our point. There is a second type of exception structure, namely the Nixon diamond case, which goes as follows: (A) Quakers are pacifists, (B) republicans are not pacifists and (C) Nixon is both a Quaker and a republican (Brewka, 1991, p. 14). The Nixon diamond case lacks specificity, and, so, no unambiguous conclusion can be drawn concerning Nixon's pacifism. We could make our point equally well with the Nixon diamond case (see Schurz, 2011, §2.).

a clear plus of probabilistic approaches in cognitive science (e.g., Griffiths et al., 2010, p. 360).

We will see that monistic Bayesianism does not fare significantly better than classical logic with respect to (w.r.t.) the Tweety case: Monistic Bayesians suffer from the same fate as proponents of a classical logic analysis, inasmuch as there is no probabilistically consistent representation of (A)-(C) in terms of a single subjective probability function. The monistic Bayesian's conundrum is, thus, the following:

(I) For any theory of epistemic states to be fully adequate, the theory must, be able to represent Tweety type beliefs as consistent.

(II) Monistic Bayesianism, however, cannot do that.

In this paper we will focus on point (II). One might also be tempted to regard point (I) as implausible. However, to deny that a fully adequate theory of epistemic states should be able to represent Tweety type cases implies that any pair of exception-admitting and instantiated conditionals or generics of type (A) and (B) are insignificant, at least from the viewpoint of a theory of beliefs. This assumption is hardly plausible, since it conflicts strongly with the fact that normic laws, as expressed by uncertain conditionals or generics, are abundant in everyday life and the sciences (cf. Schurz, 2001, p. 476f). So, the inability to represent Tweety type cases would reduce the applicability of any theory of epistemic states to a great extent, since we hardly accept or believe normic laws unless they are non-instantiated or non-conflicting with other normic laws.

1.2. Structure of Our Paper

We will follow Schurz (2011) in our use of the Tweety case against monistic Bayesianism. We will extend Schurz's argumentation in the following ways. We draw out in more detail, in which way monistic Bayesian approaches fail to solve the new Tweety puzzle. We will describe parts of the results of an empirical study that is briefly reported in Schurz (2011). We shall, furthermore, investigate in how far monistic type Bayesian approaches are tenable in cognitive science. We do so by discussing the new Tweety puzzle in the context of these approaches and supplement our argument by means of a case study of Oaksford and Chater's (2007, 2008) model of conditional inference. The latter part extends Schurz's argument against monistic Bayesianism.

The structure of the paper is as follows. In section 2 we discuss why the new Tweety puzzle cannot be handled adequately by monistic Bayesian approaches and why our findings imply that monistic Bayesianism cannot be a fully adequate theory of doxastic states. We will furthermore, describe two possible solutions to the new Tweety puzzle, based on Schurz's (2011) account. In these solutions both a subjective and a frequency-based probability function are used to account for agents' doxastic states.

In section 3 we present the results of a questionnaire study, in which we investigated whether cases such as (A)-(C) are perceived as probabilistically contradictory or not. This study serves two purposes. First, we investigated by these means whether our intuitions regarding the probabilistic consistency of (A)-(C) is not only specific to philosophers like us, but shared (at least) by other Western people. Second, the data allowed us to test – based on our argument in section 2 – whether subjects use only a subjective probability function to represent Tweety type cases or, in addition, also a frequency-based one.

Finally, we will discuss in section 4 two points. First, we criticize monistic type Bayesian approaches in cognitive science. We, then, focus on Oaksford and Chater's (2007, 2008) model of conditional inference. This serves as a case study for monistic Bayesian approaches in cognitive sciences, insofar as we aim to show that Oaksford and Chater's model of conditional inference (2007, 2008) – contrary to the authors' theoretical position – implicitly refers to frequencies and must do so. Our criticism of Oaksford and Chater (2007, 2008), hence, provides additional support for our thesis that a fully adequate (probabilistic) theory of epistemic states must refer to both a degrees of belief function and a frequency-based probability function.

Hawthorne (2005) and Schurz (2011) put forth a further argument against monistic Bayesianism. Their argument is based on the requirement that likelihoods have to be specified objectively, independent from actual degrees of belief. We will, however, not discuss the latter argument here.

2. The New Tweety Puzzle

2.1. Probabilism and Monistic Bayesianism

Before we describe the new Tweety puzzle, let us specify in more detail, what we mean by *probabilism* (including "weak" or "dualistic" Bayesianism) in contrast to *monistic* Bayesianism. Probabilism is a certain type of Bayesianism, which has the following core thesis:

1. Rational beliefs come in degrees, and rational degrees of belief obey the standard probability axioms (e.g., Joyce, 1998; Hájek, 2008).

In the context of conditionals, probabilism is usually associated with the following additional thesis:

2. At least in many cases, uncertain conditionals or generics should be understood as assertions of high conditional probabilities (Adams, 1975;

McGee, 1989; Bennett, 2003; psychology: Evans, Handley, & Over, 2003; Oberauer & Wilhelm, 2003).

While generics, such as (Ag) and (Bg), have the linguistic form '*A*s are *B*s', conditionals have the linguistic structure 'if ... then'. Despite this fact, it is almost universally assumed that we can paraphrase generics in terms of generic conditionals, i.e., 'if ... then' structures; the classical logic approach and the Bayesian and probabilist approaches do not differ herein and we will not deviate from this assumption either.

We employ a more cautious formulation of thesis 2 of probabilism, for the following two reasons. First, thesis 2 is not a core tenet of probabilism, but expresses a widely held expansion of probabilism w.r.t. uncertain conditionals. Second, we do not claim that all uncertain conditionals express high conditional probabilities, but only that many of them do, and that in particular the conditionals (A), and (B) in the Tweety case are to be understood in that way; this is sufficient for making our point against monistic Bayesianism. Leslie, for example, has objected to a universal formulation of thesis 2 (cf. Schurz, 2005; Nickel, 2009) that some generics such as 'Mosquitos carry the West Nile Virus' do not seem to imply that the majority of mosquitos carries that virus (Leslie, 2007, p. 376; 2008, p. 7). She rather argues that this generic is true, even if only a small percentage of mosquitos actually carry the disease. However, Leslie (2007, 2008) does not argue that no generics satisfy thesis 2. She only aims to show that some important subclass of generics does not have that property. It is, however, hard to argue that (A) and (B) violate thesis 2. The statements 'Birds can fly' and 'Penguins cannot fly' are not prone to an interpretation in line with the above mosquito case. The conditional or generic (Ag) [(Bg)] is true (and/or assertible), if the great majority of birds [penguins] can [cannot] fly. It is not true (and/or assertible), if only a minority of birds [penguins] can [cannot] fly. Leslie seems to accept this assumption at least for some generics (Leslie, 2008, p. 1). Moreover, there are reasons to suppose that such important classes of generics exist, which satisfy condition 2. Schurz (2005, p. 39; 2012), for example, argues that evolutionary systems generate high probabilities or frequencies associated with conditionals or generics based on their regulatory features. So, given that evolutionary systems are almost omnipresent in our environment, it seems plausible to assume that there are at least some important classes of generics, which satisfy thesis 2.

Note also that thesis 2 does not imply the Adams-Stalnaker thesis, namely that the probability of a conditional $P(A \Rightarrow B)$ equals its conditional probability $P(B|A)$ for $P(A) > 0$. Here '$A \Rightarrow B$' and '$P(B|A)$' stand for the conditional 'if *A*, then *B*' and for the 'conditional probability of *B* given *A*', respectively. This is relevant in this context, since Lewis (1976) famously showed that the Adams-Stalnaker thesis implies that $P(A \Rightarrow B)$ is equal to $P(B)$ for $P(A) > 0$, which would be hard to swallow in any case. Thesis 2, however, is a minimal requirement, that only assumes that the truth, respectively acceptibility or assertibility, of a conditional $A \Rightarrow B$ implies that its associated conditional probability $P(B\,|\,A)$ is high. Therefore, our representation of generics and conditionals on the basis of thesis 2 is not affected by Lewis' (1976) triviality result.

"Monistic" Bayesianism. Monistic (or "strong") Bayesianism strengthens probabilism, as defined above, inasmuch as it satisfies besides theses 1 and 2 also the following additional assumption:

3. Rational probabilistic reasoning is based on only *one* probability function: one's actual degrees of belief.

Note that in the new Tweety puzzle it is assumed that (A)-(C) describe a subject's beliefs in a synchronic fashion, i.e. beliefs *at the same point in time*. So, according to monistic Bayesianism a belief in (A)-(C) should be describable by a single, actual probability function. Hence, a representation of (A)-(C) by means of more than one probability function in a diachronic way – for example, by means of conditionalization – would not do, since this does not make sure that (A)-(C) can be accepted or believed at the same time.

2.2. The New Tweety Puzzle and Monistic Bayesianism

Let us now focus on probabilistic representations of the Tweety case. Again we distinguish between a singular and a generic representation:

Generic Form:	Singular Form:
(A2g) $p(Fx\,|\,Bx)$ is high	(A2s) $P(Ft\,|\,Bt)$ is high
(B2g) $p(\neg Fx\,|\,Px)$ is high	(B2s) $P(\neg Ft\,|\,Pt)$ is high

(C2) $P(Pt\&Bt)$ is high

Here, 'Fx', 'Bx', 'Py' and 't' abbreviate 'x can fly', 'x is a bird', 'x is a penguin' and 'Tweety', respectively. Note that in the Bayesian reconstruction also the factual (non-conditional) premise (C) has to be explicated in terms of a degree of belief, which may well be equal to 1 in this case. (A2g) and (B2g) are interpreted as assertions of high relative frequencies, whereas (A2s), (B2s) and (C2) refer to conditional degrees of belief in the respective proposition. Note that the generic reconstruction involves two probability functions (p and P). The singular construction involves only one probability function (P) – at least prima facie. So, the weak (dualistic) Bayesian might choose either one representation for (A) and (B) – the generic or else the singular version – whereas a monistic Bayesian has to choose a singular formalization. This is due to the fact that a monistic Bayesian uses only a single subjective probability function, which cannot provide an adequate reconstruction of generic probabilities. The reason for this is a logical one: Generic (or frequency-based) probabilities always refer to a *repeatable* kind of event, or to a property, that is expressed by an *open* formula (Fx), where the free variables (e.g., x) are bound by the probability operator. So, $p(Fx \mid Bx)$ is the frequency (limit) of flying animals among all birds in the domain. Subjective probabilities, on the other hand, always refer to a particular event, state of affairs, or proposition, by being the object of a belief that is expressed by a *closed* formula, i.e. a sentence. Only a proposition, but not a property can be the object of a graded belief. I can believe (to a high degree) that 'this bird can fly', that 'all birds can fly', or that 'most birds can fly', *but not* that 'x can fly, given x is a bird', because the latter statement is an open formula, which leaves open what the *content* of this belief really is.

A monistic Bayesian could try to defend his position in several ways. For example, she might suggest that the generic conditionals should be reconstructed as high degrees of belief in "general propositions". There are two ways to implement this suggestion, both of which are unsatisfactory. The first way would be to reconstruct '(most) Bs are Fs' as (a) 'the degree of belief in the strictly universal proposition 'all Bs are Fs' is high'. This reconstruction is inadequate, because its content differs strongly from the generic (b) '(most) Bs are Fs'. While (a) implies that one believes to a high degree that there are no exceptions, (b) does not entail this at all. (On the contrary people asserting (b) usually assume that there are exceptions). The second way would be to reconstruct '(most) Bs are Fs' as (c) 'the degree of belief in 'most Bs are Fs' is high'. This reconstruction is *not in*adequate, but it does not help the monistic Bayesian out of her problem. For this reconstruction merely repeats the generic probability statement within the scope of a 2nd order probability statement. So this reconstruction involves still *two* notions of probability, generic (at the 1st order level) and subjective (at the 2nd order level).

A further move for the monistic Bayesian would be to understand a high degree of belief in 'x can fly, given x is a bird' as a quantification over subjective probability statements, saying that *for all* x one has a high degree belief in 'x can fly, given x is a bird', or formally $\forall x(P(Fx \mid Bx) = \text{high})$. But this would be inadequate, because we assume that the agent is familiar at least with *some* birds, for example penguins, for which she believes that they cannot fly (cf. Bacchus, 1990, p. 134, on this point). Finally, if one suggests the reconstruction of 'most Bs are Fs' as 'For most x, $P(Fx \mid Bx) =$ high', then one reintroduces frequency-based probabilities at the 2nd level (apart from further inadequacies). In conclusion, there seems to be *no way* a monistic Bayesian could reduce generic probabilities to subjective probabilities without either changing their semantical content or reintroducing generic probabilities in some way or other.

We now turn to the singular form of the Tweety puzzle. The monistic Bayesian will typically express the singular form in a propositional logic framework as follows:

(A3) $P(F \mid B)$ is high

(B3) $P(\neg F \mid P)$ is high

(C3) $P(B\&P)$ is high

The problem for monistic Bayesians is that we cannot represent the belief in (A)-(C) consistently in a monistic Bayesian framework, with probability higher than .71. To show this, let us, first, ask, what the highest admissible value of r is, such that $P(F \mid B) \geq r$, $P(\neg F \mid P) \geq r$, and $P(B\&P) \geq r$. Due to the Kolmogorov axioms of probability theory the following observations hold:

1. $P(F) \geq P(F\&B) = P(F \mid B)P(B) \geq P(F \mid B)P(B\&P)$
2. $P(\neg F) \geq P(\neg F\&P) = P(\neg F \mid P)P(P) \geq P(\neg F \mid P)P(B\&P)$

Note that $P(F) + P(\neg F) = 1$. By 1 and 2 the following holds:

3. $1 = P(F)+P(\neg F) \geq P(F \mid B)P(B\&P)+P(\neg F \mid P)P(B\&P) \geq r\cdot r+r\cdot r \Rightarrow r^2 \leq 1/2 \Rightarrow r \leq 1/\sqrt{2} \approx .71$.

The calculations show that we cannot assign a higher probability value to $P(F \mid B)$, $P(\neg F \mid P)$, and $P(B\&P)$ than .71 on the basis of the Kolmogorov axioms. In conclusion, Bayesians render Tweety type cases as probabilistically inconsistent. For the monistic Bayesian has to represent the Tweety case in one way or another by probability assignments of the form (A3)-(C3). For values higher than $1/\sqrt{2}$, however, she can only do that on pains of probabilistic inconsistency.

One might object that $1/\sqrt{2}$ is a *quite* high value. However, the conditional probabilities involved in exception structures of everyday experience, such as the Tweety case, are assumed to be much higher than $1/\sqrt{2}$, typically close to 1; and nevertheless these exception structures are intuitively regarded as consistent. So, this defense against the new Tweety puzzle does not work.

Moreover, one can construct more complex exception structures, so-called multiple exceptions situations, in which the upper limit of the conditional probabilities in a monistic reconstruction sinks to arbitrary low values. Assume $O_1, \ldots, O_n$ to be a partition of possible outcomes of an experiment, for example a competition with n contestants (named by natural numbers $1, \ldots, n$). Each O_i says that contestant i wins. Let $E_1, \ldots, E_n$ be a set of pieces of evidence concerning the probable winner of the competition, where each E_i says that contestant i and only contestant i has a certain property, which makes it highly probable that i wins. Assume that for each i, $P(O_i \mid E_i)$ is high and that $P(E_1\& \ldots \&E_n)$ is high, where 'high' means 'a probability value $\geq r$'. Then, it is easy to prove that if one assumes the same probability function for conditional and factual probabilities, r cannot be greater than $1/\sqrt{n}$.[2]

[2] Proof: $\sum_{1\leq i\leq n} P(O_i) = 1$. But for each i holds that $P(O_i) \geq P(O_i \mid E_i) \cdot P(E_i) \geq P(O_i \mid E_i) \cdot P(E_1\& \ldots \&E_n) = r^2$. So, $P(O_i) \geq r^2$. Hence $n \cdot r^2 \leq 1$, $r^2 \leq \frac{1}{n}$, which implies our claim.

At this point of our argument, the monistic Bayesian has the following defense: She will deny that the singular and the generic formulations are treated on par. In the singular case, she argues, in which we deal with only one probability function, the evidence that Tweety is a penguin *lowers* our conditional degree of belief that Tweety can fly, given it is a bird. Indeed, this must be true for every coherent degree of belief function, because the following must hold for the agent's degrees of belief at a time point t when she has received evidence that Tweety is both a bird and a penguin: $P_t(Ft \mid Bt) = P_t(Ft \mid Bt\&Pt)$, provided the agent's evidence at time t contains 'Pt', i.e., $P(Pt) = 1$; moreover, because of strict specificity (i.e., $P(Bt \mid Pt) = 1$), it follows that $P_t(Ft \mid Bt) = P_t(Ft \mid Bt\&Pt) = P_t(Ft \mid Pt) = \text{low}$.

So, the monistic Bayesian *has* a solution to the new Tweety puzzle. Accordingly, we do not argue that monistic Bayesianism is an incoherent position. However, the monistic Bayesian's solution to the Tweety puzzle seems to be *intuitively inadequate*. It is counter-intuitive that the probability that a given individual can fly *conditional* on being a bird is changed by new evidence about this individual – at least if this conditional probability intends to reflect the strength of the connection between the two properties "Bird" and "CanFly", which are instantiated by Tweety. That this reconstruction is counter-intuitive is not only *our* philosophical intuition, but is supported by the empirical study that is presented in the next section.

Moreover, a monistic Bayesian reconstruction that sets the conditional probability $P_t(Ft \mid Bt)$ to a low value hides important information, namely that there are two *opposing* conditionals relevant to the consequent (this problem becomes even more unpleasant in the Nixon diamond example; see footnote 1 and Schurz, 2011, §2). That a monistic Bayesian reconstruction hides conflicting probabilistic information is a further argument against monistic Bayesianism. We conclude that an adequate probabilistic reconstruction of the singular conditionals should not interpret (As) and (Bs) as actual subjective degrees of belief, but as *objective single case probabilities*. How this works is explained in section 2.3.

Let us, now, summarize our argument against monistic Bayesianism. We saw that monistic Bayesians are forced to represent Tweety type cases in terms of probability assignments of the form (A3)-(C3), on the basis of thesis 2. Assignments of this form make (A)-(C) jointly probabilistically inconsistent. On the other hand, exception structures, such as the Tweety case, are important in everyday life and science and must, thus, be representable by any fully adequate theory of epistemic states. So, monistic Bayesianism cannot be

---

Note that to prove our generalized new Tweety puzzle it would be even sufficient to assume, instead of $P(E_1\&...\&E_n) \geq r$, that $P(E_i) \geq r$ holds for all i (in the binary case that $P(Pt) \geq r$ and $P(Bt) \geq r$). An empirical investigation of the latter version of the new Tweety puzzle is in progress.

a fully adequate theory of epistemic states, at least not in the way, in which we defined this position.

2.3. Two Possible Solutions to the New Tweety Puzzle

Do weak Bayesians run into the same problem regarding the new Tweety puzzle as the monistic Bayesians? The answer is 'no'. Two intuitively adequate solutions are available to the weak Bayesian. *First* and most straightforwardly, the weak Bayesian may use the generic probabilistic reconstruction. No inconsistency for the Tweety case arises, since two different probability functions can be used, a frequency-based probability function representing (A) and (B) and a second probability function representing actual degrees of beliefs to describe (C) relative to one's background beliefs.

In the first solution the weak Bayesian does not make literal sense of the singular version of the Tweety puzzle. He would have to argue that singular conditional probabilities are merely an "elliptic" way of speaking, in which one is implicitly referring to the existence of underlying generic conditional probabilities. The *second* solution takes the singular version of the Tweety puzzle literally: It makes the "elliptic" way of speaking explicit by means of the *statistical principal principle*, which transfers a frequency-based (statistical) probability to the single case. This principle is one of the most important connection principles between objective and subjective probabilities.[3]

The probabilities that one obtains by transferring statistical probabilities to single cases are (in their standard interpretation) still epistemic, but 'objective' in the sense of being intersubjective; they must, therefore, be different from actual probabilities (this has been pointed out by many authors; cf. Hawthorne, 2005; Schurz, 2011, §3). In what follows we denote these intersubjective epistemic probabilities by means of P_o (with 'o' for 'objective'):

Statistical Principal Principle $P_o(Ga \mid Fa) = p(Gx \mid Fx)$.[4]

The interpretation of "objective" single case probabilities in terms of subjective degrees of belief is established in terms of Pearl's interpretation of these as hypothetical degrees of beliefs (Pearl, 1988, p. 475), or in terms of Carnap's related subjective prior probabilities or "credibilities" (Carnap, 1971, 21-23). The central idea is that $P_o(Ga \mid Fa)$ equals the rational degree

[3] Strevens (2004, p. 370) calls this principle the 'probability coordination principle'. The name 'principal principle' has been coined by Lewis (1980); he suggested this principle not for statistical probabilities, but for objective single case chances.

[4] In the logically general case, instead of Fa one has a closed formula $A(a_1, \ldots, a_n)$, which contains exactly the distinct individual constants $a_1, \ldots, a_n$, and in place of Ga we have a formula B, whose individual constants are among $a_1, \ldots, a_n$. Similarly, in place of Fx one has $A(x_1, \ldots, x_n)$ (each a_i, for $1 \leq i \leq n$, is replaced by a distinct individual variable x_i), and likewise for Gx.

of belief in Ga that one *would* have if *all that one knew* was Fa and the relevant frequentistic probability estimate $p(Gx \mid Fx)$. The actual degrees of belief (given by P) are connected with these frequency-based hypothetical degrees of beliefs (provided by p) by the principle of total evidence that goes back to Carnap (1962, p. 211) and Reichenbach (1949, §72; he named it 'narrowest reference class'):

Principle of Total Evidence $P(Ga) = P_o(Ga \mid Fa)$, provided Fa contains all evidence that is relevant to the conditional probability of Ga.[5]

The class of birds and the class of penguins that are birds are identical. It follows that $p(Fx \mid Bx\&Px) = p(Fx \mid Px)$. So, by the principle of total evidence $P(\neg Ft)$ equals $p(\neg Fx \mid Bx\&Px)$, which is in turn equal to $p(\neg Fx \mid Px)$. Thus, we draw the conclusion that Tweety cannot fly.

3. Empirical Study

3.1. Introduction

In the following section we will describe an empirical questionnaire study, which we conducted to support our argument against monistic Bayesianism. Our aim was to find out whether subjects regard assertions of the type (A)-(C) from section 1, understood as assertions of high probability, as jointly consistent or not. If human beings judge such probabilistic Tweety type statements as jointly consistent, this provides evidence for our claim that monistic Bayesianism is not a fully adequate theory of cognitive belief states because it renders Tweety type statements as jointly inconsistent. If subjects, however, find those type of statements conjointly contradictory, this indicates that monistic Bayesianism adequately describes Tweety type cases from a cognitive viewpoint.

We will restrict ourselves here to those parts of the study that directly bear on Tweety type cases. We do not report rating tasks concerning the Nixon diamond case (items 6 and 8 from appendix A; see also footnote 1) and omit soundness ratings, where the sets of statements of each consistency rating described later served only as premises. We exclude these results here, since their discussion would increase the complexity of our argumentation to a high degree. We, however, saw in section 2 that Tweety type cases are quite sufficient to successfully criticize a monistic Bayesian position.

[5] In a complete Carnapian reconstruction of the total evidence principle, one assumes a prior degree of belief function P_o that one has before acquiring any evidence whatsoever, and considers $P_o(Ga \mid Fa) = r$ as an abbreviation of $P_0(Ga \mid Fa \,\&\, p(Gx \mid Fx) = r)$.

The two target items pertaining to the Tweety case were the following sets of assertions of high probability, formulated in a singular or in a generic way:

TwGen a) Tigers are most probably dangerous.
b) 3 months old tigers are most probably not dangerous.
c) This animal is most probably a 3 months old tiger.

TwSp a) This animal is most probably dangerous, given it is a lion.
b) This animal is most probably not dangerous, given it is a one-month-old lion.
c) This animal is most probably a one-month-old lion.

The expressions 'TwGen' and 'TwSing' stand for 'Tweety type case generic formulation' and 'Tweety type case singular formulation', respectively. TwGen employs a formulation of a) and b) in terms of generics, while TwSp used only propositions pertaining to the same animal ('this animal'). Note that all items were presented in German. It was indicated in the instructions that the German expression translated as 'most probably' pertained to probabilities of 90% or higher. This qualification was included to make certain that subjects understood the probabilistic statements in TwGen and TwSp as being clearly higher than .71, the point value, above which the Tweety type case cannot be described by a single probability function (see section 2).

Subjects were, then, asked to rate for each item whether assertions a), b), and c) are jointly contradictory or not in order to test the assumptions of monistic Bayesianism. Recall from section 2 that monistic Bayesians assume that beliefs must be represented by a single subjective probability function, which describes the agent's actual degrees of belief. Our instructions indicate that the probability of all assertions a), b), and c) have to have probability .9 or higher. So, if subjects rate a)+b)+c) as non-contradictory, monistic Bayesianism must be empirically inadequate, since according to it a)+b)+c) are jointly contradictory, provided a), b), and c) have a probability of .71 or higher. However, if subjects regard a)+b)+c) as contradictory, this would provide support for monistic Bayesianism, insomuch as the representation of human beings' beliefs would agree with the monistic Bayesian's tenets.

Moreover, if subjects regard a)+b)+c) as non-contradictory, by our argumentation in section 2 this presents evidence that human agents do use two types of probability functions to represent their probabilistic beliefs rather than one type of probability function, as the monistic Bayesians assume. If, however, subjects rate both TwGen and TwSp as non-contradictory, this would provide additional evidence for the solutions described in Section 2.3 and against monistic Bayesians, insofar as even in a singular formulation of the Tweety case, subjects were to use two types of probability function rather than one. Our hypothesis is, thus, the following:

HYPOTHESIS 1. *A clear majority rates the Tweety case as being non-contradictory, both in the generic and the singular formulation (i.e., TwGen and TwSp, respectively).*

One might object that the empirical study is only decisive for or against monistic Bayesianism, if human beings are able to draw correct inferences with probabilities. In the heuristics and biases literature (e.g., Tversky & Kahneman, 1983) it is often argued that naïve human beings are prone to violations of the axioms of probability theory w.r.t. their reasoning behavior. It, hence, appears as though human beings might not be able to carry out judgments of probabilistic consistency adequately.

Observe, however, that there is also ample evidence that human beings are quite capable of processing probabilities adequately (e.g., Hertwig & Gigerenzer, 1999; Gigerenzer & Hoffrage, 1995) and that human cognition can be well described by means of probabilistic and Bayesian principles (Chater et al., 2006a; Griffiths et al., 2010). In particular, some of the evidence against human probabilistic competence in the heuristics and biases literature is less conclusive than it might seem at first hand.

For example, the famous experiments of Tversky and Kahneman showed that under specific circumstances subjects rate the probability of a conjunction *A&B* higher than the probability of *A*, which contradicts the axioms of probability calculus (Tversky & Kahneman, 1983).[6] But Hertwig and Gigerenzer (1999, p. 291) found that when probabilities were described in a frequency format (e.g., "How many of the 200 women are bank tellers?") rather than a degree of belief version (e.g., "Rank the following hypotheses according to their probability.") , the conjunction fallacy essentially disappeared. Furthermore, studies in the area of uncertain reasoning (Evans et al., 2003; Oberauer & Wilhelm, 2003) clearly show that subjects' reasoning with conditional probabilities concurs with postulates of probability theory. Moreover, we are not aware of any study, which demonstrates that there exists a strong bias for consistency judgments of probabilistic statements. These types of consistency judgments are quite different from the probabilistic tasks inquired in the heuristics and biases tradition (e.g., Tversky & Kahneman, 1983; see also above).

In conclusion, subjects' competence in probabilistic reasoning is not so bad, after all. Moreover, in the evaluation of our empirical study of the new Tweety puzzle we do not need to assume that subjects are able to master complicated probabilistic calculations; all we need to presuppose is that subjects have a basic understanding of the concept of probabilistic inconsistency. In order to test the latter thesis, we included items, which represent clear-cut

[6] See Crupi, Fitelson, and Tentori (2008) and Cevolani, Crupi, and Festa (2011) for an interesting recent investigation and discussion of the conjunction fallacy from a philosophical perspective.

cases of probabilistically consistent and inconsistent sets of assertions. The clearly inconsistent sets of assertions were the following two items:

InCons1 a) Erika visits most probably an evening class.
b) Carl visits most probably an evening class.
c) Most probably neither Erika visits an evening class nor Carl visits an evening class.

InCons2 a) If Arnold teases Joseph, then Joseph will most probably get irritated.
b) If Joseph is teased by Arnold, then Joseph will most probably not get irritated.
c) Arnold most probably teases Joseph.

The expressions ‘InCons1’ and ‘InCons2’ abbreviate ‘Inconsistency item 1’ and ‘Inconsistency item 2’, respectively. The qualitative logical form of the assertions InCons1.a-c are A, B, and $\neg A \& \neg B$, respectively. InCons2.a-c have the qualitative form $A \Rightarrow B$, $A \Rightarrow \neg B$, and A, respectively. Note that A in InCons2.b is – contrary to A in InCons2.a and InCons2.c – formulated in the passive voice. The clearly consistent sets of assertions in the questionnaire are the following:

Cons1 a) Hans travels most probably by train to Munich.
b) Peter most probably does not travel by train to Munich.

Cons2 a) If Suzy goes shopping, then Suzy is most probably happy.
b) If Lena goes shopping, then Lena is most probably not happy.

The expressions ‘Cons1’ and ‘Cons2’ stand for ‘Consistency item 1’ and ‘Consistency item 2’, respectively. The qualitative form of assertions Cons1.a and Cons1.b are A and $\neg B$, respectively. Assertions Cons2.a and Cons2.b have the qualitative forms $A \Rightarrow B$ and $C \Rightarrow D$, respectively. Our second hypothesis is as follows:

HYPOTHESIS 2. *The subjects have a basic probabilistic competence, insofar they rate clear cases of probabilistically contradictory sets of assertions (i.e., InCons1 and InCons2) as being contradictory and clear cases of non-contradictory sets of assertions (i.e., Cons1 and Cons2) as being non-contradictory.*

3.2. Method

3.2.1. *Participants*

Twenty-seven persons (59% female) participated in the study. We required all participants to be between 18 and 50 years of age, to be native speakers of German and not to have participated in a study we conducted earlier.[7] The average age was 25.67 years (the standard deviation was 6.4 years). 96% were students, 82% of these from the humanities. The remaining 18% of students came from the sciences, medicine, management, and art. Only two psychology students and four philosophy students participated in the questionnaire study. 40.7% and 55.6% of the subjects indicated that they had in some form or other exposure to logic and statistics or probability theory (in high school, at the university, etc.), respectively.

3.2.2. *Test Material*

The questionnaire consisted of two parts. The first part included the consistency-ratings described in the previous section and the second part a probability estimation task (see previous section). All testing material was administered in German. We will report here only the consistency ratings of the first part of the questionnaire. In the consistency rating task subjects were instructed to answer all items according to their intuitive understanding. We told subjects in the instructions that the phrase 'most probably' (German: 'höchstwahrscheinlich'), used throughout all items, implies a probability of 90% or higher. Subjects were, then, asked to tell for TwGen, TwSp, Incons1 and Incons2 whether assertions a)+b)+c) [or alternatively a)+b) for items Cons1 and Cons2] were contradictory, viz. whether they contradict each other. The exact formulation of the question for item Nr. X was: 'Xa), Xb) and Xc) are conjointly: □ non-contradictory □ contradictory.' The complete list of items administered in the consistency rating task can be found – in the order of presentation – in Appendix A.

3.2.3. *Procedure*

All questionnaires were administered in terms of single person testing by a psychologist. All subjects answered the questionnaire in a cubicle located in a quiet office and received 5 euros for their participation in the study.

[7] Despite our participation requirements one subject was not a native speaker of German, but spoke German fluently. We decided to include the participant in the study.

Table I. Percentages of Inconsistency Ratings for Item Types

Item Type	Percentage of Rating as 'Inconsistent'
Cons1+Cons2	9.3%
InCons1+InCons2	83.3% (96.3% without the Arnold item)
TwGen	11.1%
TwSing	22.2%

3.3. Results and Discussion

Percentages of mean ratings for all item types can be found in Table I.[8] The data clearly confirms our hypothesis 2. Subjects were able to categorize both, clearly probabilistically contradictory sets of assertions (83.3%) and non-contradictory sets of assertions (90.7%), correctly. We interpret this result as indicating that human beings have a basic understanding of the concept of probabilistic consistency and inconsistency and are in this context not prone to violations of the axioms of probability theory, as, for example, described by the conjunction fallacy (see section 3.1). Note that the percentage of inconsistency ratings for clearly inconsistent items was higher (96.3%) when we did not include the Arnold example (item InCons2). This was expected to some degree, since the antecedent was used both in the active and the passive voice. Despite this fact, the pooled percentage of correct inconsistency ratings for InCons1 and InCons2 was quite high (83.3%).

Hypothesis 1 was also clearly confirmed by our data. The great majority of subjects rated both, a singular formulation of the Tweety case (TwSing) and a generic formulation (TwGen) as probabilistically consistent (88.9% and 77.8%, respectively). So, since monistic Bayesianism renders both singular formulations and generic formulations of Tweety type cases as probabilistically inconsistent (see section 2), this study provides clear evidence against monistic Bayesianism as a fully adequate theory of epistemic states.

Furthermore, our discussion in section 2 suggests that the second probability function – besides a probability function interpreted in terms of degrees of belief – is, as Schurz (2011) suggests, best reconstructed as being frequency-based. It is interesting to observe that the singular formulations of (A) and (B) in the item TwSing did not seem to have an effect on the consistency ratings to any significant degree. This result suggests that the singular formulations of (A) and (B) were essentially read as objective single case probabilities backed up by frequencies rather than actual degrees of belief. So, the consistency/inconsistency ratings provide evidence that the human mind works

[8] We ran a separate analysis for subjects, who indicated that they had some exposure to logic. The pattern of results for this subgroup did not differ from the results for the total sample.

with these two types of probability functions, without making this necessarily explicit.

Note, however, that we do not argue that the present study is decisive. Despite the clear pattern observed in our sample, a replication of our results is needed, best with a variety of items. We, however, believe that this study provides a first confirmatory result for the inadequacy of monistic Bayesianism as a fully adequate theory of epistemic states.

4. The Tweety Problem in Cognitive Science

In this section we will, first, survey Bayesian approaches in cognitive science and describe their key tenets. As we will see, the cognitive scientists endorsing a Bayesian approach almost universally employ a type of monistic Bayesianism, which aims to rely solely on a degree of belief interpretation of probabilities. We will, first, criticize this type of approach by discussing the new Tweety puzzle in the context of these approaches. We will, then, take a closer look at Oaksford and Chater's models of conditional inference, and argue – although the authors are clear proponents of a monistic brand of Bayesianism – that their model eventually has to refer to frequencies in one form or another. The latter discussion is intended as a case study, which aims to show that also in an applied context in cognitive science, a monistic Bayesian approach is not viable.

4.1. Bayesian Approaches in Cognitive Science

In recent years Bayesian approaches in cognitive science became more and more prominent. For example, a special issue in Trends in Cognitive Sciences (2006, Vol. 10, Issue No. 7) was specifically dedicated to Bayesian and probabilistic approaches in cognitive science. Furthermore, a range of papers applied Bayesian ideas and methods to topics, such as knowledge representation (Griffiths et al., 2010; Chater et al., 2006a), language processing and acquisition (Xu & Tenenbaum, 2007; Chater & Manning, 2006), vision (Yuille & Kersten, 2006), inductive learning and reasoning (Kemp & Tenenbaum, 2009; Tenenbaum, Griffiths, & Kemp, 2006; Oaksford & Chater, 2007, 2008), sensorimotor control (Körding & Wolpert, 2006), and memory (Griffiths, Steyvers, & Tenenbaum, 2007; Steyvers, Griffiths, & Dennis, 2006).

The cognitive scientists endorsing the Bayesian framework are very explicit about their own theoretical position and seem to share the idea that only subjective probability functions are needed in their approach, much like the monistic Bayesians criticized in section 2. So, Chater et al. (2006a, p. 288f) and Chater, Tenenbaum, and Yuille (2006b, p. 292) – in their introduction and programmatic outlook of the special issue described in the previous paragraph – clearly state that they intend to use a subjective rather than a frequentistic interpretation of probabilities. Moreover, Oaksford and Chater (2007, p. 10f; 2009, p. 69) also explicitly endorse this assumption in their Bayesian approach in the area of reasoning.

Let us now see at which theoretical level Bayesian cognitive scientists locate their monistic Bayesian principles. Oaksford and Chater (2007), for example, do not presuppose that human cognition actually carries out probability calculations, but rather assume that Bayesian assumptions allow one to characterize human behavior in terms of the rational problems that human cognition has to solve (cf. Anderson's 1990 program of rational analysis). In other words Bayesian cognitive scientists, such as Oaksford and Chater (2007), use Bayesian principles to describe, which type of problems human cognition aims to solve. The human mind uses, then, heuristics to approximate these normatively correct solutions (Oaksford & Chater, 2007, p. 14f).

Note that the rational analysis, endorsed by Oaksford and Chater (2007), pertains in Marr's (1982) terminology to the computational level (Oaksford & Chater, 2007, p. 46) and differs from approaches at the representational/algorithmic and the neuronal-implementational level. Accounts at the latter levels aim to answer the questions 'what is represented in human cognition' and 'which brain mechanisms underlie these phenomena', respectively. In contrast, connectionist accounts, such as Rogers and McClelland (2004), do not start with assumptions regarding the computational/rational analysis level, but rather the representational/algorithmic level, by making assumptions regarding the neuronal implementation of human cognition, as described by connectionist models.

Observe further that Oaksford and Chater (2007) explicitly reject deductive-logical approaches for a rational level analysis and argue at length (pp. 41–65) against any deductive-logical approach, including default logic approaches (pp. 60–62), as, for example, Reiter (1980). Oaksford and Chater's alternative is a Bayesian approach, as described above. The Bayesian cognitive scientists also explicitly indicate that Marr's computational level is best understood in terms of probabilistic approaches, such as their Bayesian approach (Griffiths et al., 2010, p. 363; cf. also Chater et al., 2006a, p. 289f).

4.2. Criticism of monistic Bayesian approaches in Cognitive Science

Unlike other critics, such as McClelland et al. (2010), we do not argue against the general approach of the Bayesian cognitive scientists, in terms of an analysis of cognition at a computational level. We, rather, endorse a pluralistic view regarding investigations at all Marrian levels, including the computational and the representational/algorithmic level. In particular, we regard analyses of cognition at the computational level, as for example done by the Bayesian cognitive scientists, as being fruitful and complementary to analyses at the representational/algorithmic level, such as connectionist approaches (e.g., Rogers & McClelland, 2004).

We, furthermore, do not reject a probabilistic analysis at the computational level in favor of an alternative analysis, such as default logic approaches (e.g., Reiter, 1980). Although we do not agree with much of Oaksford and Chater's (2007) criticism of deductive-logical approaches and default logic approaches, we will not argue here – in principle – against probabilistic analyses of cognition. On the contrary, we share and appreciate the idea that Bayesian methods and ideas from mathematics and epistemology can and should be fruitfully applied in cognitive science.

We, however, criticize the particular probabilistic approach employed by Oaksford and Chater (2007) and other Bayesian cognitive scientists (Chater et al., 2006a; Griffiths et al., 2010), namely a monistic type Bayesian approach. In particular, we regard the thesis that a probabilistic theory on the computational level can solely rely on a subjective interpretation of probabilities and still be fully adequate as not tenable. We rather suggest that Bayesian cognitive scientists should weaken their monistic standpoint and explicitly incorporate frequency-based probability functions besides subjectively interpreted probability functions, also *at the computational level*.

To argue against the viability of a monistic Bayesian approach in cognitive science we will, first, strengthen our general argument against monistic type Bayesianism, based on the new Tweety puzzle. We will, then, focus on Oaksford and Chater's (2007, 2008) Bayesian model of conditional inference as a cases study. We saw earlier, that Oaksford and Chater (2007, p. 10) also endorse an account that only draws on subjectively interpreted probabilities. We will argue here that their model of conditional inference implicitly and non-trivially relies also on a frequentistically interpreted probability function. Our argumentation aims to show that also in an applied context in cognitive science a monistic type Bayesian approach is not viable.

Note that Oaksford and Chater (2008) admit that their 2007 model of conditional inference implicitly uses in addition a frequency-based probability function and propose a modification, which does not rely on relative frequencies explicitly. We will aim to demonstrate that neither their 2007 model nor their 2008 version can be adequate without employing a frequency-based probability function. We will also draw out in more detail, in which way both Oaksford and Chater's 2007 and their 2008 model refer to a frequency-based probability function, since in Oaksford and Chater (2008) it is not explicitly described, in which way their 2007 model relies on the notion of relative frequencies.[9]

Let us, now, apply our argumentation from section 2 to the monistic type Bayesian approaches in cognitive science, as described in section 4.1. We saw in section 2 that exception structures, such as the Tweety case, cannot adequately be represented by monistic Bayesian approaches. But what are Tweety type cases? They are simply cases, in which an exceptional subclass, such as penguins, of a class, such as birds, is instantiated. Now, as we saw in section 1, it is one of the motivations of probabilistic approaches that these are more flexible than, for example, classical logic in allowing for exceptional cases and classes. Griffiths et al. (2010) defends the project of monistic type Bayesianism, as described in section 4.1, against criticism by Rogers and McClelland (2004) who argue that in general symbolic models cannot handle exceptions adequately:

> "Connectionists have criticized symbolic models for failing to handle exceptions or produce graded generalizations, or to account for how representations are learned [...] Combining structured representations with probabilistic inference meets those challenges" (Griffiths et al., 2010, p. 360)

Thus, Griffiths et al. (2010) argue that one can, when combining structured representations with probabilistic inferences, also account for handling exceptions. We, however, saw in section 2, that we cannot do so on the basis of a monistic type Bayesian approach. For, whenever we use an instantiated exceptional subclass, we cannot describe this scenario adequately by means of a single subjective probability function only, but must also refer to a frequency-based probability function.

Note that Griffiths et al. (2010, p. 359) argue that monistic type Bayesians can account for exception structures. For that purpose they suggest a procedure, by which assignments of relative frequencies, such as (A2g), are replaced by assignments of subjective probabilities to the corresponding universal conditional, such as (A1g) (see section 2). Note, however, that there are strong reasons that speak against pursuing such a strategy (see section 2 for a discussion of these points): First, this contradicts thesis (2) of monistic Bayesianism. Although Griffiths et al. (2010) do not explicitly endorse this principle, their procedure amounts to a rejection of the thesis that conditionals should be probabilistically represented by conditional probabilities, which is for many Bayesians hard to swallow. Second, the logical properties of these new assignments differ from those of the frequency-based versions (see also section 2.2).

[9] Also Oaksford and Chater (2009) describe a model of conditional inference. Since it is very close to their 2007 version, not to say identical, we will rather focus on Oaksford and Chater's (2007) model.

We conclude here that Griffiths et al.'s approach is not viable. We suggest that instead of sticking to monistic tenets, Bayesians in cognitive science should rather endorse a weaker form of Bayesianism, which – unlike monistic Bayesianism – allows for the specification of instantiated exceptions.

4.3. Case Study: Oaksford and Chater's (2007, 2008) Model of Conditional Inference

4.3.1. *Oaksford and Chater's (2007) Model of Conditional Inference*

Before we describe Oaksford and Chater's model of conditional inference, let us, first, specify what they mean by 'conditional inferences'. The conditional inferences, on which Oaksford and Chater's model focuses, are of the following four types, which have a long tradition in the psychology of reasoning (Oaksford & Chater, 2007, p. 113, p. 100; Evans, 1982, Ch. 8):

From A and $A \Rightarrow B$ conclude B	(Modus Ponens, short: MP)
From $\neg A$ and $A \Rightarrow B$ conclude $\neg B$	(Denying of the Antecedent, DA)
From B and $A \Rightarrow B$ conclude A	(Affirming of the Consequent, AC)
From $\neg B$ and $A \Rightarrow B$ conclude $\neg A$	(Modus Tollens, MT)

Here '$A \Rightarrow B$' stands for the conditional 'If A then B'. The formulas in 'From ...' and 'conclude ...' are called 'premises' and 'conclusion', respectively. We also distinguish between conditional and categorical premises. The categorical premises in the MP, DA, AC, and MT inferences are A, $\neg A$, B, and $\neg B$, respectively.[10]

Oaksford and Chater (2007) propose a Bayesian framework in order to describe inferences of type MP, DA, AC, and MT. For this purpose Oaksford and Chater (2007) use their 2000 model of conditional inference (Oaksford et al., 2000) as a starting point. In the latter model, which they also call 'conditional probability model' (see Oaksford & Chater, 2007, p. 121), conditional inferences of type MP, DA, AC, and MT are described by an "update [procedure of] their beliefs about the conclusion by using the categorical premise to conditionalize on the relevant conditional probability" (Oaksford & Chater, 2007, p. 121). In other words MP, DA, AC, and MT inferences are described by conditionalization on the respective categorical premise.

Oaksford and Chater (2007, p. 120f), then, specify the relevant conditional probabilities for the conditional probability model the following way, where a, b, and c are specified by $a = P_0(B|A)$, $b = P_0(\neg B)$ and $c = P_0(A)$, respectively (Oaksford & Chater, 2007, p. 119f):

[10] Oaksford and Chater (2007, p. 122) give $\neg p$, as an example for a categorical premise of the inference DA, which is specified as 'from $p \Rightarrow q$ and $\neg p$ conclude q' in that context (Equation 5.14, p. 119).

Cond MP	$P_1(B) = P_0(B \mid A) = a$	for $c > 0$
Cond DA:	$P_1(\neg B) = P_0(\neg B \mid \neg A) = \frac{b-(1-a)c}{1-c}$	for $c < 1$
Cond AC:	$P_1(A) = P_0(A \mid B) = \frac{ac}{1-b}$	for $b < 1$
Cond MT:	$P_1(\neg A) = P_0(\neg A \mid \neg B) = \frac{b-(1-a)c}{b}$	for $b > 0$

So, in their 2000 model, Oaksford et al. presuppose that the probability of the conclusion is determined by conditionalizing on the categorical premise with probability 1, viz. the probability of a conclusion $P_1(C)$ is specified by assuming the categorical premise with certainty (see Oaksford & Chater, 2007, p. 120f).

Oaksford and Chater (2007) argue that we do not need to look at generalizations of Cond MP, DA, AC, and MT, where the probability of the categorical premise is smaller than 1, since "[f]or almost all the data we look at later on, the generalization [...] [in terms of the probability of the categorical premise not equaling 1] is not needed" (Oaksford & Chater, 2007, p. 121). We will follow here Oaksford and Chater's suggestion and focus on conditionalization in line with Cond MP, DA, AC, and MT as a starting point.

While the Oaksford et al. (2000) model provided overall good fit with the empirical data (Oaksford & Chater, 2007, p. 126), this model did not perform particularly well w.r.t. the MP-MT asymmetry (2007, p. 126; 2008, p. 100). By 'MP-MT asymmetry' we refer to the unequivocal finding in the area of conditional inferences that human beings draw more MP inferences than MT inferences (Oaksford & Chater, 2008, p. 98f; see also Evans & Over, 2004, p. 47). In a model fitting exercise Schroyens and Schaeken (2003) compared different models of conditional inference and noted that Oaksford et al.'s (2000) model underestimated the probability of MP endorsements and overestimated the probability of MT inferences (Oaksford & Chater, 2007, p. 126).

To account for this MP-MT asymmetry in a better way, Oaksford and Chater suggested a modified account, as described in the 2007 and 2008 versions of their revised model. At the heart of the revised model lies the assumption that the so-called rigidity condition is violated for MT inferences.[11] The rigidity condition specifies that the conditional probability $P_0(B \mid A)$ remains stable in the course of conditionalizing from P_0 to P_1 for all conditional inferences, as specified by Cond MP, DA, AC, and MT (Oaksford & Chater, 2007, p. 127). According to Oaksford and Chater's (2007, 2008) revised model it is highly probable that a violation of the rigidity condition occurs for MT inferences, but not so for MP inferences (Oaksford & Chater, 2007, p. 127; 2008, p. 108).

[11] Oaksford and Chater suggest that DA and AC inferences are also prone to rigidity violations (Oaksford & Chater, 2008, p. 110). We will, however, restrict our discussion here to rigidity violations for MT inferences.

Oaksford and Chater (2007, p. 127f; 2008, p. 108) explain why agents violate the rigidity condition for MT inferences. They do so by means of the pragmatics of item presentation. Oaksford and Chater use the following example to illustrate their point (2007, p. 127f; 2008, p. 108):

(D) If you turn the key (p), the car starts (q).

(E) The car did not start ($\neg q$).

Suppose you belief (D) and on that occasion also (E). Then, "[t]here would seem to be little reason to expect the car to start unless one was reasonably confident that the key had been turned" (Oaksford & Chater, 2007, p. 127). Furthermore, "the assertion of the categorical premise of MT only seems to be informative against a background where the car was expected to start [...] So this seems like a case where rigidity might be violated, i.e. it is a counter-example, and so $P_0(q\,|\,p)$ needs to be adjusted" (Oaksford & Chater, 2007, p. 127f).[12]

Before we go on with the description and criticism of the Oaksford and Chater 2007 model of conditional inference, let us, first, clarify where our general punch line is. We do not, on general grounds, argue against Oaksford and Chater's model of conditional inference. We rather believe that their model, based on rigidity violations to be explained below, represents a plausible and promising approach. We, however, believe that Oaksford and Chater's restriction to a monistic Bayesian approach, which only allows for subjective probability functions, limits the full potential of their model and we argue that their model has to acknowledge in one way or another a probability function which is interpreted in terms of relative frequencies, in addition to a subjective probability function.

Oaksford and Chater argue that in rigidity violations the original value $P_0(q\,|\,p)$ is revised and lowered to the new value $P_0^R(q\,|\,p)$ ('R' for 'revised'; Oaksford & Chater, 2008, footnote 6, p. 115; 2008, p. 128). The revised value $P_0^R(q\,|\,p)$ is, then, used to calculate the new probability function P_1, as described by Cond MT. Why, however, does in the Oaksford and Chater (2007) model the counter-example imply a lower value $P_0^R(q\,|\,p)$ compared to $P_0(q\,|\,p)$ rather than an equal or a higher value?[13]

[12] Note that Oaksford and Chater (2008, p. 108) describe the case with almost the same words.

[13] Oaksford and Chater (2007, p. 128) say that the value of $P_0(q\,|\,p)$ can be determined on the basis of a dependence-independence model comparison, in which a conditional independence model ($P_0(q\,|\,p) = P_0(q)$) is compared with a conditional dependence model ($P_0(q\,|\,p) \neq P_0(q)$). However, the authors do not explain how this model comparison is used to achieve lower values for $P_0^R(q\,|\,p)$, but only vaguely sketch it. Moreover, in their empirical evaluation of their revised model Oaksford and Chater (2007, p. 130) do not use any values predicted by the dependence-independence model, but rather estimated $P_0^R(q\,|\,p)$ by their best fit values when testing the fit of their model with the data. In Oaksford and Chater (2008)

Let us inquire how it is that Oaksford and Chater (2007) might account for a lower probability of $P_0^R(q \mid p)$. For that purpose, we shall take a closer look at the example, which was used in both the 2007 and the 2008 version of their model. In this example, an agent is described as believing (D) ('if you turn the key, the car starts') and (E) ('the car did not start'). According to Oaksford and Chater's (2007, 2008) model, agents make the following assumption, on pragmatic grounds of conversational relevance:

(F) You turned the key (p).

In Oaksford and Chater's model, the fact that you turned the key and the car did not start – that is (F)+(E) which is represented by $p\&\neg q$ – is a single counter-example (cf. Oaksford & Chater, 2007, p. 129f). Oaksford and Chater (2007) argue, then, that "participants update on the evidence of a single $p, \neg q$ counter-example" (Oaksford & Chater, 2007, p. 130).[14] But in which sense can this counter-example decrease $P_0^R(q \mid p)$ in comparison to $P_0(q \mid p)$?

The problem with Oaksford and Chater's (2007, 2008) formalization of (D) is that the probability $P_0(q \mid p)$ is not just lowered to some degree by considering the "single counter-example" $p\&\neg q$, but receives the value *zero*, if we apply Oaksford and Chater's approach. On that basis, however, one can hardly argue that $p\&\neg q$ is just a single counter-example that lowers the probability $P_0(q \mid p)$ to some degree, to the value $P_0^R(q \mid p)$. Rather $p\&\neg q$ is a "falsifying" counter-example that lowers $P_0^R(q \mid p)$ to zero.

Let us draw this out in more detail: The probability $P_1(p\&\neg q)$ is, as we saw in section 4.2, interpreted as the degree of belief in the proposition $p\&\neg q$, namely the probability that on a specific occasion you turned the key, but the car did not start. However, in order for $p\&\neg q$ to serve as a counter-example, the agent has to be certain or nearly certain that $p\&\neg q$ is in fact the case. Given that one is certain that $p\&\neg q$ is a counter-example, viz. $P_1(p\&\neg q) = 1$, it follows that $P_1(q) = 0$, and by Cond MT this implies that $P_1(q) = P_0^R(q \mid p) = 0$. So, the single counter-example $p\&\neg q$ fully determines $P_0^R(q \mid p)$ as having probability 0, without having to refer to any other positive $p\&q$ or negative $p\&\neg q$ instances.

What went amiss in Oaksford and Chater's (2007) model? The problem is that the probability associated with (D) cannot be sensibly specified by $P(q \mid p)$, when the counter-example has the form $p\&\neg q$. It rather seems that (D) is understood as a generic conditional, not a conditional, which only refers to a specific occasion. The representation of the probability associated

the authors do not comment on this idea any further; instead they argue that they employed in their 2007 version the notion of relative frequencies, without saying explicitly how and why (Oaksford & Chater, 2008, p. 108). In conclusion, the dependence-independence model-comparison *does not answer* the question why Oaksford and Chater's model implies a lower value for $P_0^R(q \mid p)$ compared to $P_0(q \mid p)$.

[14] Note here that '$p, \neg q$' in Oaksford and Chater (2007, cf. p. 130) correspond in our terminology to '$p\&\neg q$'.

with (D) by an assignment of the form $P(q \mid p) = r$, however, implies that this conditional probability only refers to the very same occasion described by the counter-example $p \& \neg q$. Note that also Oaksford and Chater (2007, p. 127) seem to imply that the counter-example, in contrast to (D), refers to a specific occasion (see also Oaksford & Chater, 2008, p. 108).

Since (D) is understood as a generic conditional, it seems, hence, best to associate it with a generic probability assignment of the following sort:

(G) $p(Sx \mid Tx) = r$

Here 'Sx' and 'Tx' stand for 'the car starts on occasion x' and 'you turn the key on occasion x', respectively. (G) does not refer to an assignment of a state of affairs relative to another state of affairs, as a degree of belief interpretation of probabilities suggests. Rather (G) specifies the relative frequency of individuals having property S among those individuals, which have property T (cf. example (A2g)). In the case of (D) both Tx and Sx range over individual occasions, on which the key has been turned, and the car has started, respectively.

By this interpretation of (D) we can also explain why an instance, on which the key was turned but the car did not start, is a single counter-example for (D). Let us assume that among four occasions before time t_0 the car always started when the key was turned. Then, the probability associated with (D) at time t_0, namely $p_0(Sx \mid Tx)$, as estimated from the observed sample frequencies, equals $\frac{4}{4}$ and, thus, 1. Suppose we encounter now on this occasion o_1, the new, pragmatically implied counter-example, which amounts to the key being turned (To_1), but the car not starting ($\neg So_1$). Then, the new counter-example $To_1 \& \neg So_1$ reduces the estimated frequentistic probability from $p_0(Sx \mid Tx)$ to $p_0^R(Sx \mid Tx)$, which equals $\frac{4}{5}$, taking into account the new counter-example.

Our frequentistic interpretation of (D), hence, provides a clear explanation of why the probability $p_0(Sx \mid Tx)$ is lowered to $p_0^R(Sx \mid Tx)$ on learning about the counter-example, while the probability value associated with (D) need not be zero. For the subjective Bayesian, however, there seems to be no way out of the problem described here. She can, hence, not explain why the pragmatically implied counter-example is just a single counter-example for (D) rather than a falsification of (D). Hence, the monistic Bayesian cannot explain why the revised probability associated with (D) is lowered by the single counter-example. Rather, the very concept of a counter-example seems to imply that (D) must be generically understood, as described by (D1). However, a monistic Bayesian cannot achieve this. Once again we see that subjectively interpreted probability functions cannot express conditionals or generics, such as (D), in a generic, rather than singular way (see also section 2).

4.3.2. *Oaksford and Chater's (2008) Model of Conditional Inference*

In this section we will describe Oaksford and Chater's (2008) modified approach. We will, then, discuss whether their modified model is describable solely in terms subjective probability functions.

In 2008 Oaksford and Chater present a revised version of their 2007 model. As in their 2007 model, Oaksford and Chater (2008) still employ the idea of a rigidity violation for MT inferences based on a pragmatically implied counter-example, as described above. We saw earlier, Oaksford and Chater (2008) even use almost the same wording as in Oaksford and Chater (2007). The authors, however, go on to argue that a frequentistic interpretation seems implausible. To make a case for this point they discuss the following example, describing a conditional promise, as given by the agent John (p. 109):

(G) If it is sunny tomorrow, then I will play tennis.

Oaksford and Chater (2008, p. 109) suppose that (G) relates to a specific occasion and argue about this case the following way:

> "Under these circumstances, it is nonsensical to suggest that the truth of this claim could be assessed by looking at the relative frequencies of whether John plays or does not play tennis when it is sunny on, say the 14th February 2007 (tomorrow[']s date as we write). Rather we must look to other sorts of evidence that bear on this claim derived from world knowledge. For example, is John generally reliable when making promises, what is the chance that something will prevent him from playing tennis even if it is sunny?" (Oaksford & Chater, 2008, p. 109).

In this way Oaksford and Chater (2008) argue that not for all types of probabilistic conditionals a frequency-based approach might be possible and, thus, make a case for a specification of belief updating based on a purely subjective probability function.

Let us, first, discuss the conditional promise example (G) before we move on to the description and discussion of the second part of Oaksford and Chater's modified model. We think that this example is unconvincing for two reasons. First, Oaksford and Chater's original model seems to be targeted for conditional inferences with "normal" conditionals rather than conditional promises. In Oaksford and Chater (2007, 2008) the premier example, namely (D), is not a conditional promise but rather a "normal" conditional in the sense of conditional assertions, such as (D). Given that framing conditionals as promises, threats, tips and warnings has a strong impact on endorsement rates of MP, DA, AC, and MT inferences (Evans & Twyman-Musgrove, 1998; Evans & Over, 2004, p. 108f), it is not clear, which implications a discussion of (G) has for conditional assertions without a promise modifier.

Second, to make their point Oaksford and Chater (2008) argue that one cannot sensibly interpret (G) – understood as pertaining to a specific point in time and space – in terms of frequencies and has to refer to specific evidence, such as the promising person's (John's) reliability regarding keeping promises. While we accept that not all conditionals can be interpreted in terms of frequencies, such as for example conditional (a) of TwSp (see section 3.1), we do not accept that all the evidence can be specified purely by a subjective probability function. For example, the evaluation of John's reliability of keeping promises and, hence, the determination of the probability associated with (G), eventually has to refer to relative frequencies, in terms of how often he has kept promises in the past. One cannot escape relative frequencies when specifying a model of conditional inference. To ignore this aspect, seems to set aside important information, as specified by probability assignments interpreted in terms of relative frequencies.

We shall now continue with our description and discussion of the modification of Oaksford and Chater's (2007) earlier model. Oaksford and Chater (2008) go on to suggest – in contrast to their 2007 model – that "[t]he classes of evidence we identified can not be regarded as individual pieces of evidence that can be used to update $P_0(q \,|\, p)$ to $P_0^R(q \,|\, p)$ by Bayesian revision as proposed by Oaksford and Chater (2007). Rather they must directly influence $P_0(p, \neg q)$ and $P_0(\neg p, q)$, while the marginals [that is $P_0(p)$ and $P_0(q)$] remained fixed" (Oaksford & Chater, 2008, p. 109). Note here that '$p, \neg q$' in Oaksford and Chater (2008) correspond in our terminology to '$p\&\neg q$'.

In a nutshell, in Oaksford and Chater's (2008) new model the counter-example $p\&\neg q$ influences $P_0^R(p\&\neg q)$ and $P_0^R(\neg p\&q)$ instead of $P_0^R(q \,|\, p)$ directly. The problem from Oaksford and Chater (2007) resurfaces in Oaksford and Chater (2008). Oaksford and Chater (2008) still have to answer why the revised probability $P_0^R(q \,|\, p)$ becomes smaller compared to the original value $P_0(q \,|\, p)$, when we take the counter-example $p\&\neg q$ into account.

If we accept the assumptions made by Oaksford and Chater (2007, 2008), then – as our discussion from section 4.3.1 suggests – in order for $p\&\neg q$ to qualify as a counter-example the agent must be certain that $p\&\neg q$, in other words that $P_1(p\&\neg q) = 1$. It follows that $P_1(p) = 1$ and $P_1(q) = 0$. On the basis of Oaksford and Chater's assumptions concerning Cond MT it holds that $P_1(\neg q) = P_0^R(\neg q \,|\, \neg p)$ and $P_1(q) = P_0^R(q \,|\, p)$. Since $P_1(q) = 0$ and $P_1(p) = 1$, this implies that $P_0^R(q \,|\, p)$ and $P_0^R(\neg p \,|\, \neg q)$ equal zero and one, respectively. So, Oaksford and Chater's (2008) modified model in fact cannot explain, why the revised probability $P_0^R(q \,|\, p)$ becomes smaller to some degree compared to its original estimate. It rather treats $p\&\neg q$ once more as a *falsifying* counter-example.

Again, no such problem arises, if we assume that (D) is interpreted in terms of a frequency-based probability function, as described by (D1). Then, the counter-example adds to the negative instances of the respective frequency-based probability function, as specified in section 4.3.1. We, thus, conclude that the model of Oaksford and Chater, as specified in Oaksford and Chater

(2007) and in Oaksford and Chater (2008), essentially has to refer to a frequency-based probability function in one way or another.

Let us make some final remarks regarding the change of frequencies when accounting for MT inferences, as suggested by Oaksford and Chater in their 2007 and 2008 model. Observe that there are empirical studies, which suggest that human beings use relative frequencies of counter-examples to estimate the probability of conditionals (Geiger & Oberauer, 2007), which supports the model of modification of frequencies by singular counter-examples, as suggested in the previous section. There exists, however, also evidence that under some circumstances human beings use classes of counter-examples rather than frequencies to estimate the probability of conditionals (Markovits, Forgues, & Brunet, 2010). Observe that the estimation of probabilities of conditionals by means of classes of counter-examples also aims to predict the prevalence or frequency of certain attributes. So, even if Markovits et al.'s (2010) finding turns out to provide the more accurate account, this does not speak against our main point, namely that in order to adequately describe Oaksford and Chater's model, one has to refer to relative frequencies in one way or another.

5. Conclusion

We, finally, give an overview over the topics discussed in this paper. The first part of our paper (sections 1 and 2) focused on the new Tweety puzzle. We argued on that basis that monistic Bayesians – Bayesians who restrict themselves to a single "actual" probability function for the specification of epistemic states – cannot account for Tweety type cases adequately. Since Tweety type cases represent an important class of situations, which an agent should be able to represent in terms of beliefs, we argue that monistic Bayesianism cannot be a fully adequate theory of epistemic states.

In the second part (section 3) of our paper we described an empirical study, which provides initial empirical support for the thesis that human beings regard Tweety type cases as probabilistically consistent. Given our theoretical argument in the first part of the paper, this result suggests that in order to account for this empirical finding we have to rely on both a subjective probability function and an objective frequency-based probability function to describe human beings' epistemic states.

In the third part of the paper (section 4) we discussed monistic type Bayesian approaches in cognitive science. We provide a general argument against monistic type Bayesian accounts in cognitive science by applying the new Tweety puzzle to these approaches. We, then, focus on a model of this brand, put forth by Oaksford and Chater (2007, 2008), namely their model of conditional inference. We argue that their model has to rely on a frequency-based probability function, in order to explain why the probability associated with a conditional should be lowered for modus tollens inferences, but not for modus ponens inferences. Since the probability functions in Oaksford and Chater's model are used to model (probabilistic) beliefs of agents, our argumentation shows that we need in Oaksford and Chater's model of conditional inference in addition to a subjective probability function also a frequency-based probability function to adequately represent agents' epistemic states.

Appendix

A. Items Used in the Empirical Study

The following items are English translations from their German originals used in our empirical study. Subjects were asked to rate whether the items a)+b)+c) or else a)+b) were jointly contradictory or not. The items are listed in order of administration. For more details see section 3.

1. a) Erika visits most probably an evening class.
   b) Carl visits most probably an evening class.
   c) Most probably neither Erika visits an evening class nor Carl visits an evening class.
2. a) This animal is most probably dangerous, given it is a lion.
   b) This animal is most probably not dangerous, given it is a one-month-old lion.
   c) This animal is most probably a one-month-old lion.
3. a) Hans travels most probably by train to Munich.
   b) Peter most probably does not travel by train to Munich.
4. a) Tigers are most probably dangerous.
   b) 3 months old tigers are most probably not dangerous.
   c) This animal is most probably a 3 months old tiger.
5. a) If Arnold teases Joseph, then Joseph will most probably get irritated.
   b) If Joseph is teased by Arnold, then Joseph will most probably not get irritated.
   c) Arnold most probably teases Joseph.
6. a) Animals, which can fly, are most probably egg-laying.

b) Mammals are most probably not egg-laying.

c) This animal is most probably a mammal, which can fly.

7. a) If Suzy goes shopping, then Suzy is most probably happy.

   b) If Lena goes shopping, then Lena is most probably not happy.

8. a) This animal is most probably egg-laying, given it is animal, which can fly.

   b) This animal is most probably not egg-laying, given it is a mammal.

   c) This animal is most probably a mammal, which can fly.

References

Adams, E. W. (1975). *The Logic of Conditionals*. Dordrecht, Netherlands: D. Reidel Publishing Company.

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Bacchus, F. (1990). *Representing and Reasoning with Probabilistic Knowledge*. Cambridge, Mass.: The MIT Press.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford: Oxford University Press.

Bovens, L., & Hartmann, S. (2003). *Bayesian Epistemology*. Oxford: Oxford University Press.

Brewka, G. (1991). *Nonmonotonic Reasoning. Logical Foundations of Commonsense*. Cambridge: Cambridge University Press.

Carnap, R. (1962). *Logical Foundations of Probability* (2nd ed.). Chicago: University of Chicago Press.

Carnap, R. (1971). Inductive Logic and Rational Decisions. In R. Carnap & R. C. Jeffrey (Eds.), *Studies in Inductive Logic and Probability* (Vol. I, pp. 5–31). Berkeley: University of California Press.

Cevolani, G., Crupi, V., & Festa, R. (2011). Verisimilitude and Belief Change for Conjunctive Theories. *Erkenntnis*, *75*, 183–202.

Chater, N., & Manning, C. D. (2006). Probabilistic Models of Language Processing and Acquisition. *Trends in Cognitive Sciences*, *10*, 335–344.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006a). Probabilistic Models of Cognition. Conceptual Foundations. *Trends in Cognitive Sciences*, *10*, 287–291.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006b). Probabilistic Models of Cognition. Where Next? *Trends in Cognitive Sciences*, *10*, 292–293.

Crupi, V., Fitelson, B., & Tentori, K. (2008). Probability, Confirmation, and the Conjunction Fallacy. *Thinking & Reasoning*, *14*, 182–199.

de Finetti, B. (1973). Foresight. Its Logical Laws, Its Subjective Sources. Reprinted in H. E. Kyburg and H. E. Smokler (Eds.), *Studies in Subjective Probability*, Wiley, New York 1964, 93-158.

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, Mass.: MIT Press.

Evans, J. S. B. T. (1982). *The Psychology of Deductive Reasoning*. London: Routledge & Kegan Paul.

Evans, J. S. B. T., Handley, S. J., & Over, D. E. (2003). Conditionals and Conditional Probability. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *29*, 321–335.

Evans, J. S. B. T., & Over, D. E. (2004). *If*. Oxford: Oxford University Press.

Evans, J. S. B. T., & Twyman-Musgrove, J. (1998). Conditional Reasoning with Inducements and Advice. *Cognition*, B11-B16.

Geiger, S. M., & Oberauer, K. (2007). Reasoning with Conditionals. Does Every Counterexample Count? It's Frequency that Counts. *Memory & Cognition*, *35*, 2060-2074.

Gigerenzer, G., & Hoffrage, U. (1995). How to Improve Bayesian Reasoning without Instructions. Frequency Formats. *Psychological Review*, *102*, 684–704.

Gillies, D. (2000). *Philosophical Theories of Conditionals*. London: Routledge.

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic Models of Cognition. Exploring Representations and Inductive Biases. *Trends in Cognitive Sciences*, *14*, 357–364.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in Semantic Representation. *Psychological Review*, *114*, 211–244.

Hájek, A. (2008). Arguments for – or against – Probabilism? *British Journal for the Philosophy of Science*, *59*, 793–819.

Hájek, A. (2012). Interpretations of Probability. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy.* Retrieved from http://plato.stanford.edu/archives/sum2012/entries/probability-interpret/.

Hawthorne, J. (2005). Degree-of-Belief and Degree-of-Support. Why Bayesians Need Both Notions. *Mind*, *114*, 277–320.

Hertwig, R., & Gigerenzer, G. (1999). The 'Conjunction Fallacy' Revisited. How Intelligent Inferences Look like Reasoning Errors. *Journal of Behavioral Decision Making*, *12*, 275–305.

Howson, C., & Urbach, P. (2006). *Scientific Reasoning. The Bayesian Approach* (3rd ed.). Chicago: Open Court.

Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, *65*, 575–603.

Kemp, C., & Tenenbaum, J. B. (2009). Structured Statistical Models of Inductive Reasoning. *Psychological Review*, *116*, 20–58.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian Decision Theory in Sensorimotor Control. *Trends in Cognitive Sciences*, *10*, 319–326.

Leslie, S.-J. (2007). Generics and the Structure of the Mind. *Philosophical Perspectives*, *21*, 375–403.

Leslie, S.-J. (2008). Generics. Cognition and Acquisition. *Philosophical Review*, *117*, 1–47.

Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities. *The Philosophical Review*, *85*, 297–315.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability* (Vol. 2, pp. 263–293). Berkeley: University of California Press.

Markovits, H., Forgues, H. L., & Brunet, M.-L. (2010). Conditional Reasoning, Frequency of Counterexamples, and the Effect of Response Modality. *Memory & Cognition*, *38*, 485–492.

Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman.

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., et al. (2010). Letting Structure Emerge. Connectionist and Dynamical Systems Approaches to Cognition. *Trends in Cognitive Sciences*, *14*, 348–356.

McGee, V. (1989). Conditional Probabilities and Compounds of Conditionals. *The Philosophical Review*, *98*, 485–541.

Nickel, B. (2009). Generics and the Ways of Normality. *Linguistics and Philosophy*, *31*, 629–648.

Oaksford, M., & Chater, N. (2007). *Bayesian Rationality. The Probabilistic Approach to Human Reasoning*. Oxford: Oxford University Press.

Oaksford, M., & Chater, N. (2008). Probability Logic and the Modus Ponens-Modus Tollens Asymmetry in Conditional Inference. In N. Chater & M. Oaksford (Eds.), *The Probabilistic Mind. Prospects for Bayesian Cognitive Science* (pp. 97–120). Oxford: Oxford University Press.

Oaksford, M., & Chater, N. (2009). Précis of Bayesian Rationality. The Probabilistic Approach to Human Reasoning. *Behavioral and Brain Sciences*, *32*, 69–120.

Oaksford, M., Chater, N., & Larkin, J. (2000). Probabilities and Polarity Biases in Conditional Inference. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *26*, 883–899.

Oberauer, K., & Wilhelm, O. (2003). The Meaning(s) of Conditionals. Conditional Probabilities, Mental Models, and Personal Utilities. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *29*, 680-693.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems. Networks of Plausible Inference*. San Francisco: Morgan Kaufmann Publishers.

(Revised Second Printing)

Reichenbach, H. (1949). *The Theory of Probability*. Berkeley: University of California Press.

Reiter, R. (1980). A Logic for Default Reasoning. *Artificial Intelligence*, *13*, 81–132.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition. A Parallel Distributed Processing Approach*. Cambridge, Mass.: MIT Press.

Schroyens, W., & Schaeken, W. (2003). A Critique of Oaksford, Chater, and Larkin's (2000) Conditional Probability Model of Conditional Reasoning. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *29*, 140–149.

Schurz, G. (2001). What is 'Normal'? An Evolution-Theoretic Foundation for Normic Laws and Their Relation to Statistical Normality. *Philosophy of Science*, *68*, 476–497.

Schurz, G. (2005). Non-Monotonic Reasoning from an Evolution-Theoretic Perspective. Ontic, Logical and Cognitive Foundations. *Synthese*, *146*, 37–51.

Schurz, G. (2011). Tweety, or Why Probabilism and even Bayesianism Need Objective and Evidential Probabilities. In D. Dieks, W. J. Gonzales, S. Hartmann, M. Stöltzner, & M. Weber (Eds.), *Probabilities, Laws and Structures* (pp. 57–74). New York: Springer.

Schurz, G. (2012). Prototypes and their Composition from an Evolutionary Point of View. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford Handbook of Compositionality.* Oxford: Oxford University Press.

Steyvers, M., Griffiths, T. L., & Dennis, S. (2006). Probabilistic Inference in Human Semantic Memory. *Trends in Cognitive Sciences*, *10*, 327–334.

Strevens, M. (2004). Bayesian Confirmation Theory. Inductive Logic, or Mere Inductive Framework. *Synthese*, *141*, 365-379.

Talbott, W. (2008). Bayesian Epistemology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy.* Retrieved from http://plato.stanford.edu/entries/epistemology-bayesian/.

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-Based Bayesian Models of Inductive Learning and Reasoning. *Trends in Cognitive Sciences*, *10*, 309–318.

Tversky, A., & Kahneman, D. (1983). Extensional versus Intuitive Reasoning. The Conjunction Fallacy in Probability Judgment. *Psychological Review*, *90*, 293–315.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.

Xu, F., & Tenenbaum, J. B. (2007). Word Learning as Bayesian Inference. *Psychological Review*, *114*, 245–272.

Yuille, A., & Kersten, D. (2006). Vision as Bayesian Inference. Analysis by Synthesis? *Trends in Cognitive Sciences*, *10*, 301–308.

Acknowledgements

We would like to thank Stephan Hartmann, Jon Williamson, James Hawthorne, Niki Pfeifer, Gernot Kleiter, Paul Thorn, Ludwig Fahrbach, Vincenco Crupi, Edouard Machery and four anonymous referees for valuable comments on an earlier version of the paper and for valuable discussions at the Tilburg workshop "Formal Epistemology Meets Experimental Philosophy". We acknowledge support by DFG-Grant SCHU1566/5-1 as part of the LogiCCC project "The Logic of Causal and Probabilistic Reasoning" (LcpR), and by DFG-Grant SCHU1566/3-2 as part of the research group "Functional concepts and frames" (FOR 600).