

CHANCE, CREDENCE AND CIRCLES

FABRIZIO CARIANI

f-cariani@northwestern.edu

ABSTRACT

I target Pettigrew's application of the accuracy framework to derive chance-credence principles. My principal contention is that Pettigrew's preferred version of the argument might in one sense be circular and, moreover, that Pettigrew's premises have content that goes beyond that of standard chance-credence principles.

Accuracy and the Laws of Credence (henceforth *ALC*) is a special achievement. It is, at once, a manifesto for the epistemic utility research program, a comprehensive introduction to the subject, and a trail-blazing investigation of its frontier. In this note, I investigate one segment of that frontier: Pettigrew's attempt to justify a family of principles related to Lewis's Principal Principle (Lewis 1980).¹ These principles constrain the relationship between an agent's credences about chancy events and her credences in propositions about objective chances. Following custom, and despite some worries about the accuracy of the label, I refer to the family as *chance-credence principles*.

The book's central goal is to identify and evaluate accuracy arguments in support of several epistemic norms on rational credences. Part I offers Pettigrew's rendition of the accuracy argument for probabilism (originally from Joyce 1998, 2009). Then, part II presents two broad strategies for extending this argument to chance-credence principles. These strategies involve different interventions on the original.

To appreciate the difference between the strategies, it helps to recall the main ingredients of that argument. Start with the idea that, to assess the rationality of credal states, we must deploy decision theoretic techniques. What puts the 'epistemic' in 'epistemic rationality' is not a special notion of rationality, but a distinctive kind of utility – epistemic utility – to be plugged into general-purpose decision-theoretic constraints. The epistemic utility of having credence c in world w is the inverse of the inaccuracy of c in w . The inaccuracy of c in w is the sum of the 'local' inaccuracies of $c(p)$ in w for each proposition p that c is defined on; in turn, these are measured by the squared Euclidean distance of $c(p)$ from the truth-value of p in w . That this is the correct notion of epistemic utility stems from a commitment to *veritism*, the thesis that the sole source of epistemic value is accuracy (see Goldman 2002, and *ALC*, pp. 6–9).

Arguments for probabilism typically invoke the rationality constraint that dominated credences are irrational. Pettigrew's version appeals to an even weaker constraint: credence c is irrational if it is strongly dominated by a credence that is not even weakly dominated and that expects itself, and only itself, to be maximally accurate (see *ALC*, ch. 2 for discussion).² The last component of the argument is a theorem showing that for every

1 This thread of *ALC* extends Pettigrew's previous work in (2012, 2013).

2 It is very easy to get confused when using the language of 'weakness' and 'strength' in relation to norms, especially ones that have a conditional form. So let's be totally explicit: say that rationality constraint N_i

non-probabilistic credence c there is a probabilistic credence c^* that dominates c in the above sense. Putting these together, we get that non-probabilistic credences are irrational.

The first strategy for extending this argument to chance-credence principles, inspired by Hájek (ms), renounces the letter of veritism in favor of the thesis that credences are vindicated by chances. Accordingly, we ought to measure the accuracy of a credence held at t in w not against the distribution of truth-values (at t in w), but against the objective chances (as they are at t in w). This strategy does not require any amendments to the decision rule, as it relies on standard dominance constraints. The second strategy holds on to veritism but appeals to a stronger constraint than dominance (see below). Pettigrew ultimately endorses this second strategy as his official justification of chance-credence principles, because of the first strategy's deviation from veritism.

After presenting some background (§1) and providing details on the second strategy (§2), I will challenge two aspects of this strategy. §§3–4 question whether Pettigrew has successfully dealt with concerns about the circularity of the argument (these concerns arise because of the strength of the decision-theoretic rule). In §5, I argue that the second strategy proves more than we should expect to fall out of chance-credence principles. In particular, while chance-credence principles are most naturally interpreted as coherence constraints (regulating the coherence between beliefs of one sort and beliefs of another sort), the new decision-theoretic principle entails narrow-scope verdicts that go well beyond requirements of coherence.

One last note before diving in: in part IV of the book (§14.2), Pettigrew also runs an accuracy argument for the reflection principle (van Fraassen 1984). This argument closely mirrors the structure of the arguments for chance-credence principles. Many of the comments I will make about the justification of chance-credence principles have close correlates within that particular argument for reflection. Since these correlates can nonetheless be evaluated differently, I will occasionally remind the reader of this parallel (though I will also lack the space to provide much additional detail).

I. THE VARIETY OF CHANCE-CREDENCE PRINCIPLES

Stating chance-credence principles requires modeling languages with unusual expressive capacities. It is not enough for the language L to talk about chancy events. That is, it is not enough if L features sentences like “the coin will land heads”, or “most of the polonium on the victim's body will decay within a year”. In addition, L must express what I will call *chance hypotheses*: these are propositions that characterize the state of objective chances in a world (or perhaps, in a world at a time). Chance hypotheses come in two varieties:

entails constraint N_2 just in case every act (or state) that is classified as irrational by N_2 is also classified as irrational by N_1 . Furthermore, say that N_1 is stronger than N_2 iff N_1 asymmetrically entails N_2 (in which case N_2 is weaker). This means that the principle that says that being strongly dominated is irrational is weaker than the principle that says that being weakly dominated is irrational. This might be confusing, but it is as it is supposed to be. So to say that Pettigrew bases his argument on a weakening of dominance is to say that he uses a principle that classifies strictly fewer acts as irrational.

Ur-chance hypotheses are propositions, here denoted by ‘ C_{cb} ’, that are true in world w just in case w ’s *ur*-chance function is cb (the *ur*-chance of w is the chance function at the beginning of w ’s history, assuming that w ’s history has a beginning).

Temporal chance hypotheses are propositions, here denoted by ‘ T_{cb} ’, that are true in w at t just in case the chances in w at t are provided by cb .

The key difference is that the former are time-invariant, while the latter may have truth-values that vary across times (Caie 2015; ALC, 9.2).

Among the many side contributions of ALC is a useful taxonomy of chance-credence principles. These are classified along three main dimensions. Given principle P and agent α , consider these questions:

- Q1 What kind of chance hypothesis does P appeal to? (*ur*-chance hypotheses vs. temporal chance hypotheses).
 Q2 Is P formulated as a constraint on α ’s initial credence function or as a constraint on α ’s current credence function?
 Q3 Does P include a restriction to admissible evidence?

Lewis’s Principal Principle constrains initial credence functions conditional on *ur*-chance hypotheses.³ More recently, however, Caie (2015) has defended chance-credence principles that constrain current credences conditional on temporal chance hypotheses. Pettigrew agrees that these are the crucial principles to justify, and thus focuses on:⁴

Evidential Temporal Principle (ETP) If an agent has credence function c and total evidence E , then rationality requires that

$$c(X|T_{cb}) = cb(X|E)$$

for all propositions X in \mathcal{F} and all possible chance functions cb such that T_{cb} is in \mathcal{F} and $c(T_{cb}) > 0$.⁵

The two main accuracy arguments in part II of ALC seek to establish ETP.⁶

2. TWO ARGUMENTS FOR CHANCE-CREDENCE PRINCIPLES

I sketched two routes to the conclusion that an agent who violates etp must be irrational. An example will help flesh them out:

3 Notoriously, Lewis gave two formulations of the Principal Principle, only one of which made reference to admissible information (Meacham 2010). The principles I discuss here do not appeal to admissible evidence.

4 If one prefers something closer to the original version of the Principal Principle, Pettigrew (2013) offers a very similar accuracy argument, as does the introduction to chapter 10 of ALC. While in footnote-land, I might add that Pettigrew’s reasons for focusing on ETP are different from Caie’s.

5 ‘ \mathcal{F} ’ is Pettigrew’s label for the algebra over which the agent’s credences are defined.

6 I will operate under the simplifying assumption that there are no self-undermining chances. That is, there is no chance cb such that $cb(T_{cb}) < 1$. Most of ALC (with the exception of chapter 11, the last of part II) operates under this assumption as well.

Suppose that Diana is about to take a three point shot, and assume for the sake of argument that three point shots are chancy events. Steph has credences over the outcome of Diana’s shot and over the relevant temporal chance hypotheses. Steph’s total evidence is E . Steph is certain that a particular temporal chance hypothesis T_{cb} is true (i.e. $c(T_{cb}) = 1$). That is, he is certain that the current chance function is given by cb . According to cb , the chance of Diana’s making the shot (conditional on E) is .4. Steph, however, is very confident that Diana will make the shot: he assigns credence .9 to that proposition.

Plausibly, something is defective in Steph’s credal state, and ETP tells us what it is: Steph is much more confident in X than is warranted by a chance hypothesis he is certain of. If ETP admits of a deeper justification, we should be able to characterize the defectiveness of Steph’s credal state directly in terms of more fundamental principles.

According to the first strategy, Steph is irrational because his credence is dominated. But, in the context of this strategy, the appropriate epistemic utility function at t in w tracks the squared Euclidean distance from the chance function at t in w – not the squared Euclidean distance from the truths at t in w . Equipped with this alternate notion of epistemic utility, Pettigrew proves that there is a credence c^* that strongly dominates Steph’s c , and that c^* is itself not weakly dominated. Steph’s credence is irrational after all.

Pettigrew finds this strategy objectionable because he thinks that the idea that chances vindicate credences is not truly veritistic. This is because he maintains that, on virtually every theory of objective chance, the chances in a world are “information-losing summaries of the truths” at that world. From a veritistic standpoint, the claim that our credences ought to match those information-losing summaries seems unmotivated (ALC, § 9.4).

The second strategy – ALC’s official strategy – reverts back to the idea of alethic vindication: local inaccuracy is to be measured as distance from truth-value. Because vindication is purely alethic, Steph’s probabilistically coherent credence is not dominated in the standard sense. Hence, any argument for ETP must go by way of some stronger decision-theoretic constraint.

Pettigrew proposes a different rationality constraint. Momentarily ignoring some nuances, this constraint states that option o is irrational if there is an option o^* such that all the possible current chance functions expect o to be less accurate than o^* . As far as I can tell, the notion of ‘possible current chance function’ C is not explicitly defined in ALC. Pettigrew (personal communication) suggests interpreting this as the set of functions cb such that T_{cb} is consistent with the agent’s total evidence. For later discussion, I want to offer another, more subjective interpretation: C consists of the set of functions cb such that the agent assigns non-zero credence to T_{cb} .

Working towards a more precise statement of the rationality constraint, let ‘ $Exp_u(o|pr)$ ’ denote the expected utility of o calculated relative to utility function u and probability function pr ; let ‘ cb_E ’ denote the function that inputs a proposition p and outputs $cb(p|E)$. Next, define some auxiliary concepts:

o^* **bests** o relative to cb and E iff

$$Exp_u(o|cb_E) < Exp_u(o^*|cb_E)$$

(when the sign is ‘=’ we say that o^* *equals* o)

Letting C be the set of possible current chance functions, define:

- (i) o^* **strongly chance dominates** o relative to C and E iff for all $ch \in C$, o^* bests o relative to ch and E .
- (ii) o^* **weakly chance dominates** o relative to C and E iff
 - (a) for all $ch \in C$, o^* bests or equals o relative to ch and E and
 - (b) for some $ch \in C$, o^* bests o relative to ch and E .

In terms of these concepts we can state Pettigrew's principle, a new sufficient condition for irrationality:

Current Chance Dominance (CCD) Credence c is irrational for an agent with total evidence E if

- (a) c is strongly chance dominated by a probabilistic c^* conditional on E
- (b) c^* is not weakly chance dominated conditional on E and
- (c) c^* is not extremely modest.

Pettigrew proves that an agent who violates probabilism or ETP must meet that sufficient condition, and hence is irrational. As anticipated in the introduction, a parallel principle of 'future credence dominance' is involved in the justification of the reflection principle (ALC, p. 194).

Before assessing this second version of the argument, pause to note that there is a level of ambition that is only available to proponents of the first strategy. A proponent of the first strategy might hope to obtain a simultaneous justification of probabilism and ETP. The second strategy cannot share this ambition. This is because it is plausible that a justification of probabilism should not presuppose overtly probabilistic constraints (Easwaran 2014). Arguably, CCD's use of expectations is one such problematic presupposition.⁷

This does not damage ALC's overall argument. Instead of seeking to justify both requirements by a single argument, we might think of the justification as coming in two separate steps (appropriately corresponding to the first two parts of the book). Step one is the dominance-based justification of probabilism. Step two is the CCD-based justification of chance-credence principles. As long as the rationality constraints do not conflict and the utility function remains veritistic, the epistemic value monism at the center of Pettigrew's program remains unassailed.

3. THE TWO CIRCULARITY OBJECTIONS

However, there are other important questions concerning the viability of the second strategy. Pettigrew identifies and discusses the worry that CCD might pack too much to serve as a justification of ETP.

Pettigrew calls this the *Circularity Objection*. I think, however, that it is useful to distinguish at the outset between two different kinds of circularity worries—corresponding to two ways in which arguments might be circular (Sinnott-Armstrong 1999; Rips 2002).

⁷ Leitgeb and Pettigrew (2010) might have taken exception to this claim, but it is notable that ALC relies on weakenings of dominance for all its main arguments.

Very roughly, we can say that arguments are *structurally* circular when their conclusion is one of the premises or must figure in any reasonable justification of the premises. By contrast, they are *dialectically* circular when they rely on premises that would not (or should not) be accepted by the opponent, because they are too close to the conclusions that they are to support.⁸ An important difference is that dialectical circularity does not require that the conclusion itself appear in the justification chain. For an extreme example of dialectical circularity, imagine trying to justify probabilism by assuming, among other things, that credences are additive. This attempt would fail not because any justification of the additivity constraint must appeal to probabilism, but because additivity packs too much of what needs to be justified.

Dialectical circularity shows up in ALC as the worry that CCD might be “no more plausible than ETP itself” (p. 129). Structural circularity as the worry that ETP might itself be necessary to justify CCD. Though both versions of the objection are voiced, Pettigrew’s discussion is heavily focused on the structural variant of the objection. After sketching an argument purporting to justify CCD on the basis of a strengthening of ETP, he gives two reasons to reject it: first, the premises in this argument seem to be less general in their application than CCD; second, they seem more substantive in content. This suggests that the argument is not structurally circular, after all.

Even granting all that, this line of reasoning only deals with structural circularity. This means that, as far as the discussion of ALC goes, the threat of dialectical circularity is still looming, since dialectical circularity is perfectly compatible with the claim that CCD is more general and less substantive in content than ETP.

To investigate the dialectical circularity objection, we need to get traction on the vague-sounding question whether CCD is too close to ETP to justify it. The first thought might be to ask: what do we need, in addition to CCD, to obtain the basic verdicts that the chance-credence principles are meant to systematize? If we need to add relatively little, that will count as defeasible pressure to accept the claim that CCD is ‘too close’ to ETP to justify it. If we need extremely substantive assumptions, then the charge of dialectical circularity seems misplaced. The problem is that this approach does not seem very conclusive in this case: although CCD does do much of the work in deriving ETP, it does not do all of it.⁹ And it is hard to have an independent assessment of whether the required auxiliary assumptions are substantive enough.

4. CHOOSING DECISION RULES

There is another, and I think more productive, way of advancing the worry that CCD is dialectically circular. This focuses not on what we must add to it to derive etp, but on what we must sacrifice to take it on board.

Start by asking: why should the rationality of an agent’s preferences be constrained by CCD? Suppose that, as most Bayesians do, we accept EU. Rationality requires that an agent with credence c , total evidence E and utility u prefer o to o^* iff $Exp_u(o^*|c_E) < Exp_u(o|c_E)$.

8 Sinnott-Armstrong (1999) reserves the word ‘circular’ for the first type of argument. Instead, he labels the second type of argument as ‘question-begging’.

9 In fact, a first step in this direction is already in Pettigrew (2013), who proves that the general theorem goes through for any strictly proper scoring rule.

If anything deserves to be called a standard constraint on preference, it is EU.¹⁰ Now, we might ask a broad methodological question: when are we justified in deviating from EU in setting up an accuracy argument for some rational constraint?

One reason for deviating is concern that the full strength of EU would make our decision-theoretic justifications circular. In such cases, we might need a weaker rule that is nonetheless compatible with EU. This is why, as noted, standard accuracy arguments for probabilism do not deploy expectation-based rules, and instead use dominance requirements.

The other possible reason not to use EU is if it is not applicable in the circumstances our agents find themselves in. Pettigrew resorts to this kind of consideration in part III of *ALC*. Part III lays out accuracy arguments for the principle of indifference. These arguments appeal to generalizations and strengthenings of MAXIMIN – the rule according to which it is rational to prefer o^* to o just in case the minimum outcome guaranteed by o^* exceeds the minimum outcome guaranteed by o . But Pettigrew is careful to formulate MAXIMIN and its variants as principles that apply only when the agent has no evidence whatsoever. It is plausible to rule such circumstances out of the domain of applicability of EU. If we do that, there is no tension between maximin-like principles and EU. However, neither of these reasons for deviating from EU applies to CCD. The only requirement for applying CCD is that the agent assign credences to chance hypotheses.

I think, then, that a more convincing way of advancing the dialectical circularity worry is based on this line of thought.

Once we are done justifying probabilism, we ought to default back to EU, unless we have a reason to think that EU does not apply. But in the vast majority of cases, we have no such reason. So, in the vast majority of cases, we ought to default to EU. Now, in some of these cases EU and CCD make incompatible requirements. This can happen because EU constrains rational preferences on the basis of the agent's credences while CCD constrains them on the basis of the unanimous agreement of the possible current chance functions.

The threat of dialectical circularity arises here because, to uphold CCD, we need a reason to think that it ought to take priority over EU when they conflict. It is hard to see what would push us to sacrifice such a central part of the Bayesian picture, other than a desire to vindicate chance-credence principles.

5. NARROW SCOPE ENTAILMENTS OF CCD

There is a passage in *ALC* that speaks to a worry in this general ballpark. Pettigrew notes that the second strategy “will not satisfy someone who is not already convinced that we should defer to chances *in some way*. Doing more is beyond the scope of this project” (p. 131). In other words: if we start with some mild attitude of deference to chance, accuracy arguments will help us squeeze out more robust chance-credence principles. Perhaps (here comes an additional step that is not explicitly taken by Pettigrew), that mild attitude involves prioritizing CCD over EU when they conflict.

¹⁰ I recognize of course that someone like Buchak (2014) would resist EU. But the points made in this section survive even if we prefer a different account of rational preference, such as Buchak's.

I think we should resist this line of thought. To motivate my resistance, I want to offer a more direct criticism of CCD and of the second strategy as a whole. The second strategy delivers verdicts that go well beyond the content of chance-credence principles – and, what is more, verdicts that do not strike me as mild ways of deferring to chance. Which verdicts these are depends on how we characterize the set of possible current chance functions. As I noted earlier, there are a couple of possible ways of characterizing this set. Consequently, there are a couple of ways of running this argument. I consider each in turn.

Start with the evidential construal suggested to me by Pettigrew (personal communication): the possible current chance functions are those ch such that T_{ch} is consistent with one's total evidence. Suppose that Margot's total evidence E is only compatible with chance function ch_1 (i.e., E entails T_{ch_1}). Suppose that, despite this, Margot is certain that the current chances are given by ch_2 (i.e., $c_{Margot}(T_{ch_2}) = 1$). Given that E entails T_{ch_1} , she is either not probabilistically coherent or not logically omniscient. Importantly, she might be non-omniscient without being incoherent, provided we treat logical omniscience with techniques such as those advocated in Garber's (1983).¹¹ Finally (and crucially), suppose that all of Margot's other credences harmonize with ch_2 , so when X is a chancy event, $c_{Margot}(X) = ch_2(X|E)$.

My judgment here is that, although Margot is in one sense irrational, she is not violating a chance-credence principle. Chance-credence principles regulate the coherence of Margot's credences about chancy events and her credences about chancy hypotheses. Margot's irrationality does not stem from this kind of conflict: her attitudes concerning these propositions are coherent even though she is failing to recognize what her evidence supports.

This is not the verdict we get if we follow the second strategy. The formal reason behind this is proven (you guessed it!) by Pettigrew, who notes that the second strategy actually entails a strengthening of ETP:

ETP* If an agent has credence c and total evidence E then rationality requires that c be in the closure of the convex hull of the set C (where C is the set of possible current chance functions).

If C is a singleton then the closure of the convex hull is that same singleton. When, as in Margot's case, $C = \{ch_1\}$, rationality requires Margot to prefer ch_1 to her own credence. My reason for concern, to repeat, is that this goes beyond the content of ETP. Because ETP is a coherence requirement, it is silent on how the agent ought to resolve the incoherence in her state. It is also silent on whether the agent should prefer a credence that alters her belief in chance hypotheses as opposed to a credence that alters her beliefs in propositions describing chancy events. If that is right, CCD would appear to embody more than a mild commitment to defer to chances.

The other construal involves a subjective interpretation of possible current chance functions as those functions ch such that $c(T_{ch}) > 0$. This avoids the problem of the previous construal: when $C = \{ch_2\}$, ETP* requires Margot to have ch_2 as her credence, which she does. If she is irrational, it is because of a principle other than ETP.

¹¹ Since the issue of logical omniscience does not arise at all in ALC, I assume that Pettigrew intends to stay consistent with the main Bayesian treatments of logical omniscience.

However, on this construal, the problematic narrow scope verdicts arise for agents who violate ETP. Consider an agent who, like Steph in my example from § 2, has $c(X) = .9$, $c(T_{cb}) = 1$, with $cb(X|E) = .4$. Once again we have that $C = \{cb\}$, so, according to ETP⁺, epistemic rationality requires Steph to prefer cb to his own credence. But that too is too strong: given that chance-credence principles are coherence constraints, they are silent on whether the defect in Steph's state lies with his credence in X or with his certainty in T_{cb} . He should not be mandated by epistemic rationality to prefer cb to his own credence.

Taking stock, no matter how we characterize the set of possible current chance functions, the second strategy yields narrow scope requirements of rationality that go beyond the content of chance-credence principles. Moreover, checking-in once again on the parallel argument for the reflection principle, we can make similar considerations. Pettigrew's 'future credence dominance' principle yields similar narrow scope verdicts: if one is certain that one will have a particular credence, then one is presently required to prefer that credence to one's current credence if they disagree. But all the standard formulations of reflection, including the one that Pettigrew uses, merely require that one's present credence and one's credences about one's future credences harmonize in a particular way.

Of course, my contention that chance-credence and reflection principles are just coherence constraints might be controversial. For instance, it does not sit well with the fact that these principles are often glossed as 'deference' principles, or described as having an evidential source. But nothing corresponds to these glosses in the formal content of the principles themselves. Perhaps, the glosses are best understood as concerning what happens when the principles are applied to model certain types of situations. If so, they would not be part of what needs to be vindicated by a justification of those principles.¹²

REFERENCES

- Buchak, L. 2014. *Risk and Rationality*. Oxford: Oxford University Press.
- Caie, M. 2015. 'Credence in the Image of Chance.' *Philosophy of Science*, 82: 626–48.
- Easwaran, K. 2014. 'Decision Theory without Representation Theorems.' *Philosophers Imprint*, 14: 1–30.
- Garber, D. 1983. 'Old Evidence and Logical Omniscience in Bayesian Confirmation Theory.' In J. Earman (ed.), *Testing Scientific Theories*, pp. 99–131. Minneapolis, MN: University of Minnesota Press.
- Goldman, A. I. 2002. *Pathways to Knowledge: Private and Public*. New York, NY: Oxford University Press.
- Hájek, A. ms. 'A Puzzle about Partial Belief.' Manuscript, Australian National University.
- Leitgeb, H. and Pettigrew, R. 2010. 'An Objective Justification of Bayesianism I: Measuring Inaccuracy.' *Philosophy of Science*, 77: 201–35.
- Joyce, J. 1998. 'A Nonpragmatic Vindication of Probabilism.' *Philosophy of Science*, 65: 575–603.
- . 2009. 'Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.' In F. Huber and C. Schmidt-Petri (eds), *Degrees of Belief*, pp. 263–97. Berlin: Springer.
- Lewis, D. 1980. 'A Subjectivist's Guide to Objective Chance.' In R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Volume II, pp. 83–112. Berkeley, CA: University of California Press.
- Meacham, C. 2010. 'Two Mistakes Regarding the Principal Principle.' *British Journal for the Philosophy of Science*, 61: 407–31.

¹² Thanks to Mike Caie, Kenny Easwaran, and Richard Pettigrew.

- Pettigrew, R. 2012. 'Accuracy, Chance, and the Principal Principle.' *Philosophical Review*, 121: 241–75.
- . 2013. 'A New Epistemic Utility Argument for the Principal Principle.' *Episteme*, 10: 19–35.
- Rips, L. 2002. 'Circular Reasoning.' *Cognitive Science*, 26: 767–95.
- Sinnott-Armstrong, W. 1999. 'Begging the Question.' *Australasian Journal of Philosophy*, 77: 174–91.
- van Fraassen, B. 1984. 'Belief and the Will.' *Journal of Philosophy*, 81: 235–56.

FABRIZIO CARIANI is an Associate Professor of Philosophy at Northwestern University. He has published papers in epistemology (especially concerning applications of formal methods in social epistemology) and philosophy of language, formal semantics and philosophical logic (modality, especially deontic modality, and conditionals).
